Pacific Graphics (2026)
Y. He, N. Thürey, and L. Liu (Editors)

Poster

A Spatio-Temporal Acceleration Framework for Order-Independent Transparency

Grigoris Tsopouridis [1], Christos Georgiou-Mousses[1], Aris Panagiotidis[1], Andreas A. Vasilakis[2], Tobias A. Franke[3], Aleksei Gorbonosov[3], Andrei Astapov[3] and Ioannis Fudos[1]

[1]Department of Computer Science and Engineering, University of Ioannina, Greece
[2]Department of Informatics, Athens University of Economics and Business, Greece
[3]Huawei Ireland Research Center, Ireland

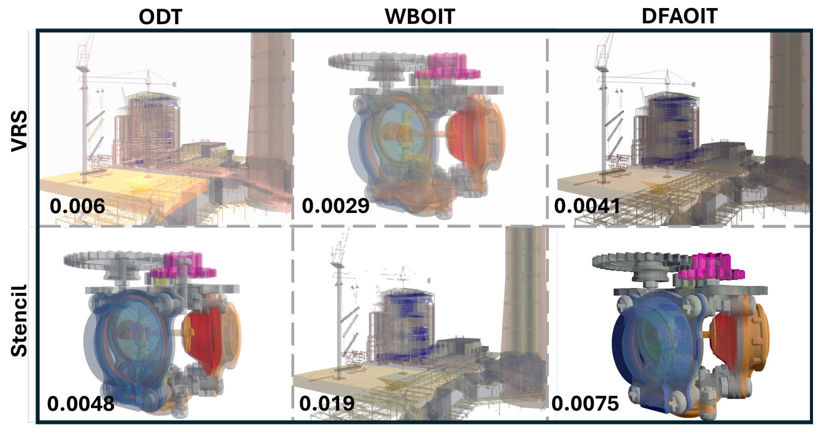

Figure 1: *ꟻLIP perceptual error comparison across ODT, WBOIT, and DFAOIT pipelines using VRS and Stencil acceleration, measured against non-accelerated references. All errors remain below 0.02, with VRS consistently achieving lower error. Stencil shows higher error due to its simple bi-linear interpolation.*

Abstract

Order-independent transparency (OIT) remains costly, particularly on mobile and legacy hardware. We present a spatio-temporal acceleration suite that exploits spatial and temporal coherence to reduce rendering cost while preserving visual quality. Our method combines adaptive tiled screen-space sampling or hardware-accelerated variable rate shading (VRS) where available, and modular temporal reprojection. Designed as a plug-in acceleration layer, it can be integrated with existing OIT techniques to significantly improve real-time performance with minimal integration overhead.

CCS Concepts

• ***Computing methodologies*** $\rightarrow$ ***Visibility;*** *Rasterization;*

1. Introduction

Rendering order-independent transparency (OIT) with high visual fidelity remains challenging for real-time applications, particularly on memory- and power-constrained devices. Exact methods such as the A-buffer [Car84] guarantee correct compositing but incur high bandwidth, memory, and execution costs at medium-to-large depth complexity. Approximate alternatives, from heuristic blending (WBOIT [MB13]) to neural approaches (DFAOIT [TVF24]), reduce cost but remain expensive at high quality, resolution, and hardware coverage.

We introduce a spatio-temporal tool suite that accelerates transparency by combining spatial adaptivity and temporal reuse, while staying agnostic to the underlying OIT solver. Our approach integrates with existing pipelines—from classical multifragment techniques to recent neural methods—and delivers consistent performance gains with minimal impact on visual quality.

Early OIT work emphasized exact multifragment schemes. The A-buffer [Car84] stores and sorts all per-pixel fragments before compositing, ensuring correctness but demanding unbounded memory and bandwidth. Depth peeling [BM08] iteratively extracts and blends depth layers to achieve exact ordering, with cost growing linearly in visible layers. To improve scalability, order-dependent approaches such as Object-Sorted Transparency (ODT) [TGMF*25] pre-sort objects per view, while heuristic blended methods like WBOIT [MB13] trade strict correctness for GPU-friendly accumulation. Learning-based OIT has recently emerged as a promising alternative. DFAOIT [TVF24] uses a compact network to predict final transparency from per-pixel fragment statistics, achieving bounded per-pixel cost suited to interactive use. While such neural methods can pair with A-buffer or depth peeling to improve quality, they still evaluate transparency at each pixel and frame. In contrast, our framework is orthogonal to the OIT formulation of exact, approximate, or neural approaches, and focuses on reducing spatial and temporal workload, providing complementary performance gains while preserving the underlying quality.

2. Spatiotemporal Tools for Accelerating OIT Rendering

Our framework accelerates OIT by exploiting spatial and temporal coherence. Adaptive tile subdivision reduces shading via stencil masking or variable rate shading (VRS) [YZK*19], and modular temporal reprojection reuses results across frames. Components integrate with current methods and work independently or combined.

Adaptive quadtree tile subdivision. We divide the screen into 16×16 pixel tiles, each encoded by a four-level quadtree bitmask ($16\times16 \rightarrow 2\times2$, 21 nodes). Each bit marks whether the corresponding node is split, so a tiles entire subdivision fits in a 21 bits (32KB at 1080p). Per frame, a compute pass estimates per-node color variance from prior transparency output using the four corner samples. Nodes split when variance exceeds level thresholds and merge when it drops below lower thresholds.

Tile-based sampling. The finalized quadtree writes a stencil mask before the transparent pass. A fullscreen triangle samples only leaf-node corners and marks them in stencil, yielding a sparse set of required pixels. Transparent geometry shades only at marked pixels; a compute pass then reconstructs the full-resolution layer via bilinear interpolation from leaf corners.

Tile-driven variable rate shading (VRS). Alternatively, the quadtree drives hardware VRS: coarse nodes shade at 2×2 or 4×4, fine nodes at 1×1. This yields smoother quality than stencil discard but requires VRS support and produces approximate results in coarse regions.

Temporal interpolation. We reuse prior-frame transparency via depth-based reprojection using depth and camera motion. This module is plug-in: any temporal reconstruction can replace depth-based warping without changing the spatial pipeline. As with any reprojection-based reuse, artifacts can appear when the real frame rate is too low or camera motion too fast, i.e. when inter-frame displacement exceeds what the warp can reliably recover.

Table 1: *Rendering cost (ms) with stencil-driven discard and hardware VRS, alone and combined with temporal re-projection. Baseline cost shown in parentheses next to each scene.*

ODT (ms)	Stencil	+Temporal	VRS	+Temporal
Powerplant (1.80)	1.57	1.18	1.58	1.21
Engine (1.10)	1.01	0.74	0.98	0.8
WBOIT (ms)	Stencil	+Temporal	VRS	+Temporal
Powerplant (1.96)	1.64	1.22	1.75	1.28
Engine (1.12)	1.04	0.78	1.03	0.8
DFAOIT (ms)	Stencil	+Temporal	VRS	+Temporal
Powerplant (2.50)	2.10	1.35	2.25	1.4
Engine (1.49)	1.36	1.05	1.39	1.1

3. Results

We evaluate our framework on two scenes of varying depth complexity (Powerplant and Engine) across three OIT pipelines: ODT, WBOIT, and DFAOIT. As shown in Table 1, both spatial accelerators reduce rendering cost relative to each pipeline's baseline: stencil-based resolution reduction by roughly 12% and VRS by roughly 10%. Coupling either with temporal reprojection nearly doubles the speedup, cutting cost by up to 46% for ODT, 38% for WBOIT, and 46% for DFAOIT on Powerplant, with similar or larger relative gains on the lighter Engine scene. These savings are consistent across all pipelines, confirming that our suite is agnostic to the underlying transparency method. Visual quality remains high: in Figure 1, FLIP error stays below 0.02 for all pipeline/acceleration combinations, with VRS consistently lower than stencil due to its smoother, shading-rate-based degradation, whereas discard-plus-bilinear reconstruction used in stencil buffer yields slightly higher error. Overall, the spatio-temporal suite delivers substantial, broadly applicable performance improvements with minimal perceptual quality loss.

References

[BM08] Bavoil L., Myers K.: Order independent transparency with dual depth peeling. In *NVIDIA OpenGL SDK.* (2008). 2

[Car84] Carpenter L.: The A-buffer, an antialiased hidden surface method. *SIGGRAPH Comput. Graph. 18*, 3 (jan 1984), 103–108. doi:10.1145/964965.808585. 1, 2

[MB13] McGuire M., Bavoil L.: Weighted blended order-independent transparency. *Journal of Computer Graphics Techniques (JCGT) 2*, 2 (December 2013), 122–141. 1, 2

[TGMF*25] Tsopouridis G., Georgiou-Mousses C., Fudos I., Corrigan D., Franke T. A.: Traditional and Neural Order-Independent Transparency. In *Eurographics 2025 - Tutorials* (2025), Mantiuk R., Hildebrandt K., (Eds.), The Eurographics Association. doi:10.2312/egt.20251001. 2

[TVF24] Tsopouridis G., Vasilakis A. A., Fudos I.: Deep and fast approximate order independent transparency. *Computer Graphics Forum 43*, 6 (2024), e15071. doi:10.1111/cgf.15071. 1, 2

[YZK*19] Yang L., Zhdan D., Kilgariff E., Lum E. B., Zhang Y., Johnson M., Rydgård H.: Visually lossless content and motion adaptive shading in games. *Proc. ACM Comput. Graph. Interact. Tech. 2*, 1 (June 2019). doi:10.1145/3320287. 2