Fast Fluid Up-Res

Xavier Martin Ramirez
Industrial Light & Magic
London, United Kingdom
xmartin@ilm.com

Richard Jones
Industrial Light & Magic
London, United Kingdom
richardj@ilm.com

Abstract

We present a production-proven method for up-resing volumetric fluid simulations, allowing creation of high-resolution results efficiently from fast-to-compute, low resolution sims. Existing approaches often rely on secondary simulation and supplementary turbulence to inject missing details at higher resolutions. At Industrial Light and Magic (ILM), we have devised a simpler approach, one that is frame independent and therefore parallelisable across frames. Our combination of gridless advection, anti-aliased sampling and camera-frustum aligned grids has been used with great success on many recent productions and is a mainstay of our fluids toolkit.

ACM Reference Format:
Xavier Martin Ramirez and Richard Jones. 2025. Fast Fluid Up-Res. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH Talks '25), August 10–14, 2025, Vancouver, BC, Canada.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3721239.3734123

1 Introduction

Simulating volumetric FX, e.g. smoke, fire and explosions at the required level of detail for realism in visual effects can be computationally prohibitive. Up-resing offers a solution, allowing artists to generate detailed visuals from less expensive lower-resolution simulations. This leads to shorter turnaround times and lower storage requirements.

At ILM, a significant proportion of our simulated volumes leverage up-resing techniques to meet production demands. Usually up-resing requires injection of smaller scale features using procedurally generated turbulences/noises such as Wavelet Turbulence [Kim et al. 2008], and/or guided re-simulation at a higher resolution. Our method here differs in that it simply tries to recover details lost to numerical diffusion using the initial simulation's velocity fields. Combining gridless advection with anti-aliased sampling on frustum-aligned grids, allows us to up-res quickly, easily, and in a render farm-friendly (frame independent) way.

2 Gridless Advection

To perform our up-res, we use N steps of gridless advection (or more accurately, SELMA) [Tessendorf and Kowalski 2016], to trace the position x_t of each voxel from a naive upsampling of the volume backwards in time through the velocity field by N frames. Then, by sampling the input volumes f at these positions and returning this value for their new position, we get the up-res'ed value at frame t as $F_t(x_t) = f_{t-N}(x_{t-N})$. This effectively recovers details lost to numerical diffusion over the last N frames of advection and smoothes out grid artifacts in our higher resolution grid. For this process, each new higher resolution frame only requires the N prior velocity frames to backward advect the sample position, and frame $t - N$ of the volume we wish to up-res to sample from. Thus, up-resing can be evaluated as soon as a frame is available from the input simulation, and/or in parallel across all frames as a post-process. Simple single step advection schemes are effective for this, though higher order methods can be more temporally stable at the cost of higher compute times.

Figure 1: The Balrog of Moria was created with heavy use of fluid up-resing in "The Lord of the Rings: Rings of Power". ©Amazon MGM Studios. All Rights Reserved.

Figure 2: Smoke simulation with 0, 2, 4, 8 and 16 iterations of up-res (left to right). More iterations lead to increased detail recovery. ©Industrial Light & Magic. All Rights Reserved.

3 Anti-Aliased Volume Sampling

Using the above method alone can introduce aliasing artifacts, noise and flickering when neighbouring output voxel samples span large distances in the input volume. This is usually in areas of high visual interest e.g. turbulence. To mitigate this, we use 3D mip-map anti-aliasing.

Mip-maps are a hierarchy of pre-calculated, decreasing resolution versions of the original volume used to prevent aliasing when sampling far apart locations. We calculate the mip level of the sample positions output from the gridless advection. This level indicates

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGGRAPH Talks '25, Vancouver, BC, Canada
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1541-9/25/08
https://doi.org/10.1145/3721239.3734123

which volume in our hierarchy represents a corresponding wide enough sampling width. As this can introduce blurring we often apply a sharpening filter to either input or output grid to restore lost high-frequency details.

4 Frustum Grids

We commonly use camera frustum-aligned VDBs, which offer significant advantages for rendering efficiency. Here, one voxel corresponds to one pixel in screen space. These grids offer multiple possible accelerations for path-tracing ray traversal processes. With frustum VDBs we can compute and store our up-res'ed volume before the rendering process, hugely reducing the computational cost of the rendering vs attempting the process during the render itself. Additionally, frustum grids have the added benefit of allocating more voxels based in screen-space rather than world space, so volume resolution is optimal for the rendering target. Frustum volumes can sometimes experience banding so we add a random offset to the z axis to favour noise over banding artifacts.

5 Implementation and Production Use

This technique is implemented in Houdini with OpenVDB using a combination of VEX and custom SOP operators. A high-level tool is exposed to artists as an HDA with a minimal interface. ILM has successfully deployed this technique in numerous productions including the following. **"Avatar: The Way of Water"**: Up-resing was crucial for the 4K, high-frame-rate (HFR) requirements. This technique benefited fire, pyroclastic, and smoke simulations, allowing for efficient iterations and detailed final renders. **"The Lord of the Rings: Rings of Power"**: Season 2 gave us a closer look at the Balrog of Moria. All of the creature's volumetrics and flames were enhanced with up-resing. **Disney's Castle Logo** and 100th-anniversary production Logo appearing at the beginning of films released by Disney. The 6K, stereo, 120FPS delivery required up-resing all volumetrics to ensure a quality of detail that would future-proof the visuals for years to come.

6 Discussion

This technique consistently delivers high-quality results while significantly reducing both simulation and render times. It roughly costs as much as N advection steps at higher resolution, but can also be parallelised across frames. It is robust to lossy compression on the input simulation fields and exhibits temporal stability on both fast and slow-moving fluids. However, there are some limitations, requiring additional post-processing e.g. sampling back in time makes new fluid sources disappear. Similarly, it is unaware of effects such as density dissipation and temperature cooling. Frustum grids, whilst greatly beneficial are not strictly necessary for our technique and in some cases cartesian grids are preferred for their flexibility.

This technique has become an integral part of ILM's FX pipeline due to its simplicity and ease-of-use. It has enabled the creation of high-resolution fluid FX while maintaining efficient workflows and manageable data sizes. The novel combination of gridless advection, anti-aliased sampling and frustum grids allows for a successful balance between visual fidelity and computational efficiency, keeping creative control in the hands of the artist whilst maintaining the high quality of work that ILM is known for.

References

Theodore Kim, Nils Thürey, Doug James, and Markus Gross. 2008. Wavelet Turbulence for Fluid Simulation. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 1–6. doi:10.1145/1360612.1360649

Jerry Tessendorf and Michael Kowalski. 2016. *Resolution Independent Volumes.* Technical Report. Clemson University.