Auto-Adaptivity: An Optimization-Based Approach to Spatial Adaptivity for Smoke Simulations

Michael B. Nielsen
michael.nielsen@autodesk.com
Autodesk

Konstantinos Stamatelos
konstantinos.stamatelos@autodesk.com
Autodesk

Morten Bojsen-Hansen
morten.bojsen-hansen@autodesk.com
Autodesk

Robert Bridson
robert.bridson@autodesk.com
Autodesk

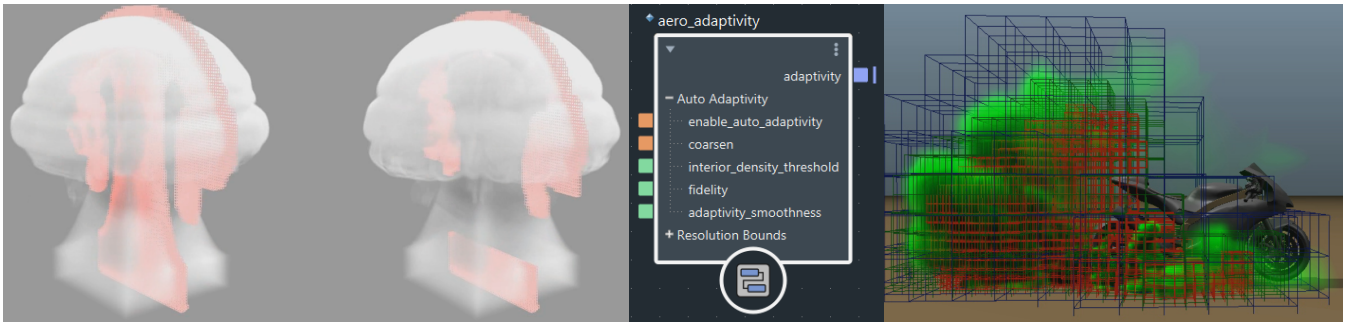

(a) Non-adaptive (left) and adaptive (right) with slice of finest-level voxels. (b) Auto-adaptivity inputs (c) Adaptive dirt bike simulation

Figure 1: Our new approach to spatial adaptivity enables the user to run adaptive simulations of smoke that are visually close to identical to their sparse non-adaptive counterpart (a) with the benefit of a reduction in computation-time and memory. A few input parameters (b) are fed into our new auto-adaptivity algorithm that retains the voxels which — subject to the constraint of a user-specified computation-budget (fidelity) — globally maximizes the quality of the simulation according to criteria such as distortion-rate, detail and interpolation error. The auto-adaptivity algorithm frees the user from explicitly managing and combining adaptivity controls by automatically determining which voxels to coarsen and refine as shown in the dirt bike simulation which represents the dust at four different levels of resolution (c). Dirt bike courtesy of Kerosene VFX.

CCS CONCEPTS

• **Computing methodologies** → **Physical simulation**.

KEYWORDS

fluid simulation, smoke simulation, spatial adaptivity

ACM Reference Format:
Michael B. Nielsen, Konstantinos Stamatelos, Morten Bojsen-Hansen, and Robert Bridson. 2020. Auto-Adaptivity: An Optimization-Based Approach to Spatial Adaptivity for Smoke Simulations. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH '20 Talks), August 17, 2020*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3388767.3407320

SIGGRAPH '20 Talks, August 17, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7971-7/20/08.
https://doi.org/10.1145/3388767.3407320

1 INTRODUCTION

Grid-based simulations of gaseous phenomena are immensely popular in visual effects for commercials and movies as a means to faithfully capture effects such as smoke, dirt, spray and fire. However, such simulations quickly become compute-intensive as the voxel-resolution is increased to match the demands and quality required for high-end VFX. One approach to reducing compute-time and memory that has gained increasing popularity as simulation packages have matured is spatial adaptivity where the voxel-resolution varies throughout the computational domain. Spatial adaptivity for gaseous simulations was first introduced to graphics by Losasso et al. [2004] and follow-up work has contributed with parallel and scalable algorithms designed for modern CPU-architectures [Nielsen et al. 2018; Setaluri et al. 2014]. Several criteria for dynamically coarsening and refining the computational grid over time have been proposed such as truncation error, optical depth, vorticity, proximity to smoke and geometric primitives. However, directly exposing many adaptivity controls leads to a wealth of input parameters that are difficult for the user to grasp, control and combine meaningfully (see video). As an example of this, Ando et al. [2013] recently proposed a local mapping from adaptivity criteria to desired spatial resolution called *sizing functions*. However, in our experiments we found that the scale-factors involved need to be

exposed and hand-tuned on a per-simulation basis, thereby making it difficult to use an approach like this in production. In this talk we describe and discuss a new workflow and algorithm for spatially adaptive gaseous simulations which, based on a few input parameters, automatically combines any number of adaptivity controls while delivering and optimizing for the simulation quality specified by the user. Our method has been implemented in Autodesk Bifrost and is already in the hands of users.

2 WORKFLOW DESIGN CHALLENGES

Volumetric simulations already suffer from workflow complexity, and burdening the artist further for performance gains should be minimized — at least, that is our assumed mindset when considering how to present spatial adaptivity capabilities to the user. The hope is that depending on the nature of the simulated body of smoke, the user can improve performance by allowing certain voxels to coarsen (or remain coarse) with at most moderate repercussions on the aesthetics and behavior of the simulation. The task then is to find an easy way to specify the criteria by which to identify these voxels, and also what level of sacrifice they are willing to make. We decided on three useful criteria to be presented to the user in a drop-down menu as options for a *coarsen* property (figure 1.b): the solver can afford to coarsen interiors of dense plumes, areas of low disturbance, areas where there is low fog detail, or all of the above. Determining the extent of the coarsening is the next challenge. Each criterion has its own set of parameters and thresholds for controlling when, where and by how much to coarsen, the totality of which would be unnecessarily confusing to manage for the user. Thus, the idea of a *fidelity* control was introduced, a single parameter with a range from zero to one, letting the user decide how much they are willing to deviate from the original simulation as a tradeoff for more coarsening. The algorithm described in the next section then sets internal parameters accordingly; in most cases, even a *fidelity* of zero produces results with barely noticeable visual or behavioral differences from the original simulation while providing performance advantages (see video). In summary, the parameters exposed are (figure 1.b):

- *coarsen*: the adaptivity criteria used to rank all voxels. Currently we support *disturbance* (the norm of the velocity gradient), *fog-detail* (the absolute value of the second derivatives of the visible gas), *interior* and *all*.
- *interior density threshold*: the visible gas in a voxel is classified as *interior* if above this value.
- *fidelity*: a value between zero (discard as many voxels as possible) and one (retain as many voxels as the equivalent sparse but non-adaptive simulation).
- *adaptivity smoothness*: a value between zero and one that maps to the maximum average $L1$ interpolation-error allowed by coarsening a voxel.

Some challenges remain, such as not having any obvious way to fold in the *interior density threshold* with the *fidelity* parameter along with not being able to detect whether interior coarsening produces visual artifacts in less opaque simulations. Perhaps decreasing *fidelity* should automatically increase *adaptivity smoothness*, or perhaps a user could directly specify an outer thickness. Other considerations include concepts such as optical depth and camera position to identify areas for coarsening, but this also has dependencies on external data which might cause a deviation from the desired workflow simplicity, unless somehow automatically detected. Nonetheless, this consolidated and simplified workflow has made auto-adaptivity practical for our users at last.

3 OUR AUTO-ADAPTIVITY ALGORITHM

The auto-adaptivity algorithm determines which voxels to respectively keep, coarsen and refine based on the user-provided input parameters listed in the previous section and shown in figure 1.b. A target voxel-budget B is first computed from the *fidelity* parameter f as $B = (f(1-p)+p)\,B_{\text{sparse}}$ where $p = 0.3$ is the fraction of voxels kept with a $fidelity$ of zero, and B_{sparse} is the number of voxels in the corresponding sparse simulation that would result from refining all voxels containing smoke to the finest level. Next, we compute a normalized quality measure q from the adaptivity criteria for each voxel that is a candidate for refinement as well as tile (4^3 voxels) that is a candidate for coarsening. Now let both types of candidates be denoted by $x_i \in \{0, 1\}$. In particular, if x_i is a tile candidate for coarsening, $x_i = 0$ implies coarsen and $x_i = 1$ implies keep. If x_i is a voxel candidate for refinement, $x_i = 0$ implies keep and $x_i = 1$ implies refine. Assuming we have a total of N candidates, a global optimization to coarsen, keep and refine the most optimal set of voxels can be formulated as:

$$\arg\max_{x_i} \sum_{i}^{N} x_i q_i \quad \text{s.t.} \quad \sum_{i}^{N} x_i n_i \leq B$$

where n_i is the number of voxels the resulting coarsening or refinement contributes to the total voxel-count. This is an instance of the *0-1 knapsack* problem which is weakly NP-hard and can be solved by dynamic programming in pseudo-polynomial time and space $O(NB)$. However, since both N and B can be large this is impractical and we opt to simplify the problem in the search for a faster algorithm. By introducing the restriction that a voxel (tile) can at most be refined (coarsened) by one level at a time, we can simplify the problem because $\forall i, j : n_i = n_j = n$ where $n = 4^3$ is a constant. In particular, coarsening a tile will reduce the voxel-count by 4^3, and refining a voxel will increase the voxel-count by 4^3. The simplified problem is solvable in time and space $O(N)$ as follows. Set $k = B/n$ and find the k'th largest quality q_m. This is an instance of the *selection* problem for which $O(N)$ time algorithms are known. Next define the solution as $x_i = 1$ if $q_i \geq q_m$ and $x_i = 0$ otherwise.

REFERENCES

Ryoichi Ando, Nils Thürey, and Chris Wojtan. 2013. Highly Adaptive Liquid Simulations on Tetrahedral Meshes. *ACM Trans. Graph.* 32, 4, Article Article 103 (July 2013), 10 pages. https://doi.org/10.1145/2461912.2461982

Frank Losasso, Frédéric Gibou, and Ron Fedkiw. 2004. Simulating Water and Smoke with an Octree Data Structure. In *ACM SIGGRAPH 2004 Papers (SIGGRAPH '04)*. Association for Computing Machinery, New York, NY, USA, 457–462. https://doi.org/10.1145/1186562.1015745

Michael B. Nielsen, Konstantinos Stamatelos, Morten Bojsen-Hansen, Duncan Brinsmead, Yannick Pomerleau, Marcus Nordenstam, and Robert Bridson. 2018. A Collocated Spatially Adaptive Approach to Smoke Simulation in Bifrost. In *ACM SIGGRAPH 2018 Talks (SIGGRAPH '18)*. ACM, New York, NY, USA, Article 77, 2 pages. https://doi.org/10.1145/3214745.3214749

Rajsekhar Setaluri, Mridul Aanjaneya, Sean Bauer, and Eftychios Sifakis. 2014. SPGrid: A Sparse Paged Grid Structure Applied to Adaptive Smoke Simulation. *ACM Trans. Graph.* 33, 6, Article Article 205 (Nov. 2014), 12 pages. https://doi.org/10.1145/2661229.2661269