



Real-time landscape-size convective clouds simulation and rendering

Prashant Goswami, Fabrice Neyret

► To cite this version:

Prashant Goswami, Fabrice Neyret. Real-time landscape-size convective clouds simulation and rendering. [Research Report] RR-8919, INRIA. 2016, pp.17. hal-01325905

HAL Id: hal-01325905

<https://inria.hal.science/hal-01325905>

Submitted on 7 Jun 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution| 4.0 International License



Real-time landscape-size convective clouds simulation and rendering

Prashant Goswami, Fabrice Neyret

**RESEARCH
REPORT**

N° 8919

6 February 2016

Project-Teams Maverick



Real-time landscape-size convective clouds simulation and rendering

Prashant Goswami^{*†}, Fabrice Neyret^{‡†}

Project-Teams Maverick

Research Report n° 8919 — 6 February 2016 — 17 pages

Abstract: This paper presents an efficient, physics-based procedural model for the real-time animation and visualization of cumulus clouds at landscape size. We couple a coarse Lagrangian model of *air parcels* with a procedural amplification using volumetric noise. Our Lagrangian model draws an aerology *i.e.*, the atmospheric physics of hydrostatic atmosphere with thermodynamics transforms, augmented by a model of mixing between parcels and environment. In addition to the particle-particle interactions, we introduce particle-implicit environment interactions. In contrast to the usual fluid simulation, we thus do not need to sample the transparent environment, a key property for real-time efficiency and scalability to large domains. Inheriting from the high-level physics of aerology, we also validate our simulation by comparing it to predictive diagrams, and we show how the user can easily control key aspects of the result such as the cloud base and top altitude. Our model is thus fast, physical and controllable.

Key-words: clouds, natural phenomena, real-time, animation, amplification, GPU

* Now at Blekinge Institute of Technology – prashant.goswami@bth.se

† INRIA & Grenoble University

‡ fabrice.neyret@inria.fr

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Simulation et rendu temps-réel de nuages convectifs à l'échelle du paysage

Résumé : Cet article présente un modèle temps-réel procédural basé sur la physique pour l'animation et la visualisation temps réel de cumulus à l'échelle du paysage. Nous couplons un modèle Lagrangien basse-résolution de *parcelles atmosphériques* avec un modèle d'amplification procédurale à base de bruit volumique. Notre modèle Lagrangien s'appuie sur l'aérologie *i.e.*, la physique de l'atmosphère en équilibre hydrostatique avec transformations thermodynamiques, augmentée d'un modèle de mélange entre les parcelles et de l'environnement. En plus des interactions entre particules, nous introduisons les interactions entre les particules et l'environnement implicites. Contrairement à la simulation de fluides habituel, on n'a donc pas besoin d'échantillonner l'environnement transparent, une propriété clé pour l'accessibilité au temps réel et la scalabilité aux grandes étendues. Héritant de la physique à haut niveau de l'aérologie, nous validons notre simulation en la comparant aux diagrammes prédictifs, et nous montrons comment l'utilisateur peut contrôler facilement les principaux aspects du résultat comme l'altitude de la base et du haut des nuages. Notre modèle est donc rapide, physique et contrôlable.

Mots-clés : nuages, phénomènes naturels, temps-réel, animation, amplification, GPU

1 Introduction

Clouds are ubiquitous in outdoor naturalistic scenes, especially cumulus clouds which make interesting and conspicuous patterns. Thus they form an important component of many outdoor graphics applications, such as video games and flight simulators. Since these applications are becoming increasingly realistic, covering large and distributed landscapes, a similar trend is expected from their clouds which is currently not the case. Moreover, clouds could be one of the secondary features in these applications. Therefore, their animation and rendering should take only a small fraction of the total frame time.

The requirements are contradictory. Purely procedural models like [1, 2, 3] cannot convey the animated look of cumulus clouds which is a complex and evolving volumetric phenomena. On the other hand, traditional fluid simulation methods like computational fluid dynamics (CFD) simulations are not affordable in real-time, especially at landscape scale or with thin 3D details.

The evolving cloud simulation through huge Eulerian solvers is challenging even for physicists, due to the fact that small-scale and large ranges are required at the same time causing the numerical blending to affect the thermodynamics. In the scope of real-time graphics, cloud simulations do exist, for example based on Eulerian physics [4]. However, they address only very limited volumes or show very coarse resolutions. One of the big disadvantages both with CFD and Eulerian methods is that most of the sky volume is filled with invisible air that needs to be stored and computed. Cloud simulation has also been proposed using Lagrangian framework [5] but such methods are dependent on a high particle count to achieve realistic resolution.

We propose a hybrid *physics-inspired procedural* method for the real-time simulation and rendering of detailed clouds. The main contributions of our approach are:

- We introduce *smoothed-particle hydrodynamics* (SPH) based Lagrangian framework for the cloud physics, operating at the macro-level on large particles called *thermals* making it computationally inexpensive.
- To avoid sampling the ubiquitous clear sky around clouds, we extend the Lagrangian formalism to account for the hybrid particle-implicit environment which is the key for real-time and scalability. For closure, we also introduce a new formulation of cloud-environment mixing (entrainment/detrainment).
- The use of Lagrangian formalism, aerology physics and separation of environment from the dynamics brings more robustness, offers meaningful high-level handles for the user control on environment profile and evolution, cloud base and top.
- For visualization, we treat these thermals as procedural units containing hypertexture. The volumetric texture inside a thermal operates at the micro-level and is produced based on its macro physical state.

The model lightness makes possible the first real-time, landscape-size convective cloud interactive simulation, with rendering capturing small-scale details. It is at the same time physically validated and controllable, which is important for applications aiming at the contradictory requirements of natural look, predictable aspect and performance.

2 Related Work

Much work exists on fluid simulation in the graphics literature, most of which is out of reach of real-time constraints for cloud animation or rendering. Moreover, these methods often account for a very limited volume since transparent regions generally cost as much as the visible ones. There do exist a few multiscale simulation schemes which make use of multigrid and finite elements but the performance is still far from real-time. Moreover, cloud physics includes thermodynamics




Figure 1: Cloud life cycle demonstrated using 75 parcels resulting from the space-time distributed emission from the ground with probabilities depending on the local wetness and temperature. Clouds condense when reaching the condensation level altitude, rise up to equilibrium altitude, then decay and die due to the effects of entrainment and detrainment.

which is a crucial ingredient (and constraint) since advected scalars are reactive; clouds are not just smoke.

The existing approaches for cloud animation in computer graphics follow different strategies. In [6] the cloud evolution is simulated using cellular automaton where the dynamics are specified using simple transition rules. [7] proposes a method for animating clouds surrounding the earth at planetary level based on high and low pressure regions. Mostly [8, 9, 4] account for thermodynamics and are thus directly relevant to us. However, they come with all the problems of Eulerian simulations like limited domain, coarse resolution, harsh initial conditions, energy damping, numerical dissipation, lack of control, etc. Eulerian approaches also suffer a lot from low resolution since artificial blending modifies the physics, damping energy, etc. misrepresenting important characteristics such as vorticity and free boundaries.

Lagrangian approaches like SPH [10] are considered more reliable on moving features, even at low resolutions. [5] recently proposed a particle-based method for cloud simulation based on position-based fluids [11]. Their approach is based on physics and thermodynamics and relies on sampling air in a limited domain which involves merging and splitting SPH particles in transparent and cloudy regions respectively. Consequently, a high particle count is necessitated wherein physics computation takes up the bulk of computational time. and visualization is achieved as a secondary, offline step.

A key issue with sampling the transparent air is that equations are only a part of the problem; initial conditions, state of the whole atmosphere and evolving boundary conditions like humidity and temperature on the floor are equally crucial. It is therefore, advisable to separate inputs comprising the tropospheric environment (for which the hydrostatic hypothesis is very reasonable), from simulated outputs *the clouds*.

In terms of the control of fluids, the existing works aim at artistic control of fluid animation or pure proceduralism [12, 13]. This is totally different to controlling a few high-level, visible characteristics emerging from the simulation such as nebulosity, size and altitude range of clouds. Contrary to these methods, we rely on the high-level physics to simply govern these emerging characteristics. Some work on real-time visualization of clouds is also based on 2D textural cloud layer or sprite-based clouds. The rendering part in our technique could benefit by using one or more of these methods. The reader is referred to [14] for a more detailed insight into the rendering methods for clouds.

3 Atmospheric Physics

Rising air parcels are subjected to the laws of thermodynamics, buoyancy and mixing with the surrounding environment, which we briefly describe here.

Thermodynamics: Humidity is characterized by the *mixing ratio* ω which is the weight of vapor in grams per kilogram of dry air. The *saturated mixing ratio* ω_s is the limit beyond which the extra vapor condenses into liquid water, releasing *latent heat* $L (= 2260 \text{ kJ/kg})$ per gram of condensed mass. The temperature change in a parcel with mass m due to condensation (or vaporization) of Δm of vapor is

$$\Delta T = \alpha \frac{L}{C_p} \frac{\Delta m}{m} \quad [\text{Kelvin}] \quad (1)$$

considering adiabatic conditions. Here $C_p = 1.006 \text{ kJ/kg.K}$ is the specific heat of air and α is a constant (similar to [5]) which is empirically set to 0.3.

The saturated mixing ratio is obtained from the pressure P (which changes with altitude h) and the *saturated vapor pressure* e_s (which depends only on the temperature T in Kelvin):

$$P = P_0 e^{-0.143h} \quad [\text{hPascal}] \quad (2)$$

$$e_s = 6.11 \times 10^{\frac{7.5(T-273.15)}{T-35.45}} \quad [\text{hPascal}]$$

$$\omega_s = 621.97 \frac{e_s}{P - e_s} \quad [\frac{g}{kg}] \quad (3)$$

where $P_0 \approx 1013.25 \text{ hPascal}$.

Temperature, buoyancy and stability: Two key ingredients here are the atmospheric profile (temperature and humidity with altitude) and the initial conditions on the ground where the thermal bubbles start from. With the help of these parameters, one can obtain the altitude of cloud base known as the condensation level where the saturation starts ($\omega = \omega_s$) and the top where the buoyancy vanishes and the parcel reaches a stable equilibrium. To ease the comparisons and computations, atmospheric physicists introduced several alternate notions of temperature. While *potential temperature* $\Theta = T \left(\frac{P_0}{P} \right)^{0.286}$ and *virtual temperature* $T_v = T(1 + 0.61\omega)$ account for the change of pressure and moisture respectively on parcel, *virtual potential temperature* Θ_v accounts for the combined effect of both.

$$\Theta_v = \Theta(1 + 0.61\omega) \quad [\text{Kelvin}] \quad (4)$$

With these definitions, Archimedes' buoyancy force per unit mass can be simplified to

$$\mathbf{f}_{buoy} \approx \left(\frac{\Theta_v^{parcel} - \Theta_v^{atm}}{\Theta_v^{atm}} \right) g \quad [\frac{\text{Newton}}{kg}] \quad (5)$$

(by using ideal gas law $P \propto \rho T$).

Parcel-environment interactions: Mixing has a profound effect on the parcels since it modifies water and air content as well as the temperature. In this context, *entrainment* is engulfing relatively dry environment air in the cloudy parcel while *detrainment* is leaking of cloud mass into the environment. A commonly accepted formulation often assumes the form $\frac{\delta m}{m \delta z} = \varepsilon - \delta$ where δm is the change in the mass of parcel, z its altitude and ε and δ are entrainment and detrainment constants respectively [15]. Parcels are also subjected to friction or drag force $\mathbf{f}_{airDrag} = -\frac{1}{2} C_D \rho v^2 \mathbf{A}$ where v is the relative velocity of parcel, A its reference surface and C_D is the drag coefficient.

From Laplace's law $PV^\gamma = \text{constant}$ ($\gamma \approx 1.4$ for air) and accounting for possible mass flux and pressure change, the parcel's volume is estimated using

$$V_2 = V_1 \left(\frac{P_1}{P_2} \right)^{\frac{1}{\gamma}} \left(\frac{m_2}{m_1} \right) \quad [m^3] \quad (6)$$

SPH: Since our method is based on the discretization of fluid physics with particles, we need to introduce the SPH formalism. The particles influence each others' movements through pressure and viscosity forces. In SPH a weighting kernel W within a given support radius sr is associated to the particles and pressure is solved through a state law around equilibrium. Density is computed using the neighborhood set over support radius (summation over j) for each particle

$$\rho_i = \sum_j m_j W(\mathbf{x}_i - \mathbf{x}_j, sr) \quad (7)$$

where m_i refers to the mass of i^{th} particle and \mathbf{x}_i its position. Pressure follows the spring mass formulation, where fluid rest density is given by ρ_0 and K is the stiffness constant.

$$p_i = K(\rho_i - \rho_0) \quad (8)$$

The pressure force is derived from the pressure using Navier-Stokes equation

$$\mathbf{f}_{pressure_i} = - \sum_j m_i m_j \left(\frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} \right) \nabla W(\mathbf{x}_i - \mathbf{x}_j, sr) \quad (9)$$

and the viscosity force (or the drag between particles) using relative velocity between neighboring particles (μ is viscosity coefficient)

$$\mathbf{f}_{viscosity_i} = \mu \sum_j m_j \frac{\mathbf{v}_j - \mathbf{v}_i}{\rho_j} \nabla^2 W(\mathbf{x}_i - \mathbf{x}_j, sr) \quad (10)$$

We refer the reader to [16] for the weighting kernels employed in our method and for more in-depth survey of SPH.

4 Model Overview

We adopt a two-scale model, with large particles called *parcels* acting as the physics units and simultaneously containing volumetric texture. Through this decoupling between the macro and micro levels, a physics-inspired procedural method is made possible. We introduce these levels in the following and cover them in detail in section 5 and section 6 respectively.

1. *Macro level* : On the physics front, our particle solver inherits most of the SPH fluid simulation formalism with the difference that our parcels are large and correspond to the thermals in [12], with very little overlapping. Unlike existing methods, the environment profile is modelled implicitly and consists of piecewise linear potential temperature and humidity. This enables us to have particles only for thermals and clouds and not for the whole sky, thereby avoiding expensive physics and value storage for transparent regions. For closure, we introduce a model of particle-implicit environment interaction, hybridised with SPH.
2. *Micro level* : On the visualization front, our amplification relies upon on-the-fly procedural hypertexture generation within the parcels and a GPU volume ray-tracer. We evaluate a per-parcel *filling ratio* factor used to clamp a portion of the volume to account for the internal heterogeneity occurring based on the amount of condensed water content. This is inspired by the fact that condensation first occurs on the top of the parcel and at cloud decay lower average droplet density corresponds to a fractal rather than a continuous haze.

Algorithm 1 Parcel physics computation

```

1: while (animating) do
2:   update neighborhood grid
3:   for all parcels  $p$  do
4:     – get neighbors( $p$ )
5:     – compute density( $p$ ), pressure( $p$ ) 7,8
6:     – compute pressure force  $\mathbf{f}_{pressure}$ , viscosity force  $\mathbf{f}_{viscosity}$  9,10
7:     – compute ambient atmospheric pressure force  $\mathbf{f}_{atmPres}$  2
8:     – compute entrained mass  $\Delta m_e$  and update potential temperature  $\Theta'$  13,14
9:     – eject detrained mass  $\Delta m_d$  15
10:    – compute water mass exchange between neighboring parcels  $\Delta m_p$ 
11:    – update volume  $V$  6
12:    – update mixing ratio (liquid/vapor content) and potential temperature change  $\Delta\Theta$  3,1

13:   – compute buoyant force  $\mathbf{f}_{buoy}$  5
14:   – apply  $\mathbf{f}_{parcel}$  and wind to move the parcel 17

```

5 Macro scale

We rely on the classical Lagrangian model for the particle-particle interaction. The inter-parcel interaction is captured by applying repulsive pressure force $\mathbf{f}_{pressure}$ (Equation 9) that prevents parcels from getting too close to each other and the viscosity force $\mathbf{f}_{viscous}$ (Equation 10) that accounts for the drag between the parcels. Besides the SPH forces (in our model), cloud parcels are also subjected to the thermodynamic rules. Thus in the following, we describe the interaction of parcels with the implicit environment, inter-parcel mixing and the buoyant forces acting on parcels, thereby completing the total parcel physics leading to cloud formation.

5.1 Interactions with the implicit environment

In our model, particles cover only a part of the atmosphere which still must represent a continuum of all physical interactions. We also therefore, need to estimate the contribution of the implicit (*i.e.*, unsampled) environment on the area of parcels exposed to the atmosphere. Thus to these particle-particle interactions, we have to add forces corresponding to the equivalent terms in particle-environment interactions, $\mathbf{f}_{atmPres}$ and $\mathbf{f}_{airDrag}$ respectively.

Atmospheric force $\mathbf{f}_{atmPres}$ is applied to the exposed surface area of each parcel which is in direct contact with the atmosphere. The exposed area (denoted by \mathbb{S}) is determined by coarsely tessellating each parcel into quadrilaterals and determining the set of quadrilaterals that do not lie inside any other parcel. Instead of querying all the parcels, a quadrilateral needs to perform this test only on the SPH neighbors of the parcel it belongs to. Atmospheric pressure is obtained using Equation 2 and is applied to each tessellated quadrilateral along its normal weighted by the area. $\mathbf{f}_{atmPres}$ needs an additional weighting factor, which we obtain through calibration based on our experiments.

5.2 Mixing with the environment

The parcels interact with the surrounding environment through the process of *entrainment* and *detrainment*. The total amount of entrained air in a parcel during time interval Δt is a function of both its exposed area \mathbb{S} and the velocity, which in essence gives the swept volume of the

parcel through the external relatively dry air. Even if a parcel is globally motionless, small-scale turbulence pursues the mixing; otherwise clouds would not decay. To account for this, we thus add to the macroscopic drag a small turbulence-related term $\delta \mathbf{v}_t$. Furthermore, fractal nature of clouds makes the real cloud-air contact surface larger than one for the macroscopic sampling primitives. For this we assume proportionality coefficients, K_e for entrainment and K_d for detrainment, which is the surface multiplier to account for subgrid fractal surface.

Entrainment entails the mixing of external air Δm_e^a (Equation 11), which also carries with it the water vapor Δm_e^w (Equation 12) already present in it (computed using mixing ratio), with the parcel. The total mass gain Δm_e is obtained by summing these two quantities (Equation 13). To incorporate it in our formulation, we modify the standard entrainment and detrainment constants to $\varepsilon' = \varepsilon r$ and $\delta' = \delta r$, where r is the radius of the parcel.

$$\Delta m_e^a = K_e \varepsilon' A \rho^{atm} \|\mathbf{v} + \delta \mathbf{v}_t\| \Delta t \quad (11)$$

$$\Delta m_e^w = \omega_s \Delta m_e^a \quad (12)$$

$$\Delta m_e = \Delta m_e^a + \Delta m_e^w \quad (13)$$

Here ε is the entrainment rate taken from [15] and A is the exposed surface area obtained from the computed \mathbb{S} . The entrained mass comprising both dry air and vapor is at a different temperature than the parcel. The equilibrium temperature Θ' is computed by mixing the influx with the existing parcel contents m , see Equation 14.

$$\Theta' \approx \frac{m \Theta^{parcel} + \Delta m_e \Theta^{atm}}{m + \Delta m_e} \quad (14)$$

Subsequently, the quantity of detrained content Δm_d (Equation 15) is determined using a similar expression as for entrainment albeit with a few changes. ρ^{atm} in Equation 11 is substituted with the average density ρ in the parcel to compute Δm_d^a . Assuming a uniform mixing of vapor in the parcel, an equal proportion of water mass Δm_d^w is ejected together with Δm_d^a . Furthermore, ε is also replaced by δ , the detrainment constant. The total mass detrained from a parcel is given by Equation 15.

$$\Delta m_d = \Delta m_d^a + \Delta m_d^w \quad (15)$$

5.3 Inter-parcel mixing

Much as the parcels exchange mass with the external environment through their exposed surface, a similar process takes between the neighboring parcels through the contact surface. The inter-parcel mass exchange Δm_p is obtained using a simple diffusion process from high concentration parcel to lower ones surrounding it.

With this, the combined mixing occurring to a parcel could be summed up as:

$$\Delta m = \Delta m_e - \Delta m_d + \Delta m_p \quad (16)$$

5.4 Buoyant force

In addition to aforementioned forces, buoyant force \mathbf{f}_{buoy} is added to each parcel which is a key component in parcel rising and cloud formation. With the computation of updated potential temperature Θ' (Equation 14) inside the parcel, buoyant force follows from Equation 5 based on the temperature difference with the surrounding environment.




Figure 2: *Left*: hypertexture in the shell (between the blue and the red curves) of sphere \cup half-space above condensation level. *right*: with filling ratio $f=50\%$.

With this, the total force acting on a parcel is:

$$\mathbf{f}_{\text{parcel}} = \mathbf{f}_{\text{pressure}} + \mathbf{f}_{\text{viscous}} + \mathbf{f}_{\text{atmPres}} + \mathbf{f}_{\text{airDrag}} + \mathbf{f}_{\text{buoy}} \quad (17)$$

6 Micro Scale

The micro level deals with the procedural generation of cloud details on top of the existing macro structure of parcels. The hypertexture is computed on-the-fly, and is controlled by a few macro-level handles such as particle radius, condensation level plane and filling ratio which are themselves derived from the ongoing parcel physics. This is detailed in subsection 6.1. We employ GPU-based ray-tracing to efficiently render our clouds as explained in subsection 6.2.

6.1 Volumetric cloud texture

From the local list of parcels intersecting a given point in space, we employ the hypertexture framework of [17] to compute the union of clamped hypertextured parcels (see Figure 2):

- The pseudo-distance is computed as $\cup_i(\text{sphere}_i \cap \text{halfSpace}_i)$ where the sphere corresponds to the parcel and the half-space corresponds to clamping by the condensation level point plane. Since our SPH kernels have low super-imposition, we use spheres larger than these kernels for nice looking unions. Note that we can stop the \cup_i loop as soon as it saturates to 1.
- The sphere and half-space shell thickness is a user-tunable constant, but is modulated by the parcel filling ratio f to account for the cloud decay so that the noise texture gets sponge-like and erodes with the average density proportional to f . The density is then computed based on the noise and distance function (where k is a constant):

$$\rho(P) = \text{smoothstep}(f \cdot (\text{dist}(P) + 1) - 1 + k \cdot \text{noise}(P)) \quad (18)$$

The factor f estimates the percentage of a parcel filled with condensed (opaque) material. The idea is that the filling is not homogeneous. Condensation starts at condensation level (denoted by half-space) and grows making the parcel more opaque while decaying carves a sponge-like




Figure 3: Attributes of a parcel varying with time in our simulation, with mixing enabled and disabled: (*left*) parcel position, (*mid*) parcel water mass ratio relative to its initial value, (*right*) parcel radius ratio relative to its initial value.

texture. Therefore, in practice an average droplet density corresponds to a mix of clear and opaque fractions. This relative volume of opaque fraction is the filling ratio. We estimate it as $f = \frac{\omega_{cloud} - \omega_s}{V_{1kg}^{air} \cdot q_{l0}}$ where V_{1kg}^{air} is the volume of 1kg of air at parcel altitude and q_{l0} is the droplet density in the opaque fraction.

6.2 Ray-tracing clouds

The real-time volumetric details generated from subsection 6.1 are visualized using volume ray casting. Our simple physics engine can easily be coupled to a reasonably minimal volume ray-casting fragment shader. To this end, we have used the *uniform buffer object (UBO)* which allows us to pass arrays to the fragment shader. We transfer the entire grid together with the list of condensed parcels and the related attributes (parcel size, filling ratio, plane position) as *UBO* to the GPU.

7 Results and discussion

The proposed method was implemented using C++, OpenGL version 4.0 and tested on an Intel Core 3.4 GHz machine with 16 GB RAM and an NVidia GeForce GTX 970 graphics card on a display window of 730×545 pixels.

The number of parcels is user chosen and can vary from single to many. Our model allows realistic, landscape-scale cloud effects by using just few macro primitives. The physics part therefore, is a lightweight computation which can be carried out on the CPU itself. However, this is not a limitation of our method as in case where larger number of macro primitives are desired, physics computation could be easily and efficiently transferred to GPU as laid out in [18]. We also cast simplified shadow of clouds on the OpenGL terrain. This is achieved by shooting rays from the light source towards the ground to collect cloud opacities in a 2D texture in the first pass and blending them with the scene at high resolution in the second pass.

The parcel tessellation for atmospheric pressure and mixing computation is kept coarse (18×9) in our case which provides a good balance between accuracy and computational effort. We set the environment dimensions to $10km \times 10km \times 10km$ where the simulation and visualization is carried out, although it is flexible and user modifiable (and could as well be irregular). In terms of Lagrangian parameters, density of air is taken to be around $\rho_0 = 1.28kg/m^3$, stiffness constant $K = 100$ (since air is compressible) and a low viscosity coefficient $\mu \approx 0.01$ is chosen.

7.1 Aerology validation

Single-parcel kinematics is verified against aerology diagram for the following important parameters:

- *Lapse rates*: As computed directly from Figure 4a, the *dry adiabatic lapse rate* comes out to be around -10 C/km (slope between points A and B), whereas the *moist adiabatic lapse rate* is about -5 C/km (average slope between points B and D). These values are close to the actual figures, as given in [19]. Air parcels that get saturated as they rise cool at a smaller rate than the dry adiabatic lapse rate due to the heat released by the condensation of water vapor (B onwards). The wet adiabatic lapse rate varies with the altitude which is why we take into account the average value.
- *Equilibrium altitude*: In the case of no mixing with the atmosphere, the parcel reaches to a maximum altitude and comes to an equilibrium height after a series of damping oscillations, as is visible in Figure 4a, see also Figure 3a. The maximum altitude lies close to the point where parcel temperature intersects the atmospheric temperature profile. The parcel continues moving further up till its velocity is reversed by the negative buoyancy force. However, mixing makes a parcel lose its water content, thereby changing the internal dynamics. During this decaying process the parcel undergoes oscillatory motion too as it alternates between rising and sinking buoyant force [20], which is also validated by our simulation as shown in Figure 4b and Figure 3a. Note that while the parcel follows the same temperature trajectory during its oscillations in absence of mixing (Figure 4a), the temperature profile is altered with each oscillation if the mixing is enabled (Figure 4b).
- *Mass and radius*: The variation of parcel water mass and radius relative to its initial values is as demonstrated in Figure 3b and Figure 3c respectively. The initial mass gain with mixing enabled is explained by the fact that as per the model from [15], detrainment begins at condensation level while entrainment occurs throughout. It is for the same reason that the parcel begins to gain mass towards the end of its lifetime as it reaches the condensation level again. Eventually, the parcel exhausts its water content significantly by leaking it to the atmosphere. The parcel radius grows as it ascends due to the decrease in atmospheric pressure. However, in the case of mixing with atmosphere the parcel radius shrinks due to the combined effect of water loss and descent.

7.2 User control

The condensation level altitude and the nature of clouds obtained from parcels is dependent on the initial mixing ratio and the temperature profile. The cloud top is governed by the *temperature inversion*, the altitude at which the temperature begins to increase and hence the clouds parcel cease to rise further up. Though the default implicit environment is predefined in our system, our tool (Figure 5) gives the user an easy and intuitive handle to:

- alter the temperature profile graphically
- select condensation level
- choose mixing ratio of the atmosphere

The system simulates the physics with the user requirements and the clouds are generated if the state is physically plausible.

7.3 Cloud life cycle

Our parcel generator automatically generates parcels with a frequency proportional to the heat and humidity content of the region. Several time-consistent plumes of parcels could be emitted one after the other piling up on top of each other and giving rise to interesting cloud patterns.

# parcels	with FDM normals			without FDM normals		
	physics	render	fps	physics	render	fps
8	6%	89%	128	4%	78%	235
14	6%	89%	77	6%	91%	160
34	8%	87%	24	8%	89%	63
85	9%	82%	7	9%	88%	21

Table 1: Performance split-up between the physics part on the CPU and the rendering part on the GPU together with the frame rates per second (fps).

Figure 1 illustrates the complete life cycle of clouds as parcels from their birth to death. Clouds often die out in order of the sequence of their creation and also depending on their exposure to the outside dry air.

7.4 Performance analysis

We have tried upto 85 dynamically produced parcels where the simulation together with the rendering continues to be interactive. The performance analysis of our system is summarised in Table 1, where the percentages are given relative to the total time spent in the system per frame. In there we also compare our performance with and without computation of normals (for voxel shading) using finite difference method (FDM) for rendering. Our experiments suggest that rendering without finite differences normal is significantly faster than with normals. The only expensive part in parcel physics is the computation of area exposed to implicit environment on the tessellated sphere. On the other hand, hypertexture production and raycasting based visualization takes a larger share of overall computational time as compared to physics.

In Figure 6, a realistic and distributed cloud cover is obtained with just 14 parcels. Further, interesting cloud patterns are obtained by taking the union of density parcels in the overlapping regions as shown in Figure 7. In addition to the above mentioned experiments, being a Lagrangian-based technique our method can easily deal with the effects of wind, parcel-obstacle collision which could influence the cloud movement and formation. Furthermore, the atmospheric temperature and humidity profile can be varied as a function of time of the day to influence cloud formation. These aspects of our method are demonstrated in the accompanying video.

8 Conclusions and future work

We have presented an efficient, physics-based procedural hybrid method for real-time landscape-size simulation and visualization of cumulus clouds. The use of Lagrangian framework at the macro parcel level makes the physics computationally inexpensive. The introduction of the implicit atmosphere eliminates the need to sample transparent air and accounts for the parcel-air interaction and mass exchange. The real-time visualization of clouds is achieved by raytracing the procedural amplification of volumetric noise on the GPU. Our model is controllable by the user and predictive, and can be validated against the areology diagram for basic parameters.

For future work, we would like to add more complex physics in order to permit billowing of cumulus clouds, and study hierarchical Lagrangian simulation for even larger scale scalability.

References

- [1] David S. Ebert. Volumetric modeling with implicit functions: A cloud is born. In *ACM SIGGRAPH 97 Visual Proceedings: The Art and Interdisciplinary Programs of SIGGRAPH*, SIGGRAPH, pages 147–, New York, NY, USA, 1997. ACM.
- [2] Andrew P. Schneider, Trevor G. Thomson, and Mathew S. Wilson. Clouds in the skies of Rio. In *ACM SIGGRAPH Talk*, 2011.
- [3] Brett Miller, Ken Museth, Devon Penney, and Nafees Bin Zafar. Cloud Modeling And Rendering for "Puss In Boots". In *ACM SIGGRAPH Talk*, 2012.
- [4] Mark J Harris, William V Baxter, Thorsten Scheuermann, and Anselmo Lastra. Simulation of cloud dynamics on graphics hardware. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, pages 92–101, 2003.
- [5] Charles Welton Ferreira Barbosa, Yoshinori Dobashi, and Tsuyoshi Yamamoto. Adaptive cloud simulation using position based fluids. *Computer Animation and Virtual Worlds*, 26(3-4):367–375, 2015.
- [6] Yoshinori Dobashi, Kazufumi Kaneda, Hideo Yamashita, Tsuyoshi Okita, and Tomoyuki Nishita. A simple, efficient method for realistic animation of clouds. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH, pages 19–28, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.
- [7] Yoshinori Dobashi, Tsuyoshi Yamamoto, and Tomoyuki Nishita. A controllable method for animation of earth-scale clouds. In *Proc. of CASA*, pages 43–52, 2006.
- [8] James T. Kajiya and Brian P Von Herzen. Ray tracing volume densities. *SIGGRAPH Comput. Graph.*, 18(3):165–174, January 1984.
- [9] Fabrice Neyret. Qualitative simulation of cloud formation and evolution. In *8th Eurographics Workshop on Computer Animation and Simulation (EGCAS)*, pages 113–124, sep 1997.
- [10] Matthias Müller, David Charypar, and Markus Gross. Particle-based fluid simulation for interactive applications. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA, pages 154–159, 2003.
- [11] Miles Macklin and Matthias Müller. Position based fluids. *ACM Trans. Graph.*, 32(4):104:1–104:12, July 2013.
- [12] Joshua Schpok, Joseph Simons, David S. Ebert, and Charles Hansen. A real-time cloud modeling, rendering, and animation system. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA, pages 160–166, 2003.
- [13] Yoshinori Dobashi, Katsutoshi Kusumoto, Tomoyuki Nishita, and Tsuyoshi Yamamoto. Feedback control of cumuliform cloud formation based on computational fluid dynamics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 27(3), 2008.
- [14] Zhaoxin Ye. Volumetric Cloud Rendering: An Animation of Clouds. Master’s thesis, Clemson University, the Netherlands, 2014.
- [15] P. M. M. Soares, P. M. A. Miranda, A. P. Siebesma, and J. Teixeira. An eddy-diffusivity/mass-flux parametrization for dry and shallow cumulus convection. *Quarterly Journal of the Royal Meteorological Society*, 130(604):3365–3383, 2004.

- [16] J. MONAGHAN. Smoothed particle hydrodynamics. *Rep.Prog. Phys.*, 68:1703–1759, 2005.
- [17] K. Perlin and E. M. Hoffert. Hypertexture. *SIGGRAPH Comput. Graph.*, 23(3):253–262, July 1989.
- [18] Prashant Goswami, Philipp Schlegel, Barbara Solenthaler, and Renato Pajarola. Interactive sph simulation and rendering on the gpu. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA, pages 55–64, 2010.
- [19] Lazaridis Mihalīs. *First Principles of Meteorology and Air Pollution*, volume 19. Springer Netherlands, 2011.
- [20] Jay S. Hobgood. Fundamentals of weather and climate. robin mcilveen, chapman and hall (london), 1992. no. of pages: xxii + 497. price: £19.95 (p/b). isbn 0-442-31476-0. *International Journal of Climatology*, 13(4):461–461, 1993.

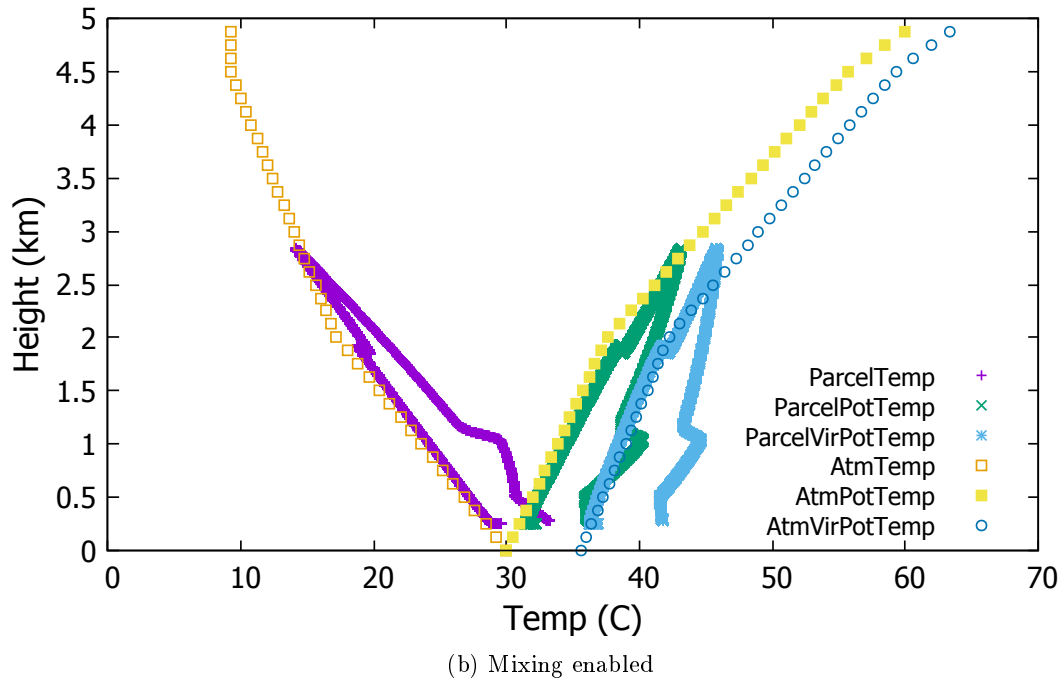



Figure 4: Parcel vs. environment temperatures (T, Θ, Θ_v) during single parcel simulation with (a) mixing disabled (b) mixing enabled.




Figure 5: Our interactive tool allows the user to experiment with intuitive parameters to generate plausible evolving fields of cumulus clouds.




Figure 6: Parcel physics and hypertexture generated with just 14 parcels leading to fast and yet, realistic visualization.



Figure 7: Interesting patterns resembling cumulus clouds are obtained by taking union of the densities in the overlapping parcel regions.



**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399