

Fast SPH simulation for gaseous fluids

Bo Ren¹ · Xiao Yan² · Tao Yang² · Chen-feng Li³ · Ming C. Lin⁴ · Shi-min Hu²

Published online: 28 April 2015
© Springer-Verlag Berlin Heidelberg 2015

Abstract This paper presents a fast smoothed particle hydro-dynamics (SPH) simulation approach for gaseous fluids. Unlike previous SPH gas simulators, which solve the transparent air flow in a fixed simulation domain, the proposed approach directly solves the visible gas without involving the transparent air. By compensating the density and force calculation for the visible gas particles, we completely avoid the need of computational cost on ambient air particles in previous approaches. This allows the computational resources to be exclusively focused on the visible gas, leading to significant performance improvement of SPH gas

simulation. The proposed approach is at least ten times faster than the standard SPH gas simulation strategy and is able to reduce the total particle number by 25–400 times in large open scenes. The proposed approach also enables fast SPH simulation of complex scenes involving liquid–gas transition, such as boiling and evaporation. A particle splitting and merging scheme is proposed to handle the degraded resolution in liquid–gas phase transition. Various examples are provided to demonstrate the effectiveness and efficiency of the proposed approach.

Electronic supplementary material The online version of this article (doi:[10.1007/s00371-015-1086-y](https://doi.org/10.1007/s00371-015-1086-y)) contains supplementary material, which is available to authorized users.

✉ Bo Ren
renboeverywhere@gmail.com

Xiao Yan
shiotoli@gmail.com

Tao Yang
yangtao9009@gmail.com

Chen-feng Li
c.f.li@swansea.ac.uk

Ming C. Lin
lin@cs.unc.edu

Shi-min Hu
shimin@tsinghua.edu.cn

¹ College of Computer and Control Engineering, Nankai University, Tianjin, China

² Department of Computer Science and Technology, Tsinghua University, Beijing, China

³ College of Engineering, Swansea University, Swansea, UK

⁴ Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, USA

Keywords Smoothed particle hydrodynamics · Gas simulation · Adaptive particle splitting and merging · Phase transition

1 Introduction

The smoothed particle hydrodynamics (SPH) method has become increasingly popular for simulating fluid motion for a wide range of natural phenomena and special effects. Yet, despite its popularity in liquid simulation, much less research has been reported in the literature on SPH simulation of gaseous fluids. Liquid is often considered to be incompressible in simulation, though gas usually behaves far more compressible. The time stepping condition in SPH methods requires much smaller time step simulating highly incompressible fluids than compressible fluids. This makes the SPH approach potentially more promising for gas simulation in a performance perspective. However, SPH simulation of gaseous fluids has not been sufficiently explored.

The challenge associated with fast SPH gas simulation can be better understood by examining a simple example, a plume of smoke rising in the air. Using a standard SPH simulator, the air flow is simulated within a fixed domain and the visible

plume is represented by a density field that flows with the transparent air. The simulation domain is usually much larger in its extent than the space occupied by the actual evolving plume; therefore, the majority of the computational cost is spent on simulating the invisible air. The computational cost required to compute the air flow in a large simulation domain can be prohibitively high, making these approaches difficult, if not impossible, to use in simulating large, complex scenes with moving plumes, e.g. a steam train moving through a field or plumes of chimney smoke billowing in a village.

To address this issue, we propose a novel SPH scheme for gaseous fluids in which computation is performed only on the visible gaseous fluids, thus minimizing the simulation cost on ambient air gas particles. The interpolation error for the gas density is compensated for by estimating the average influence of the absent ambient air particles. The missing interaction between the simulation particles and the absent ambient air particles is also compensated for by a virtual pressure that is locally determined based on the current fluid dynamics and the atmosphere pressure. We demonstrate that such an approach is able to minimize the unnecessary computation on ambient gas particles, while significantly improving the runtime performance of SPH gas simulation. The new method also enables fast SPH simulation of complex scenes that involve liquid–gas phase transition. A robust particle splitting and merging strategy is proposed for liquid–gas transition situations; we show that phase transition phenomena can be efficiently captured in a unified SPH framework. We highlight the performance benefits (up to two orders of magnitude) of this approach on several large-scale phenomena as shown in Figs. 5, 4, 6, and 7.

2 Previous work

Recent SPH methods [20,22,27], have focused primarily on liquid simulation. [29] proposed an SPH scheme to create and preserve vortical details near moving objects in gas simulation. However, they fill the whole simulation domain with particles. In some hybrid methods, SPH gas particles are generated in selected regions for various reasons, e.g. to reduce the computational cost of solving high-speed gas flow [10] and to help track the interface and compute gas-liquid interaction [4]. Nonetheless, the entire gas domain needs to be simulated by the grid-based simulator in these hybrid approaches to achieve high-accuracy simulation results.

SPH simulation generally simulates the entire domain because particle deficiencies lead to the erroneous interpolation of particle properties, such as density and pressure. To overcome this degradation of accuracy at rigid boundaries, SPH simulations often add solid particles, both to avoid particle penetration and to improve density computation of the liquid [3,7,19]. For water–air boundaries, ambient air parti-

cles have been added through a serialized sampling process to improve the density and surface tension calculation of the liquid [15,25]. Compared to the interaction between water particles, the force between air and water is largely negligible and is not included in their liquid-simulation model. However, for gas simulation, the interaction between the visible plume and the surrounding ambient air is important and cannot be ignored.

Adaptive SPH simulations have previously used the split-and-merge technique for SPH particles. In [1,13,28] SPH particles were bisected into different levels and particles at the same level were merged on demand. These methods achieve good results in liquid simulation, but require serialized adaptive sampling to maintain a uniform particle distribution and improve stability. More recently [26] proposed a two-scale particle scheme extending the processable density ratio, and [23] blended different levels of particles to achieve a smoother transition during splitting and merging. Their methods involve blending of different levels of particles and relaxing the clustered particles after splitting during a certain time period. This leads to better stability, but comes at an extra computational cost in the splitting process. Although these adaptive approaches adjust particle sizes to reduce the number of ambient air particles being simulated, the simulation cost still remains untenably high for very large simulation domains. Our approach is considerably faster than the previous methods that simulate ambient air particles, because it focuses the computing resources only on the visible gaseous fluids and estimating the average influence of the surrounding air. Such estimation is performed by compensating for the lost gas density and interaction forces that would be exerted by the unsimulated ambient air particles.

There are other works that are also partially related to our approach. [12] and [14] achieved large-scale liquid simulation with Lagrangian method. In [2], sparse particles were allowed in a vortex-based simulation scheme. [5] also tried to achieve similar goal of sparse representation using vortex sheet to represent smoke. Virtual particles for interface handling are also used in works involving solid simulation, such as in [21]. In [16,17] optimization methods are used to reduce simulation errors. In [11] the velocity field is decomposed and only a small part of it is simulated using high-resolution grid, achieving faster simulation. Recently [18] proposed a unified representation focusing on simulating all types of fluids, in which constrains and weak cohesive forces are used to maintain particle density in cases of particle deficiency; however, they involved neither long-ranged movement of smoke in large scenes nor phase transition in the work.

3 Fast SPH gas simulation

In gas simulation, standard SPH simulators (as shown in Fig. 1a) fill the entire simulation domain with ambient air




Fig. 1 **a** The standard approach fills the simulation domain with ambient air particles (*blue*). **b** Our approach removes these ambient air particles and instead adds compensation forces to the gas particles (*red*)

particles. This is impractical for large or unbounded scenes, where large numbers of particles are needed to fill the simulation domain. Yet simply removing some of these particles causes particle deficiency. This is not an issue for free-surface liquids (liquid–air systems), where air particles are usually ignored without consequence, since the liquid and air densities differ so greatly. Atmospheric forces are negligible compared to the liquid particle interaction, and the missing air particles do not cause noticeable errors in calculation. These particle deficiencies do, however, pose serious problems for gaseous simulation with SPH. The forces between the plume and the surrounding air are not negligible as in liquid simulations, but are crucial in determining the movement of simulated gas particles.

In order to achieve high-fidelity visual effects without filling the entire space with ambient air particles, we must compensate for the accuracy degradation from particle deficiency. Particle deficiency has two main effects on accuracy: (1) erroneous density calculation of SPH particles and (2) missing interaction forces from the ambient air particles.

Following the same work flow as the standard SPH method, we derive below a fast SPH scheme (illustrated in Fig. 1b) which only simulates the visible gas and does not model the surrounding ambient air. We have developed appropriate compensation formulations for both gas-density and particle-force calculations. Instead of using a secondary density field to represent the distribution of the visible gas plume, we use the SPH gas particles to directly represent the visible gas; this is also convenient for phase transformation of particles in liquid–gas phase transitions. Symbols used in the following derivation of gas simulation can be found in Table 1.

3.1 Gas density correction

Assuming the ambient air particles are included in the SPH simulation, the following relations hold for a given particle i :

Table 1 Definition of symbols in gas simulation

Symbols	Meaning
$\bar{\rho}_i$	Uncorrected density of particle i
ρ_i	Corrected density of particle i
ρ_0	Rest density of particle i
m_j	Mass of particle j
$W(\mathbf{r}, h)$	Smoothing kernel function
∇W_{ij}	Short for $\nabla_i W(\mathbf{r}_i - \mathbf{r}_j, h)$
$\mathbf{r}_i, \mathbf{r}_j$	Position of the i -th, j -th particle
V_0	Intermediate variable
p_i	Pressure of particle i
T_i	Temperature of particle i
\mathbf{n}_i	Local normal vector at the particle i
P_a, C_d	User defined strength factors
C_b, D_c, D_r	User defined coefficients for buoyancy
C_N	User defined threshold value

$$\rho_i = \sum_j m_j W_{ij} + \sum_{j'} m_{j'} W_{ij'} \quad (1)$$

$$0 = \sum_j \frac{m_j}{\rho_j} \nabla W_{ij} + \sum_{j'} \frac{m_{j'}}{\rho_{j'}} \nabla W_{ij'}, \quad (2)$$

where the index j refers to the SPH particles of interest, the index j' refers to the ambient air particles, ρ_s denotes the density of particle s , m_s denotes the mass of particle s , and $\nabla W_{st} = \nabla_s W(\mathbf{r}_s - \mathbf{r}_t, h)$ is the gradient of the smoothing kernel function with support h . We use the standard smoothing kernel functions as in [20]. Equations (1) and (2) both come from the standard SPH formulation, where Eq. (1) is the density interpolation and Eq. (2) is a first-order identity that holds exactly for ideal particle distributions.

If the ambient air particles are simply removed from the simulation, the second term in Eqs. (1) and (2) will disappear, leading to erroneous results. To correct this error, it is assumed that a virtual “average” particle k can be placed nearby to satisfy Eqs. (1) and (2), i.e.:

$$\rho_i = \sum_j m_j W_{ij} + m_k W_{ik} \quad (3)$$

$$0 = \sum_j \frac{m_j}{\bar{\rho}_j} \nabla W_{ij} + \frac{m_k}{\rho_k} \nabla W_{ik} \quad (4)$$

In Eqs. (3) and (4) the effect of ambient air particles surrounding the particle i is compensated for by a virtual “average” particle k paired with it. We use $\bar{\rho}_i = \sum_j m_j W_{ij}$ to denote the uncorrected density value, i.e. the density interpolation following Eq. (1) without the second term. To determine the properties of the virtual particle k , it is first assumed that densities in adjacent particles change gradually in the gas simulation; the density of particle k is then considered to be

the same as the interpolated density of particle i , i.e. $\rho_k = \rho_i$. Then the distance from any virtual particle k to its paired real particle i is set as a fixed value, which will be discussed later. Equation (4) contains two vector-valued functions that cancel each other. The first term in Eq. (4) can be computed directly and is known for a given particle i . On the right-hand side of Eq. (4), the direction of the second term must be in line with the first term. Therefore, we only need to consider the norms of the two terms in Eq. (4), which can be rewritten as

$$m_k = V_0 \rho_k = V_0 \rho_i, \quad (5)$$

where $V_0 = \frac{|\sum_j \frac{m_j}{\rho_j} \nabla W_{ij}|}{|\nabla W_{ik}|}$ can be predetermined for particle i .

Substituting Eq. (5) into Eq. (3), the corrected density of particle i can be computed as

$$\rho_i = \frac{\bar{\rho}_i}{1 - V_0 W_{ik}} \quad (6)$$

Equation (6) corrects the density of the SPH particles in the absence of surrounding ambient air particles. However, for a more accurate evaluation, the densities of adjacent particles in Eq. (4) should also use the corrected value ρ_j instead of the uncorrected one $\bar{\rho}_j$. This leads to the following formulae, whose detailed proof is given in Appendix A:

$$m_k = V_0 \bar{\rho}_i \quad (7)$$

$$\rho_i = \bar{\rho}_i (1 + V_0 W_{ik}) \quad (8)$$

In the above calculation, the distance from a virtual particle k to its paired real particle i is heuristically set as a fixed value δh , where $\delta \in [0.4, 0.7]$ usually brings more plausible results in our experiments. There may be other methods that allow this distance to vary, but we have found simply fixing the distance produces sufficiently good results for the SPH gas simulation.

3.2 Particle force compensation

After removing the ambient air particles from the simulation, the interaction forces between the simulation particles and the ambient air particles must also be compensated for. The missing interaction forces are responsible for three physical effects: the local force balance, the atmospheric pressure, and the local gas viscosity. Each of these physical effects requires that an appropriate compensation be made.

The virtual particles generated for the gas density correction can be used to compensate the local force balance. Following the standard SPH formulation, the pressure force for a given particle i is

$$\mathbf{F}_{p,i} = - \sum_j \frac{m_j}{\rho_j} \frac{p_i + p_j}{2} \nabla W_{ij}, \quad (9)$$

where p_s is the gas pressure calculated at particle s . Using the paired virtual particle k , an extra term is added to Eq. (9):

$$\tilde{\mathbf{F}}_{p,i} = - \sum_j \frac{m_j}{\rho_j} \frac{p_i + p_j}{2} \nabla W_{ij} + \frac{m_k}{\rho_k} \frac{p_i + p_k}{2} \nabla W_{ik} \quad (10)$$

Here the interpolated density uses the corrected density computed from Eq. (8), and the pressure is computed from the ideal gas state function:

$$p = \kappa (\rho - \rho_0), \quad (11)$$

where κ is the gas constant, and ρ_0 is the gas density at rest.

The atmospheric pressure pushes the gas towards its inner direction; to simulate this effect, we introduce a normal force. First, the local normal vector at the particle i is computed as

$$\mathbf{n}_i = \sum_j \frac{m_j}{\rho_j} \nabla W_{ij} \quad (12)$$

The normal vector defined above is identical to the first term in Eq. (2). For those internal simulation particles whose neighborhoods do not contain any ambient air particles, the second term in Eq. (2) is zero and, therefore, the norms of their normal vectors are zero or approximately zero, depending on the particle distribution. For those simulation particles that are adjacent to the missing ambient air particles, the second term in Eq. (2) is non-zero and, therefore, the norms of their normal vectors are non-zero as defined by the first term. Based on this analysis, the atmospheric pressure force is defined, as

$$\tilde{\mathbf{F}}'_{p,i} = P_a \mathbf{n}_i = P_a \sum_j \frac{m_j}{\rho_j} \nabla W_{ij}, \quad (13)$$

where P_a is a user-defined strength coefficient reflecting the ambient air pressure. The pressure force defined above automatically ensures that the atmospheric pressure is added only to the particles on the gas boundary since the inner gas particles have zero norms for their normal vectors.

To account for the effect of the local gas viscosity, we introduce a damping effect proportional to the particle velocity as follows:

$$\tilde{\mathbf{a}}_{\text{damp},i} = -C_d \mathbf{v}_i, \quad (14)$$

where $\mathbf{a}_{\text{damp},i}$ is the acceleration of the i -th particle due to the damping force, C_d is the damping coefficient, and \mathbf{v}_i denotes the velocity of particle i . This damping force is applied only if the norm of the particle normal exceeds a given threshold value, i.e. $|\mathbf{n}_i| > C_N$. This is to ensure that the damping effect is only added to those boundary particles with neighborhood deficiency.

3.3 Buoyancy and boundary treatment

The buoyancy force is related to the temperature of SPH gas particles. Similarly to [6], the temperature of each gas particle is evolved as

$$\frac{dT_i}{dt} = \sum_j \frac{m_j}{\rho_i \rho_j} D_c (T_i - T_j) \frac{(\mathbf{r}_i - \mathbf{r}_j) \cdot \nabla W_{ij}}{(\mathbf{r}_i - \mathbf{r}_j)^2 + \gamma^2}, \quad (15)$$

where T_s is the temperature of particle s , D_c is the thermal conductivity, and $\gamma^2 \ll 1$ is a small positive number to avoid computational singularity. For those particles with $|\mathbf{n}_i| > C_N$, a radiation term is also applied to simulate the heat transfer to the missing particles:

$$\frac{dT_i}{dt} = -T_i/D_r, \quad (16)$$

where D_r denotes the radiation half time, and lower value of D_r leads to faster cooling of the corresponding particles. Then the buoyancy of the gas particle is set to be proportional to the temperature, i.e.

$$\tilde{\mathbf{a}}_{b,i} = C_b T_i \mathbf{b}, \quad (17)$$

where $\tilde{\mathbf{a}}_{b,i}$ is the acceleration of the i -th particle due to the buoyancy force, C_b is the buoyancy coefficient, and \mathbf{b} is a unit vector pointing to the buoyancy direction.

In this work, the rigid boundaries are represented with boundary particles. The positions of boundary particles are not affected by gas particles, but they participate in all SPH calculations in the same way as normal simulation particles, as in [27].

4 Phase transition between liquid and gas

Using the proposed fast SPH scheme for gas simulation, it becomes possible to develop an efficient and unified SPH scheme for complex scenes that involve liquid–gas phase transition. The density of liquid is often 1000 times higher than the gas density. When the phase transition occurs, this high-density ratio causes a drop in simulation resolution and robustness and ultimately loss of visual plausibility. To prevent this loss of plausibility, this situation requires a proper particle splitting and merging strategy, which is explained in the following sections. We combine a pattern-based splitting scheme originated from [9] and a merging strategy based on local mass decentration [8] to achieve this goal, and a practical approach is described in the following subsections to make these strategies fit for phase-transition simulation.

4.1 Keeping simulation resolution in phase transition

Due to the large density ratio between liquid and gas phases, the effective volume of an SPH particle will be significantly




Fig. 2 2D examples of particle splitting. Based on a predefined pattern, e.g. the hexagon pattern shown in **a** and the square pattern shown in **b**, the mother particle (the red node) is split into a set of evenly distributed daughter particles (the blue nodes). The distance between the mother particle and the daughter particles is ϵ . The smoothing radius of the daughter particles is αh , where h denotes the smoothing radius of the mother particle

larger in the gas phase if it is directly vaporized from the liquid phase. This leads to degraded simulation resolution. To overcome this problem, we split the involved liquid particles before transferring them into the gas phase. To be truthful to the underlying physics, the change of local mass distribution due to this refinement process must be minimized, and the global mass and momentum conservation should be preserved.

As shown in Fig. 2, the mother particle (denoted by the red node) is split into a set of daughter particles (denoted by the blue nodes), which are evenly distributed around the original position. Depending on the simulation need (to be explained in Sect. 4.2), a daughter particle can also be placed at the center, i.e. the position of the mother particle. The distance from the mother particle to each daughter particle is denoted by a scalar parameter ϵ . As in [1, 9], we allow variable smoothing lengths in the simulation. Thus, another scalar parameter α is introduced to denote the change of smoothing radius, i.e. $h_{\text{new}} = \alpha h_{\text{origin}}$. Finally, the mass ratio between the daughter particles and the mother particle is denoted by a vector $\lambda = [\lambda_1, \dots, \lambda_M]^T$, where $i = 1, 2, \dots, M$ denotes the daughter particles and M denotes the number of daughter particles. Thus, the mass of each daughter particle is $m_i = \lambda_i m_N$, where N denotes the mother particle and m_N is the mass of the mother particle.

It is proved in [9] that, given the values α and ϵ , the refinement error can be written in terms of λ as

$$E[\alpha, \epsilon](\lambda) = \frac{m_N^2}{h_N^d} \left(\bar{C} - 2\lambda^T \bar{\mathbf{b}} + \lambda^T \bar{\mathbf{A}} \lambda \right), \quad (18)$$

where each $\bar{A}_{ij} = \frac{1}{\alpha^d} \int_{\Omega} W_i(\mathbf{r}, 1) W_j(\mathbf{r}, 1) d\mathbf{r}$, each $\bar{b}_j = \int_{\Omega} W_N(\mathbf{r}, 1) W_j(\mathbf{r}, 1) d\mathbf{r}$, $\bar{C} = \int_{\Omega} W_N^2(\mathbf{r}, 1) d\mathbf{r}$, and d is the simulation dimension. $W_i(\mathbf{r}, 1) = W(\mathbf{r} - \mathbf{r}_i, 1)$ denotes the

smoothing kernel function at the i -th particle position \mathbf{r}_i with the unit kernel radius.

Therefore, given the refinement parameters (ϵ, α) , the optimal values for λ^* can be obtained by minimizing the error $E[\alpha, \epsilon]$ with the constraint of $\sum_{j=1}^M \lambda_j = 1$. Note that in a given pattern, for each choice of (ϵ, α) , the coefficients \bar{C} , $\bar{\mathbf{b}}$, $\bar{\mathbf{A}}$ can be calculated numerically, which allows for the pre-computation of the best λ^* . This process can also be reversed, and the corresponding (ϵ, α) can be found from a desired λ^* pattern. For a certain pattern, at the precomputation stage we need to calculate corresponding values of (E, λ) from (ϵ, α) pairs and save them into a two-dimension chart. This provides the necessary information to determine better patterns that give a good balance between refinement error and mass distribution.

4.2 Boiling and evaporation

Liquid can transform into gas through boiling or vaporizing. The former is a vibrant process occurring at the boiling point, and the latter is a quiet process taking place slowly at lower temperatures. The pattern-based splitting scheme in Sect. 4.1 provides a natural way to model these two different phenomena. For boiling we choose a splitting pattern that does not place a daughter particle at the center. Patterns without a central node usually divide the mass of the original particle evenly into new particles, especially when the pattern is symmetric. Typically we use a regular icosahedron pattern with $(\epsilon, \alpha) = (0.3, 0.9)$, but other patterns can also be chosen when necessary. After splitting, the new particles are transferred into the gas phase, experiencing density drop and volume expansion in the boiling phenomena. The smoothing radius of these gas particles is also enlarged in this process, matching their expansion in effective volume. For the relatively slower vaporization, we choose a splitting pattern with a daughter particle in the center. In these types of patterns, the majority mass of the original particle remains in the new particle at the center. For example, in a cubic pattern with $(\epsilon, \alpha) = (0.4, 0.9)$ and with a relatively low refinement error, the mass of the new particle is $\lambda_i = 0.992$ for the particle at the center and is $\lambda_i = 0.001$ for others. After particle splitting, the center particle remains in the liquid phase and the other particles are transformed into the gaseous phase. The smoothing radius of these transformed particles is also re-adjusted.

The above patterns are chosen with balanced consideration. The cubic pattern, for example, provides a moderate phase transition rate between liquid and air, with only a small fraction of liquid mass turned into gas phase in each splitting, and keeps the effective volume of the gas particles similar to that of the liquid particles. For the boiling process the regular icosahedron pattern provides a good balance between sim-

ulation efficiency (i.e. avoiding particle number explosion) and resolution. These patterns also provide a relatively small refinement error under the above mass distribution. There may be better patterns; however, the symmetric and easy-to-compute property of these patterns make the calculation both in precomputation and actual simulation significantly easier and faster. The effect of refinement error related to different patterns and corresponding (ϵ, α) values and the effectiveness of this method comparing to original distribution are thoroughly analyzed and validated in [9].

The pattern-based splitting scheme described in [9] is mainly developed for one-time splitting during the simulation and should be further extended to be applicable in simulations of certain phenomena as evaporation. In evaporation, a liquid particle could be split several times. It is possible to pre-compute different splitting schemes with different (ϵ, α) values for each split, making the mass of the new gas particles similar to each other. However, after several splits, the mass of the liquid particle reduces and the splitting error discussed in Sect. 4.1 becomes gradually larger. To overcome this problem, we split liquid particles only up to a threshold number of splits. Once a liquid particle reaches this splitting threshold, we merge the particle with adjacent similar liquid particles using a merging strategy originating from [8].

Specifically, for each candidate particle M , we compute its inertia matrix as

$$\mathbf{I}_M = \begin{pmatrix} I_{xx} & -I_{xy} & -I_{xz} \\ -I_{xy} & I_{yy} & -I_{yz} \\ -I_{xz} & -I_{yz} & I_{zz} \end{pmatrix} \quad (19)$$

where $I_{xx} = \sum_j m_j ((y_j - y_M)^2 + (z_j - z_M)^2)$, $I_{yy} = \sum_j m_j ((x_j - x_M)^2 + (z_j - z_M)^2)$, $I_{zz} = \sum_j m_j ((x_j - x_M)^2 + (y_j - y_M)^2)$, $I_{xy} = \sum_j m_j (x_j - x_M)(y_j - y_M)$, $I_{yz} = \sum_j m_j (y_j - y_M)(z_j - z_M)$, and $I_{xz} = \sum_j m_j (x_j - x_M)(z_j - z_M)$. Here, the summation is performed over the neighborhood of the candidate particle M defined by its smoothing radius. Then, following [8], we define the mass decentrations as

$$\eta_M = \left| \left(\frac{\text{trace}(\mathbf{I}_M)}{3} \right)^3 - \det(\mathbf{I}_M) \right| \quad (20)$$

The smaller the η_M value, the mass distribution within the neighborhood of particle M is more similar to a uniform sphere. The mass decentrations η_M vanishes when the mass distribution around the candidate particle M is an ideal sphere. In our implementation, the particle with the smallest mass decentrations is merged with its neighborhood. The merged particle, denoted by M' , has the total mass of the entire neighborhood before merging, and it is placed at the same position as M . After merging, the merged particle M' is split again into several particles of proper size, following




Fig. 3 **a** The candidate particle M to be merged with its neighborhood. **b** A larger particle M' is generated after merging. **c** Then the large particle M' is split into a set of daughter particles with desired size and distribution

the splitting algorithm described in Sect. 4.1. As shown in Fig. 3, this merge–split process ensures the quality of SPH particles during merging and splitting operations.

It is noted that, right after the splitting or merging process, partially “overlapping” daughter particles can temporarily appear. However the new particle position, mass, etc. are carefully pre-computed so that particle interactions in the new distribution are maintained almost the same as before even in the case of partial overlapping. Thus stable simulation can be achieved without additional calculation like some other methods, e.g. Poisson disc sampling.

5 Implementation

In the standard SPH algorithm, the motion of a particle i is determined by

$$\rho_i = \sum_j m_j W_{ij} \quad (21)$$

$$\mathbf{a}_i = \frac{d\mathbf{v}_i}{dt} = \frac{1}{\rho_i} \mathbf{F}_{p,i} + \frac{1}{\rho_i} \mathbf{F}_{v,i} + \mathbf{g}, \quad (22)$$

where $\mathbf{F}_{p,i}$ denotes the pressure force, $\mathbf{F}_{v,i}$ the viscosity force, \mathbf{g} the gravity acceleration, \mathbf{a}_i the acceleration of particle i , \mathbf{v}_i the velocity of particle i , and t the time variable. Our approach first compensates for the density of each particle using Eq. (8). Then the local pressure balance and the atmospheric pressure are combined to form the pressure force term, i.e.

$$\mathbf{F}_{p,i} = \tilde{\mathbf{F}}_{p,i} + \tilde{\mathbf{F}}'_{p,i}, \quad (23)$$

where $\tilde{\mathbf{F}}_{p,i}$ and $\tilde{\mathbf{F}}'_{p,i}$ are defined in Eqs. (10) and (13).

The damping and buoyancy accelerations defined in Eqs. (14) and (17) are also added to the right-hand side of Eq. (22). The complete equation of our approach is

$$\mathbf{a}_i = \frac{1}{\rho_i} \left(\tilde{\mathbf{F}}_{p,i} + \tilde{\mathbf{F}}'_{p,i} \right) + \frac{1}{\rho_i} \mathbf{F}_{v,i} + \mathbf{g} + \tilde{\mathbf{a}}_{\text{damp},i} + \tilde{\mathbf{a}}_{b,i} \quad (24)$$

To obtain more visual variance of the simulated gas, it is possible to artificially adjust the atmospheric pressure $\tilde{\mathbf{F}}'_{p,i}$ according to the particle’s normal vector. Specifically, $\tilde{\mathbf{F}}'_{p,i}$ is multiplied by a constant factor $\beta \in [0, 1]$ if the norm of its normal vector is less than a given threshold $|\mathbf{n}_i| < C_N$, where near-zero β values will offer more visual variation.

If desired one can also add artificial vortex to the simulation, such as Lagrangian vortex methods adopted in [18, 24]. Specifically, the vortex carried by each particle is evolved by the following equation:

$$\frac{D\omega_i}{dt} = \omega_i \cdot \nabla \mathbf{v} + \beta (\mathbf{n}_i \times \mathbf{g}), \quad (25)$$

where ω_i is vortex carried by the i -th particle, and β is a user-defined constant used to control the overall strength of the effect. Then a particle is driven by an extra force due to the neighbors:

$$\mathbf{F}_{\text{vort},i} = \sum_j (\omega_j \times (\mathbf{r}_i - \mathbf{r}_j)) W_{ij} \quad (26)$$

The above vortex method adds more visual variance to the result with about 15 % increase in computational cost.

To deal with situations when a particle moves too far away from other particles and be isolated, we give them a life-strength value l_s and subtract $\Delta l_s + \epsilon$ from it each step it remains isolated, where ϵ is a small random number. If the particle becomes no longer isolated later, its life-strength value is restored. A particle is removed whenever its life-strength value is less than 0. Other calculations are not changed under such situation.

6 Results

The performance gain for the proposed approach will depend on the size of the simulation domain in that larger domains require more ambient air particles to fill the entire space. Our approach is most efficient dealing with scenes that have smoke plumes evolving across a large empty space with open boundaries, which we will demonstrate in the examples. For rendering, the fluid particles are directly rendered, which is convenient for liquid–gas phase transition, and it can easily cope with long-range movement of gaseous fluids in very large scenes. Post-processing rendering techniques, similar to those in [18], can also be adopted to enhance diffusion effects, which we use in the steam train example.

To demonstrate the effectiveness and speed up of our approach, a boiling example is presented in Fig. 4. At the center of the scene, the water is boiled in a small pan producing vapor. In Fig. 4a, using our approach described in Sects. 3 and 4, the vapor phase is successfully simulated, showing plausible movement and shape. The simulation uses only about




Fig. 4 Boiling. **a** Our approach. Using only about 9000 vapor particles in total, the simulation runs at 67.3 steps per second. **b** Standard approach with ambient air particles. 640,000 ambient air particles are required to fill the empty space, reducing the simulation speed to 7.4 steps per second. **c** Removing the correction and compensation from

a, the simulation speed is comparable with that in **a** at 84.6 steps per second; however, the vapor does not hold its shape, scattering, and becoming thin. **d** An implementation of the two-scale particle simulation method. Up to 305,000 particles are generated in the boundary region of the method. The simulation speed is 4.9 steps per second

9000 vapor particles in total, plus about 40,000 liquid particles, and with artificial vortex added the simulation runs at 67.3 steps per second on a Nvidia GeForce GTX 780 GPU. The rendering volume grid is $160 \times 120 \times 400$. The simulation can run stably at 2 steps per frame (step size $1/60$ s). For comparison, a vapor plume with the same resolution, but simulated using fully filled ambient air particles, is shown in Fig. 4b. In this case, no correction or compensation is needed, but this small, simple scene requires about 640,000 ambient air particles to fill the empty space, reducing the simulation speed to 7.4 steps per second. In Fig. 4c, we demonstrate what happens when we remove the correction and compensation from the case shown in Fig. 4a. The erroneous calculation causes the vapor phase to scatter and become thin. The simulation speed is comparable with (a) at 84.6 steps per second.

Previous SPH methods using adaptive+-sampling have mainly dealt with liquid simulation; it is not straightforward to extend previous methods to GPU-based gas simulation without adjustment. We tested the method of [26] on the above example. The vapor particles are used to define the active region, and the low-resolution simulation in their method is carried on by a simulation of half resolution. The result is shown in Fig. 4d. Their method still needs to fill

the entire simulation domain with ambient air particles, and a large amount of particles are also needed in the boundary region. In total, the low-resolution simulation contains about 88,000 particles (including liquid particles), and up to 305,000 “boundary particles” are added into the simulation around the vapor particles; the simulation speed is 4.9 steps per second. This speed is slower than the standard algorithm, possibly because the splitting and relaxing algorithm for the generation of “boundary particles” are not as well parallelizable on GPU as other parts of the implementation, and because the number of “boundary particles” needed does not give much performance improvement space in gas simulations.

Figure 5 shows an outdoor bath using a rendering volume grid of $260 \times 260 \times 400$, where the water is evaporating and producing a misty fog. The gas phase has a rest density that is $1/1000$ of the liquid phase. Filling this scene would need 9.9 million ambient air particles besides liquid particles. With our approach, up to 329,000 particles are used to simulate the vapor plus about 78,000 liquid particles, achieving approximately 25 times speed up. The simulation runs at 13.9 steps per second.

Figure 6 shows a steam train running across open terrain and passing through a tunnel in a light-yellow hill.




Fig. 5 Evaporation. A misty fog rises from an outdoor bath. The simulation includes about 407,000 particles in total, which is approximately 25 times faster than filling the scene using an estimated 9.9 million ambient air particles with standard approach




Fig. 6 Steam train running across an open terrain through a tunnel in a light-yellow hill. The smoke rising from its smokestack is simulated with up to 433,000 gas particles using our approach. If ambient air par-

ticles were used, this example would have required an estimated 119.3 million ambient air particles to fill the simulation space. Our approach achieves more than 200 times speed up on this scene



Fig. 7 Plumes rising from an ancient city towards a magic attractor located above the center building. Up to 1.3 million particles are used with our approach, saving computational cost on more than 600 million ambient air particles. More than 400 times speed up is achieved in this scene

We simulate the smoke rising from its smokestack in a $200 \times 420 \times 1920$ region measured by volume data grid size in rendering. If the standard approach were used, the example would have required an estimated 119.3 million ambient air particles to fill the simulation space. In this case the artificial vortex is added. With our approach only 433,000 gas particles are used in total, and the simulation is able to run at 17.0 steps per second, achieving more than 200 times speed up.

Figure 7 shows an ancient city, from which a number of plumes are rising into the sky towards a magic attractor located above the center building. The plumes then converge and mix around the magic attractor. Based on the proposed approach, up to 1.3 million particles are used in this example

and the simulation runs at 3.0 steps per second. The rendering volume grid of this example reaches $1320 \times 560 \times 1240$. If the standard approach were used, the example would have required more than 600 million ambient air particles in order to fill the simulation space. For such a large scene it would be too costly for normal GPUs to simulate ambient air particles, and our approach brings more than 400 times speed up to its simulation.

7 Conclusion

We have developed a novel fast SPH simulation approach for gaseous fluids. The proposed approach is able to completely

avoid simulating ambient air particles, bringing significant performance benefit (up to two orders of magnitude) to SPH gas simulation, especially for large open scenes. Liquid–gas phase transition phenomena, such as boiling and evaporation, can also be efficiently captured following the fast simulation approach.

Unlike grid-based solvers, the SPH schemes do not have built-in diffusion, and it is difficult to handle the diffusion effect using the proposed approach. Although it can be partially alleviated by rendering, this is an intrinsic limitation related to single-phase SPH gas simulation, and using multi-fluid models can be a possible research direction in the future. Removing ambient air particles also makes it difficult to analyze the momentum conservation property. Although little disobedience of momentum conservation is visually observed for the overall gas movement, the global momentum conservation is not accurately ensured in the proposed approach. Theoretically this could lead to error accumulation, making the simulation possibly biased from the natural movement in very long simulations. Though the interactions between particles due to pressure is reproduced in our approach, the immediate viscous interactions between simulated particles and ambient air particles are not fully recovered, losing vortical motions to some degree. This can be partially compensated by adding artificial vortex; however, additional compensation mechanisms can be desirable. In reality, in relatively narrow spaces sometimes it is possible for the unseen air to “carry over” interactions between separated parts of visible gas through propulsion, viscosity or other mechanisms. However, since these ambient air particles are removed in our approach, these indirect interactions are no longer captured; our approach is thus less effective for simulating narrow surroundings.

Acknowledgments The authors would like to thank the anonymous reviewers for their constructive comments, and Prof. Jianjun Zhang and Prof. Jian Chang for their suggestions. This work is supported by the National Basic Research Project of China (Project Number 2011CB302205), the Natural Science Foundation of China (Project Number 61120106007) and the National High Technology Research and Development Program of China (Project Number 2012AA011503). The authors would also like to acknowledge the financial support provided by the National Science Foundation (NSF IIS-1320644) and UNC Carolina Development Foundation, and Sêr Cymru National Research Network in Advanced Engineering and Materials.

Appendix A: Derivation of Eqs. (7–8)

This appendix shows the detailed derivations of Eqs. (7) and (8). Since densities in adjacent particles are assumed to change gradually in the gas simulations, we can further consider that for a particle i , the densities of its neighboring particles $\rho_{i,j}$ are corrected with the same proportion as itself and in a similar form as Eq. (6):

$$\rho_{i,j} = \frac{\bar{\rho}_{i,j}}{1 - V W_{ik}} \quad (27)$$

Here we want to determine the value of the scalar V , which for a given particle i does not vary with adjacent particle j . Using the form in Eq. (27), Eq. (4) can be rewritten as

$$0 = (1 - V W_{ik}) \sum_j \frac{m_j}{\bar{\rho}_j} \nabla W_{ij} + \frac{m_k}{\rho_i} \nabla W_{ik} \quad (28)$$

Utilizing V_0 in Eq. (5), the above equation can be further rewritten as

$$m_k = (1 - V W_{ik}) V_0 \rho_i \quad (29)$$

Note that it has been assumed in the first place that the final correction form should be in the form of Eq. (27), which is equivalent to requiring $m_k = V \rho_i$. Comparing this with Eq. (29), the following equation can be obtained:

$$(1 - V W_{ik}) V_0 = V \quad (30)$$

Solving Eq. (30) yields


$$V = \frac{V_0}{1 + V_0 W_{ik}} \quad (31)$$

Substituting Eq. (31) into Eqs. (27) and (29) leads to Eqs. (7–8).

References

1. Adams, B., Pauly, M., Keiser, R., Guibas, L.J.: Adaptively sampled particle fluids. In: ACM Transactions on Graphics (TOG), vol. 26, p. 48. ACM (2007)
2. Angelidis, A., Neyret, F., Singh, K., Nowrouzezahrai, D.: A controllable, fast and stable basis for vortex based smoke simulation. In: Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on computer animation, SCA '06, pp. 25–32. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2006). <http://dl.acm.org/citation.cfm?id=1218064.1218068>
3. Bender, J., Erleben, K., Teschner, M., et al.: Boundary handling and adaptive time-stepping for pcisph. In: Workshop on virtual reality interaction and physical simulation VRIPHY (2010)
4. Boyd, L., Bridson, R.: Multiflip for energetic two-phase fluid simulation. ACM Trans. Graph. **31**(2), 16 (2012)
5. Brochu, T., Keeler, T., Bridson, R.: Linear-time smoke animation with vortex sheet meshes. In: Lee, J., Kry, P.G. (eds.) Symposium on Computer Animation, pp. 87–95. Eurographics Association (2012)
6. Cleary, P.W.: Modelling confined multi-material heat and mass flows using SPH. Appl. Math. Model. **22**(12), 981–993 (1998). doi:[10.1016/S0307-904X\(98\)10031-8](https://doi.org/10.1016/S0307-904X(98)10031-8)
7. Colagrossi, A., Landrini, M.: Numerical simulation of interfacial flows by smoothed particle hydrodynamics. J. Comput. Phys. **191**(2), 448–475 (2003)
8. Desbrun, M., Cani, M.P.: Space-time adaptive simulation of highly deformable substances. Rapport de recherche RR-3829, INRIA (1999). <http://hal.inria.fr/inria-00072829>

9. Feldman, J., Bonet, J.: Dynamic refinement and boundary contact forces in sph with applications in fluid flow problems. *Int. J. Numer. Methods Eng.* **72**(3), 295–324 (2007)
10. Gao, Y., Li, C.F., Hu, S.M., Barsky, B.A.: Simulating gaseous fluids with low and high speeds. *Comput. Graph. Forum* **28**(7), 1845–1852 (2009). doi:[10.1111/j.1467-8659.2009.01562.x](https://doi.org/10.1111/j.1467-8659.2009.01562.x)
11. Gao, Y., Li, C.F., Ren, B., Hu, S.M.: View-dependent multiscale fluid simulation. *Vis. Comput. Graph. IEEE Trans.* **19**(2), 178–188 (2013). doi:[10.1109/TVCG.2012.117](https://doi.org/10.1109/TVCG.2012.117)
12. Gerszewski, D., Bargteil, A.W.: Physics-based animation of large-scale splashing liquids. *ACM Trans. Graph.* **32**(6), 185:1–185:6 (2013)
13. Hong, W., House, D.H., Keyser, J.: Adaptive particles for incompressible fluid simulation. *Vis. Comput.* **24**(7–9), 535–543 (2008)
14. Ihmsen, M., Cornelis, J., Solenthaler, B., Horvath, C., Teschner, M.: Implicit incompressible sph. *IEEE Trans. Vis. Comput. Graph.* **20**(3), 426–435 (2014). doi:[10.1109/TVCG.2013.105](https://doi.org/10.1109/TVCG.2013.105)
15. Keiser, R., Adams, B., Dutre, P., Guibas, L.J., Pauly, M.: Multiresolution particle-based fluids. ETH Zurich, NAKinci (2006)
16. Kim, S.T., Hong, J.M.: Visual simulation of turbulent fluids using mls interpolation profiles. *Vis. Comput.* **29**(12), 1293–1302 (2013)
17. Li, X., Liu, L., Wu, W., Liu, X., Wu, E.: Dynamic bfcc characteristic mapping method for fluid simulations. *Vis. Comput.* **30**(6–8), 787–796 (2014). doi:[10.1007/s00371-014-0969-7](https://doi.org/10.1007/s00371-014-0969-7)
18. Macklin, M., Müller, M., Chentanez, N., Kim, T.Y.: Unified particle physics for real-time applications. *ACM Trans. Graph. (TOG)* **33**(4), 104 (2014)
19. Monaghan, J.J.: Simulating free surface flows with sph. *J. Comput. Phys.* **110**(2), 399–406 (1994)
20. Müller, M., Charypar, D., Gross, M.: Particle-based fluid simulation for interactive applications. In: *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on computer animation, SCA '03*, pp. 154–159. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2003). <http://dl.acm.org/citation.cfm?id=846276.846298>
21. Müller, M., Schirm, S., Teschner, M., Heidelberger, B., Gross, M.: Interaction of fluids with deformable solids: research articles. *Comput. Animat. Virtual Worlds* **15**(3–4), 159–171 (2004). doi:[10.1002/cav.v15:3/4](https://doi.org/10.1002/cav.v15:3/4)
22. Müller, M., Solenthaler, B., Keiser, R., Gross, M.: Particle-based fluid-fluid interaction. In: *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on computer animation, SCA '05*, pp. 237–244. ACM, New York, NY, USA (2005). doi:[10.1145/1073368.1073402](https://doi.org/10.1145/1073368.1073402)
23. Orthmann, J., Kolb, A.: Temporal blending for adaptive sph. In: *Computer Graphics Forum*, vol. 31, pp. 2436–2449. Wiley Online Library (2012)
24. Pfaff, T., Thuerey, N., Gross, M.: Lagrangian vortex sheets for animating fluids. *ACM Trans. Graph.* **31**(4), 112:1–112:8 (2012). doi:[10.1145/2185520.2185608](https://doi.org/10.1145/2185520.2185608)
25. Schechter, H., Bridson, R.: Ghost sph for animating water. *ACM Trans. Graph. (TOG)* **31**(4), 61 (2012)
26. Solenthaler, B., Gross, M.: Two-scale particle simulation. In: *ACM Transactions on Graphics (TOG)*, vol. 30, p. 81. ACM (2011)
27. Solenthaler, B., Pajarola, R.: Density contrast sph interfaces. In: *Proceedings of the 2008 ACM SIGGRAPH/Eurographics symposium on computer animation, SCA '08*, pp. 211–218. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2008). <http://dl.acm.org/citation.cfm?id=1632592.1632623>
28. Yan, H., Wang, Z., He, J., Chen, X., Wang, C., Peng, Q.: Real-time fluid simulation with adaptive sph. *Comput. Animat. Virtual Worlds* **20**(2–3), 417–426 (2009)
29. Zhu, B., Yang, X., Fan, Y.: Creating and preserving vortical details in sph fluid. *Comput. Graph. Forum* **29**(7), 2207–2214 (2010)



Bo Ren is currently a Ph.D. candidate in the Department of Computer Science and Technology, Tsinghua University. His research interest is in physical simulation and rendering.



Xiao Yan is currently a Ph.D. student in the Department of Computer Science and Technology, Tsinghua University. His research interests include physically based simulation and rendering.



Tao Yang is now a Ph.D. candidate of the Department of Science & Technology in Tsinghua University. His research interest is computer graphics.



Chen-feng Li is currently an Associate Professor at the College of Engineering of Swansea University, UK. He received his BEng and MSc degrees from Tsinghua University, China in 1999 and 2002, respectively, and received his Ph.D. degree from Swansea University in 2006. His research interests include computational solid mechanics, computational fluid dynamics, and computational graphics. He is on the editorial board of *Engineering Computations* (Emerald).



Ming C. Lin received the BS, MS, and Ph.D. degrees in electrical engineering and computer science from the University of California, Berkeley. She is currently John R. & Louise S. Parker Distinguished Professor of computer science at the University of North Carolina at Chapel Hill and a distinguished visiting professor at Tsinghua University, China. Her research interests include physically based modeling, haptics, real-time 3D graphics for virtual environments,

robotics, and geometric computing. She has (co)authored more than 240 refereed publications, coedited/authored four books. She has served on more than 100 program committees of leading conferences on virtual reality, computer graphics, robotics, haptics, and computational geometry, and co-chaired more than 25 international conferences and workshops. She is the editor-in-chief of the *IEEE Transactions on Visualization and Computer Graphics* and an editorial board member of several scientific journals. She has also served on steering committees and advisory boards of international conferences, as well as technical advisory committees for government agencies and industry. She has received several honors and awards, including IEEE VGTC VR Technical Achievement Award 2010 and several best paper awards at international conferences on computer graphics and virtual reality. She is a fellow of the ACM and the IEEE.



Shi-min Hu received his Ph.D. degree from Zhejiang University in 1996. He is currently a professor in the Department of Computer Science and Technology, Tsinghua University, Beijing, and a professor at the School of Computer Science and Informatics, Cardiff University, United Kingdom. His research interests include digital geometry processing, video processing, rendering, computer animation, and computer-aided geometric design. He is an associate editor-

in-chief of *The Visual Computer* (Springer), and on the editorial boards of *Computer-Aided Design* (Elsevier) and *Computer & Graphics* (Elsevier). He is a member of the IEEE. He is the corresponding author of this paper.