

Volcanic Skyscapes

P. Cilliers Pretorius

Supervisor: James Gain



Minor dissertation presented in partial fulfillment
of the requirements for the degree of
MASTER OF SCIENCE
in the Department of Computer Science,
Science Faculty, University of Cape Town.

22 July 2024

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

12 Abstract

13 Explosive volcanic eruptions are some of the most visually impressive natural phe-
14 nomena, and are frequently represented in visual media such as films or games,
15 usually with a bespoke, one-off solution. Existing plume models treat the atmo-
16 sphere as a static boundary, while existing atmosphere models are not designed for
17 the extremes present in a plume. We present the first general-purpose model with
18 bi-directional interaction between a volcano plume and the atmosphere. The atmo-
19 sphere is modelled as sparse layers of 2D Eulerian grids, while the plume dynamics
20 are modelled with Lagrangian disks and spheres. We enable the transfer of physi-
21 cal quantities such as ash, moisture, temperature, and wind velocity between these
22 submodels. Volumetric animations are generated through procedural upsampling
23 using noise functions, specifically tuned to the physical processes in play, to create
24 a complete volcanic skyscape. Our model depicts most of the relevant facets arising
25 from plume and atmosphere interactions, such as ash rain, embedded bell, cap, and
26 skirt clouds, shockwave effects, and volcanic lightning.

27 Plagiarism Declaration

28

29 I know the meaning of plagiarism and declare that all of the work in the disser-
30 tation, save for that which is properly acknowledged, is my own.

31

Pieter Cilliers Pretorius
22 July 2024

32 **Acknowledgements**

33 I would like to firstly thank my supervisor, Professor James Gain, for his advice,
34 patience, and support throughout the entire process. Without his advice and insight,
35 this work would not be anywhere near its current standard. I would also like to
36 thank the many friends and family who supported me over the many long nights
37 and weekends, serving as soundboards and providing assistance with comprehending
38 the complexities of fluid dynamics, Sané, Leighton, and Alec in particular. This has
39 been a long journey which would not have been completed without the support from
40 too many people to name.

41 Contents

42	1 Introduction	1
43	2 Background	4
44	2.1 Movement in the atmosphere	4
45	2.2 Cloud formation	7
46	2.2.1 Water cycle	7
47	2.2.2 Precipitation	8
48	2.2.3 Cloud types	9
49	2.3 Explosive volcanic eruptions	9
50	2.3.1 Volcanic weather phenomena	11
51	2.4 Fluid dynamics of the atmosphere	15
52	2.4.1 Dynamics of incompressible fluids	15
53	2.4.2 Navier-Stokes	17
54	2.4.3 Advection	18
55	2.4.4 Semi-Lagrangian advection	19
56	2.4.5 Correcting for divergence	19
57	2.4.6 Boundary conditions	22
58	3 Related work	24
59	3.1 Cloud modelling and animation	25
60	3.1.1 Procedural	25
61	3.1.2 Simulation	27
62	3.1.3 Data-driven	28
63	3.1.4 Hybrid models	29
64	3.1.5 Combined models	29
65	3.1.6 Synoptic scales	31
66	3.2 Volcanic plume	31
67	3.2.1 Geophysical simulations	33
68	3.3 Summary	34
69	4 Weather modelling	36
70	4.1 Model structure	37
71	4.1.1 Horizontal air layers	37
72	4.1.2 Vertical exchange layers	40
73	4.1.3 Terrain model	40
74	4.2 Physics within the atmosphere	42
75	4.2.1 Terrain heating	42
76	4.2.2 Atmospheric heating	43
77	4.2.3 Advection	44

78	4.2.4	Terrain friction	44
79	4.2.5	Pressure correction	45
80	4.2.6	Phase transition	45
81	4.2.7	Moisture uptake	47
82	4.3	Interlayer communication	48
83	4.3.1	Precipitation	48
84	4.3.2	Layer exchange	49
85	5	Plume Modelling	53
86	5.1	Plume model	53
87	5.1.1	Plume slices	53
88	5.1.2	Plume spheres	54
89	5.2	Sphere Enhancements	55
90	5.3	Lightning	55
91	5.4	Embedded Clouds	56
92	5.5	Transient condensation clouds	58
93	6	Combining submodels and amplification	61
94	6.1	Model combination	61
95	6.2	Procedural amplification	63
96	6.2.1	Cloud amplification	65
97	6.2.2	Ash rain	66
98	6.2.3	Plume	66
99	6.2.4	Lightning	67
100	7	Results	68
101	7.1	Validation	70
102	7.1.1	Ash rain	71
103	7.1.2	Lightning	71
104	7.1.3	Windfield ablation	74
105	7.2	Performance	74
106	7.3	Limitations	78
107	8	Conclusion	80
108		Appendices	94
109	A	2D pressure solve when $\Delta x \neq \Delta y$	95
110	B	Derivation of solid interface in pressure solve	96
111	C	Example calculation of layer heights	97

112 List of Figures

113	1.1	Combining the atmosphere and plume submodels with an offline rendering system.	3
114			
115	2.1	Average atmospheric temperature as a function of altitude.	6
116	2.2	Cloud classification according to height and convective activity.	9
117	2.3	Phases present in a typical Plinian eruption.	10
118	2.4	Transient condensation cloud formed by US Navy test shot.	12
119	2.5	Mount Tavurvur transient condensation clouds.	13
120	2.6	Hunga Tonga-Hunga Ha'apai skirt clouds.	14
121	2.7	Sarychev Peak pileus cloud.	14
122	2.8	Shinmoedake ash rain.	15
123	2.9	Eyjafjallajökull lightning around volcanic plume.	16
124	2.10	Comparing Lagrangian and Eulerian approaches.	19
125	2.11	Semi-Lagrangian advection on a given grid point.	20
126	2.12	A Marker-and-Cell grid representation.	21
127	3.1	Examples from previous works on cloud modelling and animation.	26
128	3.2	Example renders from two combined models.	30
129	3.3	Examples of volcanic plume animation in previous work.	33
130	4.1	Simplified structure of our atmospheric model.	37
131	4.2	Different weather components and the quantities each affect.	38
132	5.1	Plume model divided into slices and spheres.	54
133	5.2	Skirt uplift dependent on distance from plume.	58
134	5.3	Various types of embedded clouds in volcanic plumes.	59
135	5.4	Simulating a pressure wave and subsequent condensation cloud.	59
136	6.1	Combining the atmosphere and plume submodels with an offline rendering system.	62
137			
138	6.2	Samples of different noise functions.	65
139	7.1	Transient condensation cloud timelapse and air layer's moisture map.	69
140	7.2	Timelapse of skirt clouds deformed and uplifted by plume.	69
141	7.3	Timelapse of cap cloud being pushed ahead of and then through by plume.	69
142			
143	7.4	Timelapse of ash rain falling from the plume in umbrella phase.	70
144	7.5	Volcanic lightning.	70
145	7.6	Qualitative validation of ash precipitation in atmosphere.	72
146	7.7	Ash deposition depth across transects of the terrain.	73

147	7.8	Ash deposition depth of real volcanic eruptions as a function of isopach	
148		area.	73
149	7.9	Incremental and accumulated electrical charge in plume.	74
150	7.10	Comparing a uniform wind field to an advected wind field.	75
151	7.11	Computation time per time step as a function of simulation time. . .	78

152 List of Tables

153	2.1	Symbols used in Navier-Stokes equations.	17
154	3.1	Cloud generation systems classified according to scale and method.	25
155	4.1	List of components in each horizontal air layer.	39
156	4.2	List of components of each vertical exchange layer.	40
157	4.3	List of quantities used in the planet model.	41
158	4.4	List of quantities stored for each cell of the terrain.	41
159	4.5	List of quantities present in each terrain type.	41
160	4.6	List of quantities advected by each horizontal layer.	44
161	4.7	Scaling factors used for the precipitation modes.	49
162	5.1	Cap values for calculating charge accumulation for each sphere.	57
163	7.1	Performance for complete system across three scenarios.	75
164	7.2	Performance cost for each component of atmosphere system	76
165	7.3	Impact of layer count and resolution on performance.	76
166	7.4	Total computation time for each scenario.	77
167	7.5	Performance for each sub-model across three scenarios.	77

Chapter 1

Introduction

“The sea seemed to roll back upon itself, driven from its banks by the convulsive motion of the earth. On the other side, a black and dreadful cloud, broken with rapid, zigzag flashes, revealed behind it variously shaped masses of flame: these last were like sheet-lightning, but much larger. Soon, the cloud began to descend and cover the sea.” Such were the words of awed horror, written by Pliny the Younger to the Roman historian Tacitus, describing the eruption of Mount Vesuvius in 79 CE.

Volcanic eruptions have captured the human imagination, across history and across cultures. As much as 10% of the world’s population live in close proximity to active or recently active volcanoes[SN01]. Volcanoes can dramatically alter the landscape and ecosystems surrounding them, or even affect global weather patterns for months or years. Given this, it is perhaps no surprise that volcanic eruptions are often depicted in digital media, such as films or games. Beyond the visually spectacular lava flows, the volcanic plume is the core component of any scene including an erupting volcano. Visualising the plume is also of significant importance in both public education about volcanoes and risk forecasting, given the population potentially at risk in volcanic eruptions.

However, depicting a plume in digital media is usually done through an expensive, bespoke, one-off solution. The greatest challenge in depicting these plumes lie in the complicated interactions between the plume, the atmosphere, the terrain, and the effusive ejecta, such as lava flows and lahars (a violent slurry of pyroclasts, mud, and debris moving at high speeds). These require very expensive simulations to accurately model, which, while prevalent in the geophysical sciences community, are not suitable for graphical applications.

Existing graphical simulations of the plume[MDN02b; Miz+03; Las+22] focus on the plume dynamics and treat the atmosphere as a static boundary condition. For instance, Lastic et al.[Las+22] assumes a constant wind velocity at a given altitude. While this method can result in visually appealing plumes, it is not physically accurate with regards to the influence of the atmosphere’s advective and convective forces on the plume. Similarly, it ignores the effect that the plume has on the atmosphere through the transfer of heat, moisture, and ash.

While graphical simulations of the plume are limited in number, cloud modelling

205 and animation in computer graphics is a mature and rich field of research. Despite
206 covering a wide variety of cloud types[Vim+20] and energy dynamics[Häd+20], these
207 are not capable of accurately modelling the plume due to accuracy constraints re-
208 sulting from modelling typical atmospheric conditions. Volcanic eruptions present
209 extremes in temperature at the vent and in the plume, while also extending high
210 above the usual altitudes allowed for by these atmospheric models. They also fail to
211 take any ash content in the atmosphere into account, which is crucial for key plume
212 features such as ash rain. Given the extremely high energy dynamics present in a
213 volcanic eruption, the time steps employed by atmospheric models are usually too
214 large to capture the fine detail arising from extreme turbulence in the plume.

215
216 In this dissertation, we focus on specific phenomena arising from the bi-directional
217 interaction of the plume and atmosphere. We focus on explosive eruptions since
218 these eruptions have the most visually striking phenomena. In particular, cold at-
219 mospheric air entrained (taken up and mixed into the plume) into the plume can
220 give rise to embedded clouds, while ash, moisture, and ice within the plume lead to
221 lightning strikes. Shockwaves arising from the initial explosion can lead to transient
222 condensation clouds forming if atmospheric conditions are suitable. Ash sedimenting
223 out of the plume and subsequently mixing with moisture in the atmosphere leads
224 to the formation of ash rain that precipitates downwind of the plume. This ash
225 rain often has the most prominent effect on civilisation, disrupting transport and
226 damaging buildings and infrastructure on large scales. At the same time, it greatly
227 enriches the soil, leading to very fertile terrain for agriculture.

228
229 We aim to simulate the phenomena described above with a model combining two
230 sub-models. To address the specific challenges that arise from plume and atmosphere
231 interaction, we combine a Lagrangian model of plume dynamics[Las+22] with an at-
232 mospheric model using Eulerian advection and procedural convection[Vim+20]. We
233 find it necessary to extend these models to handle extreme temperatures and alti-
234 tudes, as well as to enable a variety of phenomena that were previously unsupported.

235
236 The plume is approximated as a stack of cylinders, with movement determined
237 by Lagrangian dynamics arising from ejection and air entrainment forces, as well
238 as temperature, moisture, and wind velocity provided by the atmospheric model.
239 Lagrangian spheres of various sizes are created on the cylinders to approximate a
240 coarse shape for the plume, while also transferring temperature, moisture, and ash
241 to the atmospheric model.

242
243 The atmospheric model is structured as sparse two-dimensional layers, with each
244 layer being a Eulerian grid with pertinent quantities such as temperature, mois-
245 ture, water, ash, and wind velocity. Horizontal wind is simulated through a semi-
246 Lagrangian advection scheme, while vertical uplift due to convective forces and oro-
247 graphic effects is modelled pseudo-physically for computational efficiency. We utilise
248 packet transport calculated by the moist static energy gradient. The moist static
249 energy is an approximation of the total energy in an air packet, including heat en-
250 ergy, potential energy due to gravity, and kinetic energy. Transport will occur from
251 packets of high energy to packets with lower energy. We calculate a pressure ad-
252 justment resulting from interactions with the terrain. Heat, moisture, and ash is

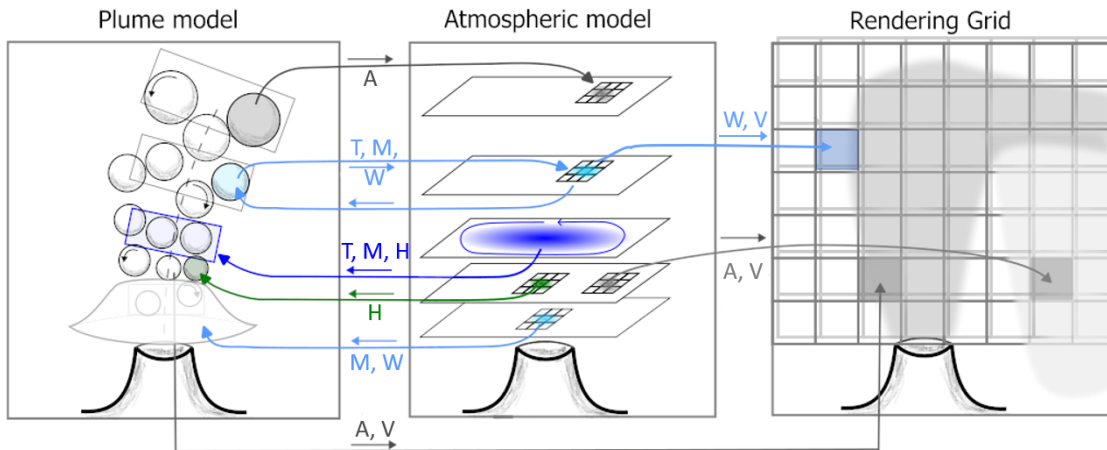


Figure 1.1: We transfer physical quantities between the plume and atmospheric submodels, before outputting the results to the final rendering system. We transfer ash A , moisture density M , water content W , temperature T , 3D velocity V , and horizontal 2D velocity H .

253 transferred from the plume to the atmospheric model and is used to induce cloud
 254 formation, convection, and ash rain once the plume has developed sufficiently.

255

256 The output of this coupled simulation is fed to a procedural upsampling pro-
 257 cess. This results in a finely-sampled three-dimensional grid for the atmosphere,
 258 and a collection of ellipsoids with variable density for the plume. This can then be
 259 used for offline volumetric rendering, or exported for use as part of a scene in films,
 260 games, or educational programs. By utilising an upsampling approach, we generate
 261 a plausible, fine-scale animation while only requiring a relatively coarse simulation,
 262 resulting in computational efficiency gains. The coarse simulation does result in the
 263 loss of some of the complex turbulence, which is then replaced by the output from
 264 procedural noise functions. Figure 1.1 provides an overview of the system and the
 265 quantities transferred between the atmospheric and plume submodels.

266

267 In this dissertation, we present our framework, which couples efficient and in-
 268 teracting sub-models for the volcanic plume and atmosphere, and combines it with
 269 procedural amplification. Our model enables the generation of several phenomena
 270 previously unavailable in computer graphics, such as ash rain, embedded clouds in
 271 the plume, and volcanic lightning. Chapter 2 provides the background for atmo-
 272 spheric physics and plume dynamics that give rise to the phenomena we model.
 273 Chapter 3 provides an overview of previous work in this field, as well as related
 274 work in other fields. The atmospheric sub-model is described in chapter 4, while the
 275 plume dynamics are described in chapter 5. The coupling scheme between the two
 276 sub-models and procedural amplification processes are described in chapter 6. We
 277 present our results, performance timings, and validation in chapter 7.

278 Chapter 2

279 Background

280 In this chapter, we provide a primer on the movement of air within the atmosphere,
281 its causes and the factors that influence it, and an overview of the water cycle and
282 the resulting cloud formation. We restrict our overview to the troposphere, which
283 contains more than 75% of the atmosphere's total mass and nearly all of its water.
284 We also provide a primer on explosive volcanic eruptions, the dynamics present in
285 the volcanic plume, and factors that influence the appearance of the plume.

287 2.1 Movement in the atmosphere

288 Our work is focused on a limited domain, both in space and in time. Explosive
289 volcanic eruptions are comparatively short-lived, rarely lasting more than 24 hours.
290 Similarly, while these eruptions can have effects lasting several years (such as the
291 eruption of Mount Tambora in 1815), the visually salient features rarely last more
292 than a few hours after the eruption comes to an end. Therefore, we will cover only
293 the salient points of weather formation for our scenario. For more detail on weather
294 processes as a whole, we refer the reader to works by Barry et al.[BC09; Bar13] and
295 Pielke et al.[Pie+98]. While the same processes govern movement in all atmospheres
296 regardless of composition, we will restrict this work to Earth-like atmospheres with
297 a water cycle (see section 2.2).

298
299 To describe movement in an atmosphere, an understanding of its composition
300 and nature is required. An atmosphere consists of a mixture of gases bound to a
301 planet's surface by the planet's gravity. The atmosphere can carry material other
302 than gases, such as liquid water (water clouds) or particulates (dust or ash clouds).

303
304 All movement in the atmosphere (collectively referred to as wind) is a result of
305 a pressure-gradient force. This is the force applied to air molecules as a result of a
306 difference in pressure within a volume, with the force pushing from an area of high
307 pressure to an area of low pressure. As per the general gas equation:

$$pV = nR^*T, \quad (2.1)$$

308 pressure p is proportional to the amount of substance n and temperature T , and
309 inversely proportional to volume V , where R^* is the universal gas constant. If we
310 discretise the atmosphere, then each parcel or packet of air under consideration will

311 have a volume that remains constant. Thus, the only differences that can lead to
 312 a pressure differential are differences in the temperature T , or in the amount of
 313 substance n .

314

315 The atmosphere on Earth (at speeds below Mach 1) behaves as an incompressible
 316 fluid. Thus, modelling the movement of air is done through the equations for
 317 fluid dynamics. It also means that the shape of the underlying surface can have a
 318 significant effect on the wind field. In our case, the underlying surface is the terrain.
 319 See section 2.4 for a detailed explanation of the fluid dynamics required for our
 320 model.

321

322 Due to the effect of gravity, the pressure at a given altitude remains relatively
 323 constant, with only minor variations. This is expressed by the barometric equation:

$$P_h = P_b \left(1 - \frac{L}{T_b} h \right)^{\frac{gM}{R^*L}}, \quad (2.2)$$

324

325 where P_b refers to the base pressure at sea level (which is 101325 pascal on
 326 Earth), L refers to the lapse rate, the change in temperature as a function of alti-
 327 tude ($6.5 \text{ K} \cdot \text{km}^{-1}$ on Earth), T is the average temperature at sea level (288.15 K
 328 on Earth), and h is the altitude in question. In the exponent, g is the gravitational
 329 acceleration ($9.8 \text{ m} \cdot \text{s}^{-2}$ on Earth), M is the molar mass of air ($0.029 \text{ kg} \cdot \text{mole}^{-1}$
 330 on Earth), R^* is the universal gas constant, and L is once again the lapse rate.
 331 However, there are processes that can create sufficiently steep horizontal pressure
 332 gradients for weather phenomena to arise despite the effects of gravity.

332

333 The density of an air parcel can be approximated as $\rho = n/V$, following the
 334 general gas equation (equation 2.1). It follows that an air parcel with a higher tem-
 335 perature will be less dense compared to a parcel with a lower temperature. Since
 336 the atmosphere is in the planet's gravitational field, buoyancy forces act on parcels
 337 of differing densities. Therefore, warm air will tend to rise above cold air.

338

339 The majority of an atmosphere's energy budget comes from the photons emitted
 340 by a star. Depending on the wavelength of the photons and atmospheric composi-
 341 tion, not all of the energy is directly absorbed by the atmosphere. In the case of
 342 Earth's atmosphere, the vast majority of the photons are in wavelengths that the
 343 atmosphere is transparent to. Some of the photons are reflected back out to space
 344 (depending on the albedo of the surface). The rest of the photons are absorbed by
 345 the planet surface, where the photon's energy adds to the temperature of the surface.

346

347 The surface radiates this energy as greybody radiation peaking in the infrared,
 348 and the atmosphere is opaque to infrared wavelengths. Thus, some of the radi-
 349 ated energy is absorbed by the lower layers of the atmosphere. Some energy is also
 350 exchanged through direct contact and conductance from the surface and the air
 351 touching it, but since air is a terrible conductor and very low density, this conduc-
 352 tance is negligible compared to the radiative heat transfer.

353

354 As a consequence of the heating mechanisms of the atmosphere, warm air will
 355 form at the surface and then rise due to buoyancy. Such a convective column will

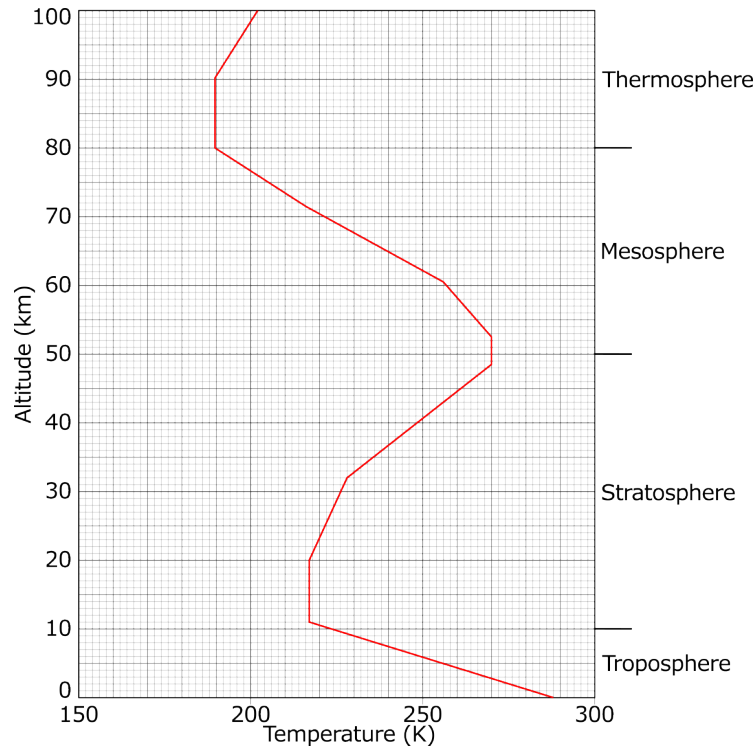


Figure 2.1: Average atmospheric temperature as a function of altitude, with temperature on the x-axis and altitude on the y-axis[OA76]. Note the lapse rate tending to 0 between the altitudes of 10 and 20 km.

356 tend to have a low pressure zone at its base relative to the average pressure at that
 357 altitude. Similarly, the top of the convective column will tend to have a high pres-
 358 sure zone relative to the average pressure at that altitude. This convective cycle
 359 creates pressure gradients in the horizontal and thereby drives wind currents.

360

361 Another facet of convective columns to take into account is the *tropopause*. As
 362 a hot air parcel rises, it exchanges some of its thermal energy for potential energy,
 363 as a result of adiabatic cooling. This is cooling as a result of a drop in pressure
 364 in the parcel as it rises. As per equation 2.1, if pressure reduces, but volume and
 365 amount of substance remain constant, then temperature must reduce as well. At
 366 some point, the air parcel temperature matches the temperature of the surrounding
 367 air, and therefore reaches the same density. At this point, buoyancy effects no longer
 368 apply and the air parcel will not rise any further.

369

370 The average lapse rate in the atmosphere is sharply negative for the lower parts
 371 of the atmosphere, averaging around 6.5 K per kilometre of altitude. This changes
 372 at around 11 kilometres of altitude, where the lapse rate suddenly transitions to zero.
 373 See figure 2.1. This boundary is known as the tropopause, and weather phenomena
 374 rarely cross it. Explosive volcanic eruptions and the subsequent plumes are one of
 375 the few natural phenomena that can penetrate the stratosphere (the layer above the
 376 tropopause). Very energetic convective columns that produce cumulonimbus clouds
 377 can, in exceptional cases, do this as well. Barring those exceptions, convective
 378 columns will tend to terminate near the tropopause, restricting surface weather
 379 effects to the lower 11 kilometres of the atmosphere.

2.2 Cloud formation

Water clouds (often referred to as simply clouds) are collections of liquid water (droplets) or solid ice suspended in the atmosphere (technically referred to as an *aerosol*) in sufficient density that the majority of light passing through the area is scattered such that the collection becomes visible. Clouds can consist of material other than water, such as other liquids like carbon dioxide or methane or particulates like dust or ash. In this work, we only consider clouds consisting of liquid water and ash.

Water vapour (gaseous H_2O) makes up a small, but highly variable portion of the Earth's atmosphere[RLB19]. Typically, it makes up around 1% of the atmosphere at sea level, though it varies from negligible to 3% of the atmosphere. Higher altitudes and latitudes will generally have lower concentrations of water vapour, since the amount of water vapour that a given parcel of air can hold is dependent on the temperature of the air. Higher temperatures mean a higher capacity for water vapour. Some of the water vapour will condense out into its liquid state. Similarly, if there is liquid water in the atmosphere, some of the liquid water may evaporate to its gaseous state.

The rate at which the two states of water change depends on the *saturation vapour pressure* for the parcel. This is the point where the quantities of liquid water and gaseous water vapour are such that there are equal numbers of molecules changing to and from each state. Specifically, the water vapour will tend to be in thermodynamic equilibrium with the liquid state. The saturation vapour pressure increases with temperature, which is why air at a higher temperature can hold more water vapour. For reference, the saturation vapour pressure is equal to one bar (101325 Pascal) at a temperature of 100 degrees Celsius. See section 4.2.6 for a more detailed look at this phenomenon and its implementation in this work.

2.2.1 Water cycle

Water vapour (hereafter referred to as moisture) enters the atmosphere by evaporating from the Earth's surface in some fashion. This can be from open water expanses like oceans, lakes, and rivers, or it can be from moist soil that dries out, or from transpiration from plants. In this work, we focus only on the abiotic sources, though for a fully realistic simulation, the biotic sources and influences cannot be ignored.

Once the moisture is in the atmosphere, it changes state to remain in thermodynamic equilibrium depending on the temperature of the air. Air is heated largely through infrared radiation from the ground, and is cooled largely through radiation to space, conducting to colder surfaces (such as land at night), and adiabatic cooling, that is, cooling due to a drop in pressure (see equation 2.1). This pressure drop is a result of the air parcel rising higher, whether because of convection or being forced up by a collision with the terrain (known as orographic uplift). The air molecules exchange kinetic energy, i.e., heat, for potential energy in the gravitational field. If moisture and liquid water changes state, that also has an effect on the temperature, but it is less influential than adiabatic cooling. See section 4.2.6 for an examination

425 of the temperature changes that result from changing states.

426

427 Moisture is indirectly removed from the atmosphere and returned to the surface
428 through precipitation which removes liquid water from the atmosphere, and thus
429 moisture will condense out to liquid water to restore the thermodynamic equilibrium.
430 Moisture cannot condense out without some sort of condensation nucleus. This can
431 be ice, dust, or other particles. In the atmosphere of the Earth, it is rare for
432 conditions to exist where condensation nuclei are not present in sufficient quantities
433 to satisfy all condensation, and we do not consider those circumstances in this work.

434 2.2.2 Precipitation

435 Precipitation refers to the various forms of non-gaseous water that fall from clouds
436 due to the gravitational attraction of the Earth. Precipitation can be solid (hail or
437 snow), or it can be liquid (rain or drizzle). For this work, we consider only liquid and
438 ash precipitation. Liquid precipitation makes up the vast majority of precipitation,
439 while ash precipitation is of particular visual interest in the case of volcanoes.

440 Liquid precipitation

441 Rain falls when some of the water droplets that make up a cloud grow too large
442 to be kept in the air through electrostatic forces or air resistance. Droplets grow
443 through *coalescence*, which is where two or more droplets combine to form a larger
444 droplet. This is a result of collision, and turbulent conditions within the cloud lead
445 to faster coalescence[Voh+99]. Once a droplet has grown large enough to overcome
446 air resistance, it starts falling. Depending on the conditions below the droplet, it
447 might continue to grow through collision with other droplets (such as in strong
448 convective clouds), or it might evaporate before it hits the surface (known as *virga*).
449 Once the rain has reached the surface, it can be taken up by open water, or it can be
450 absorbed by the soil (turning the soil moist). It thus becomes available for uptake
451 into the atmosphere again, starting the water cycle afresh.

452 Ash precipitation

453 Ash rain is a common sight in the vicinity of volcanoes undergoing an eruption. It
454 is formed when fine ash particles mix with liquid water in the atmosphere, or serve
455 as condensation seeds for water to condense around. These particles of mixed water
456 and ash will grow through collision similar to the coalescence of liquid precipitation,
457 though it is called *aggregation* in the case of ash rain[BBD12; RD11]. Many factors
458 influence the size and granularity of the ash particles in a volcanic plume[Van+15],
459 such as the chemical composition of the volcanic magma, the geology of the terrain,
460 and the explosivity of the eruption, but we do not consider those intricacies in this
461 work.

462

463 This ash rain can pose a hazard to people in the vicinity of an eruption, often
464 smothering existing plant life, and can suffocate animal life if the ash fall is heavy.
465 It also presents a significant logistical challenge when clogging roads and railways,
466 or when it has to be removed from buildings[Jen+15]. However, this ash also greatly
467 enriches the soil, leading to highly fertile ground surrounding volcanoes.

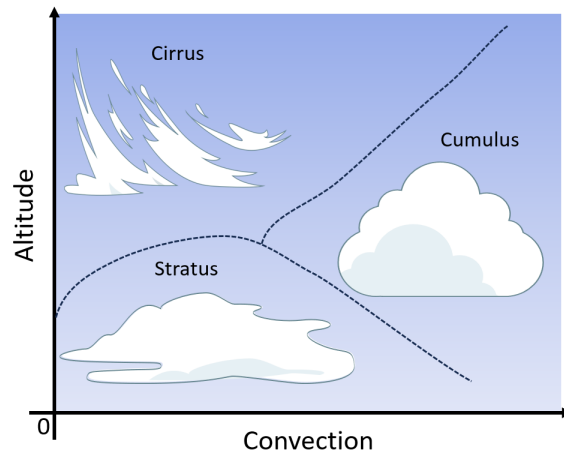


Figure 2.2: Cloud classification according to height and convective activity, with activity increasing to the right.

2.2.3 Cloud types

Clouds are classified into various different types according to the altitude (or level) of the clouds (which generally dictate the composition, whether majority liquid or majority ice), as well as the convective structure within the cloud (which dictates the shape of the cloud). The level of a cloud is generally indicated with a prefix, with no prefix for low level clouds, *alto-* for mid level clouds, and *cirro-* for high level clouds. Multilevel clouds usually get their own designation depending on the convective type.

In this work, we consider only a simplified cloud classification scheme, as shown in figure 2.2. We refer the reader to the International Cloud Atlas[Org17] for specifics on cloud types, as well as more information on cloud classification in general.

Cumulus clouds are clouds with a high degree of convection, characterised by a "fluffy" appearance. These clouds are almost entirely liquid, leading to the fine edge due to the strong electrostatic bonds between liquid water molecules. *Cirrus* clouds are wispy, filamentary clouds, with little convective activity, and occur only at high altitudes (and therefore do not get a prefix). These generally consist of ice crystals, which have weaker electrostatic bonds compared to liquid water, and this results in less defined edges compared to other types of clouds. *Stratus* clouds are flat sheets of cloud, with little to no convective activity.

2.3 Explosive volcanic eruptions

A short overview of explosive volcanic eruptions is required for an understanding of our model. In this work, we focus on the evolution of the volcanic plume and place less emphasis on the terrain altering effects, such as the formation of calderas or magma chambers, or the presence of lava or lahars. For a more complete treatment of volcanic eruptions, we refer the reader to works by Cas and Wright[CW12] and Parfitt and Wilson[PW09].




Figure 2.3: The typical progression of a Plinian eruption, divided into the jet, convective, and umbrella phases. The salient visual features resulting from an eruption and the phases where they occur are indicated as well.

497 There are many types of volcanic eruptions. We focus on a subset of magmatic
 498 eruptions, i.e., eruptions that involve the expulsion of molten rock and ash. In
 499 particular, we focus on the so-called Plinian or ultra-Plinian eruptions, since these
 500 eruptions have the most visually striking interactions with the atmosphere. This
 501 is compared to Peléan, Vulcanian, Strombolian, or Hawaiian eruptions. Magmatic
 502 eruptions are categorised according to the explosivity of the eruption, with Plinian
 503 eruptions being the most explosive and Hawaiian eruptions the least explosive.

504

505 Plinian eruptions are caused by excessive pressure within the volcano's magma
 506 chamber, often exacerbated by a plug, usually consisting of solidified lava, sealing
 507 the volcano vent. Pressure builds as solubles such as water vapour and other gases
 508 separate out of the magma, or as additional rock in the chamber is melted and ex-
 509 pands. Eventually, the plug gives way and the pressure is released catastrophically.
 510 This leads to an eruptive column of gas, lava, and tephra (the fragmented, rocky
 511 material erupted by the volcano) travelling at high velocities (50 to $200 \text{ m} \cdot \text{s}^{-1}$).
 512 This column is typically very dense compared to the air, often hundreds of kilograms
 513 per cubic metre, despite being extremely hot (in excess of 1200 K).

514

515 A Plinian eruption typically consists of three phases, demarcated by the primary
 516 forces that drive the movement of the eruptive column. In the initial *jet* phase, the
 517 column is lifted by the momentum from the explosion. Once expelled from the vol-
 518 cano vent, the column starts to expand due to the immense heat it has. The densest
 519 components of the jet, such as lava, volcanic bombs, and coarse ash, sediments out,
 520 further reducing the density of the eruptive column. Additionally, air surrounding
 521 the jet is dragged in and entrained in the plume, where it heats up. As described
 522 in section 2.1, hot air will expand and rise up due to convective forces. This is the
 523 *convective* phase of the eruption. If the initial explosion is not powerful enough for

524 entrainment to occur, or if the initial jet has a significant horizontal component,
525 the plume can collapse and spread across the terrain. This is typically considered a
526 Peléan eruption.

527

528 At some altitude, the volcanic plume will reach a neutral buoyancy, and enter
529 the *umbrella* phase. Depending on the initial force and temperature of the erup-
530 tion, this might be at an extremely high altitude, up to 58 kilometres above sea
531 level[PPS22]. This is well past the tropopause (see section 2.1), and is the reason
532 that the sulphur and other particulates ejected by volcanic eruptions can have global
533 long-term effects, such as a volcanic winter (a sustained period of lower temperatures
534 across the globe – due to the ash reflecting and blocking sunlight – lasting months
535 to years). Figure 2.3 displays the different phases of the eruption and some of the
536 visual features that occur in the plume and atmospheric interaction.

537 2.3.1 Volcanic weather phenomena

538 Each phase of the eruption has distinct weather phenomena associated with it. We
539 describe some of the most visually striking examples, which we focus on in this work,
540 and provide reference images.

541 Transient condensation clouds

542 In the initial explosion, the rarefaction phase of the shockwave causes a drop in
543 pressure and temperature. In humid atmospheres, this can lead to a very short-
544 lived imprint of condensation as the colder air cannot sustain the previous levels of
545 moisture[YI07]. This imprint showcases the composition of the atmosphere around
546 the volcano, highlighting sections of the atmosphere that are more humid. The
547 atmosphere rapidly returns to normal pressures and temperatures, leading to the
548 condensation cloud evaporating again. This phenomenon rarely lasts more than a
549 handful of seconds.

550

551 This phenomenon is seldom captured in relation to volcanoes, mainly due to the
552 remote location of many volcanoes and the lack of visual monitoring equipment in
553 close proximity to volcanic explosions. However, it is often captured with man-made
554 explosions, such as munitions tests or accidents like the 2020 Beirut explosion. Fig-
555 ure 2.4 shows a TNT explosion with an associated transient condensation cloud.

556

557 Figure 2.5 shows two transient condensation clouds forming after an explosion
558 of Mount Tavurvur in Papua New Guinea. While this explosion is of a Strombolian
559 nature, Plinian or ultra-Plinian eruptions will be more explosive and have stronger
560 shockwaves. The first transient condensation cloud forms at $t = 0.5$ s (figure 2.5(4))
561 around the plume and dissipates by $t = 1.5$ s (figure 2.5(10)). The second, more
562 distant cloud forms at $t = 2.83$ s (figure 2.5(14)) in the air near the top of the frame,
563 and disappears from the camera view by $t = 4.5$ s (figure 2.5(24)).

564 Embedded clouds

565 Along with air being entrained into the plume, air layers can also be forced upwards
566 by the rising plume. Eventually, these air layers mix with the plume. As with



Figure 2.4: A photo of 500 tons of TNT exploding shows the shockwave progress visible on the water and the condensation cloud forming behind it in the air. *Operation Sailor Hat, Hawaii, 1965. Courtesy of the US Navy.*

567 normal convective columns, the air in these layers cools as it rises. If it rises high
 568 enough (or has sufficient moisture), the moisture in the layers can condense to form
 569 clouds. Depending on the structure of the atmosphere at the time of the explosion,
 570 these layers can take the form of layered “bell” or “skirt” clouds that appear to
 571 hang from the plume[Bar82]. Figure 2.6 shows examples of this. It can also form
 572 saucer-like pileus or “cap” clouds on top of the plume, as happens occasionally with
 573 strong convective clouds. Figure 2.7 demonstrates an excellent example of a cap
 574 cloud forming and the plume breaking through it.

575 **Ash rain**

576 While the denser components of the eruptive column will fall out soon after the
 577 eruption, the less dense, finer ash can be carried up with the plume, often for dozens
 578 or hundreds of kilometres. In the umbrella phase, this fine ash can mix with water
 579 droplets to form ash rain. This is a common feature of most volcanic eruptions,
 580 though the sheer volume of ash erupted by Plinian eruptions make the subsequent
 581 ash rain more striking and long-lasting. Figure 2.8 shows a plume advected by
 582 the wind in one direction at its top, but the ash rain is subject to different wind
 583 conditions and therefore follows a different direction.

584 **Lightning**

585 The strong convection within a volcanic plume results in friction between particles
 586 of ash, ice, and tephra. This friction often leads to electric discharges in and around
 587 the plume as volcanic lightning. While singular strikes are brief, many strikes can
 588 occur in a short span of time. In addition, these strikes are often spatially extensive,
 589 taking the appearance of a web or a sheath enclosing the plume. Figure 2.9 gives
 590 an example of volcanic lightning in nature, with the glow from the lava adding a



Figure 2.5: A timelapse showing two condensation clouds forming near the vent of a volcano following an explosion. The first occurs at $t = 0.5 \text{ s}$ around the plume, while a second forms at $t = 2.83 \text{ s}$ in the air near the top of the frame. The full video is available on [YouTube](#). *Mount Tavurvur, Papua New Guinea, 2014.* © Phil McNamara.



Figure 2.6: Skirt clouds are formed by the strong convective uplift around the plume. *Eruption of Hunga Tonga-Hunga Ha'apai, Tonga, 2022. Courtesy of the Tonga Geological Services.*



Figure 2.7: A pileus cloud can form at the peak of the plume, before the plume breaks through the cap cloud. *Eruption of Sarychev Peak, Russia, 2009. Courtesy of NASA.*



Figure 2.8: Ash carried up by the plume mixes with water droplets in the upper atmosphere, and then falls to the ground as ash rain. During the fall, it is subject to the local wind field that will dictate the dispersal. *Eruption of Shinmoedake, Japan, 2009. Courtesy of the U.S. Geological Survey.*

591 striking contrast to the plume.

592 2.4 Fluid dynamics of the atmosphere

593 At the heart of our simulation lies the fluid dynamics engines. These are responsible
594 for transport within the layer and ensuring that the wind field remains consistent
595 with the properties of an incompressible fluid. In this section we explain why the
596 atmosphere can be treated as an incompressible fluid, provide the mathematical
597 grounding on the Navier-Stokes equations for incompressible fluids, and go over the
598 implementation of these continuous functions on a 2D discretised grid (see section
599 4.1 for more details on our implementation).

600 2.4.1 Dynamics of incompressible fluids

601 The atmosphere behaves as a fluid, and whether or not a fluid is considered incom-
602 pressible depends on whether the density changes as a result of the fluid moving
603 through a flow field. This depends on the elastic modulus and density of the fluid,
604 which can be approximated using the speed of sound within the fluid. A higher elas-
605 tic modulus (i.e., a stiffer material) will have a higher speed of sound, while a higher
606 density will lead to a lower speed of sound. All fluids are compressible under certain
607 circumstances, e.g., high speed movement or confined spaces. The atmosphere is an
608 open system and thus we only have to concern ourselves with the speed at which
609 the fluid is moving.

610
611 If the fluid is moving much slower than the speed of sound (the rate at which
612 vibrations within the fluid is transmitted), then any thrust applied to the fluid has
613 time to disperse to neighbouring parcels, and any compression will be negligible. If
614 the fluid is moving closer to the speed of sound within the fluid, then the applied
615 thrust does not have time to disperse, and the parcels within the fluid will be forced
616 closer together, i.e., be compressed.

617



Figure 2.9: A sheath of lightning surrounding the plume, with a long night time exposure highlighting the multiple strikes. *Eruption of Eyjafjallajökull, Iceland, 2010.* ©Terje Sørgerd.

618 The speed of sound within air at standard temperature and pressure is $343 \text{ m} \cdot \text{s}^{-1}$.
 619 Since the temperature (and thus density) of the air varies with altitude, the speed
 620 of sound also varies with altitude. The speed of sound will generally decrease with
 621 altitude, reducing to about $300 \text{ m} \cdot \text{s}^{-1}$ at the tropopause[BHR07]. In comparison,
 622 the fastest wind speed ever recorded on Earth was $113.3 \text{ m} \cdot \text{s}^{-1}$ in the eye wall
 623 of the category 4 tropical cyclone Olivia in 1996[Cou+12], while the average wind
 624 speed across the globe is merely $4.9 \text{ m} \cdot \text{s}^{-1}$ [AJ05]. Even the most extreme case we
 625 can expect is well below the speed of sound for air.

626
 627 While the pressure waves produced by volcanic eruptions can approach the speed
 628 of sound[Wri+22], the wind speed resulting will be less given the short-lived nature
 629 of the pressure waves. Wind resulting from displaced terrain could potentially be
 630 faster, but we do not model terrain displacement and can therefore ignore it. Thus,
 631 we can safely model the atmosphere as an incompressible fluid.

632 2.4.2 Navier-Stokes

633 The Navier-Stokes equations for incompressible fluids are given by the following
 634 equations:

$$\frac{\partial \vec{U}}{\partial t} + \vec{U} \cdot \nabla \vec{U} + \frac{1}{\rho} \nabla p = \vec{g} + \nu \nabla \cdot \nabla \vec{U} \quad (2.3)$$

635 and

$$\nabla \cdot \vec{U} = 0 \quad (2.4)$$

636 Equation 2.3 describes the change in the velocity vector field (\vec{U}), and is de-
 637 rived from Newton's Second law of motion, $\vec{F} = m\vec{a}$. Equation 2.4 describes the
 638 conditions that have to be satisfied if the fluid is incompressible, which is derived
 639 from calculating the change in volume with the Divergence Theorem, and having
 640 the constraint that $\frac{dV}{dt} = 0$ for all t in our system (where V is the volume of the
 641 entire system, i.e., the total volume remains unchanged). See table 2.1 for an ex-
 642 planation of the symbols used in these equations and their units. For the details of
 643 the derivations, we refer the reader to *Fluid Simulation for Computer Graphics* by
 644 Robert Bridson [Bri15], on which much of this section's work was based.

645

Symbol	Description	Unit
\vec{U}	2D vector velocity	$\text{m} \cdot \text{s}^{-1}$
t	time	s
ρ	fluid density	$\text{kg} \cdot \text{m}^{-3}$
p	(modified) pressure	Pa
\vec{g}	Vector sum of body forces	N
ν	kinematic viscosity	$\text{Pa} \cdot s$

Table 2.1: Symbols and their units used in equations 2.3 and 2.4

646 In equation 2.3, the acceleration for a given fluid parcel is calculated with three
 647 forces affecting it: the vector sum of body forces $-\vec{g}$ (i.e., forces such as gravity);
 648 the viscosity $-\nu \nabla \cdot \nabla \vec{U}$ (the fluid's resistance to deformation during flow); and most

649 importantly, the force exerted by the pressure gradient $\frac{1}{\rho}\nabla p$. This last term is the
 650 corrective term that ensures that our velocity field remains non-divergent and sat-
 651 isfies equation 2.4.

652

653 We model the atmosphere as a closed system of 2D layers. The only body force
 654 applicable to our system is gravity, since we solve the transport equations separately
 655 for each layer, neglecting the forces that can arise from vertical transport. These
 656 effects are considered in a procedural manner, described in section 4.3.2. We model
 657 only the horizontal velocity in our fluid simulation, that is, the velocity parallel
 658 to a flat terrain. The acceleration due to gravity is perpendicular to this velocity,
 659 and therefore the force as a result of gravity will always be zero in our 2D layers.
 660 Therefore, we can discard all body forces in our calculation. The kinematic viscosity
 661 ν of air is very low, with magnitudes on the order of 10^{-5} . Therefore, we can discard
 662 the viscosity term, though it can cause errors. Interestingly, the numerical diffusion
 663 that arises from discretising equation 2.3 will often have similar results to including
 664 the viscosity term [Bri15]. These factors combine to give us the simplified equation
 665 below:

$$\frac{\partial \vec{U}}{\partial t} + \vec{U} \cdot \nabla \vec{U} + \frac{1}{\rho} \nabla p = 0 \quad (2.5)$$

666 This, combined with equation 2.4 which still holds, is the final equation we use for
 667 modelling movement within layers.

668 2.4.3 Advection

669 For transporting the quantities in the layer accurately, we need to solve the **advect-**
 670 **tion** equation, equation 2.6. The quantities need to be moved without changing their
 671 values, that is, the quantities need to be *advected*. For a quantity q , the following
 672 constraints hold:

$$\begin{aligned} \frac{\partial q}{\partial t} + \vec{U} \cdot \nabla q &= 0 \\ \frac{Dq}{Dt} &= 0 \end{aligned} \quad (2.6)$$

673 There are two approaches to solving the advection equation, which also tracks
 674 and models the movement within a continuous field such as a fluid. See figure 2.10
 675 for a comparison between the two methods. *Lagrangian* methods are particle-based,
 676 with each particle having its own position, velocity, and scalar values (see table 4.6).
 677 At each timestep, each particle calculates its acceleration as per equation 2.3 and
 678 updates its position. This method is conceptually simple and has little or no issue
 679 with resolution, but it is difficult to accurately calculate the pressure gradient to
 680 solve equation 2.4.

681

682 *Eulerian* methods discretise the field into cells, with the velocity and scalar
 683 quantities sampled at these fixed cell points in the space. This method is easier
 684 when calculating the pressure gradient, but can suffer from accuracy and resolution
 685 issues if the distance travelled by a parcel of air is less than the size of a cell, or
 686 significantly greater than the size of a cell.

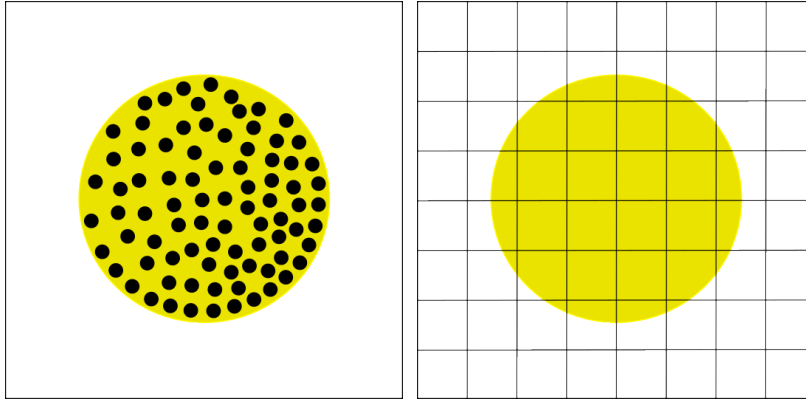


Figure 2.10: On the left, a Lagrangian approach. On the right, a Eulerian approach. In both cases, the fluid volume is indicated in yellow, with the outer square indicating the total volume.

2.4.4 Semi-Lagrangian advection

Solving the advection equation for a Eulerian method is unconditionally unstable[Bri15]. That is, there is no Δt that will not lead to a divergent result after some number of steps. Smaller values for Δt can increase the number of steps that are stable, but not indefinitely and this becomes impractical in short order. In comparison, the Lagrangian approach is unconditionally stable and solves the advection equation automatically as the particle is moved. However, it is possible to combine these two approaches, as was first done by Stam[Sta99].

Let us say we are examining the value of a quantity q at a grid point \vec{x}_G at a time $t + \Delta t$. We can create a virtual particle P that is exactly at \vec{x}_G at time $t + \Delta t$. To determine the value of q_P , we can trace P back to time t , when it would have been at position \vec{x}_P . We determine the position of \vec{x}_P with the velocity field \vec{U} at \vec{x}_G . By sampling the value of q at \vec{x}_P at time t , we know the value of q_P , since equation 2.6 holds and is automatically solved as part of the Lagrangian treatment. Therefore, for each grid point \vec{x}_i at time $t + \Delta t$, we can determine the value q_i by creating a virtual particle and tracing it back to time t . This method benefits from the spatial completeness and simpler solutions to equation 2.4 offered by the Eulerian approach, while still maintaining the unconditional stability of the Lagrangian approach.

2.4.5 Correcting for divergence

Applying the Navier-Stokes momentum equation (equation 2.3) on a discretised grid will violate the incompressibility condition (equation 2.4). This necessitates a correction step, which is the pressure term in equation 2.3 ($\frac{1}{\rho}\nabla p$). For a grid point \vec{x}_i , we can calculate the pressure gradient from the pressure values at neighbouring points using a central difference method:

$$\left(\frac{\partial p}{\partial x}\right)_i \approx \frac{p_{i+1} - p_{i-1}}{2\Delta x} \quad (2.7)$$

Marker-and-Cell Grid

However, a naive implementation on a uniform grid can violate equation 2.6 if $p_{i-1} = p_{i+1}$ in equation 2.7, since the derivative could be 0 despite p not being




Figure 2.11: An example of semi-Lagrangian advection of a grid point at position p at time $t + \Delta t$.

715 constant across the grid. Alternatively, using p_i as one of the points in equation 2.7
 716 will cause a bias towards one side, which is not desirable.

717

718 A solution to this is to use the Marker-and-Cell (MAC) discretisation, first pro-
 719 posed by Harlow and Welch[HW65]. This is a staggered grid where the velocity
 720 vector \vec{U} is separated into its components and the components are placed on the
 721 edge of the cell, while scalar quantities including pressure are placed in the middle
 722 of the cell. See figure 2.12 for an example.

723

724 This grid discretisation is useful since it allows a pressure gradient calculation
 725 without discarding any values (as equation 2.7 does for p_i). The central differences
 726 method for a point \vec{x}_i will then be:

$$\left(\frac{\partial p}{\partial x}\right)_i \approx \frac{p_{i+1/2} - p_{i-1/2}}{\Delta x}, \quad (2.8)$$

727 which ensures that the derivative will only be 0 if $p_{i+1/2} - p_{i-1/2}$ is 0, i.e., $p_{i-1/2} =$
 728 $p_{i+1/2}$.

729 Pressure projection

730 Once the pressure gradient (∇p) is calculated, we project the effect of pressure onto
 731 the velocity vector field, with the constraint that equation 2.4 must hold afterwards.
 732 This gives us equation 2.9 for velocity at step $n + 1$.

$$\vec{U}^{n+1} = \vec{U} - \frac{\Delta t}{\rho} \nabla p \quad (2.9)$$

733 When applied to a two-dimensional MAC-grid (such as shown in figure 2.12),
 734 the velocity decomposes into u and v :




Figure 2.12: A standard Marker-and-Cell (MAC) grid representation, with velocity components (u and v) on the edges of the cells and pressures in the centre of the cell.

$$u^{n+1} = u - \frac{\Delta t}{\rho} \frac{p_{i,j+1/2} - p_{i,j-1/2}}{\Delta x} \quad (2.10)$$

$$v^{n+1} = v - \frac{\Delta t}{\rho} \frac{p_{i+1/2,j} - p_{i-1/2,j}}{\Delta y} \quad (2.11)$$

735 Calculating divergence

736 The pressure gradient is used to enforce equation 2.4. In two dimensions, the diver-
737 gence of \vec{U} is expressed as:

$$\nabla \cdot \vec{U} = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \quad (2.12)$$

738 When applied to a MAC grid where the cells are square (i.e., $x \equiv y$), the diver-
739 gence of the velocity in a cell i, j can be approximated using a central differences
740 method as the following:

$$\begin{aligned} x &\equiv y \\ (\nabla \cdot \vec{U})_{i,j} &= \frac{u_{i,j+1/2} - u_{i,j-1/2}}{\Delta x} + \frac{v_{i+1/2,j} - v_{i-1/2,j}}{\Delta y} \\ (\nabla \cdot \vec{U})_{i,j} &= \frac{u_{i,j+1/2} - u_{i,j-1/2} + v_{i+1/2,j} - v_{i-1/2,j}}{\Delta x} \end{aligned} \quad (2.13)$$

741 Calculating pressure

742 At this point, no pressure value has been calculated yet. But by substituting equa-
 743 tions 2.10 and 2.11 into equation 2.13, we obtain a system of linear equations (equa-
 744 tion 2.14) where the pressures are unknown. See appendix A for the equivalent
 745 derivation when $\Delta x \neq \Delta y$, the square case is demonstrated here for legibility and
 746 comprehensibility.

$$\begin{aligned}
 & \frac{u_{i,j+1/2} - u_{i,j-1/2} + v_{i+1/2,j} - v_{i-1/2,j}}{\Delta x} = 0 \\
 & \frac{\left(u_{i,j+1/2} - \frac{\Delta t p_{i,j+1} - p_{i,j}}{\rho \Delta x}\right) - \left(u_{i,j-1/2} - \frac{\Delta t p_{i,j} - p_{i,j-1}}{\rho \Delta x}\right)}{\Delta x} + \\
 & \frac{\left(v_{i+1/2,j} - \frac{\Delta t p_{i+1,j} - p_{i,j}}{\rho \Delta x}\right) - \left(v_{i-1/2,j} - \frac{\Delta t p_{i,j} - p_{i-1,j}}{\rho \Delta x}\right)}{\Delta x} = 0
 \end{aligned} \tag{2.14}$$

747 We can rearrange equation 2.14 to have the divergence of the cell (with inverse
 748 sign) on the right hand side:

$$\begin{aligned}
 & \frac{\Delta t (4p_{i,j} - p_{i+1,j} - p_{i-1,j} - p_{i,j+1} - p_{i,j-1})}{\rho \Delta x^2} = \\
 & - \left(\frac{u_{i,j+1/2} - u_{i,j-1/2} + v_{i+1/2,j} - v_{i-1/2,j}}{\Delta x} \right)
 \end{aligned} \tag{2.15}$$

749 This forms a sparse linear system of the form $Ap = D$, where A is an $N \times N$
 750 square matrix, where N is the number of cells in the system, and contains the coef-
 751 ficients of the various pressures in equation 2.15. Here, p is a column vector with N
 752 entries, as is D , where $D_i = -(\nabla \cdot \vec{U})_i$.

753

754 We then solve for p using any matrix solver of our choice. However, equation
 755 2.15 is a symmetric positive definite matrix and this makes it very well suited to the
 756 conjugate gradient solver, as first described by Hestenes and Stiefel[HS+52]. Since
 757 we constrain our Δt such that air parcels do not travel more than a handful of cell
 758 widths at a time, the pressures in our system tend not to vary wildly between steps.
 759 This allows us to use the pressures calculated in the previous time step as the first
 760 guess for the next, which significantly speeds up convergence.

761

762 The pressures calculated can then be used with equations 2.10 and 2.11 to ensure
 763 \vec{U} is nondivergent.

764 2.4.6 Boundary conditions

765 Domain boundaries

766 We model the domain boundaries using a wraparound scheme, creating a toroidal
 767 domain for each layer. For a domain of $N_x \times N_y$, the first column ($j = 1$) has as
 768 its neighbours the second column ($j = 2$) and the last column ($j = N_x$). Similarly,
 769 the first row ($i = 1$) has the second row ($i = 2$) and the last row ($i = N_y$) as its
 770 neighbours.

771 **Solid boundaries**

772 The equations presented in section 2.4.5 assumes that all cells in the system are
773 fluid cells. However, a system with solid obstacles present require an additional
774 constraint to the pressure calculation. Specifically, the velocity tangential to a solid
775 interface must be equal to the velocity of the solid interface (equation 2.16, where
776 \vec{u}_S is the velocity of the solid interface).

$$\vec{u}^{n+1} \cdot \hat{n} = \vec{u}_S \cdot \hat{n} \quad (2.16)$$

777 In our case, the obstacles are in the form of terrain (e.g., mountains, cliffs, or
778 hills impinging on the air layer). These are stationary, and therefore \vec{u}_S will always
779 be 0.

780
781 Since we use the MAC grid, we will always have a velocity stored on the exact
782 grid point where a solid interface might be (since cells can either be fluid or solid).
783 By replacing the velocity of some boundary in equation 2.13 with \vec{u}_S , a few general-
784 isations become apparent. The coefficient in front of $p_{i,j}$ is the number of non-solid
785 cells that border cell i, j , and any neighbour that is solid has its pressure term
786 removed from that row of the matrix. Also, the velocity of the point in question on
787 the right hand side (in the divergence calculation) is replaced by \vec{u}_S . See appendix
788 B for an example derivation.

789 Chapter 3

790 Related work

791 Ever since technology progressed sufficiently to display digitally generated images,
792 the ability to generate and display clouds and other weather phenomena has been
793 considered a key part of digital landscapes on multiple scales. Like many advances in
794 technology, the first steps in this direction were fuelled by military or pseudo-military
795 applications. The first published work on visual cloud modelling is a report for the
796 National Atmospheric and Space Administration (NASA) by Barnes et al.[Bar+68]
797 in 1968, which used data from meteorological records to generate probabilistic cloud
798 cover distributions on a global scale, aimed to assist in planning and training for the
799 Apollo space missions.

800

801 Research in meteorology had been using numerical simulations for a few years by
802 then, but the first report using a visual representation was by Mancuso et al.[MSB71]
803 in 1971, incorporating a simple command line representation of the clouds from be-
804 low and from a side-on perspective. A lot of research in this area followed, though
805 the researchers' aims were generally to create a model that was as physically accu-
806 rate as possible, giving less concern to the appearance of the clouds.

807

808 As late as 1983, the state of the art graphics used for NASA's Shuttle Mission
809 Simulator did not have clouds for the in-atmosphere scenes, and only one config-
810 uration for clouds as seen from orbit[Hen83]. But with the proliferation of digital
811 computer games during the 1980s, an increased focus was put on the visual appear-
812 ance of clouds and skylscapes, especially if it could be done at a lower computational
813 cost. Gardner[Gar85] built on the work of several contemporaries, such as Ka-
814 jiya[KV84], to create a model for cloud generation focused on appearing as realistic
815 as possible without having an exorbitant computational cost. The rate of progress
816 has only increased since, with digital games, the entertainment industry, and mete-
817 orologists all requiring clouds of varying realism and physical accuracy, buoyed by
818 the exponential increase in computing power available. However, to date, no one
819 has addressed the extremes of temperature (up to 1300 K) and altitude (>20 km)
820 found within volcanic plumes in a computer graphics application.

821

822 In contrast with the glut of work on clouds, volcanic plumes have languished,
823 with little attention given to this topic within the computer graphics community.
824 The methods presented have limited physical accuracy, especially with regards to
825 the effects that interaction with the atmosphere present.

Scale	Procedural	Simulation	Data-driven	Hybrid
Micro-	[Ney97; Sch+03]	[KV84; MDN02a] [Har+03; Bi+16]	[Yua+14]	[Xu+09a]
Meso-	[Gar85; MB15] [Web+18]	[Dob+00; Miy+01] [FDY15]	[Sat+22] [Gos+23]	[Xu+09a; Xu+09b] [DG17; GN17; Gos19] [Häd+20; Vim+20]
Synoptic	[WM81]	[DYN06]	[Dob+99]	[Xu+09a]

Table 3.1: Cloud generation systems for skylscapes can be classified according to the scale at which the methods operate, and the broad category that the method falls under. While a lot of research has been devoted to micro- and mesoscale skylscapes, synoptic scales have received little attention. Data-driven and hybrid methods have become prevalent in recent years.

826

827

828

In this chapter, we examine the previous body of work in cloud animation and volcanic plume modelling, and highlight the shortcomings we address with this work.

829

3.1 Cloud modelling and animation

830

831

832

833

834

835

There are several methods of generating and animating clouds, each with their own advantages and disadvantages. We divide the methods into procedural, simulation, data-driven, and hybrid techniques. We also divide the scale on which proposed methods operate into the microscale (<1 km across), the mesoscale (<100 km across), and synoptic (global) scales. See Table 3.1 for a classification of proposed methods into these categories.

836

837

838

839

840

841

Once the clouds are generated and animated, care must be taken when rendering to achieve a realistic appearance. However, the particulars of cloud rendering is beyond the scope of this dissertation, and we refer the reader to Kallweit et al. [Kal+17] and Wilkie et al. [Wil+21] for a more comprehensive look at the subject.

842

843

844

Because cloud modelling and animation is a mature research topic, a full account is impossible and this section will therefore cover the essentials only. See Goswami [Gos20] for a recent survey of the field.

845

3.1.1 Procedural

846

847

848

849

850

851

852

853

Procedural techniques generally focus on speed and quality rather than the physical accuracy of the clouds. These techniques generally aim to create a density field or function, which is then used to generate and render clouds. Given the origins of cloud modelling coming from climate simulation, most of the initial work in the field was focused on accuracy rather than computational cost. Procedural methods are generally fast, but they generally struggle to model clouds over time and do not scale well.

854

855

Procedural noise functions are generally used to create the textures after the clouds are generated in some fashion. Gardner [Gar85] modelled cumuliform clouds

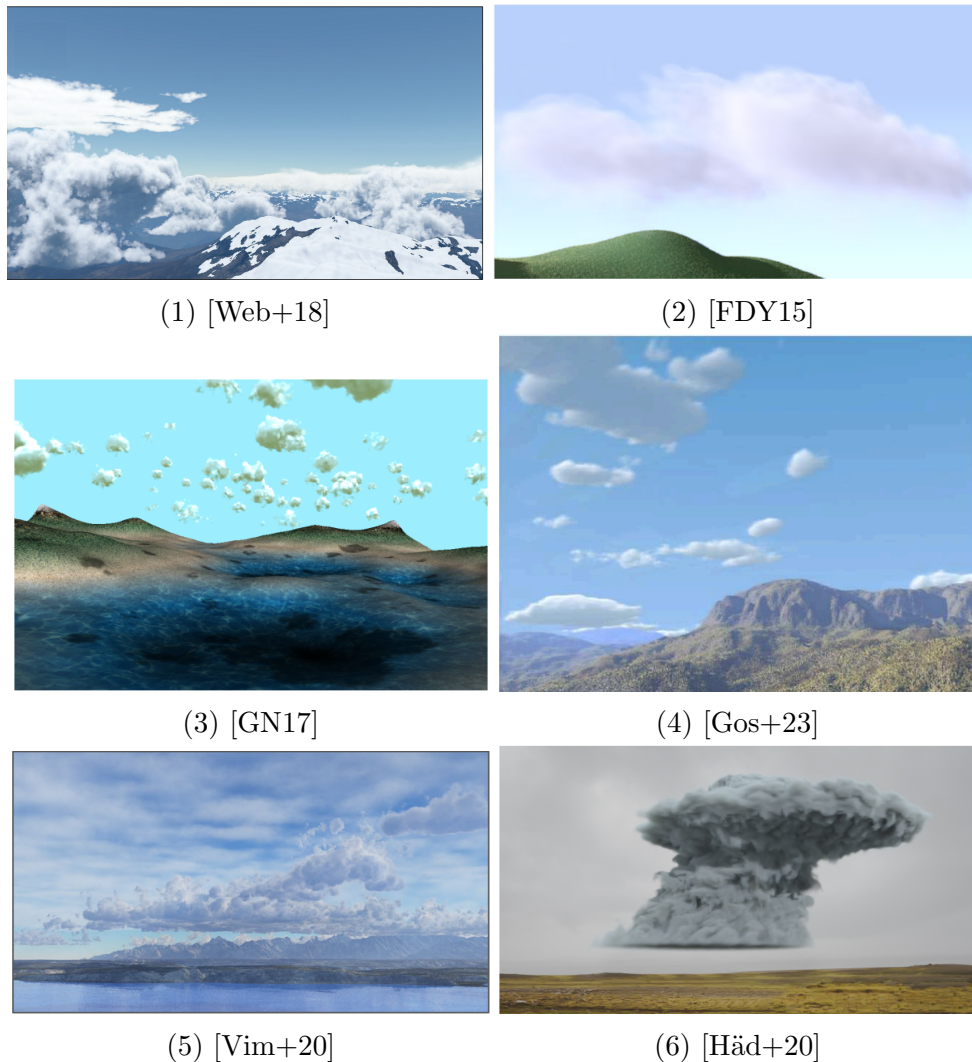


Figure 3.1: Examples from several works on cloud modelling and animation. (1) shows a procedural method from Webanck et al.[Web+18], while (2) and (3) show Lagrangian simulation methods from Ferreira Barbosa et. al[FDY15] and Goswami and Neyret[GN17], respectively. (4) shows a data-driven method from Goswami et al.[Gos+23]. (5) and (6) are hybrid methods from Vimont et al.[Vim+20] and Hädrich et al.[Häd+20], both employing simulation methods with procedural amplification.

856 using ellipsoids and blending multiples to form a cloud, while Neyret[Ney97] used
857 bubbles to similar effect. Schpok et al.[Sch+03] used a combination of ellipsoids and
858 particles textured with perlin noise to produce volumetric clouds that had realistic
859 animation.

860

861 Mukhina and Bezgodov[MB15] used perlin noise to generate 2D stratocumulus
862 clouds, using multiple layers to give the appearance of realistic, volumetric clouds.
863 With rendertimes under 10 milliseconds per frame even on low-end hardware, it
864 is very fast, albeit limited to only one type of cloud. Webanck et al.[Web+18]
865 combined two procedural functions, one for shape and one for cloud presence, to
866 compute density functions for multiple types of clouds, using primitives to model
867 the cloud. They also allowed for users to define shapes for the clouds using a control
868 field. However, their system encounters challenges with terrain interaction given the
869 procedural nature of their cloud map. Figure 3.1(1) shows an example image from
870 their system.

871

872 Procedural methods have the benefit of being very fast and very scaleable. How-
873 ever, because procedural functions have some degree of randomness, temporal and
874 spatial consistency can be difficult to achieve over small timesteps. In addition,
875 procedural methods are usually tuned to a specific cloud type, requiring different
876 functions or parameters for different cloud types. The clouds generated by procedu-
877 ral methods have no guarantee of being physically accurate, which prevents coupling
878 a procedural system on its own with other circulation models for different phenom-
879 ena where bi-directional feedback is required. These methods are of significant use
880 to enhance the results of other cloud generation methods, which we cover in section
881 3.1.4.

882 3.1.2 Simulation

883 Most simulations for cloud modelling rely on solving systems of partial differential
884 equations (PDEs) based on the theories of fluid dynamics and Navier-Stokes equa-
885 tions. Simulations have the advantage of generally appearing realistic and being
886 trivial to model dynamically over time and scalable to large volumes, at the cost
887 of very long computation times, as well as offering little or no control. Simula-
888 tion methods can usually be divided into Lagrangian (particle-based) and Eulerian
889 (grid-based) methods.

890 **Eulerian** Eulerian methods have been used in a number of approaches. Miyazaki
891 et al.[MDN02a] and Harris et al.[Har+03] both used a grid of voxels and a simplified
892 model of fluid dynamics to model wind, humidity, and temperature. Dobashi et
893 al.[DYN06] used a voxel-based grid to model a global cloudscape with user-provided
894 pressure maps according to fluid dynamics. Xu et al.[Xu+09a] used a combination
895 of grid-based fluid dynamics and metaballs to have realistic images at both global
896 and small scales.

897 **Lagrangian** Kajiya and von Herzen[KV84] defined a system of PDEs to create
898 a convective atmosphere where the humidity determines the formation of clouds,

989 using 'air parcels' advected over time. Ferreira Barbosa et al.[FDY15] created a La-
990 grangian model for cloud simulation, with particles splitting and merging depending
991 on the interest level of an area for increased efficiency, building on the work of Adams
992 et al.[Ada+07]. Goswami and Neyret[GN17] used a more traditional particle model
993 to create clouds, with procedural generation of textures for a realistic look, but is
994 limited in resolution and accuracy by the number of packets chosen. Figure 3.1(3)
995 shows an example render from this system.

996
997 Cellular automata were used by Dobashi et al.[Dob+00], using discrete val-
998 ues, to efficiently simulate clouds and generate a density distribution. Miyazaki
999 et al.[Miy+01] used real numbers in cellular automata to generate multiple cloud
1000 shapes that developed accurately over time. More recently, Bi et al.[Bi+16] used a
1001 combination of cellular automata and particle simulations to simulate clouds.

1002
1003 Simulation methods have the advantage of being physically accurate, temporally
1004 and spatially consistent, and capable of reproducing nearly all cloud types at suffi-
1005 cient resolution. This comes at a cost of performance, with simulation usually being
1006 in the interactive range or slower, only rarely being faster than real time. It also
1007 carries the cost of complexity, usually requiring a full computational fluid dynamics
1008 treatment. Comparing different simulation methods, Lagrangian methods generally
1009 result in better peak resolution than Eulerian methods. However, this comes at a
1010 cost of some areas of the domain potentially having a very poor resolution.

1011 3.1.3 Data-driven

1012
1013 Data-driven techniques rely on real data to generate or modify a model. Dobashi
1014 et al.[Dob+99] used satellite images to create a metaball-based model. Dobashi et
1015 al.[DSY10] modelled cumulus, altocumulus, and cirrus clouds by extracting features
1016 from photos of real clouds. Yuan et al.[Yua+14] refined this work for cumulus clouds,
1017 generating a more detailed three-dimensional representation of the cloud. More re-
1018 cently, Satilmis et al.[Sat+22] used a convolutional neural network (CNN) trained
1019 on 630 images of real skies to create skymaps from user provided cloud maps and
1020 sky settings. Goswami et al.[Gos+23] utilised the growing trend of convolutional
1021 generative adversarial networks (CGANs) to animate cloud maps generated from a
1022 model trained on ground-based videos of real cloudscapes. However, these anima-
1023 tions do not have the fine detail present in real clouds. Figure 3.1(4) depicts an
1024 example from Goswami et al.'s system.

1025
1026 Data-driven models suffer from a number of drawbacks. Firstly, acquiring a suf-
1027 ficiently large dataset that is labelled is challenging at best. Satilmis et al.[Sat+22]
1028 used a very limited number of images for their method, much less than what is
1029 standard in machine learning literature. This can lead to temporal artefacts, or
1030 limiting the length of animation to avoid said artefacts. To boot, acquiring volu-
1031 metric data for datasets is both technically challenging and expensive. With only
1032 one planet's example of real data, it is challenging to generalise to different planetary
1033 bodies. Data-driven models do have the benefit of being fairly fast (after the model
1034 is trained), and requiring fairly little in terms of initial conditions and parameter
1035 tuning (once the model is trained).

3.1.4 Hybrid models

It has become increasingly common for modelling systems to incorporate a hybrid method combining two or more of the techniques mentioned above. Xu et al.[Xu+09b] combined a simulation using cellular automata to determine the density field and procedural generation with Brownian motion functions for the fine shape of the clouds. Goswami and Neyret[GN17] used a Lagrangian fluid dynamics simulation at a large scale and procedural hypertextures on smaller scales for computationally inexpensive yet realistic-looking clouds. This work was extended by Goswami[Gos19] for more realistic cloud distributions and improved performance, albeit limited to only a single layer.

Duarte and Gomes[DG17] used a simplified Lagrangian simulation combined with thermodynamic diagrams from real data to simulate clouds in real-time. Vimont et al.[Vim+20] used a layered Eulerian meso-scale simulation to simulate weather at a coarse scale, combined with procedural noise amplification at finer scales to generate a wide array of cloud types at interactive rates. However, this system uses a procedural method to couple the layers, which can lead to unrealistic behaviour. Figure 3.1(5) shows an example render from Vimont et al.’s method. Hädrich et al.[Häd+20] combined a detailed physical simulation with procedural amplification of heat and wind maps to create intricate clouds of many different types, including several variations of cumulonimbus seen in nature. However, the underlying heat maps require parameters set by the user, and this requires in depth knowledge to accurately create the desired cloud types. Figure 3.1(6) shows an example render of a cumulonimbus cloud in Hädrich et al.’s system.

Hybrid models have a benefit over procedural by generally being more accurate over time, as well as modelling a wider range of cloud types. They tend to require less time than pure simulation methods for the same fidelity. However, this comes at the cost of significant complexity, and an increased overhead for tuning the parameters and setting the correct initial conditions for the desired scenario.

3.1.5 Combined models

Of particular interest to us are methods where two or more submodels for different phenomena interact for a combined simulation, in contrast to hybrid models which utilise two or more methods for generating clouds (a single output). Global circulation models (GCMs) are common in climate science, and combine atmospheric, ecological, and oceanographic models to cover as many aspects of the global climate as possible. However, these models are extremely complex, and require immense computation power. For example, Wedi et al.[Wed+20] simulated four months of global circulation using 5 days of compute time on OLCF-4/Summit[Hin18]. This is not feasible for graphical applications (with some exceptions, discussed in section 3.1.6), and we therefore focus on smaller scale simulations for graphics applications.

Hädrich et al.’s Fire in Paradise[Häd+21] combines a module-based ecosystem model with a hybrid atmospheric system drawn from their previous work on stormscapes[Häd+20] and extended to include rain vaporisation. This method models the spread of fire, smoke, and clouds to depict forest fire dynamics. Complex feedback




Figure 3.2: Example renders from two combined models. (1) depicts a flammaginetus cloud from Hädrich et al.’s *Fire in Paradise*[Häd+21], where a forest fire leads to convection and subsequent condensation. (2) is an example render from Pałubicki et al.’s *Ecoclimates*[Pał+22], which combine terrain and vegetation submodels with an atmospheric model.

991 effects arise from the two-way interaction between the fire sub-model and the atmo-
 992 sphere. Flammaginetus clouds form in strong convective zones above the hot fires,
 993 leading to uplift. This uplift can cause water vapour to condense, leading to rainfall
 994 which can extinguish fires again. Hädrich et al. successfully model forest fires under
 995 different conditions, including conditions that lead to flammaginetus clouds forming
 996 in convective columns. The authors do not provide detailed computation times,
 997 but claim interactive rates in all scenarios. They are limited in resolution to the
 998 microscale as per the definitions at the start of this section. However, the authors
 999 refer to their system as operating on the mesoscale, based on the processes modelled
 1000 not being at particle size.

1001

1002 Ecoclimates from Pałubicki et al.[Pał+22] provide a system capturing the bi-
 1003 directional feedback cycles between the atmosphere and the ecosystem that result
 1004 in microclimates and specific vegetation/biome patterns. The authors use temper-
 1005 ature and moisture maps to exchange values from the atmospheric and the soil and
 1006 vegetation sub-models. Reducing precipitation leads to reduced vegetation cover,
 1007 which leads to reduced moisture uptake through transpiration. The system runs
 1008 at interactive rates, at between 0.5 and 8 seconds per step (where each step is one
 1009 simulated year). The resolution is fairly small at 20 metres on a side for the voxels
 1010 of the atmospheric grid. They do not take diurnal or seasonal changes into account,
 1011 and it suffers from a poor time resolution (each time step is a simulated year) for
 1012 the purpose of cloud animation.

1013

1014 Like our system, both Hädrich et al. and Pałubicki et al.’s simulations are tai-
 1015 lored to very specific phenomena. Similar to hybrid models, combined models have
 1016 potential for increased physical accuracy, as well as modelling phenomena that are
 1017 not available with single-system models. However, this comes at the cost of signifi-
 1018 cant complexity with combining multiple models, as well as increasing the overhead
 1019 for tuning parameters and initial conditions. These combined models generally also
 1020 increase the computation time required.

3.1.6 Synoptic scales

Most research on cloud modeling in computer science has focused on micro and meso-scales, sufficient for many applications, but synoptic (global) scale modeling has distinct advantages, such as for planetary visualization or providing boundary conditions for smaller-scale models. While this work focuses on meso-scale skiescapes, it has the potential to be expanded to synoptic scales and a mention of salient work in that field is warranted. Early work by Wetherald and Manabe[WM81] modeled global clouds using atmospheric parameters, though it faced limitations with seasonality and calibration. Dobashi et al.[DYN06] created an Eulerian simulation for global cloud generation and later improved it with a rendering system[DYN10] for multiple scales. Xu et al.[Xu+09a] further combined large-scale simulations with smaller-scale cloud amplification, though this method lacked fine detail and terrain interaction. Microsoft’s Flight Simulator (2020)[SM20] uses real-time weather data and global circulation simulations for high resolution skiescapes across the globe, but its heavy reliance on data and cloud computing limits its application to Earth’s weather modeling.

3.2 Volcanic plume

In stark contrast to the body of work dedicated to cloud generation in graphics, there has been little focus on the generation of volcanic plumes. All previous work for plumes in the graphics domain focused solely on the volcanic plume, with no regard to atmospheric phenomena arising from the plume interactions. In addition, all previous work utilised a uniform wind field that could, at best, be varied as a function of altitude, along with a static atmosphere for entrainment interactions. Mizuno et al.[MDN02b] provided the first simulation of the volcanic plume for graphical applications, using a full 3D Eulerian grid combined with semi-Lagrangian advection. Using a grid of $100 \times 100 \times 100$ voxels, they perform a simulation of the volcanic plume based on the Navier-Stokes equations for fluid dynamics.

By applying procedural buoyancy, "vorticity confinement", and wind forces before advection, their system produces plumes that appear relatively realistic. This procedural basis allows their system to retain a fair degree of control over the plume by varying the parameters of the eruption, such as a narrower, taller cone if the eruption is more powerful. However, their resolution is coarse, with little or no fine detail being produced. Despite using an efficient Coupled Map Lattice solver for the Navier-Stokes pressure projection, performance is slow. However, a more modern implementation on a GPU would result in a significant performance increase. The system is limited to the plume, ignoring related atmospheric phenomena, and the uniform wind field is not physically accurate. The procedural nature of the buoyancy forces leads to artefacts and behaviour that isn’t physically accurate, nor does it account for turbulent effects arising from the interaction of the plume with the atmosphere. Figure 3.3(1) depicts an example render from Mizuno et al.’s system, utilising a rendering system created by Dobashi et al.[DYN02].

To address some of the issues arising from the procedural nature of Mizuno et al.’s[MDN02b] plume dynamics, Mizuno et al.[Miz+03] created a model that simu-

1066 lates the air and the hot ejecta from the eruption as two separate fluids. The hot
1067 ejecta (referred to as “magma” by Mizuno et al.) is the mix of pyroclasts and hot
1068 gas that is erupted. This magma has an initial velocity and entrains atmospheric
1069 air as it rises. A buoyancy force is calculated based on the density in a voxel, which
1070 in turn is determined by the temperatures and masses of the atmosphere and the
1071 magma components in the voxel. For performance, Mizuno et al. make the sim-
1072 plifying assumption that the temperatures remain constant, since that allows the
1073 density calculation to be reduced to a preprocessed lookup table. Using a grid of
1074 $150 \times 150 \times 150$ voxels, Mizuno et al. simulate volcanic plumes in greater variety than
1075 the earlier work ([MDN02b]) merely by increasing or decreasing the temperature dif-
1076 ference between the two fluids. This approximates the force of the eruption, leading
1077 to noticeably different behaviour in the plume. This work can generate physically
1078 accurate and realistic shapes for volcanic plumes across a range of eruptions. How-
1079 ever, the performance is slow despite the simplification, and this limits resolution.
1080 Like with the earlier work, this system cannot generate fine detail due to the coarse
1081 resolution of the voxels (each voxel being 100 m on a side). The system is limited to
1082 the volcanic plume and does not take atmospheric phenomena into account. Also,
1083 similar to the earlier work, Mizuno et al. use a uniform, procedural wind field, which
1084 is not physically accurate. Figure 3.3(2) shows a plume generated by this system.

1085

1086 In a departure from earlier work, Lastic et al.[Las+22] proposed a model util-
1087 ising cylinders following Lagrangian dynamics stacked along the centre line of the
1088 volcanic plume. These cylinders are enhanced by a hierarchy of rotating spheres
1089 and subspheres, which are used in a procedural upsampling step to introduce fine
1090 resolution detail. Our model is partly inspired by this work and we direct readers to
1091 chapter 5 for more detail on the implementation of this model. Lastic et al.’s model
1092 can generate a variety of volcanic eruptions, and is the first to explicitly consider
1093 pyroclastic flows and subsequent minor plumes for graphics. Combined with a fair
1094 degree of control available to the user and impressive simulation performance, it can
1095 be considered the state of the art. However, like its predecessors, it relies on a uni-
1096 form windfield, and does not consider any interaction with a dynamic atmosphere.

1097

1098 While not directly related to volcanic plumes, work on simulating explosion
1099 plumes[YOH00; KK14] and smoke requires similar examinations of entrainment as
1100 occurs in volcanic plumes. A full review of the field is beyond the scope of this
1101 work and we direct the reader to a review by Thuerey et al.[TKP13].The majority
1102 of work on smoke in computer graphics used Eulerian grid-based methods[KV84;
1103 FM97; Sta99; FSJ01; WB07]. These have much the same characteristics as the
1104 work on volcanic plumes by Mizuno et al.[MDN02b; Miz+03], with computationally
1105 intense simulations at fairly coarse resolutions, aiming to model the physical
1106 processes resulting in smoke. Shi and Yu[SY05], following up on work by Treuille
1107 et al.[Tre+03] and Fattal and Lischinski[FL04], provide an interesting contrast with
1108 a system to constrain smoke (generated by a fluid simulation) to transform into
1109 specific, user-selected shapes.

1110

1111 Several extensions aimed at reducing the computation required for realistic smoke
1112 have been devised, such as moving the simulation grid to follow the smoke[Sha+04;
1113 Zar20], or using dynamic grids[Pat+05]. As with cloud simulation, using a coarse

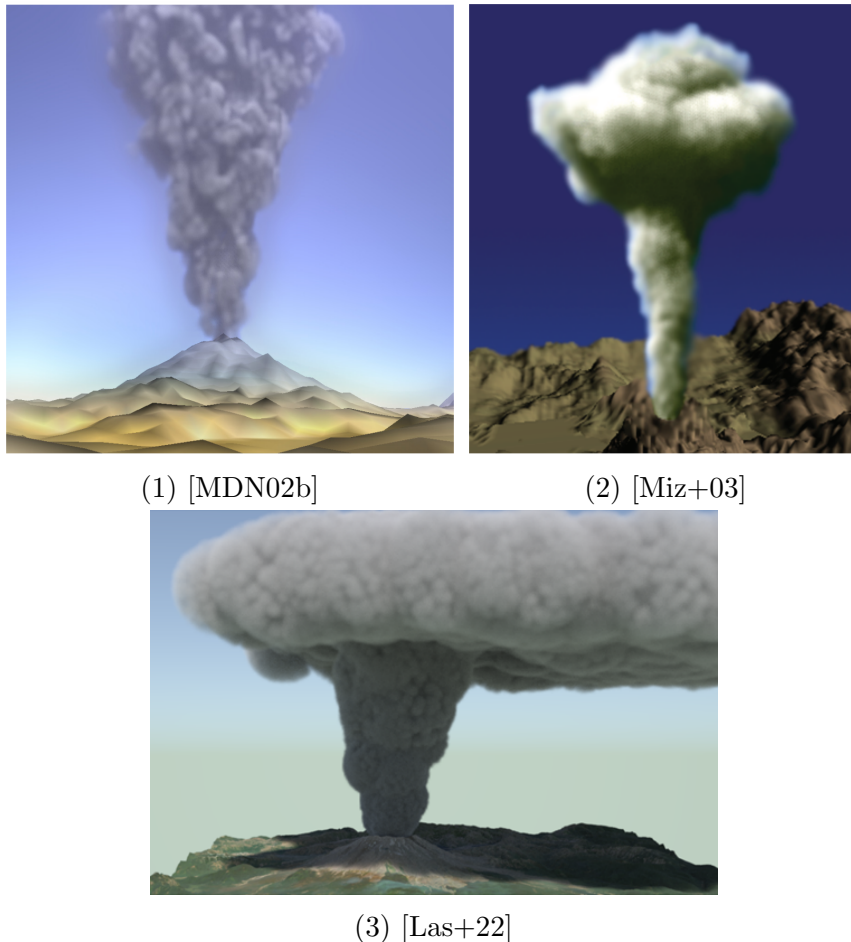


Figure 3.3: Example renders from three models of the volcanic plume. (1) shows a render from Mizuno et al. [MDN02b] using a 3D simulation with procedural buoyancy forces. (2) shows a render from Mizuno et al.’s [Miz+03] modelling the air and hot ejecta as two different fluids. (3) gives an example from Lastic et al.’s [Las+22] system using Lagrangian slices and a hierarchy of spheres and subspheres for procedural upsampling.

1114 simulation with procedural upsampling for detail has been used for smoke [SB08]. Re-
 1115 cently, using a machine learning approach for both solving and upsampling fluid sim-
 1116 ulations for smoke animation has become an extremely active field [CT17; Xie+18;
 1117 EUT19; Kim+19; Wer+19; Dry+22; Den+23]. This is by no means an exhaust-
 1118 ive review, and reviews of this narrow field by Chen et al. [Che+21] and Giraldi et
 1119 al. [Gir+23] can be consulted.

1120 3.2.1 Geophysical simulations

1121 Compared to the relative dearth within the graphics research, volcanic eruptions
 1122 and the associated plumes has been an area of intense research in the geophysical
 1123 sciences. These works typically model the plume as a multiphasic flow, with the
 1124 domain discretised as a regular 3D grid [DNM93; Obe+98; DR06; Ong+07; CEB16;
 1125 Ner+21]. This is akin to the system described by Mizuno et al. [Miz+03], though
 1126 with a much larger domain and significantly more accurate physics models. This
 1127 extends to the inclusion of a much greater range of effects, including but not limited

1128 to volcanic jets, collapsing plumes, ballistic ejecta, and convective plumes. For a
1129 comprehensive review of the field, we direct the reader to a review by Kavanagh et
1130 al.[KEM18].

1131

1132 In addition to plume models, related atmospheric phenomena has been studied
1133 in detail. Volcanic eruptions are responsible for the transport of significant quanti-
1134 ties of water[GBW97; Mil+22; Kha+22], which can have effects on cloud formation
1135 in the local area, as well as long term effects on a global scale. The entrainment of
1136 ambient air into the plume, the crucial component that dominates the behaviour of
1137 all plumes, is examined in detail. Morton et al.[MTT56] provided the first model
1138 of entrainment, which has informed all plume and most smoke models since. Work
1139 is continuing in this field, with recent models focusing on microscale interactions
1140 within the plume[SK10; Bon+15; CON16].

1141

1142 The aggregation of ash and other particles within the plume, which we use to in-
1143 form our ash rain implementation, remains an active area of research[GL94; BBD12;
1144 Hof+23]. This aggregation contributes significantly to the formation of volcanic
1145 lightning, which remains under investigation both in modelling and in laboratory
1146 conditions[Jam+08; Cim+14; Van+16; McK+21; Smi+21; CG22]. Many of these
1147 models aim to characterise the relationship between the amount of lightning and
1148 the amount of ash ejected by the volcano.

1149

1150 The models produced in this field, while as physically accurate as possible,
1151 are not suitable for use in computer graphics. The computation required for the
1152 physical accuracy in these models is exorbitant, often requiring days or weeks to
1153 execute, even on hundreds of CPU cores[Suz+16]. With a large domain (up to
1154 $350km \times 350km \times 50km$), the grid sizes are still large, ranging from 300 metres to
1155 4 000 metres on a voxel side. While this is sufficient for chemical composition and
1156 thermodynamic analysis[Tex+05; Suz+16], it cannot replicate the fine-scale turbu-
1157 lence that is the defining visual characteristic of volcanic plumes.

1158

1159 However, these studies serve as a rich source for validating our contribution,
1160 as described in section 7.1. In particular, evaluating the accuracy of the emergent
1161 properties of our system, such as ash deposition, plume shape, and vertical velocity,
1162 benefits greatly from the work in the geophysical sciences.

1163 **3.3 Summary**

1164 Significant effort has been dedicated to the modelling and animation of clouds within
1165 computer graphics, at multiple scales and with a wide variety of methods. Within
1166 the last 10 to 15 years, hybrid methods have become particularly prevalent with
1167 results notable for the computational performance, and the visual accuracy and va-
1168 riety of clouds generated. With the advent of machine learning and a variety of
1169 generative models, the future of cloud animation will likely be with hybrid models
1170 incorporating some degree of generative upsampling. Importantly for this work, no
1171 previous cloud modelling method has dealt with the extremes of temperature and
1172 altitude that volcanic plumes generate.

1173

1174 In contrast, volcanic plumes on their own have received little attention, with only
1175 three previous methods in computer graphics. Notably, all three previous methods
1176 utilised a static atmospheric boundary, something that this work aims to address.

1177 Chapter 4

1178 Weather modelling

1179 In this chapter we explain our model of the atmosphere and its physics, how trans-
1180 port within the atmosphere is modelled, and how the different layers in the atmo-
1181 sphere exchange quantities.

1182

1183 There were a number of different options available for modelling the atmosphere,
1184 as detailed in section 3.1. While the choice of plume model was simple given the
1185 minimal options presented, the weather and cloud model presented more variety. We
1186 required a system that was performant, produced realistic clouds, and that could be
1187 coupled with a volcanic plume solution. This last requirement immediately elimi-
1188 nated the procedural methods (section 3.1.1), and severely constrains the potential
1189 for data-driven methods (section 3.1.3).

1190

1191 Examining the options for simulation (section 3.1.2) or hybrid models (section
1192 3.1.4), Vimont et al.’s system[Vim+20] stood out – it is fast, it produces realistic
1193 clouds, and utilises a coarse simulation with concrete values for temperature and
1194 moisture, making it suitable for a coupling with Lastic et al.’s volcanic plume sim-
1195 ulation[Las+22].

1196

1197 Examining the source code provided by Vimont et al., we found that much of
1198 their interactions are dependent on a hard-coded grid structure. However, given
1199 that most operations within their system occur on a per cell basis, this dependency
1200 is unnecessary and hinders adaptations of their code. Therefore, we reimplemented
1201 their system to remove this dependency. As part of the reimplementation, a major
1202 issue with their fluid dynamics treatment was encountered. While very fast and
1203 efficient, it does not take solid boundaries into account. This results in significant
1204 inaccuracies.

1205

1206 Therefore, we implemented our own fluid dynamics treatment (described in sec-
1207 tion 2.4), which adds a significant computation cost, but does provide much more
1208 realistic fluid behaviour. Despite the added cost from our fluid dynamics treatment,
1209 combined with optimisations and parallelisation in other components our imple-
1210 mentation still showed a two to three times speedup compared to that of Vimont et
1211 al.




Figure 4.1: The layered structure of the atmospheric model, with two air layers and cells within these layers. The cells coloured green and yellow indicate cells that are vertically adjacent to each other. The cells in air layer 2 that are coloured brown indicate some of the cells immediately above the terrain, where the terrain impinges on air layer 1.

1212 4.1 Model structure

1213 We model the troposphere (that is, the lowest layer of the atmosphere) as several
 1214 2D fluid layers, recognising the tendency for the atmosphere to be separated into
 1215 distinct horizontal layers. Each cell in a layer consists of a velocity vector and several
 1216 scalar quantities as detailed in table 4.1. These quantities need to be transported
 1217 accurately according to the local wind field, represented by the 2D velocity vector.

1218
 1219 To complete our atmospheric model, we have vertical exchange layers between
 1220 each horizontal layer to model the interaction between the different horizontal lay-
 1221 ers, and a single surface layer representing the terrain. Figure 4.1 shows the layered
 1222 nature of our model. The cells coloured green and yellow indicate cells that are
 1223 vertically adjacent to each other. The cells in air layer 2 that are coloured brown
 1224 indicate some of the cells immediately above the terrain, where the terrain impinges
 1225 on air layer 1.

1226
 1227 A full overview of all components and quantities transferred by the components
 1228 can be found in figure 4.2, while algorithm 1 shows the substeps within a single time
 1229 step of the complete weather subsystem.

1230 4.1.1 Horizontal air layers

1231 Each horizontal air layer consists of a velocity field, separated into components, a
 1232 temperature scalar field, a moisture scalar field, a condensed water (cloud) density
 1233 scalar field, a particulate density field, a (modified) pressure correction scalar field,
 1234 and a mask dictating where the terrain impinges on the air layer. The velocity,
 1235 pressure correction, and terrain mask fields are required for advection and describ-
 1236 ing the movement of the atmosphere, with the component-wise split of the velocity
 1237 described and motivated in section 2.4.5. The temperature, moisture, cloud density,
 1238 and particulate density fields dictate the real-time appearance of the atmosphere.




Figure 4.2: The weather subsystem consists of several different components. The different quantities are transferred between cells within a layer, or between vertically neighbouring cells in different layers, with the arrows indicating the direction of transfer. Temperature T is red, moisture M is dark blue, water W is white, ash density ρ is textured grey, and the 2D velocity U is dark green, as indicated by the legend at the top of the figure. The terrain model is brown, while the horizontal air layers are light blue and vertical exchange layers are light green. From left to right, in the order that the engines operate within a full time step, we have the terrain heating engine (§4.2.1), the atmospheric heating engine (§4.2.2), the moisture uptake engine (§4.2.7), and the precipitation engine (§4.3.1). The advection engine (§4.2.3), and obstacle friction engine (§4.2.4) are independent per cell in a layer, while the pressure correction engine (§4.2.5) calculates prerequisite values independently per cell before the matrix solve. The vertical coupling engine (§4.3.2) and the phase transition engine (§4.2.6) are the final components of the subsystem.

1239

1240 We simplify all particulate matter in the atmosphere to be either liquid water or
 1241 particulates, neglecting both ice and the wide array of sizes that particulate matter
 1242 can take with resulting differences in visual appearance. We use this simplification
 1243 as a compromise between visual accuracy and computational performance. A full
 1244 atmospheric simulation requires prohibitive amounts of computation while adding
 1245 details on the micro-scale, which would be lost in our grid discretisation.

1246

1247 The pressure correction field is used only in the pressure correction step (see
 1248 section 4.2.5) and will not be referred to again in this section. Both the pressure
 1249 correction field and the obstacle mask are dimensionless quantities. See table 4.1
 1250 for the components of each air layer and the symbols used to refer to them in the
 1251 rest of the chapter.

1252

Name	Symbol	Unit
North-south velocity	u_x	metres per second ($m \cdot s^{-1}$)
East-west velocity	u_y	metres per second ($m \cdot s^{-1}$)
Temperature	T	kelvin (K)
Moisture	M	kilogram per cubic metre ($kg \cdot m^{-3}$)
Condensed water	W	kilogram per cubic metre ($kg \cdot m^{-3}$)
Particulate density	ρ	kilogram per cubic metre ($kg \cdot m^{-3}$)
Obstacle mask	O	–
Pressure correction	P	–
Height above reference	h	metres (m)
Thickness	δ	metres (m)

Table 4.1: A list of components of each horizontal air layer, their symbols, and their units. Since both the obstacle mask and the pressure correction field are dimensionless quantities, no unit is applicable.

1253

1254 The fluid dynamics are applied only to these horizontal layers, where they are
 1255 modelled as 2D planes. In the greater model, however, they are 3D volumes with a
 1256 thickness δ . The height above (or below) a reference height (which is the average sea
 1257 level, in our case) refers to the middle of the air layer's thickness. The layers do not
 1258 have a uniform thickness, instead the thickness depends on where the neighbouring
 layers are with the distribution of layers being determined by the user.

1259

1260 The user provides a maximum number of layers, a maximum height, and a list
 1261 of heights that are of interest, i.e., would benefit from having a horizontal air layer
 1262 at that height. For the scenarios described in chapter 7, the maximum number of
 1263 layers is 20, the maximum height is 20000 metres, and the list of heights provided
 1264 is $\{400, 1200, 2000, 2700, 3000, 3500, 4100, 4700, 5500, 6500\}$, with those heights being
 1265 in metres above sea level. The final position and thickness of air layers is determined
 1266 at initialisation, placing layers as close to the heights prescribed by the user as
 1267 possible. In order to have air layers that are equally thick above and below the
 1268 layer height, some adjustment might be required if the user's list of heights are not
 1269 equally spaced. Appendix C shows an example initialisation, based on the values
 1270 used in the scenarios described in chapter 7.

Name	Symbol	Unit
Vertical velocity	uz	metres per second ($m \cdot s^{-1}$)
Rainfall	η	kilogram per cubic metre ($kg \cdot m^{-3}$)
Particulate fall	ζ	kilogram per cubic metre ($kg \cdot m^{-3}$)
Lower height	H_b	metres (m)
Upper height	H_t	metres (m)

Table 4.2: A list of components of each vertical exchange layer, their symbols, and their units. The first three components are per cell, while the lower two components are the same across the layer.

1271 4.1.2 Vertical exchange layers

1272 For complete physical realism, a full 3-dimensional grid with advection and pressure
 1273 solves is required. However, this comes at a prohibitive computation cost, with a
 1274 single solve of a $512 \times 512 \times 20$ grid requiring upwards of 100 times the compute
 1275 time of 20 solves of a 512×512 grid. In addition, for optimal accuracy, the reso-
 1276 lution in the z direction would need to be significantly increased, further increasing
 1277 the computation cost. To avoid this substantial cost, we instead transfer quantities
 1278 between horizontal layers and model the interaction between them in a procedural
 1279 manner.

1280

1281 The vertical exchange layers are used to transport quantities between horizontal
 1282 layers. This is done through the precipitation engine (section 4.3.1) and vertical
 1283 coupling engine (section 4.3.2). Each vertical layer stretches from the middle of the
 1284 horizontal layer below it to the middle of the horizontal layer above it. Since the
 1285 horizontal layers can have differing heights, the reference height of the exchange
 1286 layer does not have to lie on the border between the horizontal layers. See figure
 1287 4.2 for a diagrammatic representation, with the horizontal air layers in blue and the
 1288 vertical exchange layer in green.

1289

1290 Each vertical exchange layer has three scalar fields. Firstly, the vertical velocity
 1291 scalar field, as calculated by the vertical coupling engine (section 4.3.2), and used
 1292 by the precipitation and vertical coupling engine. Then there are the rainfall and
 1293 particulate fall scalar fields, which are used to track the precipitation of water and
 1294 particulates respectively (see section 4.3.1). Table 4.2 lists the components of each
 1295 vertical exchange layer and their symbols used in the rest of this chapter.

1296 4.1.3 Terrain model

1297 The terrain is treated as a single layer with varying elevation, ranging from the
 1298 lowest depth below the ocean to the highest peak on land. Each terrain cell has an
 1299 elevation, a temperature, a moisture content, and a terrain type. The terrain type
 1300 informs a number of values such as albedo, emissivity, and heat capacity, with the
 1301 different types being dictated by the planet model. In our case, we utilise the values
 1302 of the Earth/Sun system. See tables 4.4 and 4.5 for details. In our implementation,
 1303 terrain cells only interact with the air layer immediately above it, which depends
 1304 on the elevation of the cell. This can be extended by other models, such as an

Name	Symbol	Value for Earth/Sun system	Unit
Radius	R	6 371 000	metres (m)
Sidereal rotation period	Ω	86 164.1	seconds (s)
Average temperature	\bar{T}	288.15	kelvin (K)
Reference height	h_0	0	metres (m)
Molecular mass of dry air	M_{air}	0.02896	kilogram (kg)
Lapse rate of dry air	Γ	0.098	kelvin per metre ($K \cdot m^{-1}$)
Specific heat capacity of dry air	c_A	1006	joule per kilogram per kelvin ($J \cdot kg^{-1} \cdot K^{-1}$)
Specific heat capacity of moisture	c_M	≈ 4000	joule per kilogram per kelvin ($J \cdot kg^{-1} \cdot K^{-1}$)
Latent heat of vaporisation	ΔH_{vap}	2257	joule (J)

Table 4.3: A list of quantities and constants for the planet, their symbols, their units, and their values in the Earth/Sun system.

Name	Symbol	Unit
Terrain type	X_i	See table 4.5
Elevation	h	metres (m)
Moisture	M	kilogram per cubic metre ($kg \cdot m^{-3}$)
Temperature	T_i	kelvin (K)

Table 4.4: A list of quantities and constants carried by each cell of the terrain, their symbols, and their units if any.

1305 ecosystem model.

1306

1307 The planet model contains all the planetary constants, such as the lapse rate
 1308 (that is, the averaged rate of change in temperature as a function of altitude), the
 1309 gravitational acceleration at the surface, or radius. These can be changed by the
 1310 user on initialisation. Table 4.3 contains a list of the planetary constants used in
 1311 the system, while table 4.5 lists the values associated with each terrain type. The
 1312 user can define any number of terrain types at initialisation, and a given cell can
 1313 change its terrain type over time.

Name	Symbol	Unit
Terrain index	X	–
Density	ρ	kilogram per cubic metre ($kg \cdot m^{-3}$)
Specific heat capacity	c_{ter}	joule per kilogram per kelvin ($J \cdot kg^{-1} \cdot K^{-1}$)
Albedo	α	–
Emissivity	ϵ	–
Is dry	d_X	–

Table 4.5: A list of quantities for each terrain type, their symbols, and their units if any.

1314 4.2 Physics within the atmosphere

1315 All energy on Earth derives either from the Sun, or in rare cases, from left over heat
 1316 from the formation of the planet or from the decay of radioactive elements. While
 1317 the latter two are negligibly small in comparison to the former, they are generally
 1318 thought to be a contributing cause to volcanic eruptions, and are therefore of inter-
 1319 est as well.

1320

1321 Energy modelling can be divided into two broad categories. First is the addi-
 1322 tion and removal of energy in the entire system, i.e., energy from the sun added to
 1323 the planet and atmosphere, and energy removed through dispersion and dissipation.
 1324 The second category is energy storage and release within the atmosphere, i.e., en-
 1325 ergy used or released during phase transition.

1326

1327 In our case, since our time domain is limited to a few hours at most, the effect
 1328 of terrain heating is limited as well. However, for any longer term simulations, this
 1329 effect is of crucial importance and we therefore include it in our system.

1330 4.2.1 Terrain heating

1331 Since the atmosphere is largely transparent to light in the visible spectrum (where
 1332 the sun radiates the majority of its energy), we make the simplification that the
 1333 atmosphere absorbs no heat from the sun directly. This leaves heating of the ter-
 1334 rain as the only avenue for energy from the sun to be added to the system. The
 1335 Earth's surface absorbs some of the energy from the sun and then radiates it out in
 1336 the infrared. The atmosphere is opaque to infrared radiation and absorbs some of it.

1337

1338 The energy that is absorbed by a given area of the terrain depends on the loca-
 1339 tion of the sun relative to the terrain, the thickness of the atmosphere relative to
 1340 the planet, how bright the sun is, the specific heat capacity of the terrain, and the
 1341 albedo of the terrain. The albedo is the ratio of incoming light that is reflected in-
 1342 stead of absorbed. For example, fresh snow has a very high albedo (0.95), reflecting
 1343 the vast majority of light that hits it. In comparison, water has a low albedo (0.2),
 1344 absorbing the majority of incident light.

1345

1346 We model the change in temperature for a terrain cell i as follows:

$$\Delta T_i = \frac{\Delta Q_i \Delta t}{c_{\text{ter}} \rho}, \quad (4.1)$$

1347 where ΔQ_i is the energy per square metre falling on i , and c_{ter} is the specific
 1348 heat capacity of the terrain. Note that c_{ter} is calculated as the proportional sum of
 1349 the actual terrain and any water it might contain. Dry terrain has a lower specific
 1350 heat capacity than moist terrain.

1351

1352 ΔQ_i is calculated as:

$$\Delta Q_i = L_i \cdot \cos Z_i (1 - \alpha_i), \quad (4.2)$$

1353 where L_i is the solar intensity at the surface, and Z_i is the zenith angle (i.e., angle
 1354 from the zenith to the sun) for i ¹. The solar intensity is approximated using the
 1355 empirical formula below[MM76], where AM is the air mass coefficient for i , λ is
 1356 the solar intensity in a vacuum at a distance of 1 astronomical unit (AU), and D
 1357 is the distance between the sun and the planet in AU. For the Earth-Sun system,
 1358 $\lambda = 1350 \text{ W} \cdot \text{m}^{-2}$ and $D = 1$ by definition.

$$L_i = 1.1 \frac{\lambda}{D} \cdot 0.7^{AM^{0.678}} \quad (4.3)$$

1359 Using the radius of the Earth as r , the air mass coefficient is calculated with the
 1360 following equation[Sch29]:

$$AM = \sqrt{(r \cos Z_i)^2 + 2r + 1} - r \cos Z_i. \quad (4.4)$$

1361 We also model heat loss through radiation to outer space, again using equation
 1362 4.1 to calculate heat lost. In this case, ΔQ_i is calculated as follows:

$$\Delta Q_i = \epsilon_i \sigma (T_i^4 - T_\Omega^4) \times R \left(1 + \frac{h_i}{h_A} \right), \quad (4.5)$$

1363 where ϵ_i is the emissivity of the terrain in cell i , σ is the Stefan-Boltzmann constant,
 1364 and T_i is the temperature of the terrain in cell i . T_Ω is the background temperature
 1365 of the universe (2.7K), and R is a user parameter determining how much of the
 1366 energy escapes, roughly analogous to the chemical composition of the atmosphere.
 1367 h_A is the height of the atmosphere, and this ratio is applied to R to model how the
 1368 atmospheric density decreases as the elevation increases.

1369 4.2.2 Atmospheric heating

1370 After the terrain is heated by absorbing the incident sunlight, it will radiate some
 1371 of this heat out in the form of infrared radiation. The atmosphere will absorb some
 1372 of it, while clouds can reflect some of the radiation back to the terrain. Some of the
 1373 radiation will escape to outer space, which we model using equation 4.5.

1374
 1375 For simplicity, we do not model cloud reflection and instead assume that a gen-
 1376 eral proportion R of the energy radiated out is absorbed. This proportion can be
 1377 tuned. For a terrain cell i , we model the change in temperature through radiation
 1378 absorbed by the air voxel i above it using equation 4.1 again. The change in tem-
 1379 perature for the air voxel uses a variant of the same equation, with the specific heat
 1380 capacity (c) and density (ρ) changed to be that of air at the given altitude, times
 1381 the constant factor R .

1382

$$\Delta T_A = \frac{\Delta Q_i \Delta t}{c_{\text{air}} \rho_{\text{air}}} R \quad (4.6)$$

1383 ΔQ_i will match in both equations, and is calculated as follows:

$$\Delta Q_i = \epsilon_i \sigma (T_T^4 - T_A^4) \quad (4.7)$$

¹Given the limited time domain, Z is fixed for all cells for the duration of a simulation.

1384 where ϵ_i is the emissivity of the terrain in cell i , σ is the Stefan-Boltzmann constant,
 1385 and T_T and T_A is the temperature of the terrain in cell i and the temperature of the
 1386 air in voxel i respectively.

1387 4.2.3 Advection

1388 To model the movement of wind and the quantities carried by it, we implement a
 1389 fluid dynamic simulation for each horizontal layer, using the equations described in
 1390 section 2.4. We iterate through each horizontal layer, and through each cell in the
 1391 layer. Each cell has a number of quantities that need to be advected, listed in table
 1392 4.6.

1393
 1394 For each point \vec{x}_i where values are stored (scalar values at the centre of the
 1395 cell, with the velocity separated and stored on the edges of the cell), we calculate a
 1396 virtual particle. Tracing the virtual particle to its position in the previous timestep,
 1397 we sample new values for the values stored at \vec{x}_i . By using a memory buffer, the
 1398 cells of the layer can be iterated through in parallel, which provides a significant
 1399 performance increase.

Name	Symbol	Unit
2D velocity	\vec{U}	$m \cdot s^{-1}$
Temperature	T	K
Moisture	M	$kg \cdot m^{-3}$
Condensed water	W	$kg \cdot m^{-3}$
Particulate density	ρ	$kg \cdot m^{-3}$

Table 4.6: The list of quantities that are advected by each 2D layer.

1400 4.2.4 Terrain friction

1401 On Earth, the movement of the air is slowed by friction with the terrain. Since
 1402 our fluid model is limited in resolution through the use of layers, an adjustment is
 1403 necessary to model this friction.

1404
 1405 We reuse the system proposed by Vimont et al. [Vim+20], with slight adjustments
 1406 to suit our grid system (see section 2.4.5). An orographic damping factor σ_i is cal-
 1407 culated for each point i where velocity is stored in the air layer grids. The elevation
 1408 at point p directly below point i ($h(p)$) is sampled and interpolated if not exactly
 1409 on a value. Since we utilise an MAC grid, velocity values are stored on the edges
 1410 of a cell, while elevation values are stored (like all scalar values) at the centre of a cell.

1411
 1412 σ_i is then calculated as follows:

$$\sigma_i = 1 - \min\left(1, \max\left(0, \frac{h(p) - h_i}{\Theta_i}\right)\right), \quad (4.8)$$

1413 where h_i is the height of the air layer, and Θ_i is the thickness of the air layer.
 1414 This damping factor is then applied to the velocity at point i . This damping leads

1415 to non-convective vertical motion, more accurately reflecting real life despite the
 1416 limited resolution. Such orographic lift effects include stacks of lenticular clouds
 1417 that form behind a mountain peak, or the long, flat clouds that form on the peaks
 1418 of long mountains, such as Table Mountain. See section 4.3.2 for a more detailed
 1419 explanation of this.

1420 4.2.5 Pressure correction

1421 After the various quantities in the horizontal layers have been advected and the
 1422 velocities dampened, the velocity vector field \vec{U} is no longer divergence-free. To
 1423 conform to the Navier-Stokes equations, a correction must be applied. Following
 1424 the method described in section 2.4.5, we iterate through each cell in a horizontal
 1425 layer and calculate the divergence for that cell, as well as the relevant coefficients
 1426 for the sparse linear system.

1427
 1428 Once all cells have been iterated through, we calculate the pressure correc-
 1429 tion terms (one per cell) using the conjugate gradient solver from the Eigen pack-
 1430 age[GJ+10]. This solver was chosen for its ease of use and high performance, being
 1431 highly parallelised. These pressures are used to calculate the corrections to the ve-
 1432 locity field as per equations 2.10 and 2.11. Each horizontal layer in the model is
 1433 iterated through sequentially.

1434 4.2.6 Phase transition

1435 A key component of the system is to accurately model the phase transitions of the
 1436 moisture in the air and the effect this has on the temperature and appearance of
 1437 the skyscape. Given the extreme conditions present in a volcanic plume, additional
 1438 care is required here compared to traditional solutions.

1439
 1440 Water should evaporate if the saturation vapour pressure is greater than the
 1441 partial pressure exerted by the moisture in the air. Conversely, if the saturation
 1442 vapour pressure is less than the partial pressure exerted by the moisture in the air,
 1443 water should condense. This phase exchange is a dynamic equilibrium that depends
 1444 largely on temperature.

1445
 1446 Within the plume, temperatures can exceed the *critical temperature* of water,
 1447 which is $647K$. This is the temperature at which water cannot remain a liquid,
 1448 regardless of the pressure. In this state, we force the liquid in a cell to immedi-
 1449 ately transition to a gaseous state, with no moderation dependent on the partial pressure
 1450 exerted by the moisture in the air.

1451
 1452 We iterate through each cell in a layer in parallel, and each air layer in turn. For
 1453 each layer, we calculate the ambient pressure at the air layer's height P_h using the
 1454 barometric formula[Atm76]:

$$P_h = P_b \left(1 - \frac{L}{T_b} h \right)^{\frac{gM}{R^*L}} \quad (4.9)$$

1455 P_b refers to the base pressure at sea level (which is 101 325 pascal on Earth), L
 1456 refers to the lapse rate (the change in temperature as a function of altitude), T is
 1457 the average temperature at sea level (288.15 K on Earth), and h is the height of the
 1458 air layer in question.

1459

1460 In the exponent, g is the gravitational acceleration ($9.8 \text{ m} \cdot \text{s}^{-2}$ on Earth), M is
 1461 the molar mass of air ($0.029 \text{ kg} \cdot \text{mole}$ on Earth), R^* is the universal gas constant,
 1462 and L is once again the lapse rate.

1463

1464 It must be noted that the pressure values generated by the advection system are
 1465 dimensionless and not physical. Therefore, those values cannot be used to calculate
 1466 the ambient pressure in a cell.

1467

1468 Instead, for each cell, we sample the temperature, the moisture, and the liquid
 1469 water content of the cell. Cells that are marked as solid cells (i.e., the terrain
 1470 impinges on the air layer in that cell) are ignored. We then calculate the saturation
 1471 vapour pressure for the given temperature using the Arden-Buck equation (equation
 1472 4.10[Buc96]), where T is the cell temperature in degrees Celsius.

$$\begin{aligned} P_s(T) &= 6.1121 \exp\left(\left(18.678 - \frac{T}{234.5}\right)\left(\frac{T}{257.14 + T}\right)\right) \text{ if } T \geq 0 \text{ }^\circ\text{C} \\ P_s(T) &= 6.1115 \exp\left(\left(23.036 - \frac{T}{333.7}\right)\left(\frac{T}{279.82 + T}\right)\right) \text{ if } T < 0 \text{ }^\circ\text{C} \end{aligned} \quad (4.10)$$

1473 This is a change from the system used by Vimont et al.[Vim+20], who used
 1474 Tetens' equation[Tet30] to calculate the saturation vapour pressure. This change
 1475 is driven by the higher temperatures encountered in the plume and its immediate
 1476 vicinity – the Arden-Buck equations prove more accurate than Tetens' equation over
 1477 a wider range of temperatures when compared to the measured values of the satu-
 1478 ration vapour pressure of water[RLB19].

1479

1480 We calculate the density of the air ρ in the cell using the ideal gas law (equation
 1481 2.1 in a molar form:

$$\rho = \frac{P_h M}{R^* T_b \left(1 - \frac{Lh}{T_b}\right)} \quad (4.11)$$

1482 Here, P_h refers to the pressure at altitude h , as calculated by equation 4.9. M is
 1483 the molar air mass, R^* is the universal gas constant, T_b is the average temperature
 1484 on the planet, and L is the lapse rate.

1485

1486 The density is used to calculate the absolute humidity H , a simple ratio between
 1487 the moisture m and the density ρ . The absolute humidity is then used to calcu-
 1488 late the partial pressure exerted by the moisture (P_m) in the cell, where P is the
 1489 atmospheric pressure in the cell:

$$P_m = P \frac{H}{0.622 + H} \quad (4.12)$$

1490 Similarly, the saturation humidity, i.e., the humidity that would be present in a
 1491 cell in equilibrium, is calculated using:

$$H_s = \frac{0.622 + P_s}{P - P_s} \quad (4.13)$$

1492 At this point, there are three possible options in the cell. The liquid water can
 1493 evaporate, the liquid water can boil, or the moisture can condense. Depending on
 1494 the atmospheric pressure, the partial pressure exerted by the moisture, and the sat-
 1495 uration vapour pressure, we determine which of these options should occur.

1496
 1497 If the saturation vapour pressure P_s is greater than the atmospheric pressure P ,
 1498 all the liquid water in the cell will *boil* and be converted to moisture. If the partial
 1499 pressure P_m is less than the saturation vapour pressure, the liquid water will par-
 1500 tially *evaporate*. Lastly, if the partial pressure is greater than the saturation vapour
 1501 pressure, the moisture will partially *condense* into liquid water.

1502
 1503 For evaporation and condensation, the difference ΔH between the saturation
 1504 humidity H_s and the absolute humidity H is calculated and used to determine the
 1505 mass of water that changes phase. This is multiplied by a user-determined parame-
 1506 ter J for more realistic behaviour. In our case, J is set to 0.1 for both evaporation
 1507 and condensation. This parameter is required given that our approximation is a
 1508 discretisation of a continuous process with a relatively large time step. Without this
 1509 modulation, a cell can oscillate between mostly gaseous vapour and mostly liquid
 1510 moisture at each time step, which is unrealistic and does not approach the equilib-
 1511 rium that is observed in nature.

1512
 1513 The final part of the phase transition engine is the heat exchange component. To
 1514 change phase from a liquid to a gas, energy is required, and this energy is subtracted
 1515 from the temperature in the cell. Similarly, to change phase from a gas to a liquid,
 1516 energy must be removed from the gas, which is added to the temperature in the cell.

1517
 1518 Specifically, the change in temperature is calculated as follows:

$$\Delta T = \frac{L_v \Delta H}{c_A \rho}, \quad (4.14)$$

1519 where L_v is the enthalpy of vaporisation for water ($2257 \text{ J} \cdot \text{g}^{-1}$), and c_A is the
 1520 specific heat capacity of dry air ($1.0035 \text{ J} \cdot \text{g}^{-1}$). This ΔT is then added to the
 1521 temperature of the cell.

1522 4.2.7 Moisture uptake

1523 Similar to the phase transition mechanics, the atmosphere is capable of taking up
 1524 moisture from the terrain below, or condensing out onto the terrain if it is saturated.

1525
 1526 Specifically, if the partial pressure exerted by the moisture in the air is less than
 1527 the saturation vapour pressure, a portion of the liquid water in the terrain can evap-
 1528 orate and be added to the atmosphere. Inversely, if the partial pressure exerted by
 1529 the moisture in the atmosphere exceeds the saturation vapour pressure, the moisture

1530 can condense onto the terrain.

1531

1532 For each terrain cell, we sample the air cell directly above it. We use equations
1533 4.9, 4.10, and 4.11 to calculate the ambient pressure, the saturation vapour pressure,
1534 and the density of the air in the air cell. Similarly, equations 4.12 and 4.13 are used
1535 to calculate the partial pressure exerted by the moisture in the atmosphere and the
1536 saturation humidity.

1537

1538 As in the phase transition engine, the saturation vapour pressure and partial
1539 pressure of the moisture in the air are compared to determine if evaporation or
1540 condensation should take place. The change in temperature resulting from this
1541 phase transition is calculated using equation 4.14, with the specific heat capacity
1542 changed to be that of the terrain instead of the air.

1543 4.3 Interlayer communication

1544 While the previous sections set out how we modelled the atmosphere within the
1545 horizontal layers, real atmospheric phenomena are rarely limited to single layers.
1546 Therefore, we need to model the transfer of quantities such as temperature or mois-
1547 ture between different layers. Several phenomena require interaction between layers
1548 or extend across multiple layers, such as convection or orographic lift. Convection
1549 is a major component in cloud formation and cloud types, along with being a key
1550 component of the chain of events that lead to wind forming. In particular, we focus
1551 on the modelling of precipitation, as a one-way phenomena from higher layers to
1552 lower, and the bidirectional exchange of quantities due to temperature and pressure
1553 differences.

1554 4.3.1 Precipitation

1555 Precipitation refers to the products of condensed moisture in the atmosphere falling
1556 from the clouds due to gravity. In our case, we focus only on two forms of precipi-
1557 tation: water, known as rain, and volcanic ash mixed with water, referred to as ash
1558 rain.

1559

1560 We model precipitation using a physical approximation of coalescence (or aggre-
1561 gation, when referring to ash rain). Small droplets coalesce together to form bigger
1562 droplets, until the droplets grow large enough to overcome air resistance and elec-
1563 trostatic forces keeping the cloud together.

1564

1565 This method contrasts with the probabilistic models common in meteorology
1566 [Søn+20; Rav+21]. This is needed to accurately transfer physical quantities be-
1567 tween the discrete layers using our procedural approach to inter-layer exchange (see
1568 section 4.3.2).

1569

1570 We use the same method for both rainfall (water droplets) and ash rain (fine
1571 ash mixed with water). For each cell in the system, we sample the horizontal and
1572 vertical velocities in the cell. We separate horizontal F_{hor} and vertical F_{ver} scaling
1573 factors, since the vertical velocity (uz) is a procedural rather than physical value.

Parameter	Rainfall	Ash rain
F_{hor}	0.5	0.5
F_{ver}	0.1	0.1
Ω_{base}	2	0.02
F_L	0.5	0.001

Table 4.7: A list of the scaling factors used for the two precipitation modes.

1574 The magnitude of this velocity is used as an approximation for turbulence, which
 1575 influences the collision rate and therefore the total precipitation.

1576

1577 We calculate potential precipitation ρ_{pot} using the current density ρ_i of precip-
 1578 itable matter in the cell scaled by a coalescence factor C_F , according to:

$$\rho_{\text{pot}} = \rho_i \times C_F, \quad (4.15)$$

$$C_F = \min(|\vec{H}_i| \times F_{\text{hor}} + |uz_i| \times F_{\text{ver}}, 1.0). \quad (4.16)$$

1579

1580

1581

If this potential precipitation exceeds a calculated resistance Ω_{fall} , a portion of
 the precipitation $\Delta\rho$ falls to the layer below, with scaling factors from the timestep
 (Δt) and the height between the layers (h) to prevent instantaneous transport:

$$\Omega_{\text{fall}} = \Omega_{\text{base}} + uz_i \times F_L \quad (4.17)$$

$$\Delta\rho = \begin{cases} \rho_{\text{pot}} \cdot F_{\text{base}} \cdot \Delta t \times \frac{1}{h} & \text{if } \rho_{\text{pot}} > \Omega_{\text{fall}} \\ 0 & \text{otherwise} \end{cases} \quad (4.18)$$

1582

1583

1584

1585

1586

If the precipitation reaches the ground (i.e., the cell below is impinged on by
 terrain), it is removed from the atmospheric system and added to the terrain. To
 differentiate between rainfall and ash rain, the scaling factors F are changed to
 suit the specific precipitation characteristics. Table 4.7 lists the values we used in
 practice for the various scaling factors.

1587 4.3.2 Layer exchange

1588

1589

1590

1591

1592

1593

1594

The bidirectional exchange of quantities between layers is divided into two com-
 ponents: the convective (driven by buoyancy and density differences) and non-
 convective (driven by pressure differences between layers). In practice there is a
 minor effect from the Coriolis effect as well, with westward moving air rising and
 eastward moving air descending. However, this occurs on massive scales, many hun-
 dreds of kilometres across, far greater than we consider in this work. Therefore, we
 focus only on the convective and pressure-driven vertical motion.

1595

1596

1597

1598

1599

In a full 3D fluid solver, this movement would arise naturally from the solutions
 to the fluid equations. Since we have discrete horizontal layers, we must explicitly
 calculate the vertical velocities and the attendant consequences. For performance,
 we do this in a procedural manner.

1600 Convective

1601

1602

The majority of vertical motion in the atmosphere is convective in nature, driven
 by the differing densities between hot and cold air.

1603

1604 For each cell in a vertical layer, we sample the cells above and below for tem-
 1605 peratures (T_A , T_B), moisture (M_A , M_B), and altitude (H_A , H_B). We calculate the
 1606 barometric pressure for the layers above and below (P_A , P_B) using equation 4.9.
 1607 Using the pressure and temperature, we calculate the quantity of air per volume in
 1608 the cell:

$$N_l = \frac{P_l}{R^* T_l}, \quad (4.19)$$

1609 where l refers to a specific layer, either above or below the vertical exchange layer,
 1610 and R^* is the universal gas constant.

1611

1612 These values are then used with the moisture values sampled to calculate the
 1613 specific humidities:

$$Q_l = \frac{M_l}{N_l} \quad (4.20)$$

1614 This brings us to the moist static energy equation, which describes the rough
 1615 amount of enthalpy present in an air parcel[WH06].

$$MSE_l = c_A \cdot T_l + g \cdot H_B + L_v \cdot Q_l \quad (4.21)$$

1616 c_A refers to the specific heat capacity of air, g is the acceleration due to gravity, and
 1617 L_v is the latent heat of vaporisation of water moisture. The difference between the
 1618 MSE of the layers above and below gives an approximation of the vertical velocity
 1619 between the two layers in that cell.

1620

1621 Therefore, the vertical velocity uz_i is calculated as the difference between the
 1622 MSEs of the layers, multiplied by a scaling factor C_{cv} :

$$uz_i = C_{cv}(MSE_B - MSE_T) \quad (4.22)$$

1623 The scaling factor provides user control over how energetic the atmosphere is, with
 1624 greater values leading to more convection.

1625

1626 Vertical movement will occur from the area with higher enthalpy to the area
 1627 with lower enthalpy. This is a dynamic equilibrium much like the phase transition.
 1628 The velocity calculated is not preserved through multiple simulation steps, instead
 1629 being calculated anew every step.

1630 **Non-convective**

1631 While the pressure values calculated by the pressure correction system are non-
 1632 physical and unitless, they are equivalent between different layers. Therefore, we
 1633 can calculate a vertical velocity induced by a pressure difference between layers.

1634

1635 For example, if a layer is impinged on by terrain such as a hill or mountain, a
 1636 high pressure area forms at the windward side. This high pressure zone leads to
 1637 orographic lift since the layer above does not have the same obstacle and thus a
 1638 lower relative pressure.

1639

1640 For each cell in the convection layer, we sample the pressure values in the hori-
 1641 zontal layers above and below. We normalise these values to ensure that the relation
 1642 between the layers can be established. We calculate the change in vertical velocity
 1643 (from the initial velocity calculated by equation 4.22) due to pressure differences
 1644 (Δuz_{iP}) as the difference between the normalised pressure in the cell above (\tilde{P}_{i_T})
 1645 and in the cell below (\tilde{P}_{i_B}) multiplied by a scaling factor C_P .

$$\Delta uz_{iP} = (\tilde{P}_{i_T} - \tilde{P}_{i_B}) \times C_P \quad (4.23)$$

1646 Smoothing

1647 Since the vertical motion is a continuous process in a dynamic equilibrium, the
 1648 discretisation we apply to simulate it can result in oscillations forming across time
 1649 steps with sharp discontinuities between neighbouring cells. To dampen this effect,
 1650 we apply a 2D moving window average smoothing to the vertical velocity field in each
 1651 layer. The moving window has a radius of 3 cells. This size is chosen as a compromise
 1652 between performance and effectiveness, given the computation requirements of $O(n)$
 1653 for n cells. Since the distance between discontinuities is small, a small window
 1654 size suffices in this case. The simple moving average was chosen for its ease of
 1655 implementation and adequately smoothed out the discontinuities.

1656 Vertical transport

1657 Once the vertical velocity has been calculated, we transport scalar values between
 1658 the horizontal layers. For a scalar S in cell i , we transport the difference between the
 1659 scalar values (see table 4.1) of the above and below cells multiplied by a coefficient
 1660 that is scaled by the simulation timestep and the height between two layers.

1661 We calculate the coefficient according to the following equation:
 1662

$$C_i = 1 - \exp\left(\frac{-|uz_i|\Delta t}{H_A - H_B}\right) \quad (4.24)$$

1663 The amount of the scalar values transported is then calculated (for each scalar
 1664 quantity S) as:

$$\Delta S = C_i \times (|S_B - S_A|) \quad (4.25)$$

1665 ΔS is added to the destination layer in the direction of the vertical velocity, and
 1666 subtracted from the source.

1667 Inflow and outflow approximation

1668 The final component of the vertical exchange system is to model the inflow and
 1669 outflow of air in convective columns. Rising air will result in a low pressure area
 1670 directly below it, and a high pressure area above it.

1671 We model this per cell in the horizontal air layers. For each cell, we sample the
 1672 vertical velocity immediately above and below the cell. We calculate the net flow
 1673 into the cell by subtracting these values from each other:
 1674

$$uz_{\text{net}} = uz_B - uz_A \quad (4.26)$$

1675 With this, we then calculate a change in horizontal velocity (Δu_x and Δu_y) using
 1676 the net flow and a scaling factor.

$$\Delta u = C_u \times uz_{\text{net}} \quad (4.27)$$

1677 If the net flow into the cell is positive, each velocity value on the cell (see section
 1678 2.4.5) will be modified by Δu for a greater flow out of the cell. Similarly, if the net
 1679 flow is negative, each velocity value will be modified by Δu for less flow out of the cell.

1680

1681 It must be noted that these modifications affect the 2D velocities of the horizontal
 1682 layers, and thus the velocity fields are no longer divergence free (violating equation
 1683 2.4). The advection scheme we use is more sensitive to instabilities than the one
 1684 used by Vimont et al.[Vim+20]. This sensitivity is overcome by running the vertical
 1685 coupling step at a lower frequency than the advection step (see algorithm 1). If the
 1686 vertical coupling step was run at the same frequency, the velocity fields would never
 1687 be in a divergence-free state. By allowing a few steps where the velocity fields are
 1688 divergence free, the inherent stability of the system (as described in section 2.4.4)
 1689 can be maintained, albeit with less guarantees than if the vertical coupling was
 1690 integrated as part of the pressure solve. In practice, we found that a ratio of one
 1691 vertical coupling step to three advection steps was enough for indefinite stability.

Algorithm 1 The steps for a full time step of the complete weather subsystem

```

N ← MaxSteps                                ▷ Maximum number of steps can be
x ← 0                                           ▷ time bound or it can be set by the user.
K ← CellCount                                ▷ Cell count is the number of cells in the grid
Δt ← TimeStep                                 ▷ The timestep used for most engines
ΔtVCE ← VCETimeStep ▷ The timestep used for the vertical coupling engine
while x < N do
  Step terrain heating engine by Δt             ▷ Section 4.2.1
  Step atmospheric heating engine by Δt         ▷ Section 4.2.2
  Step moisture uptake engine by Δt            ▷ Section 4.2.7
  Step precipitation engine by Δt              ▷ Section 4.3.1
  Step advection engine by Δt                 ▷ Section 4.2.3
  Step obstacle friction engine by Δt          ▷ Section 4.2.4
  Step pressure correction engine by Δt        ▷ Section 4.2.5
  if x × Δt mod ΔtVCE = 0 then
    Step vertical coupling engine by ΔtVCE     ▷ Section 4.3.2
  end if
  Step phase transition engine by Δt           ▷ Section 4.2.6
  x ← x + 1
end while

```

Chapter 5

Plume Modelling

To produce a visually realistic plume, we build on the proxy-geometry approach proposed by Lastic et al.[Las+22]. As mentioned in chapter 3, their approach assumes the atmosphere to be a static boundary condition. In this chapter, we provide an overview of Lastic et al.’s method, set out the changes necessary to couple this model with the atmospheric model described in the previous chapter, and describe the additions made to represent some of the visual phenomena present in volcanic eruptions.

5.1 Plume model

We only provide an overview of the model proposed by Lastic et al., and refer the reader to the original paper[Las+22] for a full account.

5.1.1 Plume slices

The core dynamics of the eruption are modelled using cylindrical slices emitted from the vent. Each slice carries attributes for position, size, temperature, ash density, and moisture. Each slice is modelled as a Lagrangian particle, with buoyancy, friction, gravity, and wind forces acting upon it. Equation 5.1 is used to compute the acceleration applied to a given slice, where Δt is the timestep and M is the mass of the slice. \vec{F}_b is the buoyancy force, \vec{F}_f is the friction force, \vec{F}_g is the gravitational force, and \vec{F}_w is the force applied by the local wind. Figure 5.1 depicts the slice and sphere model we use.

$$\Delta\vec{v} = \frac{\Delta t}{M} \left(\vec{F}_b + \vec{F}_f + \vec{F}_g + \vec{F}_w \right) \quad (5.1)$$

In addition, we model the entrainment of the ambient air into the eruptive column. This entrainment of ambient air is what causes the plume to rise (due to a reduction in density as the entrained air heats up) and is the origin of many of the visual features of the plume, such as the billowing appearance. Each slice includes a volume of the surrounding air at each timestep. The mass of air entrained is dependent on the velocity of the slice and the direction of the local wind. The entrained air changes the temperature of the slice, since the ambient air is at a lower temperature than the plume. The slice also decreases in density, to model the sedimentation of




Figure 5.1: The plume model is divided into slices, treated as Lagrangian particles, and a hierarchy of spheres that provide lower level detail. On the left, the slices and the physical dynamics influencing it are shown, along with the buoyancy (\vec{F}_b), friction (\vec{F}_f), gravitational (\vec{F}_g), and wind (\vec{F}_w) forces. On the right, a slice with attached rotating spheres (number limited for legibility) is shown. The grey spheres are the primary spheres, while the red spheres are at the second level of the hierarchy. The spheres detach from the slice after the slice reaches neutral buoyancy.

1721 ash and tephra. Using the ideal gas law (equation 2.1), we can calculate the new
 1722 mass of the slice.

1723 5.1.2 Plume spheres

1724 To model the local velocity in the plume (as compared to the global velocity of the
 1725 plume as a whole), procedural rotating spheres of multiple sizes are used. This is in
 1726 accordance with the superficial appearance of a billowing plume, with rotation from
 1727 the core of the plume to the outside due to the convective column.

1728

1729 The spheres at the largest scale are of a similar radius to the slice at that alti-
 1730 tude. In the model, spheres are placed on a slice, with their centres at a distance
 1731 from the slice centre equal to $r/2$, where r is the radius of the slice. The spheres
 1732 are spaced to ensure overlap between each other, with minor noise perturbations to
 1733 the position and size for natural variation. The spheres move with the slice, and the
 1734 rotation speed matches the speed of the air moving past on the outside of the column.

1735

1736 Additionally, subspheres are placed randomly across the surface of each sphere.
 1737 These subspheres are set to be one fifth the size of the sphere to which they are
 1738 attached to, and their motion is locked to that of the parent sphere.

1739

1740 Once the plume reaches neutral buoyancy, it will spread out laterally in response
 1741 to the high pressure zone created by the convective column. To accommodate this
 1742 different visual appearance, the large spheres are disconnected from their parent
 1743 slices once the slice reaches its peak height and allowed to spread out, advected by
 1744 the local wind field.

5.2 Sphere Enhancements

We make a number of changes compared to the original approach [Las+22] in order to accurately portray lightning and ash rain. In addition to only slices carrying temperature, ash density, and moisture values, each sphere now tracks ash and electrical charge. This is required for accurately transferring the ash to the atmosphere once detached from the parent slices. At each timestep, the values for ash, moisture, and temperature are updated by the slice and copied to the attached spheres.

The spheres and subspheres are treated differently when interacting with the atmosphere (see chapter 6). Temperature and moisture are transferred from the primary spheres to the atmosphere, while the subspheres are responsible for transferring ash to the atmosphere. This reflects the accuracy requirements, where the ash must be transferred with as much spatial precision as possible, while the moisture and temperature do not need the same level of precision.

During the umbrella phase, the plume becomes thinner vertically and elongates as it moves away from the vent. This is a result of the plume umbrella behaving like a gravity current [WK94]. This phenomena is replicated in our system by having the spheres gradually shrink in the vertical direction once detached from the slices, becoming ellipsoids. This shrinking is done as a function of time and wind speed, with more time and greater wind speed leading to a faster rate of compression.

5.3 Lightning

As one of the most striking aspects of volcanic plumes of all types and sizes, volcanic lightning has been extensively studied both in nature and in laboratory conditions [CG22]. Indeed, so prevalent is volcanic lightning across an extremely broad range of eruption types, that it is considered “an intrinsic property of volcanic ash plumes [CG22].” Lightning in any cloud system is an electrostatic discharge between two regions that are electrically charged, with the lightning temporarily neutralising these regions. To accurately mimic lightning, the mechanisms that result in electrically charged regions of the air must be explored.

While it is still not fully understood, electrification in water clouds is thought to be the result of water droplets and different forms of ice crystals colliding with each other under the strong convective turbulence in a cloud. The different types of droplets and crystals have different shapes and sizes and therefore move on different trajectories, leading to the now charged crystals accumulating in different parts of the cloud.

Several hypotheses have been accepted as likely mechanisms for electrification in volcanic plumes, and we focus on three of these in our work. First, *fractoelectrification* takes place in or near the vent as tephra forms. Tephra is formed when hardened magma or lava is shattered by the eruption. Many small fragments form and exert friction on each other, which leads to static charge formation. This mechanism is approximated by giving each slice and its child spheres an initial, constant charge as it is emitted from the vent.

1790

1791 In the jet and convective phases, friction between the different components of the
 1792 eruption column leads to a build-up of charge in *tribo-electrification*. These compo-
 1793 nents have different chemical compositions and structures, such as glass (amorphous
 1794 crystal solids) or rocks (heterogeneous aggregates of minerals). They also have dif-
 1795 ferent granularities, referring to the density, shape, and size. It must be noted that
 1796 laboratory experiments have demonstrated that moisture has a significant suppress-
 1797 ing effect on this kind of charge build-up[Ste+19; Mén+20]. We therefore add an
 1798 exponential decay based on the ambient moisture below the -20 °C isotherm¹. At
 1799 temperatures below -20 °C, very little moisture remains in the atmosphere.

1800

1801 At such higher altitudes, supercooled water droplets and ice flakes collide and
 1802 experience friction, similar to electrification in thunderstorms. This is known as *wa-*
 1803 *ter and hydro-meteor electrification*. However, due to the presence of ash particles
 1804 that serve as crystal seeds, substantially more ice can form in a plume compared to
 1805 a thunderstorm of equal temperature. Therefore, the build-up of charge is enhanced
 1806 in the volcano plume. This method of electrification is considered to be considerably
 1807 stronger than tribo-electrification[CG22], with roughly an order of magnitude more
 1808 charge.

1809

1810 To accommodate these different sources of electrification, we have each sphere
 1811 p track its own electrical charge, calculated from the sphere’s centre. This is a
 1812 dimensionless value L_p that is increased by ΔL_p at each timestep of the plume
 1813 system. We calculate ΔL_p using equation 5.2:

$$\Delta L_p = \begin{cases} k_1 \cdot \overline{A_p}(k_{hA}) \cdot \overline{\|\mathbf{V}_p\|}(k_{hV}) \cdot e^{-10\overline{M_p}(k_M)} & \text{if } T > 253.15 \text{ K} \\ k_2 \cdot \overline{W_p}(k_W) \cdot \overline{H_p + U_p}(k_{cV}) & \text{otherwise.} \end{cases} \quad (5.2)$$

1814

1815 Note that this distinguishes between the two electrification regimes below and
 1816 above the -20 °C (253.15 K) isotherm. At temperatures above the isotherm, we
 1817 include terms for ash density, sphere velocity, and the exponential decrease with
 1818 moisture content mentioned earlier. At temperatures below the isotherm, we take
 1819 into account the atmospheric velocity and the amount of liquid water in the sur-
 rounding atmosphere.

1820

1821 The bar function used in equation 5.2 denotes a capped normalisation, where for
 1822 a constant f , $\overline{f}(k) = \min(f, k)/k$. In practice, we use the values listed in table 5.1
 1823 for capping the various electrification effects.

1824

1825 The actual electrical discharge in the form of lightning is handled by the proce-
 1826 dural amplification and offline rendering system. See section 6.2.4.

1827 5.4 Embedded Clouds

1828

1829 Our atmospheric system consists of multiple horizontal layers, with a sparse vertical
 distribution. This is designed to capture the layered nature of the atmosphere in

¹This isotherm is found at 4117 m above sea level in our simulation, though its height varies in nature depending on the atmospheric composition.

Cap	k_1	k_2	k_{hA}	k_{cA}	k_{hV}	k_{cV}	k_M	k_W
Value	0.1	0.5	100	5	400	20	50	50

Table 5.1: The values used to cap various effects during the calculation of ΔL_p for each sphere p .

1830 typical weather conditions. However, this is too coarse a resolution (especially ver-
 1831 tically) to accurately model the air parcels in and immediately around the plume.
 1832 The air in or near the plume can be subject to significant uplift and extreme turbu-
 1833 lence leading to distortion.

1834

1835 The various embedded clouds described in section 2.3.1 are a consequence of the
 1836 fine scale behaviour of the plume. To mimic these features, we create deformable
 1837 sheet surfaces that are sampled vertex-wise. This sampling is done at a much finer
 1838 resolution than the atmospheric system, particularly in the vertical. Where the at-
 1839 mosphere system can have layers that are up to 2500 metres in vertical extent, the
 1840 sheets are only 100 metres thick. The number of sheets can be varied by the user,
 1841 but in practice, we found that three sheets are sufficient to replicate the appearance
 1842 of the skirt and cap clouds found in nature.

1843

1844 A sheet is created as a horizontal disc, centred above the volcano vent, and at
 1845 a specific altitude. Prior to the eruption, multiple sheets are created and spaced
 1846 vertically to represent naturally occurring layers of alternating moist and dry air.
 1847 Each vertex σ on the sheet is initialised with moisture M_σ and liquid water W_σ , both
 1848 of which are trilinearly interpolated from the nearest grid cells in the atmospheric
 1849 system.

1850

1851 To model vertical uplift, for each vertex σ , we sample the vertical velocity of all
 1852 slices and spheres in the plume system at each timestep. For each slice and sphere p
 1853 with an altitude h_p below h_σ , a weighted vertical velocity $V_p \cdot \Psi(V_p)$ is added to the
 1854 velocity of vertex σ . The weighting function $\Psi(V)$ is dependent on the distance d
 1855 between the sphere p and the vertex s , and is calculated according to equation 5.3:

$$\Psi(V) = \begin{cases} 1 & \text{if } \|V\| \leq r_s \\ S_1(1 - d) & \text{if } \|V\| > r_s \text{ \& } d = \|(\vec{V}_x/r_h, \vec{V}_y/r_h, \vec{V}_z/r_v)\| < 1 \\ 0 & \text{otherwise,} \end{cases} \quad (5.3)$$

1856 The smoothstep function is defined as $S_1(x) = 3x^2 - 2x^3$, while \vec{V} is the vec-
 1857 tor from the surface of sphere p to the sheet vertex σ . If the sample is within the
 1858 sphere's inner radius r_s , then the weight is set to unity. Outside of the inner radius,
 1859 the weight tails off according to the smoothstep function on the distance d . The dis-
 1860 tance is calculated as the normalised distance within a horizontal offset region r_h and
 1861 a vertical offset region r_v . Outside of this ellipsoid, the weight is 0. Figure 5.2 de-
 1862 picts the uplift force applied to vertices, as a 2D slice through the centre of the skirt.

1863

1864 By varying the size of the outer ellipsoid (r_h, r_v) , we can procedurally model
 1865 the differences between air parcels pushed ahead of the plume compared to parcels




Figure 5.2: The uplift force on each vertex depends on the distance from the slices and spheres of the plume. The further away from the plume the vertex is, the less the force that is applied.

1866 that are dragged along by the plume. This approximates the deformability of the
 1867 embedded cloud, where cap clouds are less deformable than bell or skirt clouds. We
 1868 continue to model the uplift until the sheet reaches a preset altitude where it is
 1869 neutrally buoyant.

1870

1871 As the sheet vertices are uplifted, adiabatic cooling is applied, using exactly the
 1872 same mechanics as the atmospheric system for condensation or evaporation (see sec-
 1873 tion 4.2.6). At higher altitude, the sheet vertices are usually at a lower tempera-
 1874 ture, with temperature determined by the ambient temperature at the vertex's loca-
 1875 tion, sampled from the atmospheric system. If sufficiently cooled given the moisture val-
 1876 ues of the vertex, some moisture condenses to add to the liquid water content. If
 1877 the liquid water is at sufficient concentration, the sample becomes visible as a cloud.
 1878 The deformation of the sheets and subsequent shape give the characteristic appear-
 1879 ance of bell, skirt, and cap clouds. Figure 5.3 depicts the three types of embedded
 1880 clouds that can be modelled by our system.

1881 5.5 Transient condensation clouds

1882 In explosions, such as bomb blasts and volcanic eruptions, the initial shock wave
 1883 can be powerful enough to cause transient condensation clouds to form in humid
 1884 conditions as the pressure wave propagates. These clouds are formed by an adia-
 1885 batic drop in air temperature as the pressure wave passes. This causes the moisture
 1886 in the humid air to condense out and form a cloud of droplets. After a few seconds
 1887 the air temperature returns to normal and the droplets evaporate again.

1888

1889 One way to model this phenomenon is to directly decrease and then restore
 1890 pressure in the atmospheric system for the affected cells. However, as mentioned
 1891 before, the pressure values we have are dimensionless and can therefore not be used
 1892 in the phase transition updates. The values calculated by the advection system are




Figure 5.3: The three variants of embedded clouds that we simulate in our framework. [1] Bell clouds form as moist air layers are lifted more in the centre of the plume and condenses out. [2] Skirt clouds are the result of the plume breaking through bell clouds and obscure the central bell. [3] Cap clouds are confined to immediately above the plume, with no layers visible around the plume as is the case for bell and skirt clouds.

1893 purely relative, with an actual magnitude dependent on the velocities of the system.
 1894 These values are not representative of the actual values found in the atmosphere, nor
 1895 does it have a comparable range to values found in the atmosphere. In contrast, all
 1896 pressure values used in calculating phase transitions require actual pressure values
 1897 in the range found in the atmosphere.

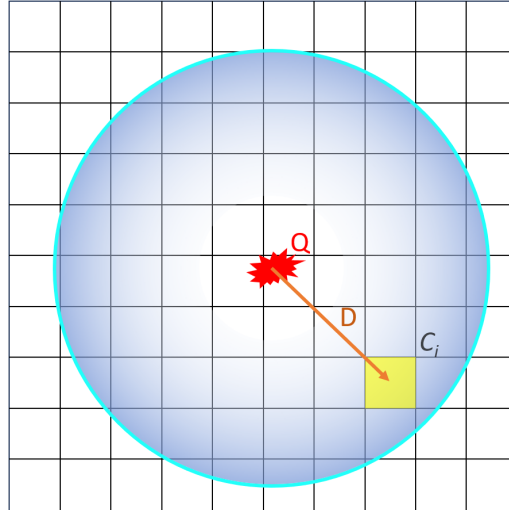


Figure 5.4: A pressure wave originating at Q (the red star) propagates outwards (represented by the cyan circle). For each cell C_i within the pressure wave, we calculate the distance D_i between C_i and Q . The temperature for C_i is reduced by N/D_i^2 to calculate T_i^* , with T_i^* increasing over time until it is restored to T_i 10 seconds after the pressure wave passes by. Darker blue indicates lower values for T^* , while white indicates ambient temperature values (T).

1898 Instead, we temporarily reduce the temperature used in the phase transition
 1899 equations (equation 4.10), again using the barometric air pressure calculated by
 1900 equation 4.9. We simulate a pressure wave centred on the volcano vent Q with an

1901 initial non-dimensional strength N . Once the wave is initiated (timing controlled by
1902 the scenario), it propagates outwards at the speed of sound. For each cell C_i that is
1903 in the spherical area encompassed by the pressure wave, we calculate the distance
1904 D_i between C_i and Q .

1905

1906 The pressure wave strength decays as the square of its radius. The adjusted
1907 temperature T_i^* is calculated by subtracting a value of N/D_i^2 multiplied by a scaling
1908 factor to fit the temperature scale. The adjusted temperature decays as time passes,
1909 and is removed completely 10 seconds after the wave first passes cell C_i , thereafter
1910 conditions are allowed to return to the normal dynamic equilibrium.

Chapter 6

Combining submodels and amplification

To accurately model the volcanic plume, a model of the atmosphere is required to capture the influence the ambient atmosphere can have on the appearance of a volcanic plume, such as embedded clouds (section 5.4) or ash rain (section 4.3.1). Lastic et al.'s original plume model used a static boundary condition, which is not accurate to reality. Similarly, many of the atmospheric characteristics present near a volcanic plume require a model of the plume to be known.

If we set our system's parameters for maximal simulation accuracy, the computation cost would be prohibitive. Therefore, we optimise and simplify as many systems as possible to achieve the best performance possible while balancing accuracy. The atmospheric system operates on a coarse scale, which does not result in a realistic appearance of the clouds from close range. Similarly, the plume, while comparatively fine-scaled, is also at too coarse a resolution to accurately capture the appearance of a volcanic plume in nature. By adding noise to the appearance dictated by the simulation models, we can achieve a significantly more realistic appearance at very little cost.

In this chapter, we describe how we combine the atmospheric and volcanic plume submodels, as well as describing the procedural amplification applied to generate the final visual results. As mentioned in section 5.3, the discharge of the electric charge is handled solely by this system and is therefore described in this chapter as well.

6.1 Model combination

In order to accurately and effectively couple the two submodels, We need to transfer physical quantities such as ash, moisture, temperature, and wind between the two submodels. Without such transfer, many processes found in nature such as entrainment or ash rain cannot be accurately modelled.

We encounter two major challenges in the combination. First, the two submodels are run at different frequencies. The volcanic plume requires smaller timesteps due to the higher velocities encountered in the eruption column. In our framework, we run the submodels at a ratio of 50 to 150 steps of the plume system ($\Delta t = 0.02$ ms)

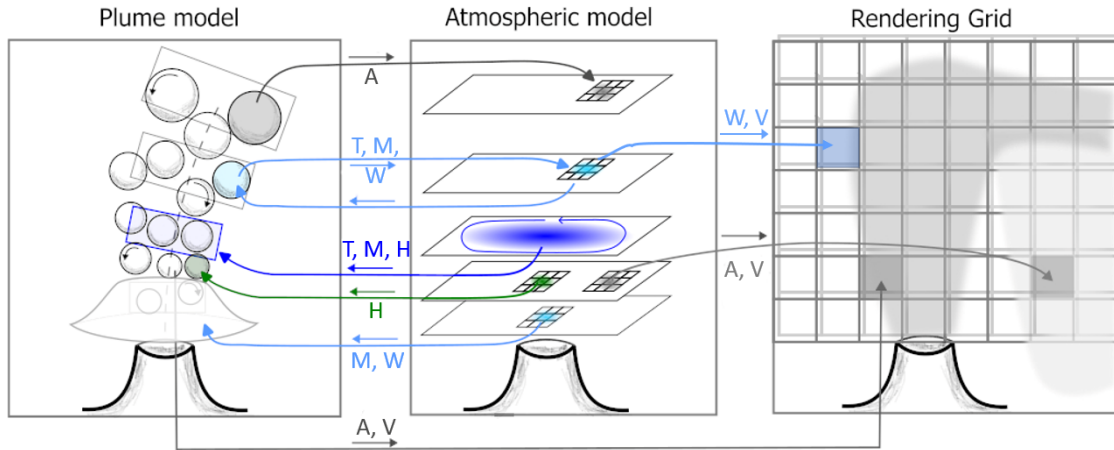


Figure 6.1: We transfer physical quantities between the two submodels, as well as the final rendering system. We transfer ash A , moisture density M , water content W , temperature T , 3D velocity V , and horizontal 2D velocity H .

1945 for each step of the atmospheric system ($\Delta t = 1 - 3$ ms). Immediately before each
 1946 step of the atmospheric system, the slower of the two, we synchronise and transfer
 1947 the physical quantities from the plume to the atmosphere. Immediately after, before
 1948 the next plume step, we transfer the physical quantities from the atmosphere to the
 1949 plume.

1950

1951 The second significant challenge deals with the different spatial structures. The
 1952 volcanic plume system, as described in section 5.1.1, is a stack of cylindrical slices
 1953 and attached spheres. In comparison, the atmosphere system is a stack of horizontal
 1954 2D layered grids.

1955

1956 To transfer quantities from the plume to the atmosphere, we iterate over each
 1957 atmosphere layer l . We find the plume slices that intersect the altitude range of
 1958 l and select the slice s that is closest to the air layer's height h_l . For each cell i
 1959 that is within $2r$ of the slice centre, we add quantities of moisture density M , water
 1960 W , and temperature T . The addition is scaled by a linear function of the distance
 1961 between the cell centre and the centre of the plume slice. When the distance is
 1962 zero, the addition is equal to the difference between the value of s and the value
 1963 of i , i.e., the cell takes the value of s . When the distance is greater than twice
 1964 the radius of s , the addition is zero, i.e., the cell remains unchanged. This trans-
 1965 fer depends only on the plume slices and does not consider the spheres or subspheres.

1966

1967 Compared with the less accurate transfer of moisture, temperature, and water,
 1968 ash needs to closely correlate with the visual appearance of the plume. Since the
 1969 plume's appearance is dictated by the spheres and subspheres, we therefore transfer
 1970 ash from the spheres to the atmosphere. At each time step of the plume system, each
 1971 sphere p calculates a projected density loss and adds it to an accumulated density
 1972 loss A_p . The ash is transferred at the same time as the other quantities mentioned.
 1973 For each sphere p , we find the closest grid cell i , and transfers A_p to i , resetting A_p
 1974 to 0. Once in the atmospheric system, the ash is subject to precipitation as described
 1975 in section 4.3.1.

1976

1977 To transfer quantities in the other direction, from the atmosphere to the plume,
 1978 we find the closest slice s to each air layer l . We sample an annular region centred
 1979 on s 's position, which roughly corresponds to the area from which air is entrained
 1980 into the eruptive column. In our implementation, this region has an inner radius
 1981 of r_s , and an outer radius of $3r_s$. We calculate the mean values for velocity V ,
 1982 temperature T , and moisture density M across each cell in the annulus, with linear
 1983 interpolation between the vertical layers above and below s . These values are then
 1984 assigned to s .

1985

1986 Figure 6.1 provides an overview of all the transfers that happen between the
 1987 different systems. Temperature T , moisture density M , and water density W are
 1988 transferred in both directions between the plume and the atmosphere. The atmo-
 1989 sphere transfers 2D horizontal velocity H to the plume, while the plume transfers
 1990 ash density A to the atmosphere. The atmosphere transfers water density W ,
 1991 ash density A , and 3D velocity V to the offline rendering system, while the plume
 1992 sends sphere ash A and 3D velocity V to the rendering system.

1993

1994 We transfer the quantities in order to mimic specific natural processes seen during
 1995 a volcanic eruption. The temperature transfer in both directions is to reproduce the
 1996 plume and atmosphere mixing that leads to localised atmospheric convection and the
 1997 convective phase of the plume. We transfer moisture and water in both directions
 1998 to model the entrainment and release of moist air leading to clouds. Horizontal 2D
 1999 velocity is transferred from the atmosphere to the plume to ensure that the plume
 2000 follows the downwind velocity field dynamically (compared to the static windfield
 2001 used by Lastic et al.[Las+22] in their original implementation). Lastly, we transfer
 2002 ash from the plume to the atmosphere to allow the ash to mix with water and model
 2003 ash rain.

2004 6.2 Procedural amplification

2005 Our simulation models are sufficiently detailed to accurately portray the dynamics
 2006 of the volcano plume and the atmosphere, as well as the interaction between the
 2007 two subsystems. However, the appearance of clouds and plumes extends to a level
 2008 of detail that the simulation cannot achieve in a reasonable amount of computation.
 2009 We therefore incorporate details of finer scales by procedurally upsampling the data
 2010 produced by the simulation.

2011

2012 We take advantage of the rich body of work in the field of procedural tech-
 2013 niques[Lew84; Per85; PN01; Ola+02; CD05; BHN07; Lag+09; Lag+10] to support
 2014 the upsampling with a wide array of auxiliary data, such as altitude, moisture, and
 2015 velocity. In particular, by keeping the velocity field as part of the upsampling, it is
 2016 possible to ensure temporal consistency in the upsampling between different frames
 2017 by backtracking the velocity field through previous timesteps.

2018

2019 We perform the upsampling through carefully directed application of procedural
 2020 noise to the data generated by the simulation. This is a field with a rich history and
 2021 we cannot cover it in detail, instead referring the reader to a survey by Lagae et

2022 al.[Lag+10]. For our purposes, we define procedural noise and give a short overview
2023 of the major methods, before describing our choice of procedural noise function and
2024 justifying the selection.

2025

2026 *Noise* as used in this context is random or pseudorandom values, usually with a
2027 probability density function that tends to be normal or Gaussian. *Procedural* noise
2028 functions refer to values that are generated by an algorithm or through code rather
2029 than from existing data sources such as an image. Some of the advantages of proce-
2030 dural noise functions, as discussed by Lagae et al.[Lag+10] and Ebert et al.[Ebe+02]
2031 make it especially suitable for our use.

2032

2033 In particular, as *continuous* functions, we can use it at any scale without re-
2034 quiring base data at a specific resolution. By being *randomly accessible*, i.e., the
2035 function can be evaluated in constant time at any point on the lattice, it is suitable
2036 for use in parallel on multi-core CPUs or GPUs, leading to increased performance.
2037 Lastly, and perhaps most importantly, procedural noise functions are *parametrised*,
2038 allowing us to tune the noise for specific use cases without having to change the
2039 underlying function. This parametrisation also ensures that a particular result is
2040 reproducible given a specific seed value.

2041

2042 While the simplest noise generation algorithm is to produce a pseudorandom
2043 number sequence, the results are not realistic and are generally considered unsuit-
2044 able for use in amplification in graphics. Instead, some form of coherent noise is
2045 preferred for its (usually) more realistic appearance. There are many methods for
2046 generating coherent noise and these are usually classified into three categories.

2047

2048 **Lattice gradient noise** functions generate noise through interpolating values
2049 or gradients at specific points on a lattice, usually the integer lattice. The defining
2050 example of this category is Perlin noise[Per85; Per02]. This was one of the first pro-
2051 cedural noise generation functions and continues to be widely used in many aspects
2052 of graphics. It calculates a pseudorandom gradient for each point based on the hash
2053 of a lattice point, with values between lattice points being the spline interpolated
2054 result of the nearest gradients.

2055

2056 There are several variants of lattice gradient noise functions. Some are cus-
2057 tomised for specific physical phenomena, such as flow noise [PN01] or curl noise
2058 [BHN07] for advective motion and turbulent flow, respectively. Other variants are
2059 based on different lattices, such as simplex noise[Ola+02], which can lower compu-
2060 tational complexity or eliminate directional artefacts.

2061

2062 **Sparse convolution noise** functions calculate noise as the sum of weighted
2063 kernels positioned randomly in the feature space. This was initially proposed by
2064 Lewis[Lew84; Lew89], who used a Poisson process convolved with a kernel to gener-
2065 ate the noise. Other variants of sparse convolution noise include spot noise[Van91],
2066 Gabor noise[Lag+09], phasor noise[Tri+19], and texton noise[GLM17]. These use
2067 different kernel types for improved control and specific use cases, such as texture
2068 mapping on parametrised surfaces.

2069




Figure 6.2: Examples of some 2D noise functions. From left to right, we show textures generated by Perlin noise[Per85], Gabor noise[Lag+09], and anisotropic noise[GZD08], covering the three major categories of noise functions.

2070 **Explicit noises** are noise functions that are calculated beforehand and used as
 2071 cached values at runtime. While not procedural techniques as defined by Lagae et
 2072 al.[Lag+10], these are often used for reducing artefacts such as detail loss or inter-
 2073 polation bias. Wavelet noise[CD05] and anisotropic noise[GZD08] are examples of
 2074 this category. It can also be used in conjunction with other methods, such as is the
 2075 case with spot noise.

2076

2077 Figure 6.2 shows 2D examples of Perlin, Gabor, and anisotropic noises.

2078

2079 Given that the movement of the plume and clouds is largely advective in nature,
 2080 we use *flow noise*[PN01] for the amplification. By varying the parameters of the
 2081 function, we can easily tune the noise for the different types of clouds we consider.
 2082 While the plume is turbulent as a whole, the scale of this turbulence is larger than
 2083 the individual spheres and subspheres (see section 5.1.2) used by the simulation
 2084 system. Therefore, the amplification does not need to portray inherent turbulence,
 2085 as that is provided by the simulation data.

2086 6.2.1 Cloud amplification

2087 The cloud data is provided as the liquid water content W in each cell in the multiple,
 2088 discrete 2D layers described in section 4.1. Any value W_i below a threshold W_{\min} is
 2089 considered to be below visibility, and therefore no clouds appear in that cell. We use
 2090 a value of $0.1 \text{ kg} \cdot \text{m}^{-3}$ for W_{\min} . This value is derived from qualitative comparisons
 2091 of different values, specific to the scenario in question. For example, if a lot of wispy,
 2092 high clouds are desired, W_{\min} can be reduced. If thick stratus clouds are the goal,
 2093 W_{\min} can be increased to make the difference in densities clearer.

2094

2095 To create a full volumetric density field, we linearly interpolate the liquid wa-
 2096 ter content between the horizontal layers. A specific cloud type (cumulus, cirrus,
 2097 or stratus) is then assigned to each cell based on the height of the layer and the
 2098 convection in the cell[Vim+20]. We use a reduced subset of the cloudtypes used by
 2099 Vimont et al., since the focus of the work is on the phenomena specific to volcanoes.

2100

2101 An issue arising from the linear upsampling of our layers is that the vertical
 2102 gradients are weaker than the comparatively dense horizontal gradients. We can
 2103 somewhat compensate for artefacts arising from this by laterally blurring the hori-
 2104 zontal layers before the upsampling.

2105

2106 To add detail to our rendering, we warp the densities with flow noise[PN01],
 2107 with parameters based on a combination of density values and the cloud type. To
 2108 ensure temporal consistency between different steps, we can calculate a displacement
 2109 $d(t)[x]$ back to the cloud source from a voxel position x . We recursively calculate
 2110 $d(t)[x]$ using the 3D wind velocity $\vec{V}(t)$:

$$d(t + \Delta t)[x] = d(t)[x - \Delta tV(t)] - \Delta tV(t) \quad (6.1)$$

2111 Over many timesteps, this backtracking can lead to discontinuities in the dis-
 2112 placement and result in rendering artefacts. We can curb these discontinuities by
 2113 relaxing the displacement, instead using a linear interpolation between $d(t + \Delta t)[x]$
 2114 and the identity. The interpolation is biased depending on the cloud type at the
 2115 voxel x , as follows:

2116

2117 We add high frequency flow noise to **Cumulus** clouds to reflect the high turbu-
 2118 lence present in these clouds. To sharpen the appearance of the cloud edges (see
 2119 section 2.2.3 for an explanation), the density is scaled up before upsampling. To
 2120 reflect the billowy appearance of cumulus clouds, the displacement interpolation is
 2121 biased towards the identity.

2122

2123 **Cirrus** clouds are windswept and wispy. The displacement interpolation is bi-
 2124 ased towards $d(t + \Delta t)[x]$ to reflect this, while the flow noise is at a lower frequency
 2125 than for cumulus clouds.

2126

2127 **Stratus** clouds are usually smooth sheets of clouds. The displacement interpo-
 2128 lation is biased toward the identity, while the flow noise is applied at low frequency
 2129 to reflect the low amount of turbulence in this type of cloud. To smooth out the
 2130 appearance, the density is scaled down before upsampling.

2131 6.2.2 Ash rain

2132 Ash rain is usually characterised by strong striation in the vertical, corresponding
 2133 to density spikes in the plume umbrella. Flow noise with a low frequency is added in
 2134 the vertical (a period of 1000 to 3000 metres) and a high frequency in the horizontal
 2135 (a period of 1 to 3 metres). This enables us to recreate the striated appearance. As
 2136 with the water clouds, we backtrack a particular voxel (equation 6.1), though the
 2137 velocity is scaled to be predominantly vertical. The horizontal velocity is matched
 2138 to the prevailing 2D wind as calculated by the simulation.

2139 6.2.3 Plume

2140 In the simulation, the spheres and subspheres (section 5.1.2) are modelled as solid
 2141 spheres of uniform density. This is difficult to represent in a volumetric rendering,
 2142 and is inaccurate. The subspheres are also small and would require an impractically

2143 dense voxel sampling to directly convert into a volumetric grid.

2144

2145 To overcome these challenges, we use a combination of a distance field, flow noise,
 2146 and vector backtracking to create a voxel grid with voxels as cubes of roughly 100
 2147 metres on a side. For this, each voxel needs to store the distance d to the nearest
 2148 plume sphere p , the ash density in the sphere A_p , and a displacement vector from
 2149 the current position to the initial position of the sphere \vec{D}_p .

2150

2151 We use ray-tracing on the distance values of each voxel to approximate the hard
 2152 transitions between the plume and the surrounding air. The opacity of this transition
 2153 is then adjusted by the ash density of the sphere. To add detail, we subtract flow
 2154 noise from the distance value of the voxel. The noise is generated with a period
 2155 of $25/2$ simulation steps to get the characteristic billowy appearance of the plume.
 2156 This frequency is chosen to ensure that the frequency of the noise is less than the
 2157 radius of the smallest sphere. As with the clouds, we need to maintain temporal
 2158 consistency despite the random nature of the noise function. We do this by seeding
 2159 the noise with the original position of the sphere derived from \vec{D}_p .

2160 6.2.4 Lightning

2161 The simulation data does not indicate when or where a lightning strike might occur.
 2162 Instead, all lightning is created during the offline rendering step. At each rendering
 2163 step, two spheres (or subspheres) in the plume with a high charge can connect to
 2164 each other. We constrain the connections to occur probabilistically, with a prob-
 2165 ability of 0.1 at each time step, though this can be adjusted to reflect a more or
 2166 less active plume. If the spheres connect, a lightning bolt is created and the charge
 2167 in an area around the endpoints of the strike is neutralised. This neutralisation is
 2168 a modification on top of the charge values provided by the simulation system (see
 2169 section 5.3), meaning that multiple strikes can occur with the same sphere if the
 2170 run continues long enough and the charge has time to accumulate again.

2171

2172 To render the final lightning strike, we use work from Reed et al.[RW94] as in-
 2173 spiration. We model lightning strikes with a branching wire structure with a single
 2174 primary branch and multiple secondary branches radiating away from the primary.
 2175 The secondary branches are randomly directed, though constrained to have an angle
 2176 of less than 45° compared to the primary branch.

2177

2178 To introduce further randomness and enhance the appearance of the lightning,
 2179 we use noise to perturb the line segments that make up the wire structure. The noise
 2180 function used is fractional Brownian motion noise[MV68], generated as a function
 2181 of the spatial position and time for each line segment. This noise function is chosen
 2182 because it generates a jagged appearance typical of lightning bolts.

2183 Chapter 7

2184 Results

2185 In this chapter, we show the results produced by our system, assess the visual real-
2186 ism of our system, and discuss the results analytically.

2187

2188 In terms of implementation, all of the simulation models are coded in C++20
2189 [Com20], and compiled at maximum optimisation using the GNU compiler for C++
2190 [Lea93], version 13.2. For previsualisation and tuning purposes we use OpenGL.
2191 The procedural upsampling and offline rendering is implemented in Houdini 19.5.
2192 Initial wind, moisture, and cloud fields were set for each layer using 2D Perlin noise
2193 [Per85] in the range $(-1, 1)$ added to a fixed value.

2194

2195 In our results, we find that our combined model can reproduce most of the vi-
2196 sually striking phenomena arising from the interaction between the atmosphere and
2197 the volcanic plume. We also perform experiments to test the correspondence be-
2198 tween the different submodels, and compare our system’s behaviour to that of real
2199 volcanic eruptions.

2200

2201 In terms of visual results, we capture the formation of transient condensation
2202 clouds, shown in figure 7.1. This shows a timelapse shortly after an explosion, rang-
2203 ing from $t = 0.3s$ to $t = 6.8s$. The condensation cloud is localised to one side of
2204 the volcano vent due to the distribution of the ambient moisture levels in the at-
2205 mosphere. This can be seen in figure 7.1(6), which shows a moisture map of the
2206 atmospheric layer at an altitude of 3000 metres, along with a visual indication of
2207 the layer’s position relative to the volcanic plume.

2208

2209 Our system successfully generates embedded clouds in the plume. Figure 7.2
2210 shows three skirt layers being driven up and then deformed by the volcanic plume.
2211 This happens over a time span of $t = 60s$. Cap clouds, which are a result of the
2212 same processes at different heights and scales, are shown in figure 7.3 over a time
2213 span of $t = 120s$. A single pileus cloud is driven ahead of the plume before reaching
2214 neutral buoyancy, where the plume then breaks through the cap cloud.

2215

2216 We also provide results that demonstrate ash rain in an eruption that extends to
2217 the umbrella phase. It must be noted that while ash fall happens during every stage
2218 of the eruption, it is usually obscured by the plume in the earlier stages. Figure 7.4
2219 shows ash rain precipitating from the plume umbrella, over a time span of $t = 510s$,

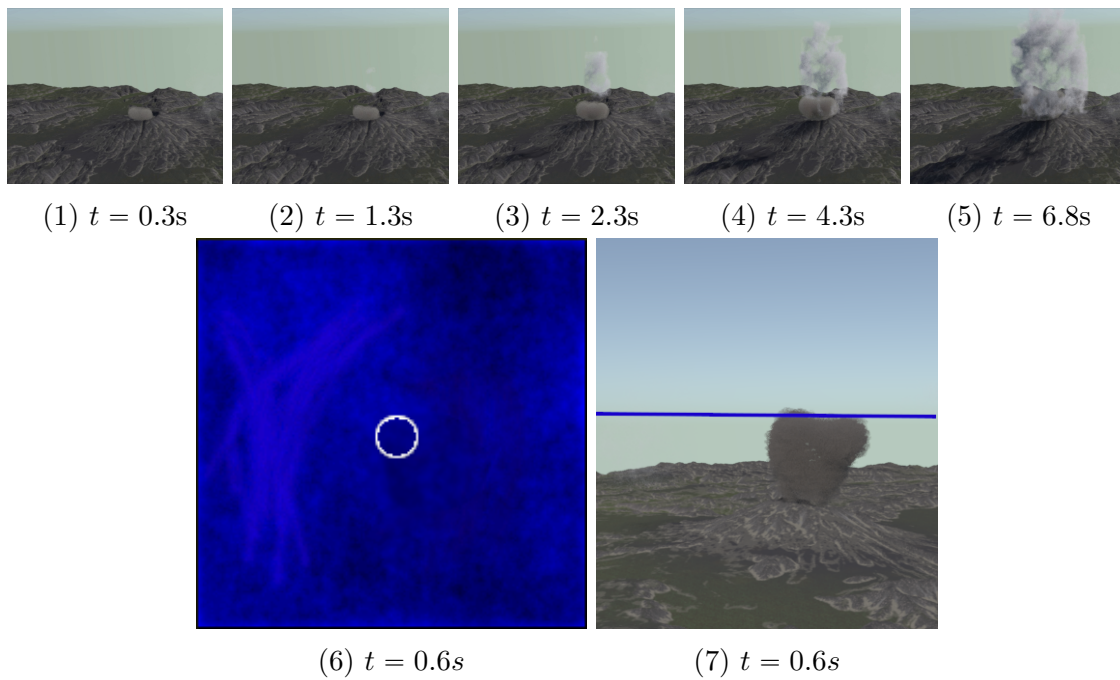


Figure 7.1: [top] A timelapse showing the formation of a transient condensation cloud. The condensation cloud is localised to one side (towards the camera) due to more ambient moisture being present in that area of the atmosphere. [bottom] On the left, a moisture map at $t = 0.6s$ with a narrow field of view around the volcano (indicated by the white circle), at an altitude of 3000 metres. Lighter blue indicates higher moisture content, particularly visible left of the vent (towards the camera in the top renders). On the right, the blue line indicates the location of the moisture layer in (6) relative to the plume.




Figure 7.2: Three skirt cloud layers deformed and driven ahead of the plume, condensing due to adiabatic cooling.




Figure 7.3: A single cap cloud forms at the head of the plume, rising with it until reaching a neutral buoyancy height. The plume then pushes through the cap cloud.




Figure 7.4: Fine ash mixes with water in the umbrella phase, before precipitating out as ash rain.




Figure 7.5: Volcanic lightning strikes.

2220 starting from $t = 450s$. The initial wind field is set to $10 \pm 1 m \cdot s^{-1}$ east to west
 2221 (left to right in the render) above an altitude of $7000m$, and $10 \pm 1 m \cdot s^{-1}$ south to
 2222 north (from front to back in the render) below that altitude.

2223

2224 Lastly, we show volcanic lightning in the plume in figure 7.5. Our system pro-
 2225 duces lightning at a peak frequency of 18 visible strikes per minute (measured across
 2226 10 frames, 30 seconds of simulation time) at time $t = 450s$. This compares favourably
 2227 with the 27 strikes per minute from real eruptions[Beh+13], which was measured
 2228 using low frequency radio waves generated by the lightning strikes and thus includes
 2229 the strikes obscured by the plume itself. This also validates our choice of 0.1 for the
 2230 connection probability per time step (see section 6.2.4).

2231 7.1 Validation

2232 To confirm that our models conform to real world data, we perform a number of
 2233 experiments. In particular, we consider ash transmission through the atmosphere,
 2234 ash deposition onto the terrain, the progressive electrical charge build up in the
 2235 plume spheres, and an ablation study comparing the advected wind field with a
 2236 uniform wind field.

7.1.1 Ash rain

We assess the accuracy of the ash rain with two methods. First, we investigate the behaviour of the ash during transfer from the plume model and in the atmospheric layers after the transfer. As part of this, we keep track of the quantity of ash deposited onto the terrain. As shown in figure 7.6(1), substantial ash accumulates around the vent from early in the eruption, even before the umbrella forms. This is to be expected since the plume undergoes a continuous density loss, regardless of the phase of the eruption. There is little or no ash accumulation on the vent itself, kept aloft by the intense uplift in the eruptive column.

The scenario that figure 7.6 is extracted from has an initial windfield that is consistent from west to east (i.e., left to right in the images). Figure 7.6(2) shows ash density in the topmost atmospheric layer below the umbrella, at a height of 17.5 km. There is a clear visual correlation between the plume in figure 7.6(3) and the distribution of ash in 7.6(2). As ash accumulates in one layer, it will fall to the layer below if sufficiently dense (see section 4.3.1). Figure 7.6(4) shows the ash density in the air layer immediately below 7.6(2), with an altitude of 16.2 km, at the same time step. Given the consistent westerly wind, the ash that has been taken up by the atmospheric system is more dispersed than the plume and further away from the plume.

The second evaluation metric is the longer term deposition of ash on the terrain. This is a fairly common metric in vulcanology, and we can compare the results of our simulation with real-world data. In figure 7.7, we show the accumulated ash depth along transects shown on the map on the right after the first 20 minutes of the eruption, with a fairly uniform wind field from the north west (i.e., the top left corner). The deposition profile is, as expected, heavily influenced by the wind strength and direction.

While eruptions in the reality can result in ash depths of multiple metres, this is the result of eruptions that continue for many days or even weeks. We find the peak ash depth to be some distance away from the vent in the general wind direction, which is a consequence of modelling only the fine ash that tend to get carried along by the convective column. Coarser grains are denser and tend to precipitate out early in the plume life cycle. Despite this, our model exhibits an exponential thinning pattern that agrees well with field studies and vulcanological models[Bra+83; Pyl89; BES98], as shown in figure 7.8. Note the logarithmic scale on the y-axis. The x-axis uses the square root of isopach area, that is, the area enclosed by a contour line of equal depth. The linear relation of $\ln(T)$ to \sqrt{A} , shown in the graphs, corresponds to an exponential relation of T to D [Pyl89]. In this case, T refers to the thickness of the ash deposition, A the isopach area, and D the distance from the vent.

7.1.2 Lightning

To evaluate the accuracy of the lightning, we visualise the charge in the spheres of the plume. Figure 7.9 shows the incremental and accumulated charge for a plume that is transitioning into the umbrella stage. We show that charge accumulation greatly increases above the -20°C isotherm, indicated by the black outline at an altitude of

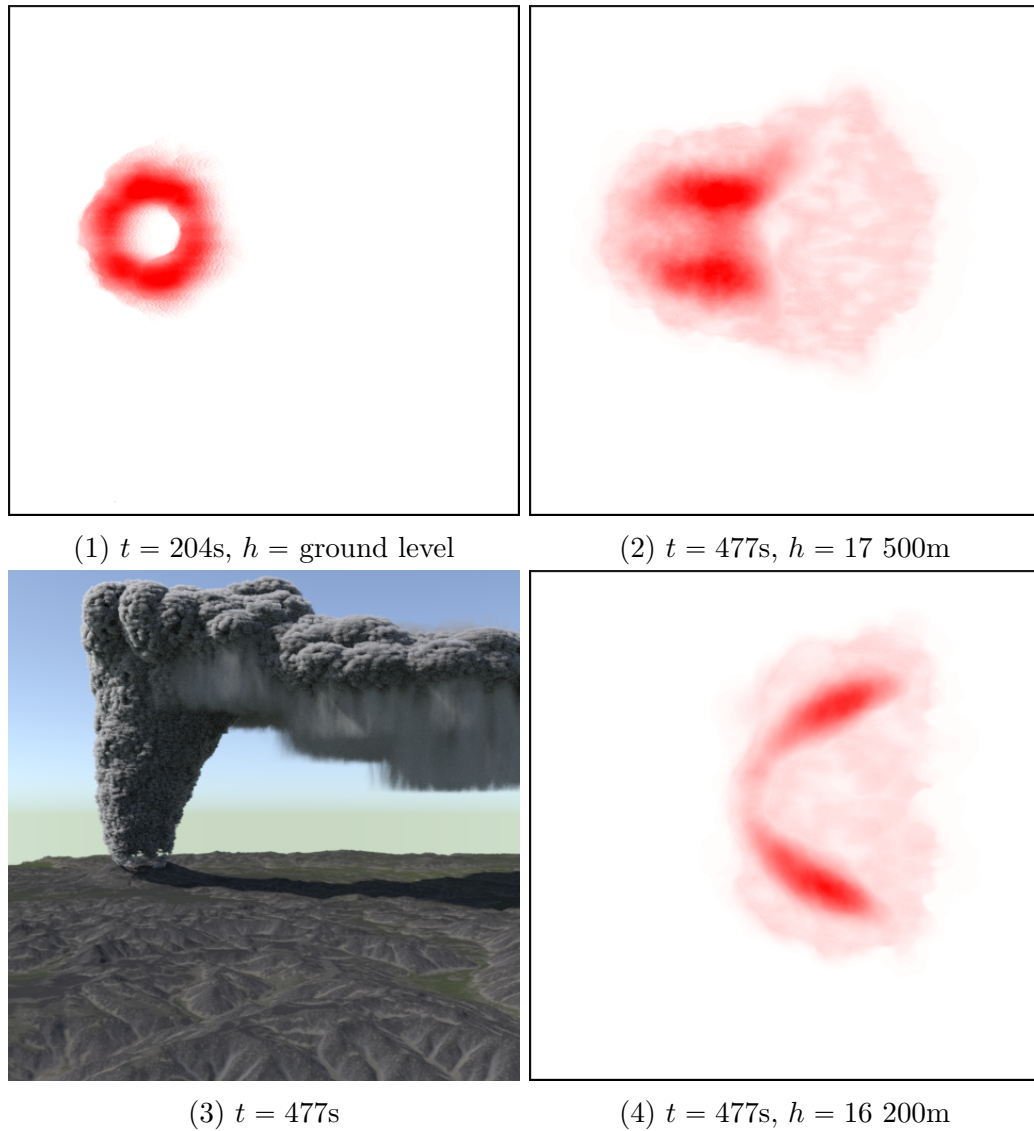


Figure 7.6: Ash density in the atmosphere with a wind from left to right. Darker red indicates a higher density. (1) Ash accumulated on the terrain prior to the plume entering the umbrella phase. (2) Ash density in the top most layer below the plume umbrella, where the majority of ash transfer from the plume to the atmosphere takes place. (3) Scene render of the plume and umbrella, showing ash rain precipitating from the umbrella. (4) Ash density in the layer immediately below (2), showing dispersal away from the plume according to the wind field.

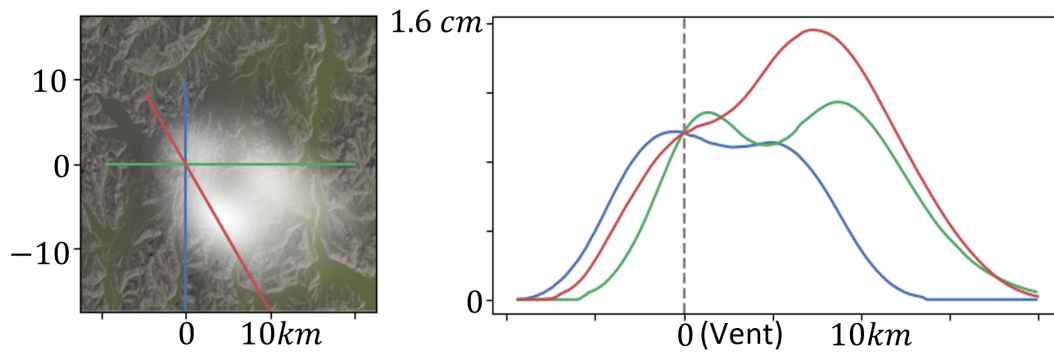


Figure 7.7: Ash deposition on the terrain 20 minutes after eruption. [left] Ash accumulation (white is deeper ash) and colour-coded transects overlaid on the landscape. [right] Ash depth along the transects graphed as a function of distance from the volcano vent.




Figure 7.8: Graphs of ash deposition from real volcanic eruptions, sourced from Pyle[Pyl89]. Note the logarithmic scale on the y-axis. The x-axis uses the square root of isopach area, that is, the area enclosed by a contour line of equal depth. This is proportional to the distance from the peak depth.




Figure 7.9: The electrical charge in each sphere visualised, with darker red indicating a greater charge and light yellow indicating no charge. The black frame is at an altitude of 4100 m above sea level, and roughly corresponds to the $-20\text{ }^{\circ}\text{C}$ isotherm. [top left] The increase in charge over a single time step. Note the significant increase above the $-20\text{ }^{\circ}\text{C}$ isotherm. [top right] Total accumulated charge in each sphere. [bottom] Expanded view of the region in the blue square. Note the regions of low charge (lighter) that result from discharge as a result of lightning strikes.

2283 4100 m. This is to be expected given the two differing regimes of charge accumulation
 2284 (see equation 5.2). We also show how the accumulated charge increases over time
 2285 for all spheres. The spheres in the vicinity of a lightning strike are discharged, and
 2286 this can be seen as areas of low charge in the visualisation.

2287 7.1.3 Windfield ablation

2288 The atmospheric sub-model, particularly the advected windfield, requires significant
 2289 computation and adds considerable design complexity. Thus, it is worth evaluating
 2290 if the benefits outweighs this performance and implementation cost. We perform
 2291 an ablation study comparing the impact of an advected windfield against a uniform
 2292 wind field. Figure 7.10 depicts the ash density in the atmosphere under a constant,
 2293 uniform wind field, and under a complex, advected wind field. Both cases utilise
 2294 a grid resolution of 512 by 512 cells per air layer. 7.10(1) shows a constant wind
 2295 field at $10\text{ m} \cdot \text{s}^{-1}$ from north to south (top to bottom), with the ash dispersing
 2296 slowly over time. In contrast, 7.10(2) shows an advected wind field, with eddies and
 2297 turbulence being reflected in the ash density. This level of detail is not possible with
 2298 a uniform wind field.

2299 7.2 Performance

2300 We perform timing tests on an Intel[®] Core[™] i9-9900K CPU @ 5GHz and an
 2301 NVidia[®] RTX 2080M (8GB). Table 7.1 shows the computational performance for
 2302 each scenario in our simulation. Table 7.5 depicts the computational performance
 2303 for each sub-model in the same scenarios. We found that, for a single simulation
 2304 step, the atmospheric model takes much longer than a step for either the plume or
 2305 embedded cloud models. However, the atmospheric model is run at a significantly
 2306 lower frequency, which partly mitigates this issue.




Figure 7.10: An ablation study comparing the ash distribution over a time span of $t = 180\text{s}$, under a uniform wind field of $10 \text{ m} \cdot \text{s}^{-1}$ north to south (top to bottom) [top]; and under an advected wind field [bottom]. The advected wind field rotates in a clockwise fashion (as set by the initial wind field), and displays eddies and turbulence that are impossible under a uniform wind field.

Scenario	Scene Dimensions	Grid size	Layers	Time step computation (in s)
Skirt (figure 7.2)	$60 \times 60 \times 20\text{km}$	512×512	20	107.2s
Cap & condensation (figures 7.1 & 7.3)	$60 \times 60 \times 20\text{km}$	512×512	20	59.6s
Ash rain (figures 7.4 and)	$60 \times 60 \times 20\text{km}$	1024×1024	20	79.4s

Table 7.1: System performance for 3 scenarios. The average performance for a simulation time of $t = 1\text{s}$ is reported, with the average calculated across 320 steps. Each scenario used the same layer structure: layers at 0.4, 1.2, 2.0, 2.7, 3.0, 3.5, 4.1, 4.7, 5.5, and 6.5 km, with the remainder evenly spaced up to 20 km.

Resolution	512 × 512		1024 × 1024	
	Percentage	Compute time	Percentage	Compute time
Advection	47.45	4.038s	25.87	13.41s
Pressure solve	39.33	3.347s	62.56	32.42s
Terrain heating	0.000002	0.0000002s	0.0000004	0.0000021s
Atmospheric heating	0.42	0.036s	0.26	0.1348s
Moisture uptake	0.57	0.049s	0.30	0.1555s
Obstacle tamping	0.0007	0.00006s	0.0003	0.0016s
Phase transition	4.7	0.4s	2.51	1.301s
Precipitation	1.01	0.086s	1.72	0.8915s
Vertical coupling	6.51	0.554s	6.78	3.514s
Total:	100%	8.51s	100%	51.83s

Table 7.2: The computation time spent on each component of the atmosphere sub-model per time step, with 20 layers on a grid of 512×512 cells. For larger grids, the pressure solve takes up a larger proportion, up to 70% on a 1024×1024 grid.

2307 Within the atmosphere sub-model, we found that the computation is dominated
2308 by the advection and pressure solve, which requires 90% of this sub-model’s compute
2309 time. At higher resolutions, the pressure solve alone requires proportionally more
2310 time, up to 63% (from 40%) of total compute for the atmospheric submodel on
2311 a 1024×1024 grid. Table 7.2 shows the average computational performance and
2312 proportion of the total for each component, using the ash rain scenario on a 512×512
2313 and a 1024×1024 grid. Adding more layers and increasing the grid resolution
2314 linearly increases computation time, as shown in table 7.3. While larger grid sizes
2315 (2048×2048) were tested, the computation times required proved infeasible and
2316 had limited visual effect on the final results. In practice, we found that around 20
2317 layers with a cell size between 60 and 100 metres on a side is enough to capture the
2318 relevant atmospheric behaviour around the volcanic plume.

		Resolution		
		256 ²	512 ²	1024 ²
Layers	10	0.86s	3.65s	21.64s
	20	1.66s	8.51s	51.83s
	30	3.12s	15.23s	98.5s

Table 7.3: Total performance requirement for each time step ($\Delta t = 1s$) of the atmospheric sub-model as a function of the number of layers and layer resolution, using the ash rain scenario.

2319 Despite the grid sizes being equal for both the skirt and cap and condensation
2320 cloud scenarios, the atmosphere sub-model requires more than double the compu-
2321 tation time for the cap and condensation cloud scenario. This is a consequence of
2322 running the phase transition component at a much higher frequency ($\Delta t = 0.1s$)
2323 compared to the rest of the sub-model, which is run with a time step of $\Delta t = 2s$. In
2324 the other scenarios, all components of the atmosphere system run at the same fre-
2325 quency, with $\Delta t = 2s$ and $\Delta t = 3s$ for the skirt and ash rain scenarios, respectively.
2326

2327 Table 7.4 shows the total time required to compute the scenarios for 320 steps.
 2328 This is sufficient to produce all the effects demonstrated in this chapter, though the
 2329 simulations can vary in length if only some effects are desired (e.g., less simulation
 2330 time is required for the skirts).

2331

Scenario	Average computation time for 1 second of simulation time	Δt for atmosphere system	Computation time for full scenario
Skirt	107.2s	2s	1143.46 minutes
Cap & condensation	59.6s	2s	635.73 minutes
Ash rain	79.4s	3s	1270.4 minutes

Table 7.4: The total computation time required for each of the three scenarios listed in table 7.1, for 320 simulation steps. The atmospheric sub-model uses different timesteps depending on the scenario, while the plume sub-model always uses a timestep of $\Delta t = 0.02s$.

Scenario	Timestep computation (in s)			
	Plume	Embedded Clouds	Atmosphere	Total
Skirt	22s	76.5s	8.7s	107.2s
Cap & condensation	5.5s	35s	19.1s	59.6s
Ash rain	26.5s	Not applicable	52.9s	79.4s

Table 7.5: Average performance for each component in the three scenarios listed in table 7.1, for a simulation time of $t = 1s$. The ash rain scenario had no embedded clouds in it, hence no cost associated with that sub-model.

2332 The computation requirements for the plume sub-model increases in a linear
 2333 fashion as the number of active slices and spheres grow as part of the eruption. When
 2334 the number of slices and spheres is low, the cost for this sub-model can be negligible.
 2335 However, when the plume is well developed (e.g., $t = 600s$), this model can require
 2336 up to 1.2 seconds of compute for each time step ($\Delta t = 0.02s$). The timings in
 2337 table 7.5 refer to the computation required for 50 steps of the plume system, for a
 2338 simulation time of $t = 1s$, averaged over a complete run of the scenario (between 20
 2339 and 40 minutes of simulated time). Figure 7.11 shows the computation time required
 2340 per time step as a function of the simulation time, showing a linear increase for the
 2341 plume subsystem while the atmospheric subsystem remains nearly constant. This
 2342 increase over time is especially noticeable in the average time required for plume
 2343 sub-model in the ash rain scenario, which contributes to the large computation time
 2344 required for the full ash rain scenario.

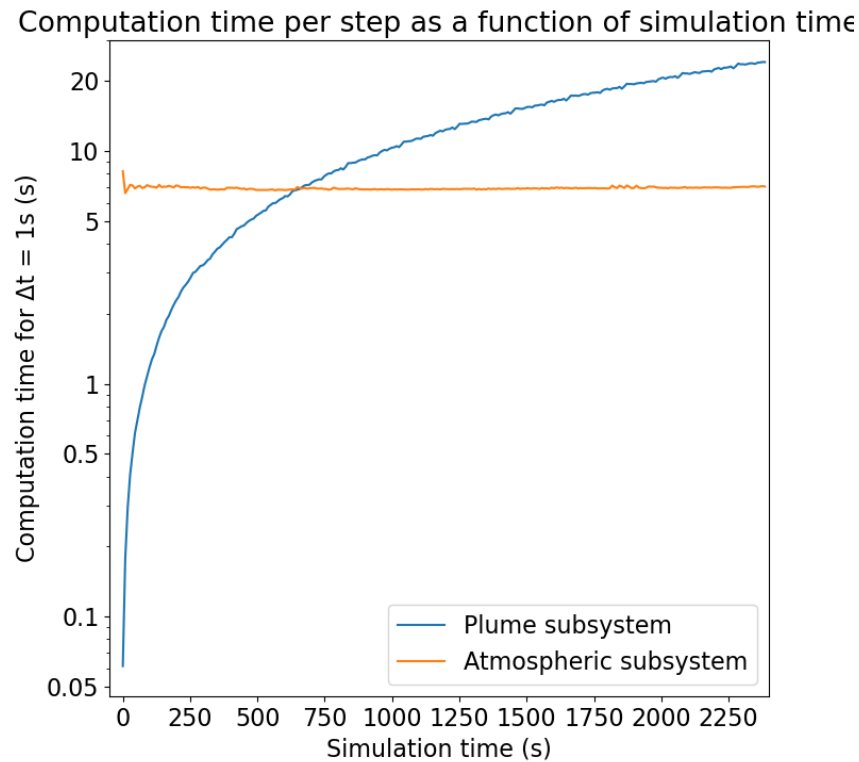


Figure 7.11: We graph the computation time required per time step as a function of simulation time. This data is extracted from the ash rain scenario (figure 7.4), on a grid of 512×512 , 20 air layers, and $\Delta t = 0.02\text{s}$ for the plume subsystem and $\Delta t = 3\text{s}$ for the atmospheric subsystem. Note the logarithmic scale on the y-axis. The plume subsystem increases linearly in computation cost as the eruption progresses, while the atmospheric subsystem remains fairly constant in its computation cost. The atmospheric system requires a few steps before the conjugate gradient solver stabilises, which explains the extra computation initially.

2345 The skirt sub-model runs in lockstep with the plume sub-model, with a time
 2346 step of $\Delta t = 0.02\text{s}$. Calculating the deformation of the skirt surfaces is particularly
 2347 expensive since it requires finely subsampling the plume trajectory. In future, a
 2348 more sophisticated scheme for calculating the intersection of the skirt sheet and
 2349 the plume could reduce the computation required. The computation cost is also
 2350 mitigated by dropping intersection testing once the skirt sheet has reached neutral
 2351 buoyancy.

2352 7.3 Limitations

2353 Our solution for coupling explosive volcanic eruptions with the atmosphere presents
 2354 a reasonable compromise between performance and accuracy. However, it remains
 2355 far slower than real time. The atmospheric sub-model remains fairly coarse, with
 2356 cells ranging in size from 60 metres to 120 metres on a side, which can lead to
 2357 some inaccuracies near the plume. While finer resolutions are possible, this incurs
 2358 a significant performance cost and is infeasible for simulation times of $t = 60\text{s}$ or

2359 more. Given that explosive volcanic eruptions range in length from tens of minutes
2360 to hours or even days, our solution is limited to a coarse resolution. A GPU imple-
2361 mentation or a multiresolution grid for the atmospheric sub-model could both serve
2362 as a potential mitigation of this issue. The algorithms in use are suitable for both,
2363 and it is likely to significantly reduce the computation time required.

2364

2365 A consequence of the largely procedural charge generation (since we do not ex-
2366 plicitly model friction) is that we do not differentiate between positive and negative
2367 charge. Highly charged areas are connected for each lightning strike, which, while
2368 visually acceptable, is not physically accurate.

2369

2370 Similarly, the coupling between the horizontal layers in the atmosphere is pro-
2371 cedural. This requires forgoing physical accuracy in a number of places elsewhere
2372 as well. The pressure values used in any of the phase transition equations (see
2373 section 4.2.6) uses a very coarse, barometric approximation, rather than the true
2374 value. Similarly, the pressure values used around the transient condensation cloud
2375 are barometric, requiring a workaround to mimic the visual appearance. The physi-
2376 cal accuracy of the layer exchange could be improved by using a more accurate, full
2377 3D model of the atmosphere, though this would come at significant computational
2378 cost.

2379

2380 Lastly, as is the case with many simulation methods that involve fluid dynamics,
2381 setting up the initial conditions and tuning the parameters requires time. This pro-
2382 cedure usually requires some experimentation and adaptation. In our case, specifi-
2383 cally, the number of air layers and their altitudes need to be carefully set to match
2384 atmospheric features such as inversion layers. Additionally, the vertical sampling
2385 near the volcano vent must be dense enough for the plume to be stable, with exact
2386 placement requiring experimentation.

2387 Chapter 8

2388 Conclusion

2389 Volcanic eruptions are visually striking phenomena and play an important role in
2390 digital media such as films or games. However, existing solutions for animating
2391 volcanic plumes are usually bespoke, one-off solutions, while existing plume models
2392 ([MDN02b; Miz+03; Las+22]) do not consider dynamic atmospheres, limiting their
2393 physical accuracy and the phenomena that can be depicted. A naive combination
2394 with existing cloud models would also not suffice considering the extremes of both
2395 temperature and altitude that are present in volcanic plumes.

2396
2397 To address this, we conceived and have developed an integrated framework to
2398 generate and animate explosive volcanic eruptions. We performed a detailed investi-
2399 gation of the physical processes that determine the appearance of skylscapes around
2400 volcanic plumes, with a focus on the movement in the atmosphere, cloud formation,
2401 and the formation of ash rain. We also did an analysis of the dynamics of incom-
2402 pressible fluids, required for implementing our treatment of fluid dynamics.

2403
2404 We compiled a thorough exploration of the state of the art in this field, with
2405 a focus on cloud modelling techniques, and the rise of hybrid models, i.e., models
2406 combining more than one modelling technique. We also conducted a survey of vol-
2407 canic plume models in both computer graphics and geophysics, finding that while
2408 little attention has been given to these phenomena in computer graphics, it is a very
2409 active field of research in the geophysical sciences.

2410
2411 We adapted Vimont et al.’s[Vim+20] cloud model for use in our system. We made
2412 changes in order to accommodate the temperature and height extremes encountered
2413 around a volcanic plume, and implement our own treatment of fluid dynamics to
2414 address the inaccuracies around solid boundaries that Vimont et al.’s system pro-
2415 duced. We also modified Lastic et al.’s[Las+22] plume model to accurately portray
2416 ash rain and volcanic lightning, while implementing our own model of lightning and
2417 embedded clouds in the plume.

2418
2419 We combined the two models through bi-directional interaction, transferring
2420 physical quantities such as ash, moisture, and temperature between the two sub-
2421 models. By running the submodels at different frequencies, we addressed some of
2422 the pitfalls that could arise from the high energy dynamics encountered in and
2423 around the plume. Through the use of procedural amplification and procedural

2424 noise functions, we enhance the final renders significantly.

2425

2426 We demonstrated the graphical depiction of several previously unsupported phe-
2427 nomena. This includes ash rain, embedded clouds, transient condensation clouds,
2428 and volcanic lightning. Moreover, the dynamic atmosphere allows the plume to track
2429 crosswinds and turbulent eddies more realistically. We validate our contributions
2430 by visually comparing our results to that of real volcanic eruptions, as well as as-
2431 sessing the ash deposition patterns of our model compared to that of real eruptions.
2432 We also justified our design decisions through an ablation study, demonstrating the
2433 necessity of our combined system.

2434

2435 Despite the high energy dynamics present, the fluid solver for the atmospheric
2436 simulation maintains indefinite stability. Computation times depend on the number
2437 of layers and the resolution of said layers, as well as the number of spheres in the
2438 plume. Generally, less than 2 minutes of computation is required for each second
2439 of simulation. Since many of the computationally intensive components (such as
2440 advection) are suitable for parallel execution[CLT07], a performance boost would
2441 be possible through GPU implementation. All of the weather components, except
2442 for the advection and the vertical inflow outflow, are independent on a per cell basis,
2443 allowing for SIMD vectorisation, which is especially suited to a GPU implementa-
2444 tion. By using a memory buffer, the advection component can be vectorised as well.

2445

2446 This work also opens up opportunities for future extensions in volcano simulation
2447 for computer graphics. Combining our model with models for terrain deformation
2448 (e.g., [CFM11] or [Pey+09]) and lava flows (e.g., [Zha+17]) would allow for a single
2449 framework simulating all aspects of a volcanic eruption. A particularly challenging
2450 facet would be to model lahar flows, a mixture of pyroclasts, water, and debris that
2451 can vary significantly in density and viscosity across its volume. Lava flows, which
2452 have been modelled in computer graphics[Sto+99; Zha+17], are fairly homogeneous
2453 in comparison and smoothly vary in temperature and density. A possible approach
2454 would be a non-homogenous Lagrangian approach, with the different components of
2455 lahars each having separate particle parameters.

2456

2457 The noise functions used in this work are largely retained from the previous sys-
2458 tems, and there are other options available as well, such as Voronoi cell noise[Wor96].
2459 A more detailed analysis and comparison of potential noise functions could be em-
2460 ployed to determine the noise functions best suited for this specific purpose.

2461

2462 Another possibility would be to incorporate ash rain and aerosolised ash in the
2463 atmosphere into weather systems, with the accompanying broad-scale effects. This
2464 would be of particular interest in games with seasonal cycles, such as *Medieval Dy-*
2465 *nasty* or *Red Dead Redemption 2*, or lightweight global climate simulations suitable
2466 for use in education.

Bibliography

- 2468 [Ada+07] Bart Adams, Mark Pauly, Richard Keiser, and Leonidas J Guibas.
2469 “Adaptively sampled particle fluids”. In: *ACM SIGGRAPH 2007 pa-*
2470 *pers.* 2007, 48–es.
- 2471 [AJ05] Cristina L Archer and Mark Z Jacobson. “Evaluation of global wind
2472 power”. In: *Journal of Geophysical Research: Atmospheres* 110.D12
2473 (2005).
- 2474 [Atm76] US Standard Atmosphere. *US standard atmosphere*. National Oceanic
2475 and Atmospheric Administration, 1976.
- 2476 [Bar+68] JC Barnes, AH Glasen, PE Sherr, and JH Willand. “World-wide cloud
2477 cover distribution for use in computer simulations Final report, 18 Jan.
2478 1967-17 Jan. 1968”. In: (1968).
- 2479 [Bar13] Roger G Barry. *Mountain weather and climate*. Routledge, 2013.
- 2480 [Bar82] Sumner Barr. “Skirt clouds associated with the Soufriere eruption of
2481 17 April 1979”. In: *Science* 216.4550 (1982), pp. 1111–1112.
- 2482 [BBD12] RJ Brown, C Bonadonna, and AJ Durant. “A review of volcanic ash
2483 aggregation”. In: *Physics and Chemistry of the Earth, Parts a/b/c* 45
2484 (2012), pp. 65–78.
- 2485 [BC09] Roger G Barry and Richard J Chorley. *Atmosphere, weather and cli-*
2486 *mate*. Routledge, 2009.
- 2487 [Beh+13] Sonja A Behnke, Ronald J Thomas, Stephen R McNutt, David J Schnei-
2488 der, Paul R Krehbiel, William Rison, and Harald E Edens. “Observa-
2489 tions of volcanic lightning during the 2009 eruption of Redoubt Vol-
2490 cano”. In: *Journal of Volcanology and Geothermal Research* 259 (2013),
2491 pp. 214–234.
- 2492 [BES98] Costanza Bonadonna, GGJ Ernst, and RSJ Sparks. “Thickness varia-
2493 tions and volume estimates of tephra fall deposits: the importance of
2494 particle Reynolds number”. In: *Journal of Volcanology and Geothermal*
2495 *Research* 81.3-4 (1998), pp. 173–187.
- 2496 [BHN07] Robert Bridson, Jim Houriham, and Marcus Nordenstam. “Curl-noise
2497 for procedural fluid flow”. In: *ACM Transactions on Graphics (ToG)*
2498 26.3 (2007), 46–es.
- 2499 [BHR07] Henry E Bass, Claus H Hetzer, and Richard Raspet. “On the speed of
2500 sound in the atmosphere as a function of altitude and frequency”. In:
2501 *Journal of Geophysical Research: Atmospheres* 112.D15 (2007).

- 2502 [Bi+16] Shuoben Bi, Shengjie Bi, Xiaowen Zeng, Yuan Lu, and Hao Zhou. “3-
2503 Dimensional modeling and simulation of the cloud based on cellular
2504 automata and particle system”. In: *ISPRS International Journal of*
2505 *Geo-Information* 5.6 (2016), p. 86.
- 2506 [Bon+15] Costanza Bonadonna, Marco Pistolesi, Raffaello Cioni, Wim Degruyter,
2507 Manuela Elissondo, and Valerie Baumann. “Dynamics of wind-affected
2508 volcanic plumes: The example of the 2011 Cordón Caulle eruption,
2509 Chile”. In: *Journal of Geophysical Research: Solid Earth* 120.4 (2015),
2510 pp. 2242–2261.
- 2511 [Bra+83] S Brazier, RSJ Sparks, SN Carey, H Sigurdsson, and JA Westgate.
2512 “Bimodal grain size distribution and secondary thickening in air-fall
2513 ash layers”. In: *Nature* 301.5896 (1983), pp. 115–119.
- 2514 [Bri15] Robert Bridson. *Fluid simulation for computer graphics*. CRC press,
2515 2015.
- 2516 [Buc96] AL Buck. “Buck Research CR-1A User’s Manual”. In: *Buck Research*
2517 *Instruments: Boulder, CO, USA* (1996).
- 2518 [CD05] Robert L Cook and Tony DeRose. “Wavelet noise”. In: *ACM Transac-*
2519 *tions on Graphics (TOG)* 24.3 (2005), pp. 803–811.
- 2520 [CEB16] Matteo Cerminara, Tomaso Esposti Ongaro, and Luigi C Berselli. “ASHEE-
2521 1.0: a compressible, equilibrium–Eulerian model for volcanic ash plumes”.
2522 In: *Geoscientific Model Development* 9.2 (2016), pp. 697–730.
- 2523 [CFM11] Justin Crause, Andrew Flower, and Patrick Marais. “A system for real-
2524 time deformable terrain”. In: *Proceedings of the South African Institute*
2525 *of Computer Scientists and Information Technologists Conference on*
2526 *Knowledge, Innovation and Leadership in a Diverse, Multidisciplinary*
2527 *Environment*. 2011, pp. 77–86.
- 2528 [CG22] Corrado Cimarelli and Kimberly Genareau. “A review of volcanic elec-
2529 trification of the atmosphere and volcanic lightning”. In: *Journal of*
2530 *Volcanology and Geothermal Research* 422 (2022), p. 107449.
- 2531 [Che+21] Qian Chen, Yue Wang, Hui Wang, and Xubo Yang. “Data-driven sim-
2532 ulation in fluids animation: A survey”. In: *Virtual Reality & Intelligent*
2533 *Hardware* 3.2 (2021), pp. 87–104.
- 2534 [Cim+14] Corrado Cimarelli, MA Alatorre-Ibargüengoitia, Ulrich Kueppers, Bet-
2535 tina Scheu, and Donald B Dingwell. “Experimental generation of vol-
2536 canic lightning”. In: *Geology* 42.1 (2014), pp. 79–82.
- 2537 [CLT07] Keenan Crane, Ignacio Llamas, and Sarah Tariq. “Real-time simulation
2538 and rendering of 3d fluids”. In: *GPU gems* 3.1 (2007), pp. 633–675.
- 2539 [Com20] C++ Standards Committee. *Programming languages – C++*. ISO/IEC
2540 14882:2020. Geneva, Switzerland: International Organization for Stan-
2541 dardization and International Electrotechnical Commission, Dec. 2020.
- 2542 [CON16] Matteo Cerminara, Tomaso Esposti Ongaro, and Augusto Neri. “Large
2543 Eddy Simulation of gas–particle kinematic decoupling and turbulent en-
2544 trainment in volcanic plumes”. In: *Journal of Volcanology and Geother-*
2545 *mal Research* 326 (2016), pp. 143–171.

- 2546 [Cou+12] J Courtney, Steve Buchan, Randall S Cervený, Pierre Bessemoulin,
2547 Thomas C Peterson, JM Rubiera Torres, John Beven, John King, Blair
2548 Trewin, and Kenneth Rancourt. “Documentation and verification of
2549 the world extreme wind gust record: 113.3 m s⁻¹ on Barrow Island,
2550 Australia, during passage of tropical cyclone Olivia”. In: *Australian
2551 Meteorological and Oceanographic Journal* 62.1 (2012), p. 1.
- 2552 [CT17] Mengyu Chu and Nils Thuerey. “Data-driven synthesis of smoke flows
2553 with CNN-based feature descriptors”. In: *ACM Transactions on Graph-
2554 ics (TOG)* 36.4 (2017), pp. 1–14.
- 2555 [CW12] Ray Cas and James Wright. *Volcanic successions modern and ancient:
2556 A geological approach to processes, products and successions*. Springer
2557 Science & Business Media, 2012.
- 2558 [Den+23] Yitong Deng, Hong-Xing Yu, Diyang Zhang, Jiajun Wu, and Bo Zhu.
2559 “Fluid Simulation on Neural Flow Maps”. In: *ACM Transactions on
2560 Graphics (TOG)* 42.6 (2023), pp. 1–21.
- 2561 [DG17] Rui PM Duarte and Abel JP Gomes. “Real-time simulation of cumulus
2562 clouds through SkewT/LogP diagrams”. In: *Computers & Graphics* 67
2563 (2017), pp. 103–114.
- 2564 [DNM93] Flavio Dobran, Augusto Neri, and Giovanni Macedonio. “Numerical
2565 simulation of collapsing volcanic columns”. In: *Journal of Geophysical
2566 Research: Solid Earth* 98.B3 (1993), pp. 4231–4259.
- 2567 [Dob+00] Yoshinori Dobashi, Kazufumi Kaneda, Hideo Yamashita, Tsuyoshi Okita,
2568 and Tomoyuki Nishita. “A simple, efficient method for realistic ani-
2569 mation of clouds”. In: *Proceedings of the 27th annual conference on
2570 Computer graphics and interactive techniques*. 2000, pp. 19–28.
- 2571 [Dob+99] Yoshinori Dobashi, Tomoyuki Nishita, Hideo Yamashita, and Tsuyoshi
2572 Okita. “Using metaballs to modeling and animate clouds from satellite
2573 images”. In: *The Visual Computer* 15.9 (1999), pp. 471–482.
- 2574 [DR06] F Dobran and JI Ramos. “Global volcanic simulation: physical mod-
2575 eling, numerics, and computer implementation”. In: *Developments in
2576 Volcanology*. Vol. 8. Elsevier, 2006, pp. 311–372.
- 2577 [Dry+22] Claudia Drygala, Benjamin Winhart, Francesca di Mare, and Hanno
2578 Gottschalk. “Generative modeling of turbulence”. In: *Physics of Fluids*
2579 34.3 (2022).
- 2580 [DSY10] Yoshinori Dobashi, Yusuke Shinzo, and Tsuyoshi Yamamoto. “Model-
2581 ing of clouds from a single photograph”. In: *Computer Graphics Forum*.
2582 Vol. 29. 7. Wiley Online Library. 2010, pp. 2083–2090.
- 2583 [DYN02] Yoshinori Dobashi, Tsuyoshi Yamamoto, and Tomoyuki Nishita. “Inter-
2584 active rendering of atmospheric scattering effects using graphics hard-
2585 ware”. In: *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS
2586 conference on Graphics hardware*. 2002, pp. 99–107.
- 2587 [DYN06] Yoshinori Dobashi, Tsuyoshi Yamamoto, and Tomoyuki Nishita. “A
2588 controllable method for animation of earth-scale clouds”. In: *Proc. of
2589 CASA* (2006), pp. 43–52.

- 2590 [DYN10] Yoshinori Dobashi, Tsuyoshi Yamamoto, and Tomoyuki Nishita. “An
2591 interactive rendering system using hierarchical data structure for earth-
2592 scale clouds”. In: *Science China Information Sciences* 53.5 (2010), pp. 920–
2593 931.
- 2594 [Ebe+02] David S Ebert, F Kenton Musgrave, Darwyn Peachey, Ken Perlin, and
2595 Steve Worley. *Texturing and modeling: a procedural approach*. Elsevier,
2596 2002.
- 2597 [EUT19] Marie-Lena Eckert, Kiwon Um, and Nils Thuerey. “Scalarflow: a large-
2598 scale volumetric data set of real-world scalar transport flows for com-
2599 puter animation and machine learning”. In: *ACM Transactions on Graph-
2600 ics (TOG)* 38.6 (2019), pp. 1–16.
- 2601 [FDY15] Charles Welton Ferreira Barbosa, Yoshinori Dobashi, and Tsuyoshi Ya-
2602 mamoto. “Adaptive cloud simulation using position based fluids”. In:
2603 *Computer Animation and Virtual Worlds* 26.3-4 (2015), pp. 367–375.
- 2604 [FL04] Raanan Fattal and Dani Lischinski. “Target-driven smoke animation”.
2605 In: *ACM SIGGRAPH 2004 Papers*. 2004, pp. 441–448.
- 2606 [FM97] Nick Foster and Dimitris Metaxas. “Modeling the motion of a hot, tur-
2607 bulent gas”. In: *Proceedings of the 24th annual conference on Computer
2608 graphics and interactive techniques*. 1997, pp. 181–188.
- 2609 [FSJ01] Ronald Fedkiw, Jos Stam, and Henrik Wann Jensen. “Visual simulation
2610 of smoke”. In: *Proceedings of the 28th annual conference on Computer
2611 graphics and interactive techniques*. 2001, pp. 15–22.
- 2612 [Gar85] Geoffrey Y Gardner. “Visual simulation of clouds”. In: *Proceedings of
2613 the 12th annual conference on Computer graphics and interactive tech-
2614 niques*. 1985, pp. 297–304.
- 2615 [GBW97] Lori S Glaze, Stephen M Baloga, and Lionel Wilson. “Transport of
2616 atmospheric water vapor by volcanic eruption columns”. In: *Journal of
2617 Geophysical Research: Atmospheres* 102.D5 (1997), pp. 6099–6108.
- 2618 [Gir+23] Gilson Antonio Giraldo, Liliane Rodrigues de Almeida, Antonio Lopes
2619 Apolinário Jr, and Leandro Tavares da Silva. *Deep Learning for Fluid
2620 Simulation and Animation: Fundamentals, Modeling, and Case Studies*.
2621 Springer Nature, 2023.
- 2622 [GJ+10] Gaël Guennebaud, Benoît Jacob, et al. *Eigen v3*. <http://eigen.tuxfamily.org>.
2623 2010.
- 2624 [GL94] JS Gilbert and SJ Lane. “The origin of accretionary lapilli”. In: *Bulletin
2625 of Volcanology* 56 (1994), pp. 398–411.
- 2626 [GLM17] Bruno Galerne, Arthur Leclair, and Lionel Moisan. “Texton noise”.
2627 In: *Computer Graphics Forum*. Vol. 36. 8. Wiley Online Library. 2017,
2628 pp. 205–218.
- 2629 [GN17] Prashant Goswami and Fabrice Neyret. “Real-time landscape-size con-
2630 vective clouds simulation and rendering”. In: *Workshop on Virtual Real-
2631 ity Interaction and Physical Simulation (VRIPHYS), Lyon*. Eurographics-
2632 European Association for Computer Graphics. 2017.

- 2633 [Gos+23] Prashant Goswami, Abbas Cheddad, Fredrik Junede, and Samuel Asp.
2634 “Interactive landscape-scale cloud animation using DCGAN”. In: *Frontiers in Computer Science* 5 (2023), p. 957920.
2635
- 2636 [Gos19] Prashant Goswami. “Interactive animation of single-layer cumulus clouds
2637 using cloudmap”. In: *STAG2019—Smart Tools and Applications in Graphics, Cagliari, November 14-15*. Eurographics-European Association for
2638 Computer Graphics. 2019.
2639
- 2640 [Gos20] Prashant Goswami. “A survey of modeling, rendering and animation of
2641 clouds in computer graphics”. In: *The Visual Computer* (2020), pp. 1–
2642 18.
- 2643 [GZD08] Alexander Goldberg, Matthias Zwicker, and Frédo Durand. “Anisotropic
2644 noise”. In: *ACM Transactions on Graphics (TOG)* 27.3 (2008), pp. 1–8.
- 2645 [Häd+20] Torsten Hädrich, Miłosz Makowski, Wojciech Pałubicki, Daniel Banuti,
2646 Sören Pirk, and Dominik L Michels. “Stormscapes: Simulating Cloud
2647 Dynamics in the Now”. In: *ACM Trans. Graph* 39.6 (2020).
- 2648 [Häd+21] Torsten Hädrich, Daniel T Banuti, Wojtek Pałubicki, Sören Pirk, and
2649 Dominik L Michels. “Fire in paradise: Mesoscale simulation of wild-
2650 fires”. In: *ACM Transactions on Graphics (TOG)* 40.4 (2021), pp. 1–
2651 15.
- 2652 [Har+03] Mark J Harris, William V Baxter, Thorsten Scheuermann, and Anselmo
2653 Lastra. “Simulation of cloud dynamics on graphics hardware”. In: *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on
2654 Graphics hardware*. 2003, pp. 92–101.
2655
- 2656 [Hen83] Thomas H Henderson. “The Shuttle Mission Simulator Computer-Generated
2657 Imagery”. In: (1983).
- 2658 [Hin18] Jonathan Hines. “Stepping up to summit”. In: *Computing in science &
2659 engineering* 20.2 (2018), pp. 78–82.
- 2660 [Hof+23] Davis W Hoffman, Larry G Mastin, Alexa R Van Eaton, Stephen A
2661 Solovitz, Raul B Cal, and John K Eaton. “A One-Dimensional Vol-
2662 canic Plume Model for Predicting Ash Aggregation”. In: *Journal of
2663 Geophysical Research: Solid Earth* 128.9 (2023), e2023JB027002.
- 2664 [HS+52] Magnus R Hestenes, Eduard Stiefel, et al. “Methods of conjugate gradi-
2665 ents for solving linear systems”. In: *Journal of research of the National
2666 Bureau of Standards* 49.6 (1952), pp. 409–436.
- 2667 [HW65] Francis H Harlow and J Eddie Welch. “Numerical calculation of time-
2668 dependent viscous incompressible flow of fluid with free surface”. In:
2669 *The physics of fluids* 8.12 (1965), pp. 2182–2189.
- 2670 [Jam+08] MR James, L Wilson, SJ Lane, JS Gilbert, TA Mather, RG Harrison,
2671 and RS Martin. “Electrical charging of volcanic plumes”. In: *Space Sci-
2672 ence Reviews* 137 (2008), pp. 399–418.
- 2673 [Jen+15] Susanna F Jenkins, TM Wilson, C Magill, V Miller, C Stewart, R Blong,
2674 W Marzocchi, Michael Boulton, Costanza Bonadonna, A Costa, et al.
2675 “Volcanic ash fall hazard and risk”. In: *Global volcanic hazards and risk
2676* (2015), pp. 173–222.

- 2677 [Kal+17] Simon Kallweit, Thomas Müller, Brian McWilliams, Markus Gross, and
2678 Jan Novák. “Deep scattering: Rendering atmospheric clouds with radiance-
2679 predicting neural networks”. In: *ACM Transactions on Graphics (TOG)*
2680 36.6 (2017), pp. 1–11.
- 2681 [KEM18] Janine L Kavanagh, Samantha L Engwell, and Simon A Martin. “A
2682 review of laboratory and numerical modelling in volcanology”. In: *Solid*
2683 *Earth* 9.2 (2018), pp. 531–571.
- 2684 [Kha+22] Sergey Khaykin, Aurelien Podglajen, Felix Ploeger, Jens-Uwe Groß,
2685 Florent Tencé, Slimane Bekki, Konstantin Khlopenkov, Kristopher Bedka,
2686 Landon Rieger, Alexandre Baron, et al. “Global perturbation of strato-
2687 spheric water and aerosol burden by Hunga eruption”. In: *Communi-*
2688 *cations Earth & Environment* 3.1 (2022), p. 316.
- 2689 [Kim+19] Byungsoo Kim, Vinicius C Azevedo, Nils Thuerey, Theodore Kim,
2690 Markus Gross, and Barbara Solenthaler. “Deep fluids: A generative
2691 network for parameterized fluid simulations”. In: *Computer graphics*
2692 *forum*. Vol. 38. 2. Wiley Online Library. 2019, pp. 59–70.
- 2693 [KK14] Genichi Kawada and Takashi Kanai. “Controlling the Shape and Mo-
2694 tion of Plumes in Explosion Simulations”. In: *Proceedings of the 11th*
2695 *Workshop on Virtual Reality Interaction and Physical Simulation*. Vol. 2014.
2696 2014, pp. 69–78.
- 2697 [KV84] James T Kajiya and Brian P Von Herzen. “Ray tracing volume den-
2698 sities”. In: *ACM SIGGRAPH computer graphics* 18.3 (1984), pp. 165–
2699 174.
- 2700 [Lag+09] Ares Lagae, Sylvain Lefebvre, George Drettakis, and Philip Dutré.
2701 “Procedural noise using sparse Gabor convolution”. In: *ACM Trans-*
2702 *actions on Graphics (TOG)* 28.3 (2009), pp. 1–10.
- 2703 [Lag+10] Ares Lagae, Sylvain Lefebvre, Rob Cook, Tony DeRose, George Dret-
2704 takis, David S Ebert, John P Lewis, Ken Perlin, and Matthias Zwicker.
2705 “A survey of procedural noise functions”. In: *Computer Graphics Fo-*
2706 *rum*. Vol. 29. 8. Wiley Online Library. 2010, pp. 2579–2600.
- 2707 [Las+22] Maud Lastic, Damien Rohmer, Guillaume Cordonnier, Claude Jau-
2708 part, Fabrice Neyret, and Marie-Paule Cani. “Interactive simulation of
2709 plume and pyroclastic volcanic ejections”. In: *Proceedings of the ACM*
2710 *on Computer Graphics and Interactive Techniques* 5.1 (2022), pp. 1–15.
- 2711 [Lea93] Doug Lea. “The GNU C++ library”. In: *The C++ Report* (1993).
- 2712 [Lew84] John-Peter Lewis. “Texture synthesis for digital painting”. In: *Proceed-*
2713 *ings of the 11th annual conference on Computer graphics and interactive*
2714 *techniques*. 1984, pp. 245–252.
- 2715 [Lew89] John-Peter Lewis. “Algorithms for solid noise synthesis”. In: *Proceed-*
2716 *ings of the 16th annual conference on Computer graphics and interactive*
2717 *techniques*. 1989, pp. 263–270.
- 2718 [MB15] Ksenia Mukhina and Alexey Bezgodov. “The Method for Real-time
2719 Cloud Rendering”. In: *Procedia Computer Science* 66 (2015), pp. 697–
2720 704.

- 2721 [McK+21] Kathleen McKee, Cassandra M Smith, Kevin Reath, Eveanjelene Snee,
2722 Sean Maher, Robin S Matoza, Simon Carn, Larry Mastin, Kyle An-
2723 derson, David Damby, et al. “Evaluating the state-of-the-art in remote
2724 volcanic eruption characterization Part I: Raikoke volcano, Kuril Is-
2725 lands”. In: *Journal of Volcanology and Geothermal Research* 419 (2021),
2726 p. 107354.
- 2727 [MDN02a] Ryo Miyazaki, Yoshinori Dobashi, and Tomoyuki Nishita. “Simulation
2728 of Cumuliform Clouds Based on Computational Fluid Dynamics.” In:
2729 *Eurographics (Short Presentations)*. 2002.
- 2730 [MDN02b] Ryoichi Mizuno, Yoshinori Dobashi, and Tomoyuki Nishita. “Modeling
2731 of volcanic clouds using CML”. In: *Proceedings of the International
2732 Computer Symposium*. 2002.
- 2733 [Mén+20] Joshua Méndez Harper, Leah Courtland, Josef Dufek, and Julian McAdams.
2734 “Microphysical effects of water content and temperature on the tribo-
2735 electrification of volcanic ash on long time scales”. In: *Journal of Geo-
2736 physical Research: Atmospheres* 125.14 (2020), e2019JD031498.
- 2737 [Mil+22] Luis Millan, Michelle L Santee, Alyn Lambert, Nathaniel J Livesey,
2738 Frank Werner, Michael J Schwartz, Hugh Charles Pumphrey, Gloria L
2739 Manney, Yuan Wang, Hui Su, et al. “The Hunga Tonga-Hunga Ha’apai
2740 hydration of the stratosphere”. In: *Geophysical Research Letters* 49.13
2741 (2022), e2022GL099381.
- 2742 [Miy+01] Ryo Miyazaki, Satoru Yoshida, Yoshinori Dobashi, and Tomoyuki Nishita.
2743 “A method for modeling clouds based on atmospheric fluid dynamics”.
2744 In: *Proceedings Ninth Pacific Conference on Computer Graphics and
2745 Applications. Pacific Graphics 2001*. IEEE. 2001, pp. 363–372.
- 2746 [Miz+03] Ryoichi Mizuno, Yoshinori Dobashi, Bing-Yu Chen, and Tomoyuki Nishita.
2747 “Physics motivated modeling of volcanic clouds as a two fluids model”.
2748 In: *11th Pacific Conference on Computer Graphics and Applications,
2749 2003. Proceedings*. IEEE. 2003, pp. 440–444.
- 2750 [MM76] A.B. Meinel and M.P. Meinel. *Applied Solar Energy: An Introduction*.
2751 Addison-Wesley series in physics. Addison-Wesley Publishing Company,
2752 1976. ISBN: 9780201047196. URL: [https://books.google.co.za/
2753 books?id=o9BSAAAAAAAJ](https://books.google.co.za/books?id=o9BSAAAAAAAJ).
- 2754 [MSB71] RL Mancuso, SM Serebreny, and RH Blackmer Jr. “A cloud simula-
2755 tion model for evaluating automatic ceilometer systems”. In: *Journal
2756 of Applied Meteorology* 10.6 (1971), pp. 1324–1330.
- 2757 [MTT56] Bruce R Morton, Geoffrey Ingram Taylor, and John Stewart Turner.
2758 “Turbulent gravitational convection from maintained and instantaneous
2759 sources”. In: *Proceedings of the Royal Society of London. Series A.
2760 Mathematical and Physical Sciences* 234.1196 (1956), pp. 1–23.
- 2761 [MV68] Benoit B Mandelbrot and John W Van Ness. “Fractional Brownian mo-
2762 tions, fractional noises and applications”. In: *SIAM review* 10.4 (1968),
2763 pp. 422–437.

- 2764 [Ner+21] Augusto Neri, Tomaso Esposti Ongaro, Mattia de’Michieli Vitturi, and
2765 Matteo Cerminara. “Multiphase flow modeling of explosive volcanic
2766 eruptions”. In: *Transport phenomena in multiphase systems*. Springer,
2767 2021, pp. 243–281.
- 2768 [Ney97] Fabrice Neyret. “Qualitative simulation of convective cloud formation
2769 and evolution”. In: *Computer animation and simulation’97*. Springer,
2770 1997, pp. 113–124.
- 2771 [OA76] National Oceanic and Atmospheric Administration. “US standard at-
2772 mosphere, 1976”. In: *Technical Report* (1976).
- 2773 [Obe+98] Josef M Oberhuber, Michael Herzog, Hans-F Graf, and Karsten Schwanke.
2774 “Volcanic plume simulation on large scales”. In: *Journal of Volcanology
2775 and Geothermal Research* 87.1-4 (1998), pp. 29–53.
- 2776 [Ola+02] Marc Olano, John Hart, Wolfgang Heidrich, and Michael McCool. *Real-
2777 Time Shading*. CRC Press, 2002.
- 2778 [Ong+07] T Esposti Ongaro, Carlo Cavazzoni, Giovanni Erbacci, Augusto Neri,
2779 and Maria-Vittoria Salvetti. “A parallel multiphase flow code for the
2780 3D simulation of explosive volcanic eruptions”. In: *Parallel Computing*
2781 33.7-8 (2007), pp. 541–560.
- 2782 [Org17] World Meteorological Organization. *International Cloud Atlas*. World
2783 Meteorological Organization, 2017.
- 2784 [Pał+22] Wojtek Pałubicki, Miłosz Makowski, Weronika Gajda, Torsten Hädrich,
2785 Dominik L Michels, and Sören Pirk. “Ecoclimates: Climate-response
2786 modeling of vegetation”. In: *ACM Transactions on Graphics (TOG)*
2787 41.4 (2022), pp. 1–19.
- 2788 [Pat+05] Sanjit Patel, Anson Chu, Jonathan Cohen, and Frederic Pighin. “Fluid
2789 simulation via disjoint translating grids”. In: *ACM SIGGRAPH 2005
2790 Sketches*. 2005, 139–es.
- 2791 [Per02] Ken Perlin. “Improving noise”. In: *Proceedings of the 29th annual con-
2792 ference on Computer graphics and interactive techniques*. 2002, pp. 681–
2793 682.
- 2794 [Per85] Ken Perlin. “An image synthesizer”. In: *ACM Siggraph Computer Graph-
2795 ics* 19.3 (1985), pp. 287–296.
- 2796 [Pey+09] Adrien Peytavie, Eric Galin, Jérôme Grosjean, and Stéphane Mérillou.
2797 “Arches: a framework for modeling complex terrains”. In: *Computer
2798 graphics forum*. Vol. 28. 2. Wiley Online Library. 2009, pp. 457–467.
- 2799 [Pie+98] Roger A Pielke, RonI Avissar, Michael Raupach, A Johannes Dolman,
2800 Xubin Zeng, and A Scott Denning. “Interactions between the atmo-
2801 sphere and terrestrial ecosystems: influence on weather and climate”.
2802 In: *Global change biology* 4.5 (1998), pp. 461–475.
- 2803 [PN01] Ken Perlin and Fabrice Neyret. “Flow noise”. In: *28th International
2804 Conference on Computer Graphics and Interactive Techniques (Tech-
2805 nical Sketches and Applications)*. SIGGRAPH. 2001, p. 187.

- 2806 [PPS22] Simon R Proud, Andrew T Prata, and Simeon Schmauß. “The Jan-
2807 uary 2022 eruption of Hunga Tonga-Hunga Ha’apai volcano reached
2808 the mesosphere”. In: *Science* 378.6619 (2022), pp. 554–557.
- 2809 [PW09] Liz Parfitt and Lionel Wilson. *Fundamentals of physical volcanology*.
2810 John Wiley & Sons, 2009.
- 2811 [Pyl89] David M Pyle. “The thickness, volume and grainsize of tephra fall de-
2812 posits”. In: *Bulletin of Volcanology* 51 (1989), pp. 1–15.
- 2813 [Rav+21] Suman Ravuri, Karel Lenc, Matthew Willson, Dmitry Kangin, Remi
2814 Lam, Piotr Mirowski, Megan Fitzsimons, Maria Athanassiadou, Sheleem
2815 Kashem, Sam Madge, et al. “Skilful precipitation nowcasting using deep
2816 generative models of radar”. In: *Nature* 597.7878 (2021), pp. 672–677.
- 2817 [RD11] William I Rose and Adam J Durant. “Fate of volcanic ash: Aggregation
2818 and fallout”. In: *Geology* 39.9 (2011), pp. 895–896.
- 2819 [RLB19] John R. Rumble, David R. Lide, and Thomas J. Bruno. *CRC handbook
2820 of chemistry and physics : a ready-reference book of chemical and phys-
2821 ical data*. eng. 2019-2020, 100th edition. Boca Raton, FL: CRC Press,
2822 2019. ISBN: 113836729X.
- 2823 [RW94] Todd Reed and Brian Wyvill. “Visual simulation of lightning”. In: *Pro-
2824 ceedings of the 21st annual conference on Computer graphics and in-
2825 teractive techniques*. 1994, pp. 359–364.
- 2826 [Sat+22] Pinar Satilmis, Demetris Marnerides, Kurt Debattista, and Thomas
2827 Bashford-Rogers. “Deep synthesis of cloud lighting”. In: *IEEE Com-
2828 puter Graphics and Applications* 42.5 (2022), pp. 8–18.
- 2829 [SB08] Hagit Schechter and Robert Bridson. “Evolving sub-grid turbulence for
2830 smoke animation”. In: *Proceedings of the 2008 ACM SIGGRAPH/Eurographics
2831 symposium on Computer animation*. 2008, pp. 1–7.
- 2832 [Sch+03] Joshua Schpok, Joseph Simons, David S Ebert, and Charles Hansen.
2833 “A real-time cloud modeling, rendering, and animation system”. In:
2834 *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on
2835 Computer animation*. Citeseer. 2003, pp. 160–166.
- 2836 [Sch29] Erich Schoenberg. *Theoretische Photometrie*. Springer, 1929.
- 2837 [Sha+04] Maurya Shah, Jonathan M Cohen, Sanjit Patel, Penne Lee, and Frédéric
2838 Pighin. “Extended galilean invariance for adaptive fluid simulation”. In:
2839 *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on
2840 Computer animation*. 2004, pp. 213–221.
- 2841 [SK10] YJ Suzuki and T Koyaguchi. “Numerical determination of the efficiency
2842 of entrainment in volcanic eruption columns”. In: *Geophysical Research
2843 Letters* 37.5 (2010).
- 2844 [SM20] Microsoft Flight Simulator and Mathias Müller. *Partnership Series:
2845 Meteoblue - Weather Forecast System*. Microsoft Flight Simulator and
2846 Meteoblue. June 2020. URL: <https://web.archive.org/web/20200612005006/>
2847 <https://www.youtube.com/watch?v=3AfQGA8i7uQ>.

- 2848 [Smi+21] Cassandra M Smith, Damien Gaudin, Alexa R Van Eaton, Sonja A
2849 Behnke, Steven Reader, Ronald J Thomas, Harald Edens, Stephen R
2850 McNutt, and Corrado Cimarelli. “Impulsive volcanic plumes generate
2851 volcanic lightning and vent discharges: A statistical analysis of Sakura-
2852 jima volcano in 2015”. In: *Geophysical Research Letters* 48.11 (2021),
2853 e2020GL092323.
- 2854 [SN01] Christopher Small and Terry Naumann. “The global distribution of
2855 human population and recent volcanism”. In: *Global Environmental*
2856 *Change Part B: Environmental Hazards* 3.3 (2001), pp. 93–109.
- 2857 [Søn+20] Casper Kaae Sønderby, Lasse Espenholt, Jonathan Heek, Mostafa De-
2858 hghani, Avital Oliver, Tim Salimans, Shreya Agrawal, Jason Hickey,
2859 and Nal Kalchbrenner. “Metnet: A neural weather model for precipita-
2860 tion forecasting”. In: *arXiv preprint arXiv:2003.12140* (2020).
- 2861 [Sta99] Jos Stam. “Stable Fluids”. In: *Proceedings of the 26th Annual Confer-*
2862 *ence on Computer Graphics and Interactive Techniques*. SIGGRAPH
2863 ’99. USA: ACM Press/Addison-Wesley Publishing Co., 1999, pp. 121–
2864 128. ISBN: 0201485605. DOI: 10.1145/311535.311548. URL: <https://doi.org/10.1145/311535.311548>.
2865
- 2866 [Ste+19] S Stern, Corrado Cimarelli, D Gaudin, B Scheu, and DB Dingwell.
2867 “Electrification of experimental volcanic jets with varying water con-
2868 tent and temperature”. In: *Geophysical Research Letters* 46.20 (2019),
2869 pp. 11136–11145.
- 2870 [Sto+99] Dan Stora, Pierre-Olivier Agliati, Marie-Paule Cani, Fabrice Neyret,
2871 and Jean-Dominique Gascuel. “Animating lava flows”. In: *Graphics In-*
2872 *terface (GI’99) Proceedings*. 1999, pp. 203–210.
- 2873 [Suz+16] YJ Suzuki, Antonio Costa, Matteo Cerminara, T Esposti Ongaro, Michael
2874 Herzog, Alexa R Van Eaton, and LC Denby. “Inter-comparison of three-
2875 dimensional models of volcanic plumes”. In: *Journal of Volcanology and*
2876 *Geothermal Research* 326 (2016), pp. 26–42.
- 2877 [SY05] Lin Shi and Yizhou Yu. “Controllable smoke animation with guiding ob-
2878 jects”. In: *ACM Transactions on Graphics (TOG)* 24.1 (2005), pp. 140–
2879 164.
- 2880 [Tet30] Otto Tetens. “Über einige meteorologische Begriffe”. In: *Z. Geophys* 6
2881 (1930), pp. 297–309.
- 2882 [Tex+05] Christiane Textor, Hans F Graf, Antonella Longo, Augusto Neri, T Es-
2883 posti Ongaro, Paolo Papale, Claudia Timmreck, and GG Ernst. “Nu-
2884 merical simulation of explosive volcanic eruptions from the conduit
2885 flow to global atmospheric scales”. In: *Annals of Geophysics* 48 (2005),
2886 pp. 817–842.
- 2887 [TKP13] Nils Thuerey, Theodore Kim, and Tobias Pfaff. “Turbulent fluids”. In:
2888 *ACM SIGGRAPH 2013 Courses*. 2013, pp. 1–1.
- 2889 [Tre+03] Adrien Treuille, Antoine McNamara, Zoran Popović, and Jos Stam.
2890 “Keyframe control of smoke simulations”. In: *ACM SIGGRAPH 2003*
2891 *Papers*. 2003, pp. 716–723.

- 2892 [Tri+19] Thibault Tricard, Semyon Efremov, Cédric Zanni, Fabrice Neyret, Jonàs
2893 Martínez, and Sylvain Lefebvre. “Procedural phasor noise”. In: *ACM*
2894 *Transactions on Graphics (TOG)* 38.4 (2019), pp. 1–13.
- 2895 [Van+15] Alexa R Van Eaton, Larry G Mastin, Michael Herzog, Hans F Schwaiger,
2896 David J Schneider, Kristi L Wallace, and Amanda B Clarke. “Hail for-
2897 mation triggers rapid ash aggregation in volcanic plumes”. In: *Nature*
2898 *communications* 6.1 (2015), p. 7860.
- 2899 [Van+16] Alexa R Van Eaton, Álvaro Amigo, Daniel Bertin, Larry G Mastin,
2900 Raúl E Giacosa, Jerónimo González, Oscar Valderrama, Karen Fontijn,
2901 and Sonja A Behnke. “Volcanic lightning and plume behavior reveal
2902 evolving hazards during the April 2015 eruption of Calbuco volcano,
2903 Chile”. In: *Geophysical Research Letters* 43.7 (2016), pp. 3563–3571.
- 2904 [Van91] Jarke J Van Wijk. “Spot noise texture synthesis for data visualization”.
2905 In: *Proceedings of the 18th annual conference on Computer graphics and*
2906 *interactive techniques*. 1991, pp. 309–318.
- 2907 [Vim+20] Ulysse Vimont, James Gain, Maud Lastic, Guillaume Cordonnier, Ba-
2908 batunde Abiodun, and Marie-Paule Cani. “Interactive Meso-scale Sim-
2909 ulation of Skyscapes”. In: *Computer Graphics Forum*. Vol. 39. 2. 2020.
- 2910 [Voh+99] O Vohl, SK Mitra, SC Wurzler, and HR Pruppacher. “A wind tunnel
2911 study of the effects of turbulence on the growth of cloud drops by
2912 collision and coalescence”. In: *Journal of the atmospheric sciences* 56.24
2913 (1999), pp. 4088–4099.
- 2914 [WB07] Roy Walmsley and Brian Butters. “Infrared smoke modelling in Coun-
2915 terSim”. In: *Technologies for Optical Countermeasures IV*. Vol. 6738.
2916 SPIE. 2007, pp. 215–223.
- 2917 [Web+18] Antoine Webanck, Yann Cortial, Eric Guérin, and Eric Galin. “Proce-
2918 dural cloudscapes”. In: *Computer Graphics Forum*. Vol. 37. 2. Wiley
2919 Online Library. 2018, pp. 431–442.
- 2920 [Wed+20] Nils P Wedi, Inna Polichtchouk, Peter Dueben, Valentine G Anan-
2921 tharaj, Peter Bauer, Souhail Boussetta, Philip Browne, Willem Decon-
2922 inck, Wayne Gaudin, Ioan Hadade, et al. “A baseline for global weather
2923 and climate simulations at 1 km resolution”. In: *Journal of Advances*
2924 *in Modeling Earth Systems* 12.11 (2020), e2020MS002192.
- 2925 [Wer+19] Maximilian Werhahn, You Xie, Mengyu Chu, and Nils Thuerey. “A
2926 multi-pass GAN for fluid flow super-resolution”. In: *Proceedings of the*
2927 *ACM on Computer Graphics and Interactive Techniques* 2.2 (2019),
2928 pp. 1–21.
- 2929 [WH06] John M Wallace and Peter V Hobbs. *Atmospheric science: an introduc-*
2930 *tory survey*. Vol. 92. Elsevier, 2006.
- 2931 [Wil+21] Alexander Wilkie, Petr Vevoda, Thomas Bashford-Rogers, Lukáš Hošek,
2932 Tomáš Iser, Monika Kolářová, Tobias Rittig, and Jaroslav Křivánek.
2933 “A fitted radiance and attenuation model for realistic atmospheres”.
2934 In: *ACM Transactions on Graphics (TOG)* 40.4 (2021), pp. 1–14.

- 2935 [WK94] Andrew W Woods and Juergen Kienle. “The dynamics and thermody-
2936 namics of volcanic clouds: theory and observations from the April 15
2937 and April 21, 1990 eruptions of Redoubt Volcano, Alaska”. In: *Journal*
2938 *of Volcanology and Geothermal Research* 62.1-4 (1994), pp. 273–299.
- 2939 [WM81] Richard T Wetherald and Syukuro Manabe. “Simulations of Cloud
2940 Cover with a Global General Circulation Model of the Atmosphere”.
2941 In: *Clouds in Climate: Modeling and Satellite Observational Studies*
2942 (1981), pp. 27–33.
- 2943 [Wor96] Steven Worley. “A cellular texture basis function”. In: *Proceedings of*
2944 *the 23rd annual conference on Computer graphics and interactive tech-*
2945 *niques*. 1996, pp. 291–294.
- 2946 [Wri+22] Corwin J Wright, Neil P Hindley, M Joan Alexander, Mathew Bar-
2947 low, Lars Hoffmann, Cathryn N Mitchell, Fred Prata, Marie Bouil-
2948 lon, Justin Carstens, Cathy Clerbaux, et al. “Surface-to-space atmo-
2949 spheric waves from Hunga Tonga–Hunga Ha’apai eruption”. In: *Nature*
2950 609.7928 (2022), pp. 741–746.
- 2951 [Xie+18] You Xie, Erik Franz, Mengyu Chu, and Nils Thuerey. “tempoGAN: A
2952 temporally coherent, volumetric GAN for super-resolution fluid flow”.
2953 In: *ACM Transactions on Graphics (TOG)* 37.4 (2018), pp. 1–15.
- 2954 [Xu+09a] Jiangbin Xu, Hanchen Song, Lingda Wu, and Chao Yang. “A seamless
2955 representation of multi-scale clouds”. In: *2009 Fifth International Joint*
2956 *Conference on INC, IMS and IDC*. IEEE. 2009, pp. 1255–1260.
- 2957 [Xu+09b] Jiangbin Xu, Chao Yang, Jian Zhao, and Lingda Wu. “Fast modeling
2958 of realistic clouds”. In: *2009 International Symposium on Computer*
2959 *Network and Multimedia Technology*. IEEE. 2009, pp. 1–4.
- 2960 [YI07] Akihiko Yokoo and Kazuhiro Ishihara. “Analysis of pressure waves ob-
2961 served in Sakurajima eruption movies”. In: *Earth, planets and space* 59
2962 (2007), pp. 177–181.
- 2963 [YOH00] Gary D Yngve, James F O’Brien, and Jessica K Hodgins. “Animating
2964 explosions”. In: *Proceedings of the 27th annual conference on Computer*
2965 *graphics and interactive techniques*. 2000, pp. 29–36.
- 2966 [Yua+14] Chunqiang Yuan, Xiaohui Liang, Shiyu Hao, Yue Qi, and Qinqing Zhao.
2967 “Modelling cumulus cloud shape from a single image”. In: *Computer*
2968 *Graphics Forum*. Vol. 33. 6. Wiley Online Library. 2014, pp. 288–297.
- 2969 [Zar20] Omar Zarifi. “Sparse smoke simulations in Houdini”. In: *Special Interest*
2970 *Group on Computer Graphics and Interactive Techniques Conference*
2971 *Talks*. 2020, pp. 1–2.
- 2972 [Zha+17] Shenfan Zhang, Fanlong Kong, Chen Li, Changbo Wang, and Hong Qin.
2973 “Hybrid modeling of multiphysical processes for particle-based volcano
2974 animation”. In: *Computer Animation and Virtual Worlds* 28.3-4 (2017),
2975 e1758.

Appendices

Appendix A

2D pressure solve when $\Delta x \neq \Delta y$

$$\begin{aligned}
& \frac{u_{i,j+1/2} - u_{i,j-1/2}}{\Delta x} + \frac{v_{i+1/2,j} - v_{i-1/2,j}}{\Delta y} = 0 \\
& \frac{\left(u_{i,j+1/2} - \frac{\Delta t}{\rho} \frac{p_{i,j+1} - p_{i,j}}{\Delta x}\right) - \left(u_{i,j-1/2} - \frac{\Delta t}{\rho} \frac{p_{i,j} - p_{i,j-1}}{\Delta x}\right)}{\Delta x} + \\
& \frac{\left(v_{i+1/2,j} - \frac{\Delta y}{\rho} \frac{p_{i+1,j} - p_{i,j}}{\Delta y}\right) - \left(v_{i-1/2,j} - \frac{\Delta t}{\rho} \frac{p_{i,j} - p_{i-1,j}}{\Delta y}\right)}{\Delta y} = 0 \\
& \frac{u_{i,j+1/2} - u_{i,j-1/2} + \frac{\Delta t}{\rho \Delta x} (2p_{i,j} - p_{i,j+1} - p_{i,j-1})}{\Delta x} + \\
& \frac{v_{i+1/2,j} - v_{i-1/2,j} + \frac{\Delta t}{\rho \Delta x} (2p_{i,j} - p_{i+1,j} - p_{i-1,j})}{\Delta y} = 0 \\
& \frac{\Delta t}{\rho} \left(\left(\frac{2x}{y} + \frac{2y}{x} \right) p_{i,j} - \frac{y}{x} p_{i,j+1} - \frac{y}{x} p_{i,j-1} - \frac{x}{y} p_{i+1,j} - \frac{x}{y} p_{i-1,j} \right) = \\
& \Delta y \cdot u_{i,j-1/2} - \Delta y \cdot u_{i,j+1/2} + \Delta x \cdot v_{i-1/2,j} - \Delta x \cdot v_{i+1/2,j}
\end{aligned}$$

2979 **Appendix B**

2980 **Derivation of solid interface in**
 2981 **pressure solve**

2982 Let's assume that cells $i, j - 1$ and $i - 1, j$ is solid, and that $\Delta x \equiv \Delta y$. We
 2983 assume that the solid interfaces are stationary (since mountains don't walk) and
 2984 thus $u_S = 0$. Therefore, the final pressure calculation for cell i, j will be

$$\begin{aligned} & \frac{u_{i,j+1/2} - u_{i,j-1/2} + v_{i+1/2,j} - v_{i-1/2,j}}{\Delta x} = 0 \\ & \frac{u_{i,j+1/2} - u_S + v_{i+1/2,j} - u_S}{\Delta x} = 0 \\ & \frac{\left(u_{i,j+1/2} - \frac{\Delta t p_{i,j+1} - p_{i,j}}{\rho \Delta x}\right) - 0}{\Delta x} + \frac{\left(v_{i+1/2,j} - \frac{\Delta t p_{i+1,j} - p_{i,j}}{\rho \Delta x}\right) - 0}{\Delta x} = 0 \\ & \frac{\Delta t}{\rho} \frac{2p_{i,j} - p_{i+1,j} - p_{i,j+1}}{\Delta x^2} = -\left(\frac{u_S - u_{i+1/2,j}}{\Delta x} + \frac{v_{i,j+1/2} - u_S}{\Delta x}\right) \end{aligned}$$

Appendix C

Example calculation of layer heights

The user provides a list of heights (in metres), a maximum altitude, a maximum number of layers, and a maximum layer thickness. More layers require more computation, but thicker layers lead to less fidelity in the output. In this example, we use [500, 1500, 2500, 2700, 3000, 3500, 4500, 5500, 6500] as the provided list, with a maximum altitude of 20000 metres, a maximum layer count of 20, and a maximum layer thickness of 1000 metres.

The initial layers (centred around the altitudes of interest to the user) are generated in a three phase process. Table C.1 demonstrates the values in practice. The first phase iterates through each height h , and finds the distance d to the nearest neighbour in the list. Initial bounds of $[h - d/2, h + d/2]$ are created. This might result in gaps between layers. The second phase iterates through the newly created list of bounds and expands each layer's bounds until the layer is in contact with at least one neighbour. The bounds are sorted with the thickest layer first before the final phase. If a layer is not in contact with both neighbours, it is expanded until all gaps are removed. The sorting is required to ensure that the loss in fidelity is minimised, and areas of interest with narrow separation (e.g., 2500, 2700, and 3000 in this example) maintain that narrow separation.

If, after the third phase, a layer exceeds the specified maximum thickness, it is split evenly to form two layers. In our example, that means the layer with bounds [750, 2400] is split into [750, 1575] and [1575, 2400], and [3850, 5000] goes to [3850, 4425] and [4425, 5000]. Finally, the space between the highest layer generated ([6000, 7000] in our example) and the maximum altitude is evenly divided into layers such that the total number of layers is equal to the maximum number of layers. That is, with 11 layers already generated, the remaining 13000m is divided into 9 layers each 1444.4m thick. Table C.2 provides the reference heights and final bounds for all layers, with the provided list [500, 1500, 2500, 2700, 3000, 3500, 4500, 5500, 6500], a maximum thickness of 1000m, a maximum altitude of 20000m, and 20 as the maximum number of layers.

User height (m)	After phase 1	After phase 2	After phase 3
500	[250, 750]	[250, 750]	[0, 750]
1500	[1000, 2000]	[750, 2250]	[750, 2400]
2500	[2400, 2600]	[2400, 2600]	[2400, 2600]
2700	[2600, 2800]	[2600, 2800]	[2600, 2800]
3000	[2850, 3150]	[2850, 3150]	[2800, 3150]
3500	[3250, 3750]	[3150, 3850]	[3150, 3850]
4500	[4000, 5000]	[4000, 5000]	[3850, 5000]
5500	[5000, 6000]	[5000, 6000]	[5000, 6000]
6500	[6000, 7000]	[6000, 7000]	[6000, 7000]

Table C.1: The bounds of the air layers generated by the user-provided list of altitudes after each phase of the initialisation, where the list provided was [500, 1500, 2500, 2700, 3000, 3500, 4500, 5500, 6500].

Reference height (m)	Bounds
375	[0, 750]
1162.5	[750, 1575]
1987.5	[1575, 2400]
2500	[2400, 2600]
2700	[2600, 2800]
2975	[2800, 3150]
3500	[3150, 3850]
4137.5	[3850, 4425]
4712.5	[4425, 5000]
5500	[5000, 6000]
6500	[6000, 7000]
7722.2	[7000, 8444]
9166.6	[8444, 9889]
10611.1	[9889, 11333]
12055.5	[11333, 12778]
13500	[12778, 14222]
14944.4	[14222, 15666]
16388.8	[15666, 17111]
17833.3	[17111, 18555]
19277.7	[18555, 20000]

Table C.2: The final layer heights and bounds after initialisation. The layers below 7000m are generated from the user-provided list [500, 1500, 2500, 2700, 3000, 3500, 4500, 5500, 6500], with the layers above 7000m evenly divided up to the maximum altitude (20000m) and maximum number of layers (20).