

MASTER THESIS

Density Driven Convection in Porous Media

A combined experimental and numerical study of CO₂ advection
diffusion in water saturated porous media

Author

Hilmar Yngvi Birgisson

Supervisors

Knut Jørgen Måløy

Eirik Grude Flekkøy

July 30th, 2020



UiO : **University of Oslo**

Abstract

Density Driven Convection (DDC) is an interesting effect which appears through a coupling between a flow field and the diffusion of some density altering component. In this project, the convective mechanics of dissolved CO_2 in water are investigated. When CO_2 comes into contact with water it partially dissolves and forms carbonic acid which gives rise to a small density increase and change in acidity level (pH).

A proprietary color indicator and analysis method was used to estimate the pH field from digital images and sample colors, which when coupled with a chemical equilibrium model can estimate the actual concentration field of dissolved carbon species. This goes beyond the more conventional single indicator methods, which are only sensitive to change in a very narrow region on the pH scale. While the strengths of the multi indicator method were clearly demonstrated, the tolerances in physical dimensions of the flow cell constructed were inadequate to guarantee a uniform permeability field, which caused some experimental artifacts to appear.

Along with flow cell experiments a numerical model was developed, which can simulate DDC. The method used is a combination of the Finite Volume Method (FVM) for the advection-diffusion equation, the Finite Difference Method (FDM) for the flow field and an Alternating Direction Implicit (ADI) method for time propagation. A linear algebra solver was optimized for the resulting matrix operations, which resulted in a computationally efficient and versatile program. The numerical results agree well with those obtained from theoretical studies and also demonstrate the complicated non-linear dynamics of DDC.

Preface

Given my interest in inter-disciplinary problems, and previous experience from the field of chemical engineering, I took great interest in tackling a problem which couples chemistry with fluid mechanics. Having gone through the entire process from countless early preliminary tests, to more systematic experimental design and execution, as well as later gaining insight in the world of computational physics, I can safely say that I've been acquainted with a strong arsenal of problem solving tools, which I'll carry with me on my future scientific endeavours.

As is often the case with experimental work, everything rarely goes as planned. Unfortunately, the COVID-19 pandemic did leave its mark on this project due to restrictions in access to campus laboratories during the main round of experiments. To make the best of an otherwise tragic situation, I seized the opportunity to delve deeper into the world of computational physics and simulations. Although not initially a part of the plan, given the circumstances, this part turned out to become a notable portion of the presented work. In a perfect world, more time and resources could have been used to further improve the experiments, but all things considered, I'm grateful for how this project turned out.

Despite this slight setback I'm very thankful for my time at the University of Oslo. I could not have accomplished this thesis without the support of my advisors, PoreLab staff, close friends and relatives, which despite distancing limitations helped me stay on track until the end of this project.

Table of contents

1	Introduction	6
2	Previous work	7
2.1	A brief introduction to carbon capture and storage	7
2.2	Theoretical work	8
2.3	Experiments	9
2.3.1	Quantitative Experiments	9
2.3.2	Experiments on convection plume patterns	9
3	Theory and modelling	10
3.1	Modeling a quasi-2D flow in porous media	10
3.2	The advection diffusion equation	11
3.3	Density Driven Convection in 2 dimensions	11
3.3.1	A simplified 1D case	12
3.4	The fundamentals of the involved chemical equilibria	14
3.4.1	The carbonic acid equilibrium	15
3.4.2	Connecting pH and carbon concentration	16
4	Simulating density driven convection	18
4.1	Developing a suitable numerical model	18
4.1.1	Spatial discretization of the concentration field via FVM	18
4.1.2	Labeling volumes and points	19
4.1.3	Constructing a linear algebra problem for the flow field	21
4.1.4	Time discretization and propagation of the C-field	22
4.2	Computational aspects	25
4.2.1	Efficient solvers for the constructed matrices	26
4.3	Testing the algorithms	28
4.3.1	A quick matrix solver benchmark	28
4.3.2	Sanity checks	30
4.3.3	The need for perturbation	31
5	Numerical results	32
5.1	Understanding the base case	32
5.1.1	A qualitative analysis of the observed patterns	32
5.1.2	A closer look at the top boundary	34
5.1.3	Estimating the characteristic wavenumber	35
5.2	Varying the Rayleigh number	36
5.2.1	The variation in plume patterns	37
5.2.2	Determining the onset time of DDC	39
5.2.3	Determining the initial wavenumber	39
5.2.4	The time dependence of characteristic wavenumbers	40
5.2.5	Total adsorbed mass	41

6	Experimental methods	42
6.1	Designing the experiment	42
6.1.1	The experiment holder	42
6.1.2	Constructing the flow cell	43
6.2	Visualizing acidity	44
6.2.1	Making a suitable indicator liquid	45
6.2.2	Image analysis	46
6.3	Refining the experiments	49
7	Experimental results	53
7.1	Patterns in the 2mm cell	53
7.2	Patterns in the 1mm cell	56
8	Conclusions	64
9	Bibliography	70

List of commonly used symbols and abbreviations

Dimensional quantities regarding DDC	
X, Y	Position
U	Velocity
P	Pressure
t	Time
g	Gravitational acceleration
ρ	Fluid density
Ψ	Stream function
D	Molecular diffusivity constant
C	Volumetric concentration
b	Hele-Shaw cell aperture
k	Permeability
μ	Dynamic viscosity
Dimensionless quantities	
x, y	Position
u	Velocity
τ	Time
Ra	Rayleigh number
ψ	Stream function
f	Permeability fluctuation function
c	Volumetric concentration
Chemical equilibrium quantities	
p_{CO_2}	Partial pressure of CO_2
K	Chemical equilibrium constant
$[A]$	Concentration of chemical species A
pH	Acidity, $-\log_{10}([H])$
pK	Shorthand for "potency of", defined as $-\log_{10}(K)$
C_T	Total carbon concentration
Mathematical operators	
$\frac{\partial}{\partial \zeta}$	Partial derivative
$\frac{D}{Dt}$	Material derivative
∇	Dimensional gradient, with respect to X, Y
$\tilde{\nabla}$	Dimensionless gradient, with respect to x,y

1 Introduction

Consider a system consisting of gaseous Carbon Dioxide (CO_2) and liquid water. At ambient conditions a chemical equilibrium exists which causes the CO_2 to dissolve and diffuse into the water and form carbonic acid. The rate of dissolution depends on the CO_2 concentration at the boundary and thus slows down and stagnates when the equilibrium is reached. However when one couples these dissolution dynamics to a flow field, something interesting happens. Since carbonic acid slightly increases the density of the liquid, the chemical dissolution at the boundary will cause a natural convective flow - which in turn brings more "fresh" liquid to the boundary. We call this phenomena Density Driven Convection (DDC). This has the effect of increasing the potential rate of CO_2 uptake of the system compared to the purely diffusive case.

Understanding these dynamics better is paramount to understanding the hydrodynamic effects involved with carbon capture and storage (CCS), since CCS is often realized by pumping CO_2 into underground geological formations sealed by some cap rock. Because of buoyancy the gas will seek to the top and the previously described dynamics take place at the gas-liquid interface [1].

To simulate this, first density driven convective flow is modelled in 2D porous media and coupled with the advection diffusion equation. Using dimensional analysis one can quantify different convective regimes and see the effect of varying the Rayleigh number, which is a number which balances the rate of convection and diffusion. Previous theory suggests the existence of a Rayleigh number threshold - below which the diffusion is fast enough to continuously stabilize the concentration front. Above the critical Rayleigh number, fluctuations in the concentration front can increase exponentially, which leads to a strong instability which causes the sinking finger-like structures to form.

The plumes which form during DDC have an interesting morphology, which can vary with system conditions. Therefore an experiment will be designed from the ground up, that is able to visualize said plumes by developing suitable experimental hardware, colorful solutions and image analysis. An example of plumes visualized with this method is shown in Fig.1. This approach is aimed at attempting to determine how the pH field of these plumes appear and vary with time. Given the pH field and certain assumptions from chemistry, this approach will also be used to try and deduce the spatial fluctuations in carbon dioxide content.




Figure 1: The liquid consists of water and a pH indicator, which is turquoise at neutral conditions and becomes yellow and orange at more acidic conditions. When gas is introduced at the top of the cell, the boundary becomes very acidic, and mildly acidic fingers begin to sink further into the media.

2 Previous work

2.1 A brief introduction to carbon capture and storage

With the current state of our energy affairs, and the climate risk associated with greenhouse gas emissions, scientists keep pursuing different ways of pushing our energy society towards a greener future with lower emissions. In tandem with a goal for reducing emissions is the relatively new approach of carbon capture and storage (CCS), which is an approach that can help achieve carbon neutrality, and although unfeasible today, could in theory be carbon negative.

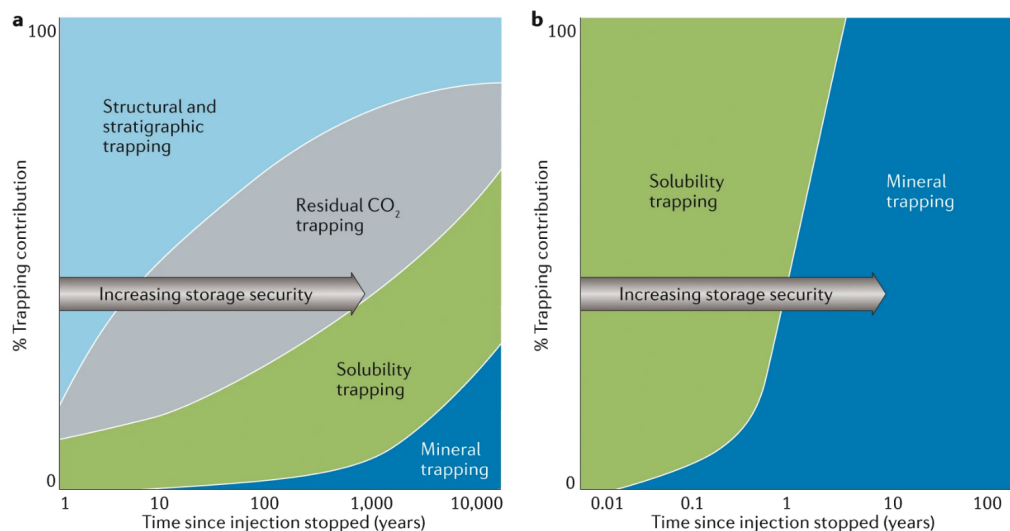


Figure 2: Change in the contribution of the carbon-trapping mechanism of CO₂ storage over time when injecting pure supercritical CO₂ into sedimentary basins (part a) and when injecting water-dissolved CO₂ for mineralization (part b), based on data from field injection experiments from the CarbFix project. Figure reprinted from [2]

CCS is most commonly realised by isolating CO₂ from high emission point sources, such as power plants, cement factories, aluminum refineries etc. and pumping it at high pressure or supercritical conditions into underground geological formations [2]. These formations need to be sealed with a relatively impermeable cap rock to prevent seepage or leakage of the buoyant CO₂. Certain oil and geothermal reservoirs have proven very useful in this regard [3].

One of the most important considerations of CCS is how the CO₂ is actually trapped for storage, as to not leak and contaminate the surrounding environment with carbonic acid. In literature, the process is often broken down into several trapping mechanics [2][3]. Initially, CO₂ is physically trapped by being injected below a caprock of low permeability. As the name suggests, this implies the injection site has a suitable geometry and low enough permeability to sufficiently contain the buoyant CO₂ from rising upwards. These traps commonly consists of folded or fractured rock.

Due to the aforementioned buoyancy of CO₂, it will migrate through the surrounding structure, and some of it will in turn be retained in pore spaces by capillary forces. The capillary forces block the formation water from entering these pore spaces again, and therefore the residing CO₂ becomes immobilized. This mechanism of trapping is typically referred to as residual trapping [1].

When the injected CO_2 comes in contact with formation water, solubility trapping will also contribute to the carbon storage process. Solubility trapping refers to trapping by dissolution into liquids already present in the geological feature. The precise details of this depends heavily on the conditions[4]. For example higher pressures and supercritical conditions speed up the dissolution process. Density Driven Convection (DDC) is also a key component of the solubility trapping mechanic, since it can accelerate the diffusive mixing that drives the dissolution process. One clear desirable aspect of solubility trapping is that the CO_2 enters the water phase, and dissolved CO_2 is thus not prone to seeping out of the reservoir due to buoyancy [1] [3]. On the contrary, solubility trapping increases fluid density and the affected liquid will instead sink. The physical and chemical modeling of this is the main subject of the work at hand, and will be discussed in further detail in the following sections.

Last but not least is the process commonly referred to as mineral trapping. Upon dissolution, the CO_2 will with the present formation water form carbonic acid. Depending on the geochemistry of the surrounding rock, a chemical reaction may take place between the ionized acidic liquid and minerals. On a longer time scale, carbonate minerals may form, which get deposited in the pore spaces of the formation rock. This is considered the most permanent form of carbon storage, but is also the slowest to take effect.

The comparison of the contributions and time scales of these individual dynamics are illustrated on Fig.2. For clarification, subplot **b** is derived from the CarbFix project, in which a mixture of CO_2 and water is injected into highly reactive basaltic formations, so the structural and residual steps do not apply. subplot **a** is a more conventional approach [2][3]. This underlines the important contribution of solubility trapping and DDC with regards to CCS. This is a clear motivation to better understand DDC, both the complicated nature of natural convection and its potential role in climate mitigation methods.

2.2 Theoretical work

The nonlinear dynamics of DDC are complicated by nature. At or around the onset of convection, linear stability analysis can be used to lay a theoretical foundation for the growth of an instability, caused by a density gradient [5]. From a mathematical standpoint, the modeling of DDC caused by CO_2 dissolution is analogous to that of density gradients due to thermal expansion, which linear stability methods have also been applied towards [6][7]. Studies like these first determine the relevant equations (and boundary conditions) of motion, diffusion of heat and/or solute and how the density varies with concentration or temperature. Then by imposing a small perturbation, one can solve for how these perturbations grow or shrink with time. These studies agree that the critical Rayleigh number at which DDC becomes possible is $Ra_c = 4\pi^2$.

Furthermore, linear stability analysis also demonstrates that different wavelengths will grow with different rates. Hassanzadeh et al. found a reciprocal relationship between fastest growing wavenumber and Rayleigh number [8]. Thus if the underlying assumptions are valid, one would expect the emerging wavelength to decrease noticeably with increasing Rayleigh number.

Additionally, the same theory can also estimate the time at which convective dynamics start to dominate [8], often referred to as the onset time or critical time. It is shortly after this time that the nonlinear behaviour takes over. Through linear stability analysis, the onset time is expected to be inversely proportional to the Rayleigh number squared, indicating that lower density gradients take significantly longer to develop convection plumes. This was reported by Ennis-King et al., who also demonstrated that the anisotropy of the porous media can play a significant role in the onset time [4].

2.3 Experiments

2.3.1 Quantitative Experiments

In this category, experiments that are concerned with the macroscopic observables are placed. Most prominently, this involves the time dependent rate of mass transfer to the system. From this one can deduce the mass flux at different times, and estimate an effective diffusion constant. These experiments are conducted in so called PVT cells. The fundamental idea is to build a an isolated chamber in which DDC is to take place. During experiments the chamber contains only a water saturated porous medium and gaseous CO_2 , which starts the DDC process. There are two approached to obtain useful data from such experiments. One can either inject a known quantity of CO_2 and then fully isolate the cell, and track the reduction in gas pressure as it dissolves into the water [9], or use a metered pump to maintain a constant pressure while measuring the flux of gas needed to maintain said pressure [10][11]. Both of these approaches show that the effective diffusion rate during density driven convection is around one order of magnitude faster than that of purely molecular diffusion. As time progresses, the water wet porous medium becomes more and more saturated. Therefore at later times, there is a much smaller density gradient driving the flow, and thus less fresh fluid is getting replenished to the liquid-gas interface, slowing down the mass flux until it approaches the flux expected by molecular diffusion. This behaviour was clearly visualized by Tang et al. who determined a time dependent effective diffusion constant which at its peak was roughly 24 times that of molecular diffusion of CO_2 [11].

2.3.2 Experiments on convection plume patterns

Much of the experimental work on plume patterns relies on the use of pH color indicators to visualize the convection plume patterns in 2D Hele-Shaw cells [12][13]. These studies can relatively easily show the CO_2 affected fluid regions, and how they propagate in time. This offers the possibility of tracking plume locations, lengths, areas etc. Compared to their relative simplicity, these experiments can give much intuitive insight into the complex plume patterns observed during DDC. This has in turn raised questions about the methodology of this visualization technique. Thomas et al. demonstrated that the resulting morphology of the convection plumes largely depends on the color indicator being used [14]. This was determined by repeated experiments in Hele-Shaw cells colored by Bromocresol Green and Bromocresol Purple. Additionally, the group compared the colored plume morphology obtained via a color indicator to fluctuations in refractive index obtained by Schlieren imaging. Among the conclusions were that the choice of indicator does in fact dictate what the imaged patterns will look like, but did not appear to have an effect on the dynamics themselves. This comes from the fact that the previous color indicator approaches have no meaningful way of estimating an amount of dissolved carbon, and as such, the collected data more resembles a binary image of an affected vs an unaffected area. This is a motivation to try and improve on the color indicator approach, in search of data more suitable to quantitative analysis. This will therefore be the main goal of this project, to design a DDC experiment with a suitable color indicator and image analysis method which is able to capture a continuous pH field and estimate the amount of dissolved carbon species from image data alone.

3 Theory and modelling

3.1 Modeling a quasi-2D flow in porous media

The system under consideration is a Hele-Shaw cell, which is characterized by confining fluid flow between two plates with a narrow spacing between plates. For the flow to be modelled as 2-dimensional, we assume the spacing between the plates is minuscule compared to the width and height of the cell. As an example of scale the plate spacing is usually about a few mm, while the width and height are typically on the scale of half a meter. For creeping flows, one typically uses Darcy's law to model fluid flow in porous media.

$$\vec{U} = -\frac{k}{\mu}(\nabla P - \rho\vec{g}) \quad (1)$$

which states that the flow follows the negative pressure gradient (in this case biased by gravity). k is the *permeability* of the medium and μ the viscosity of the liquid. For a typical Hele-Shaw cell with inner spacing b , the permeability would be $k = b^2/12$. Here the sign has been chosen such that the Y coordinate increases downwards and is 0 at the gas-liquid interface.

If one is mainly interested in the velocity field, Eq.1 can be manipulated to eliminate the pressure field - and thus reduce the amount of unknowns. For now, assume that only μ and \vec{g} is constant, but all other terms of Eq. 1 are generally 2D vector or scalar fields. One finds that:

$$\frac{1}{k}\vec{U} = -\frac{1}{\mu}(\nabla P - \rho g\hat{e}_Y)$$

Assuming 2D geometry, there is no Z -dependence, but it can still be helpful to consider these fields as a subspace in 3D. Then one can take the curl of both sides. Knowing that the curl of a field is always orthogonal to the field in question - the resulting curls will only have z -components, from which the following expression is obtained:

$$\frac{\partial}{\partial X} \frac{U_Y}{k} - \frac{\partial}{\partial Y} \frac{U_X}{k} = \frac{\partial}{\partial X} \frac{\rho g}{\mu}$$

Here the pressure field was eliminated - since the curl of any gradient field is 0. One can further reduce the amount of unknowns by introducing a vector potential:

$$\vec{U} = -\nabla \times \Psi \hat{e}_Z \quad (2)$$

Which similarly is confined to the 2D plane, since the curl of a purely Z -directed field will be orthogonal to the Z -direction. Thus instead of solving for a velocity field with 2 components, one can solve for a scalar function Ψ , which contains the velocity field as its partial derivatives:

$$U_X = -\frac{\partial \Psi}{\partial Y} \quad U_Y = \frac{\partial \Psi}{\partial X} \quad (3)$$

Such a function ψ is commonly known as a *stream function*, and is a very helpful tool for solving 2-Dimensional creeping flows. With this redefinition of the velocity field, the equation for Ψ becomes:

$$\nabla \cdot \left(\frac{1}{k} \nabla \Psi \right) = -\frac{1}{k^2} \nabla k \cdot \nabla \Psi + \frac{1}{k} \nabla^2 \Psi = \frac{g}{\mu} \frac{\partial \rho}{\partial X} \quad (4)$$

An interesting side effect of the stream function method is that the contours of the stream function correspond to stream lines, which are always tangential to the flow field. This can easily be shown by the fact that:

$$\nabla \Psi \cdot \vec{U} = -\frac{\partial \Psi}{\partial Y} \frac{\partial \Psi}{\partial X} + \frac{\partial \Psi}{\partial Y} \frac{\partial \Psi}{\partial X} = 0$$

We know that gradients of scalar fields are orthogonal to the contour lines of said fields. Thus if the velocity field is orthogonal to the gradient of the stream function, that must mean that the velocity field is tangential to the contours of ψ .

A few words on continuity

For mass conservation to be fulfilled, we need the material derivative of the density to be zero.

$$\frac{D\rho}{Dt} = \frac{\partial\rho}{\partial t} + \nabla \cdot (\rho\vec{U}) = 0 \quad (5)$$

Which tells us that any change in mass of a fluid element is caused by some net imbalance (divergence) of mass flux in and out of the element. Now, although the density itself is the driving force of the flow under consideration, the density changes are minuscule, and thus a Boussinesq approximation can be applied. This reduces the continuity equation to:

$$\nabla \cdot \vec{U} = 0 \quad (6)$$

Which is already satisfied by Eq.2, since the divergence of any curl is 0.

3.2 The advection diffusion equation

Consider some scalar parameter γ , defined in the liquid. The material derivative of γ refers to the time-change of γ as viewed from a frame of reference which moves with a given volume element.

$$\frac{D\gamma}{Dt} = \frac{\partial\gamma}{\partial t} + \vec{U} \cdot \nabla\gamma \quad (7)$$

This could in principle be any scalar quantity, like heat or a concentration of a solute material. For DDC we're interested in the transport of dissolved CO₂, or carbonic acid. For modeling purposes we will simply refer to this as *chemical*, and represent its volumetric concentration with a scalar field C . Any dissolved chemical species will diffuse through the solvent via molecular diffusion. The rate of molecular diffusion is proportional to the laplacian of the field. Thus:

$$\frac{\partial C}{\partial t} + \nabla \cdot (\vec{U}C) = D\nabla^2 C \quad (8)$$

So the time rate of change of the chemical concentration becomes:

$$\frac{\partial C}{\partial t} = \nabla \cdot (D\nabla C - \vec{U}C) \quad (9)$$

3.3 Density Driven Convection in 2 dimensions

When the aqueous system is exposed to gaseous CO₂, the gas will slowly diffuse into the liquid. If the water is exposed to a constant partial pressure of CO₂, the volumetric concentration will be determined by the carbonic acid equilibrium. Thus, assume that the liquid has a density ρ_0 at little or no chemical content C_0 , and a density ρ_e , when the chemical equilibrium concentration C_e has been reached. Assuming the density change is linear between these points, we expect:

$$\rho = (\rho_e - \rho_0) \frac{C - C_0}{C_e - C_0} + \rho_0 \quad (10)$$

From which the right hand side of Eq.4 becomes dependant on the horizontal partial derivative of the chemical concentration field. Introduce the following rescaling factors:

$$X_c = H \quad U_c = \frac{\phi D}{H} \quad t_c = \frac{H^2}{D} \quad (11)$$

Where H is the vertical height of the porous media. ϕ is the porosity of the porous media, which is the fraction of volume accessible by the liquid. Finally, D is the molecular diffusion constant. One should note that D is generally dependant on the chemical solute species, the solvent and thermodynamic conditions. The concentration field is rescaled linearly, such that dimensionless concentration is 0 initially and approaches 1 at equilibrium:

$$c = \frac{C - C_0}{C_e - C_0} \quad (12)$$

A dimensionless function f is introduced, which serves to introduce permeability fluctuations. The permeability is then $k = fk_m$ where k_m is an average permeability. For this definition to be applicable, f needs to be strictly positive and average to 1. Now the dimensionless Rayleigh number can be introduced:

$$Ra = \frac{\Delta\rho g k_m H}{\phi D \mu} \quad (13)$$

Here $\Delta\rho = \rho_e - \rho_0$ is a constant which determines the density change caused by the dissolved chemical. g is the gravitational acceleration, μ and the liquid viscosity. This number compares the rate of convective transport to diffusive transport, and a change in Ra can represent many different changes in physical parameters. An increase in Ra can for instance be achieved by making the porous medium more permeable or the liquid less viscous, increasing the fluid flow and thus convection. Similarly one can lower Ra by for instance increasing the porosity (which is the volume fraction of the porous medium accessible by the liquid) or by introducing a tilt, replacing g by $g \sin \theta$.

With these definitions of characteristic scales, the coupled system of equations can be written in dimensionless units:

$$-\frac{1}{f^2} \tilde{\nabla} f \cdot \tilde{\nabla} \psi + \frac{1}{f} \tilde{\nabla}^2 \psi = Ra \frac{\partial \tilde{c}}{\partial x} \quad (14)$$

$$\frac{\partial \tilde{c}}{\partial \tau} = \tilde{\nabla} \cdot (\tilde{\nabla} \tilde{c} - \tilde{u} \tilde{c}) \quad (15)$$

$$\psi = 0 \text{ on all boundaries} \quad c(x, y, t = 0) = 0$$

$$c(x, y = 0, \tau) = 1 \quad \tilde{\nabla} \tilde{c} \cdot \tilde{n} = 0 \text{ on the remaining boundaries}$$

In these two equations, a tilde denotes a dimensionless variable. Note that the gradient terms are also derivatives with respect to dimensionless variables. When the equations have been non-dimensionalized in this manor, the Rayleigh number is the one remaining constant which governs the set of coupled equations. For the numerical section the tildes will be dropped, and all presented values will be dimensionless as per the rescaling shown in this section, unless otherwise is specifically stated.

3.3.1 A simplified 1D case

It is known that at low Rayleigh numbers, the diffusing concentration front is stable enough to inhibit convection fingering. In this case, the flow field has neglectable effect on the propagation of

the front, and in this limit, we can solve for the time dependent concentration field using separation of variables. This simple case will of course tell us nothing about the convective transfer, but can serve as a comparison to a system in which only diffusion takes place, that is a $Ra \rightarrow 0$ limit. Consider the simplified 1D case:

$$\left(\frac{\partial}{\partial t} - \frac{\partial^2}{\partial y^2}\right) c(y, \tau) = 0 \quad (16)$$

$$c(y, 0) = 0 \quad \left.\frac{\partial c}{\partial y}\right|_{y=1} = 0 \quad c(0, \tau) = 1$$

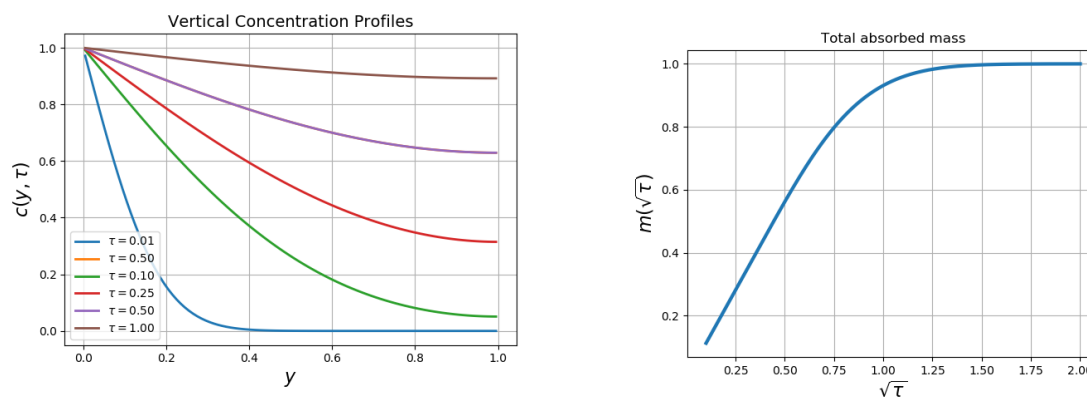
So here we assume an initial concentration field of 0, a zero-flux criteria at the bottom of the box and a constant concentration at the top. Using separation of variables, one can find a shifted sine-series expansion which satisfies the PDE and the boundary conditions. A more detailed derivation of this is given in the appendix.

$$c(y, \tau) = 1 - 2 \sum_{n=0}^{\infty} k_n^{-1} \exp(-k_n^2 \tau) \sin k_n y \quad k_n = \frac{\pi(1+2n)}{2} \quad (17)$$

Which when integrated over the domain gives an analytical expression describing the total mass of dissolved chemical in the liquid:

$$m(\tau) = \int_0^1 c(y, \tau) dy = 1 - 2 \sum_{n=0}^{\infty} k_n^{-2} \exp(-k_n^2 \tau) \quad (18)$$

These expressions are difficult to express in other terms than infinite series. They can however be plotted, if truncated after the first 1000 terms. If the mass is instead plotted as a function of the square root of time, which is often more descriptive of diffusion in open systems, one observes linear behaviour before the system gets saturated. The concentration profiles of Eq.17 and absorbed mass of Eq.18 is shown on Fig.3.



(a) Concentration fields at various times. Notice how initially the front is very sharply defined, but smoothes out over time.

(b) Total amount of dissolved chemical as a function of the square root of time. Initially the mass change with respect to \sqrt{t} is linear, with a slope of 1.126.

Figure 3: Analytical solutions for a simplified 1D case. Notice how our choice of dimensionless time (dictated by the diffusive time scale) implies that the system saturates via diffusion at $t \simeq 1$

Now if the effect of the flow field is strong enough, one would expect the rate of chemical dissolution to be quicker. The analytical solutions presented in Fig.3 can however serve as a comparison between the purely diffusive case, and the advection-diffusion case.

3.4 The fundamentals of the involved chemical equilibria

In the previous sections, we have referred to a concentration field of some chemical. That does suffice as a simple physical model, which is able to account for density driven convection of some dissolved species which alters the fluid density. However, as the chemical in question is carbonic acid, one should also familiarize with the involved chemical reactions and equilibria. Consider, without loss of generality a chemical reaction of the form:



The two sided arrow indicates that the reaction can go both ways. When A and B react to form C and D one says that the reaction goes forwards. When the rate of the forward going reaction matches that of the backwards reaction, a *chemical equilibrium* is reached. For the reaction above the associated *equilibrium constant* is:

$$K = \frac{[C]^\gamma [D]^\delta}{[A]^\alpha [B]^\beta} \quad (20)$$

Here we adopt the conventional notation that square brackets around the name of a chemical species denotes the concentration of said species once equilibrium is reached. The equilibrium constant is derived as the ratio of the rate constants of the forwards and backwards reaction of Eq.19. Thus one says that an equilibrium is reached when these rates are equal, and at this equilibrium, the concentration of all components are constrained by the equilibrium constant. Furthermore, one should mind the following [15]:

- Equilibrium constants are generally defined for reactions of aqueous solutions and/or gases. The concentration of the solvent (e.g. water) does not appear in the expressions, even though it might take part in the reaction.
- For reactions in which one or more of the components are in a gaseous phase, the concentration of said species is generally replaced by it's partial pressure.
- Granted the units of equilibrium constants depend greatly on the particular reaction, most literature unfortunately leaves the units out and assumes molar concentration (**M=mol/L**) for dissolved species and **atm** for gaseous species.
- In general, equilibrium constants are very temperature dependent (since the forwards and backwards rate laws typically include a Boltzmann term).

Table 1: Equilibrium constants relevant to the carbonic acid equilibrium at ambient conditions [16][17]. Acid pK_a values are presented for the dissociation equilibria.

Equilibrium	Definition	Eq. Constant	pK_a
Henry dissolution	$K_h = \frac{[CO_2]}{p_{CO_2}}$	$3.4 \cdot 10^{-2}$ M/atm	-
Hydration	$K_H = \frac{[H_2CO_3]}{[CO_2]}$	$1.66 \cdot 10^{-3}$	-
First Dissociation	$K_{a1} = \frac{[H^+][HCO_3^-]}{[H_2CO_3]}$	$2.67 \cdot 10^{-4}$ M	3.6
Second Dissociation	$K_{a2} = \frac{[H^+][CO_3^{2-}]}{[HCO_3^-]}$	$4.47 \cdot 10^{-11}$ M	10.3

3.4.1 The carbonic acid equilibrium


The equilibrium reactions between CO_2 and water is an example of a more involved system of equilibrium reactions. This equilibrium is of interest to various fields of the natural sciences, as it dictates processes ranging from how oceans acidify due to increased carbon emissions to how living organisms regulate their pH by breathing.

Consider a body of water in contact with the atmosphere. Assume the idealized glass of water is initially pure, containing no CO_2 or other dissolved species. Conceptually one can think of the acidifying process as the following reactions.


- Gas dissolution:



- Hydration:



- First dissociation:



- Second dissociation:




So through dissolution, hydration and two dissociations, each CO_2 molecule has the potential to yield 2 hydrogen ions in the glass of water, thus acidifying it. Each of the 4 steps have their associated equilibrium constants. It should be stressed that although presented in this order, the 4 reactions form a coupled system of simultaneous equilibria. The equilibrium constants are in general temperature dependent, and can also depend on ionic strength, i.e. the presence of other ions in solution. For the presented work, neglectable effect of other ions is assumed, so values for pure water at ambient conditions are used. The associated constants are presented in Tab. 1. These equilibria are able to account for how a body of water left for a long time in contact with gaseous CO_2 will acidify. Through a derivation, similar to the one in the following section, one can show that at equilibrium (and for $\text{pH} \ll \text{p}K_{a2}$), the equilibrium hydrogen ion concentration is governed by the following equation:

$$[\text{H}^+] = -\frac{1}{2}[\text{Na}^+] + \sqrt{\frac{1}{4}[\text{Na}^+]^2 + p_{\text{CO}_2}K_hK_HK_{a1} + 10^{-14}M^2} \quad (25)$$

This tells us exactly how a body of water in equilibrium with a known partial pressure of gas will become acidic. This equation is not explicitly used in the future experimental methods, but gives an idea about the range and nature of the chemical dynamics to be investigated. With no NaOH addition, at normal atmospheric conditions ($\simeq 350\text{ppm}$), the pH would be around 5.7. However the same body of water is in equilibrium with pure 1 atm CO_2 , the pH reaches just below 4. The exact relationship is shown on Fig.4a, both for pure water and for a small addition of NaOH matching that of the liquid used for experiments.

Similarly, the so called Bjerrum plot on Fig.4 shows the fractional composition of carbon species for a given total amount of carbon by Eq.29. The plot is derived by taking the individual components and evaluating their concentrations with respect to the total dissolved carbon C_T .



(a) The pH of a body of water in contact with specific partial pressures of CO_2 gas. Two plots are shown, one for pure water and the other for water which initially has a small concentration of NaOH base, equal to the concentration used in the presented experiments. With the added base, at 1 atm, the equilibrium is reached slightly above pH 4.

(b) A Bjerrum plot, displaying how the fractional composition of dissolved carbon species vary with pH due to the involved chemical equilibria. Note that the blue line includes both hydrated carbonic acid and unhydrated CO_2 , since hydration is independent of acidity. However, the hydration equilibrium constant in Tab.1 tells us that at any time, there is roughly 600 times as much unhydrated CO_2 as there is carbonic acid [16].

Figure 4: Some analytical chemical characteristics of the carbonic acid equilibrium. Here the partial pressure dependence on pH, and the relative compositions are shown.

3.4.2 Connecting pH and carbon concentration

The four aforementioned chemical equilibria describe how gaseous CO_2 can interact with water solutions to acidify them. Ultimately the goal with respect to experimental work is to estimate a concentration of dissolved carbon for a given observed pH value. Chemical equilibria, such as the carbonic acid system are inherently macroscopic definitions, often used in analytical chemistry in which solutions are commonly assumed to be completely homogeneous, having uniform concentrations throughout the entire solution. With that in mind, one can with relative ease deduce equilibrium concentrations and acidities of bodies of water in contact with carbon dioxide with a known partial pressure. However, in the system under consideration, the concentration fields are clearly not uniform, and thus the equilibrium models need to be applied differently.

Therefore a model is proposed, in which every small fluid element will be treated as being in pseudoequilibrium, such that the various forms of dissolved carbon are in equilibrium with each other, but not in equilibrium with the gaseous CO_2 . This is intuitive in the sense that the fluid elements under consideration are below the gas-liquid interface. Thus each fluid element will be treated as having some total carbon content (which gives rise to its density increase) and some observable pH. Of the dissolved species, only the dissociated carbonic acid forms affect pH, so for a given pH the concentration of dissociated carbonic acid can easily be found. If the underlying assumption that the dissolved species are in equilibria with each other is applied, this can in turn determine the total amount of dissolved carbon.

In order to use the equilibria to estimate a connection between pH and carbonic acid content, one needs to introduce two other useful equations, commonly used in analytical chemistry. These are the electrical charge conservation of the solution and the self-ionization equilibrium of water.

The charge conservation requirement simply states that the charge weighted total concentration of positive and negative ions must cancel each other out, thus:

$$[\text{H}^+] + [\text{Na}^+] = [\text{OH}^-] + [\text{HCO}_3^-] + 2[\text{CO}_3^{2-}] \quad (26)$$

Apply the water self-ionization condition $[\text{H}^+][\text{OH}^-] = 10^{-14}M^2$, to find the amount of carbonic acid for a given proton concentration (and thus a given pH). Also, it's helpful to introduce the second dissociation equilibria so one can solve for just one of the components:

$$[\text{HCO}_3^-] = \frac{1}{1 + 2K_{a2}/[\text{H}^+]} \left([\text{Na}^+] + [\text{H}^+] - \frac{10^{-14}M^2}{[\text{H}^+]} \right) \quad (27)$$

Now for applications within DDC, it's ultimately the total amount of dissolved carbon that's of interest. Thus, define the total amount of dissolved carbon as the following:

$$C_T = [\text{CO}_3^{2-}] + [\text{HCO}_3^-] + [\text{H}_2\text{CO}_3] + [\text{CO}_2] \quad (28)$$

Here $[\text{CO}_2]$ implies dissolved carbon dioxide (CO_{2aq}), not to be confused with the gaseous carbon dioxide above the liquid interface. Solving the equilibrium equations for these three forms of carbon species, one finds that:

$$C_T = \frac{1}{1 + 2K_{a2}/[\text{H}^+]} \left([\text{Na}^+] + [\text{H}^+] - \frac{10^{-14}M^2}{[\text{H}^+]} \right) \left(\frac{K_{a2}}{[\text{H}^+]} + 1 + \frac{[\text{H}^+]}{K_{a1}} + \frac{[\text{H}^+]}{K_{a1}K_H} \right) \quad (29)$$

where K_H , K_{a1} and K_{a2} are the equilibrium constants for the hydration of CO_2 and dissociation steps of carbonic acid respectively. Their values are given in Tab.1. The expression above, along with the fact that $[\text{H}^+] = 10^{-\text{pH}}$, gives an estimation of the total carbon concentration of a given fluid element based on its observed pH. A more detailed derivation of this result as well as Eq.25 is given in the appendix.

The prefactor in Eq.29 involving K_{a2} accounts for the behaviour at high pH. One can clearly see that if $\text{pH} \ll \text{p}K_{a2}$ that the fraction would vanish and approach 1. Therefore, as indicated on the Bjerrum plot on Fig.4b, it's safe to neglect the presence of the carbonate ion at pH values significantly lower than 10. The term can for instance be dropped in cases where one might want to solve this expression for pH rather than C_T . In such cases, dropping said fraction yields a simple quadratic formula, much simpler to solve than the otherwise cubic expression. For C_T estimations from pH however, the term is left in, although it has diminishing effect at the pH values that are of interest.

It should be clearly stated that this estimate of total dissolved carbon is rather crude by nature. In general, equilibrium reactions are most commonly used in a macroscopic sense (i.e. to represent the entire cell). The presented pseudoequilibrium approach is only valid if the underlying equilibrium reactions have rates faster than the characteristic time scale. For non-equilibrium situations, one would have to solve the full reaction kinetic equations, which is not possible by only observing pH. Therefore this shortcoming is simply stated as fact, and all estimated carbon concentrations will assume that the characteristic time scales of the convective transport is significantly slower than those of the equilibrium reactions. Also, the presence of pH indicators and NaOH will change the ionic strength and charge balance equation slightly. Considering the crudeness of the assumptions already made, these corrections have been left out.

4 Simulating density driven convection

4.1 Developing a suitable numerical model

As with any problem of a differential nature, it needs to be discretized in order to simulate it numerically. Granted we have derived two coupled differential equations, we need to find a suitable domain on which we can approximate a solution to Eq.14 and Eq.15. For the concentration field, it is advantageous to utilize a Finite Volume Method (FVM), due to the fact that it's derived from fluxes between control volumes and is thus guaranteed to be mass conserving, in such a way that the total mass only changes due to the flux at the top [18]. For the flow field, the finite difference method (FDM) is used, which yields a largely similar discretization, albeit with a simpler representation of the Dirichlet boundary conditions [19]. The methodology described in the following sections is comparable to that of Farajzadeh et. al. The previous work concerns the same dimensionless equations, but does not fully explain the computational method utilized [20]. Therefore, the following proprietary scheme is developed, suitable for the task at hand:

4.1.1 Spatial discretization of the concentration field via FVM

The FVM utilizes the integral form of the differential problem at hand and quantifies fluxes of chemical solute between so-called *control volumes*. Consider a control volume, in our case defined as a small area in the xy plane. If we consider the rate of change of chemical inside the control volume and average it over the area of the control volume, one finds that:

$$\frac{\partial c_{CV}}{\partial t} \simeq \frac{1}{A} \int_{CV} \frac{\partial c}{\partial t} dA \quad (30)$$

From Eq.15 we see that the right hand side is the divergence of a vector field. Thus we can apply Gauss' divergence theorem, to find that:

$$\int_{CV} \frac{\partial c}{\partial t} dA = \int_{\partial CV} (\nabla c - \vec{u}c) \cdot \hat{n} dS \quad (31)$$

Where ∂CV denotes the boundary of the control volume and dS an infinitesimal line segment of the boundary. Here \hat{n} is an outward facing unit normal, which the diffusive and advective current is dotted with. Until now we have made no assumptions regarding the control volume geometry. For the following computations a rectangular CV geometry will be assumed. In the FVM, each CV contributes one algebraic equation linking it to its neighbours, so for a rectangular CV we need to express across 4 line segments. In each of these equations, the CV it is derived from will be denoted P, and its neighbors N,S,E and W as per the compass directions (see Fig. 5. Lower case letters of the compass directions will denote the corresponding boundary between CV's, where as upper case letters represent the boundary line segments between them.)

Without loss of generality, consider the flux through the eastern boundary of a control volume.

$$F_e = \frac{1}{A} \int_e (\nabla c - \vec{u}c) \cdot \hat{e}_x dy$$




Figure 5: An example of a rectangular control volume with dimensions h_x and h_y . The lower case letters denote the name of the boundary segments.

Since the integration takes place on the cell boundary, the diffusive and advective currents need to be evaluated on the boundary:

$$F_e = \frac{1}{h_x} \left(\frac{\partial c}{\partial x} - u_x c \right) \Big|_e$$

Now evaluating the concentration gradient on the boundary is very easy - we simply apply a central difference approximation, since the cell centers of P and E are equal distances away from the boundary. For the advective current, we need to approximate the horizontal velocity and the concentration at the boundary. The value of C on e can similarly be determined as the average of C at P and E, or rather, linearly interpolated between P and E. For the flow velocity, recall that u_x is a y -derivative of the stream function. Since the discretized domain on which we solve for ψ will affect how and where we can apply derivatives, it is therefore helpful to choose to solve for ψ on the CV corners. Thus the chemical flux across the eastern boundary becomes:

$$F_e = \frac{1}{h_x} \left(\frac{c_E - c_P}{h_x} - \frac{\psi_{se} - \psi_{ne}}{h_y} \frac{c_P + c_E}{2} \right) \quad (32)$$

Where ψ_{se} and ψ_{ne} is the value of ψ on the south- and north-eastern corner of the cell P respectively. For any internal point, the remaining fluxes are found in a similar manor, and when summed together, we obtain the full algebraic expression for the time rate of change of the C at the CV in question:

$$\frac{\partial c}{\partial t} = F_n + F_s + F_e + F_w = \sum_i a_i c_i \quad (33)$$

Where a are coefficients and the sum is over P, and its 4 compass neighbours. Thus any internal point yields an algebraic expression which couples together 5 CVs via one equation.

A few words on boundary conditions in the FVM

Since the FVM is based on evaluating fluxes between neighbouring cells, flux-related boundary conditions are very easy to implement. On 3 of the 4 system walls, a zero concentration gradient boundary condition is imposed. Also $\psi = 0$ on the boundaries. This means that the flux on any line segment on these 3 system walls are 0, so edge and corner points just have 1 or 2 fewer couplings, but the algebra is otherwise the same.

On the top boundary however, there is clearly a non-zero diffusive flux, since $c = 1$ on the boundary. To impose this, the northern flux on the top row of CVs need some slight manipulation. By writing a taylor approximation of c in the y -direction and evaluating it at c_P and c_S , one can solve for the gradient and find that:

$$\frac{\partial c}{\partial y} \Big|_{y=0} = \frac{3}{h_y} c_P - \frac{1}{3h_y} c_S - \frac{8}{3h_y} c(x_P, 0) + \mathcal{O}(h_y^2) \quad (34)$$

which only holds true for all CVs on the top row of the system. Here $c(x_P, 0)$ can in fact be any function one wants to impose as a Dirichlet BC, but will in most cases be set to 1. A more detailed explanation of this is given in the appendix.

4.1.2 Labeling volumes and points

In order to effectively turn the FVM into a managable linear algebra problem, one needs to keep track of points, volumes and adjacencies. Since we ultimately need to construct and solve a series of matrix equations, one must find a way to efficiently convert a 2D field into a column vector and

vice versa. Thus each CV will have a single index as well as a row/column type index. To convert between the two, modular arithmetic is used. A diagram of these coordinates is shown on Fig.6.




Figure 6: An example grid geometry for $N_x = 4$ and $N_y = 3$. The grey cells are CVs, and their indices are shown as per Eq.35. Between the CVs are blue circles on which the value for ψ will be solved for. The value of ψ can also be accessed on the red circles, but is given by the Dirichlet boundary condition $\psi = 0$. The staggered grid also follows a similar indexing scheme - albeit with one fewer point count in each direction.

$$i = r * N_x + c \quad r = \lfloor i / N_x \rfloor \quad c = \text{mod}(i, N_x) \quad (35)$$

Where i is a 1 dimensional index ranging from 0 to the number of points $N_x N_y$, and (r, c) is the associated 2 dimensional row/column type index. These definitions are only used for demonstration of array indexing, and should not be confused with the concentration field itself. These definitions are mainly important to simplify accessing neighbouring values between the staggered grids. The staggered grids 1D index will have little in common, but their row/column type index will be adjacent. On Fig.6 an example is shown for $(N_x, N_y) = 4, 3$. The grey cells represent C -field CVs, and the circles represent the interlaced grid on which ψ will be approximated. The value of ψ is given on the boundary, so we need only solve for it on internal points. Thus in general this domain will be $(N_x - 1, N_y - 1)$. A general stream function index (i, j) will thus be centered around CVs (i, j) , $(i, j + 1)$, $(i + 1, j)$ and $(i + 1, j + 1)$. A similar expression can be found for which stream function indices corner a given CV. Note that when converting back and forth, ones has to be careful when using Eq.35, since the two lattices have different N_x .

When this indexing scheme has been applied, a matrix equation can be constructed. Earlier an example for one of the fluxes in Eq.33 was shown. After carrying out the remaining 3, a matrix

equation for the time rate of change can be found. In this representation, a general internal CV i will have the following equation associated with it:

$$\begin{aligned}
\frac{\partial c}{\partial t}\Big|_i &= \left(\frac{1}{h_y^2} + \frac{u_y|_{ni}}{2h_y}\right) c_{i-N_x} + \left(\frac{1}{h_x^2} + \frac{u_x|_{wi}}{2h_x}\right) c_{i-1} \\
&+ \left(-2\left(\frac{1}{h_x^2} + \frac{1}{h_y^2}\right) + \frac{u_y|_{ni} - u_y|_{si}}{2h_y} + \frac{u_x|_{wi} - u_x|_{ei}}{2h_x}\right) c_i \\
&+ \left(\frac{1}{h_x^2} - \frac{u_x|_{ei}}{2h_x}\right) c_{i+1} + \left(\frac{1}{h_y^2} - \frac{u_y|_{si}}{2h_y}\right) c_{i+N_x}
\end{aligned} \tag{36}$$

Thus an equation of the form $\partial_t \vec{c} = A\vec{c}$ can be constructed. How the time propagation will be handled is to be discussed later. To reduce cluttering the velocity terms are simply referred to by their direction and on which phase of the i -th CV it is sampled. They are found as numerical derivatives of the staggered grid, by going back and forth between i -indices and (r, c) indices of the staggered grids, as has been previously discussed. Note that CVs coinciding with edges and corners will lose couplings. For example a CV on the left edge will not have a C_{i-1} term, and only one $-h_x^{-2}$ term instead of two in the diagonal, since there is one less CV to diffuse to. Also the $u_x|_{wi}$ term would disappear from the diagonal.

4.1.3 Constructing a linear algebra problem for the flow field

In this section we shall consider how to solve Eq.14 for a known C -field. The simplest way to approach this is using the Finite Difference Method (FDM). The FDM differs slightly from the FVM by the fact that it uses the differential form of the equation to construct a linear algebra problem. That being said, the resulting matrix equations are usually very similar.

If we assume F is defined on a similar domain as C , we need some meaningful way of reducing a $N_x \times N_y$ sizes grid to a $(N_x - 1) \times (N_y - 1)$, on which the interpolated function will be denoted f . On this reduced grid one needs to be able to access the value and the gradient. As previously discussed, using Eq.35 we can easily find the 4 surrounding F and C indices associated with a given ψ index. For the terms in Eq.14 where function values are represented (f^{-1} and f^{-2}), we resort to an average over the 4 surrounding sites. Where a gradient is needed, we average the partial derivative over the orthogonal direction:

$$f_{r,c} \equiv \frac{F_{r,c} + F_{r,c+1} + F_{r+1,c} + F_{r+1,c+1}}{4} \tag{37}$$

Where capital F corresponds to function values on the $N_x \times N_y$ grid, and f is the averaged value on the interlaced $(N_x - 1) \times (N_y - 1)$ grid, the same domain as ψ will be calculated on. This sampling is shown schematically on Fig.38. Similarly, an averaged partial derivative becomes:

$$\frac{\partial f}{\partial x}\Big|_{r,c} \equiv \frac{F_{r,c+1} - F_{r,c} + F_{r+1,c+1} - F_{r+1,c}}{2h_x} \tag{38}$$




Figure 7: Sampling of values from the large grid to form averages and derivatives on the interlaced grid. The schematic here shows the weights applied to obtain Eq.37 and Eq.38.

A similar method is used for the horizontal partial derivative of C and the vertical partial derivative of f , albeit with swapped signs and step lengths. This is effectively an average of two central difference quotients. The only remaining derivatives are those of ψ . Eq.14 requires both first derivatives and the laplacian (the sum of second derivatives).

$$\left. \frac{\partial \psi}{\partial x} \right|_{r,c} = \frac{\psi_{r,c+1} - \psi_{r,c-1}}{h_x} \quad (39)$$

$$\left. \frac{\partial^2 \psi}{\partial x^2} \right|_{r,c} = \frac{\psi_{r,c+1} + \psi_{r,c-1} - \psi_{r,c}}{h_x^2} \quad (40)$$

And similar quotions for y - partial derivatives. It is these definitions that allow us to represent the differential equation as a linear algebra problem, since these derivatives can be defined as a matrix operator acting on a ψ vector. When all numerical derivatives are applied, and row/column type indexing exchanged for 1 Dimensional indices as per Eq. 35, the resulting matrix equation becomes:

$$\begin{aligned} & \left(\frac{1}{h_y^2 f_i} + \frac{1}{h_y f_i^2} \left. \frac{\partial f}{\partial y} \right|_i \right) \psi_{i-N_x} + \left(\frac{1}{h_x^2 f_i} + \frac{1}{h_x f_i^2} \left. \frac{\partial f}{\partial x} \right|_i \right) \psi_{i-1} \\ & - \frac{2}{f_i} \left(\frac{1}{h_x^2} + \frac{1}{h_y^2} \right) \psi_i + \left(\frac{1}{h_x^2 f_i} - \frac{1}{h_x f_i^2} \left. \frac{\partial f}{\partial x} \right|_i \right) \psi_{i+1} + \left(\frac{1}{h_y^2 f_i} - \frac{1}{h_y f_i^2} \left. \frac{\partial f}{\partial y} \right|_i \right) \psi_{i+N_x} = Ra \left. \frac{\partial c}{\partial x} \right|_i \end{aligned} \quad (41)$$

Which can be represented by an equation of the form $A\vec{\psi} = \vec{b}$. One should note that the matrix in question only has non-zero indices on the diagonal and bands offset by 1 and N_x from the diagonal.

4.1.4 Time discretization and propagation of the C-field

We have previously shown how the FVM leads to representing the time rate of change of the C -field as a matrix operator acting on said field. The remaining aspect of the numerical model is concerned with how to update the C -field to propagate it in time. Consider a snapshot of the spatially discretized c -field at time-step k . That instance of the C -field is represented by a column vector, which has the following time derivative associated with it:

$$\frac{\partial \vec{c}^k}{\partial \tau} = T \vec{c}^k + \vec{b} \quad (42)$$

Where T is the time rate of change matrix, which accounts for diffusion and advection and follows the general template presented in Eq.36. The vector \vec{b} represents any source term coming from the Dirichlet boundary condition via Eq.34, but is zero everywhere except the first N_x rows, which correspond to the CVs on the top of the cell (see appendix). There are now a few ways in which we can discretize the system in time:

- Explicitly, via a Forward Euler

$$\frac{\partial \vec{c}^k}{\partial \tau} \simeq \frac{1}{\Delta \tau} (\vec{c}^{k+1} - \vec{c}^k)$$

As the name suggests, this type of time discretization assumes that at each individual element at time $k + 1$ can be calculated explicitly from the information present at time k . This leads to the following time propagation rule:

$$\vec{c}^{k+1} = (I + \Delta \tau T) \vec{c}^k + \Delta \tau \vec{b} \quad (43)$$

- Implicitly, via a Backward Euler

$$\frac{\partial \vec{c}^{k+1}}{\partial \tau} \simeq \frac{1}{\Delta \tau} (\vec{c}^{k+1} - \vec{c}^k)$$

The actual discretization of the time derivative is the same here, but now we've shifted the time step at which the derivative is associated. This means that we can in principle never calculate the time derivative from present data, but we can instead construct a matrix equation for the $k + 1$ step with a right hand side dependant on step k

$$(I - \Delta \tau T) \vec{c}^{k+1} = \vec{c}^k + \Delta \tau \vec{b} \quad (44)$$

- Crank-Nicholson, via Midpoint Euler

$$\frac{\partial \vec{c}^{k+1/2}}{\partial \tau} \simeq \frac{1}{\Delta \tau} (\vec{c}^{k+1} - \vec{c}^k)$$

Here an intermediate state between k and $k + 1$ is introduced. The first half-step is done explicitly and the last half-step is done implicitly. Thus the following equation needs to be solved, to propagate the system in time:

$$\left(I - \frac{\Delta \tau}{2} T \right) \vec{c}^{k+1} = \left(I + \frac{\Delta \tau}{2} T \right) \vec{c}^k + \Delta \tau \vec{b} \quad (45)$$

The different methods have strengths and weaknesses. First of all an explicit method is computationally easy, since it is essentially repeated matrix multiplication. The computational ease however comes with a cost in terms of accuracy and numerical stability. A forward Euler difference quotient has error of order $\mathcal{O}(\Delta \tau)$ and if the time step is too large the sequence of \vec{c}^k will quickly diverge. For stability, the implicit approach is much better. Implicit methods are unconditionally stable for all $\Delta \tau$, but a pure implicit method still has a linear error term, and the computations are more demanding, since for each time-step a matrix needs to be reduced and solved, compared to a simple matrix multiplication. The Crank-Nicholson hybrid improves on the implicit approach, since the leading error term is $\mathcal{O}(\Delta \tau^2)$. Compared to an implicit approach, the computational workload is similar albeit slightly higher, since the right hand side contains a matrix multiplication. This added

increase is quite small, granted matrix multiplication is generally one order of magnitude faster than solving the left hand side matrix.

An even better alternative than Cran-Nicholson is the so called Alternating Direction Implicit (ADI) method. The fundamental idea is similar, since it relies on explicit and implicit half-steps, but comes with a much lower computational workload - since the resulting matrices are of a simpler nature. It involves splitting the advection-diffusion matrix into two operators, one for the x - and one for the y - direction. The general recipe for the two operators are the same as presented in Eq.36, albeit all the x - dependant terms are gathered into one matrix and all the y - dependant terms are gathered into another.

$$T = T_x + T_y$$

These both contain only 3 non-zero bands. The T_x matrix will be tridiagonal, having only non-zero elements on the diagonal and next to the diagonal. The T_y matrix will similarly only have non-zero elements on the diagonal and two bands offset by N_x on both sides of the diagonal. We then alternate the directions which will be treated implicitly and explicitly. We first take a half step where the y - direction is treated explicitly and the x - direction is treated implicitly, then vice versa for the last half step .

$$\left(1 - \frac{\Delta t}{2} T_x\right) \vec{C}^{k+1/2} = \left(1 + \frac{\Delta t}{2} T_y\right) \vec{C}^k + \frac{\Delta t}{2} \vec{b} \quad (46)$$

$$\left(1 - \frac{\Delta t}{2} T_y\right) \vec{C}^{k+1} = \left(1 + \frac{\Delta t}{2} T_x\right) \vec{C}^{k+1/2} + \frac{\Delta t}{2} \vec{b} \quad (47)$$

While this is clearly more linear algebra problems to solve, as will be shown later in the computational section, the structures of T_x and T_y make solving each of these much simpler. The ADI method takes an intrinsic 2D problem and converts it into a sequence of independent 1D problems, which are much easier to solve [18] [19].

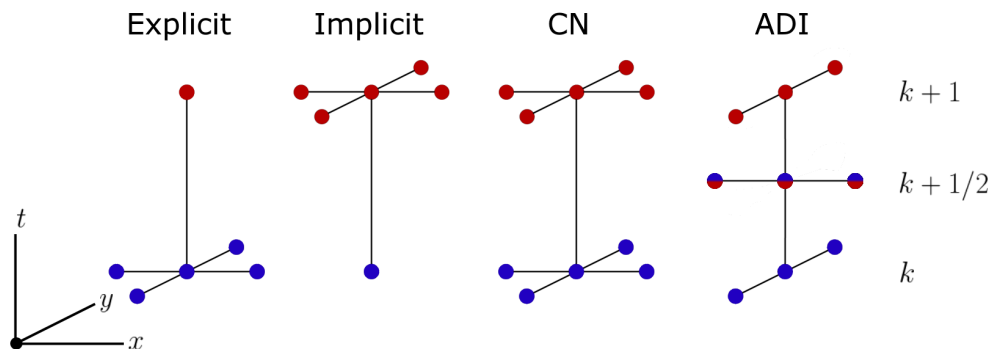


Figure 8: Schematic representation of time propagation from time step k to $k+1$ via 4 different time discretization methods. Red denotes points that are solved for implicitly. At each step a matrix equation is solved, in which red points are treated as unknowns and blue points are in the right hand side vector. In ADI, the intermediate step is first solved implicitly, then applied explicitly.

4.2 Computational aspects

All programming was done in Python 3.7. Since most matrices being constructed and solved have very few non-zero elements, the `scipy` package, as well as its sparse linear algebra modules was used. It includes highly efficient methods for constructing and solving matrix problems of a sparse nature. The presented methods were written in an object oriented fashion.

A proprietary `Field` class was written, which represents the 2D fields being worked with, and handles storing them as column vectors and most elementary operations performed on them, such as basic arithmetic, matrix multiplication and partial derivatives. These operations are vectorized by storing the row/column type indices accessed for each of these operations, reducing the amount of repeated calculations.

Additionally a `StreamSolver` class was written for the stream function FDM. Since the matrix Eq.41 is very sparse, the 5 non-zero bands are constructed as vectors and turned into a `scipy.sparse` matrix. They have built-in highly efficient LU decomposition functionality, so after constructing and decomposing the matrix for ψ it can be solved rapidly for different concentration fields. The general usage is:

```
def f(x,y):
    #Set a uniform permeability function
    return 1
#Construct a permeability field on [1,1] with 100x100 pts.
F = Field(1,1,100,100)
F.init_from_func(f)
#Construct a StreamSolver for this permeability field
SS = StreamSolver(F)
SS.factorize() #LU decompose
#Assuming C and Ra has already been defined elsewhere, solve for Psi
Psi = SS.solve(Ra*C.ddx())
```

Also a `FVM` class was written, which simplifies constructing the advection diffusion operator as per Eq.36. The class is first initialized with the `Field` geometry, from indices for each CV are precalculated to rapidly access the correct stream function values on the interlaced grid. This object constructs a diffusion operator (which never changes), and construct the advective terms (different for each time step) when updated with a stream function. The `FVM` object also has a method for applying Dirichlet boundary conditions on the top of the system built in. Assuming the `StreamSolver` has already been prepared, and `C`, `Ra` and `dt` are defined elsewhere the general usage would be:

```
#Make the FVM object, make the sparse diffusion matrix
Method = FVM(C,dt)
Method.make_pt_lists() #Allocate r,c indices
Method.make_diff_op()
for k in range(0,time_steps):
    Psi = SS.solve(Ra*C.ddx()) #solve for stream function
    #Make an implicit time prop. operator. EYE is a sparse identity matrix
    Method.make_advect_op(Psi)
    IMPLICIT = (EYE - Method.diffuse - Method.advect)
    #Apply the top BC
    Method.apply_dirichlet_top(C)
    #Update the C-field, use the built in sparse solver
    C.vector = splin.spsolve(IMPLICIT,C.vector)
```


A slightly modified `ADI_FVM` class was written and optimized for computational efficiency. Its functions are practically the same as the `FVM` class, albeit it splits the operator into two, one for x and one for y . It also returns 6 vectors instead of a sparse matrix. The vectors represent the 3 non-zero bands of the two matrices, since the solver algorithm does not require a full matrix to be constructed.

All presented code is available on GitHub [to be cited later...]. As roughly outlined above, the general solution procedure is:

1. Initialize the system, by defining a permeability field, a domain grid and construct the velocity- and concentration field objects. LU factorize the stream function operator.
2. Solve Eq.14 for the stream function.
3. Use the `FVM` or `FVM_ADI` object to construct the advection-diffusion operator.
4. Propagate the system in time.
5. Save relevant data (e.g. ψ or C fields) to an external file.
6. Repeat from (2).

4.2.1 Efficient solvers for the constructed matrices

With the chosen spatial discretization, it is now possible to delve deeper into analyzing how the structure of the matrix might affect the solution procedure. Fig.9 outlines the general structure of all matrices involved in the numerical model proposed.



(a) The general layout of a 5 point stenciled matrix operator.

(b) A tridiagonal x -component only matrix operator.

(c) A banded y -component only matrix operator

Figure 9: The different types of matrices used in the numerical methods. FDM and FVM are generally 5 point stenciled. Time propagation with ADI splits the 5 point stencil into a tridiagonal and banded matrix. The letters indicate which CV couplings the bands represent. For example E represents the eastern or right hand side neighbour.

A convenient solution procedure for tridiagonal matrices as the so called *Thomas algorithm* (sometimes simply referred to as the tridiagonal matrix algorithm or TDMA) [19]. It utilizes the fact that due to the simple structure of the matrix, each row can be reduced using only the row right above it. Consider a problem of the form $A\vec{x} = \vec{b}$, with A being a tridiagonal matrix, where

each column has non-zero indices l_i , d_i and u_i , where l represents the lower and u the upper off-diagonals. To put this matrix into echelon form, we need to get rid of the lower diagonal. This is done by forward elimination, where each row gets updated via:

$$R_i \rightarrow \frac{l_{i-1}}{l_{d-1}} R_{i-1} \quad i = 1, 2, \dots, M-2, M-1$$

Where M is the rank of A . This guarantees that the lower diagonal in line i becomes 0. Nothing changes in the upper band, but the diagonal gets updated. Similarly, the right hand side is also updated

$$s_i = \frac{l_{i-1}}{d_{i-1}} \quad d_i \rightarrow d_i - s_i u_i \quad b_i \rightarrow b_i - s_i b_{i-1}$$

Since the matrix is now in row echelon form, the solution can be found by back substitution:

$$x_{M-1} = b_{M-1}/d_{M-1}, \quad x_i = \frac{b_i - u_{i+1}x_{i+1}}{d_i} \quad i = M-2, M-3, \dots, 1, 0$$

Now for the ADI procedure, this is half the battle completed. But we also need an efficient solution procedure for banded structures as shown in Fig.9c. We will find that the algorithm is actually very similar, and has only a slight difference from the conventional Thomas algorithm. As before we assume the notation that every column i has non-zero elements l_i , d_i and u_i , and that u and l are offset by N indices from the diagonal. Thus the first N rows are already in echelon form. The forward elimination is performed in a similar fashion, the difference being that row operations for row i are now based on row $i - N$.

$$s_i = \frac{l_{i-N}}{d_{i-N}} \quad d_i \rightarrow d_i - s_i u_i \quad b_i \rightarrow b_i - s_i b_{i-N}$$


After the forward eliminations, the last N elements can be found right away. The rest are found by back substitution:

$$x_j = b_j/d_j \quad j = M-1, M-2, \dots, M-N \quad x_i = \frac{b_i - u_{i+N}x_{i+N}}{d_i} \quad i = M-N-1, \dots, 1, 0$$

As one can see, the Thomas algorithm is actually just a special case of this procedure for $N = 1$. These methods are much faster than brute force Gaussian elimination, since the amount of floating point operations is only proportional to M , not M^3 [19].

One can speed up these methods even further by identifying patterns in the sequences above. In the Thomas algorithm (used for x -component only operators), notice that $s_i = 0$ if $l_{i-1} = 0$. In such cases the dependency of earlier matrix elements are lost, and the method practically restarts at this i . This is very beneficial, because in the FVM, we can predict the occurrence of such zeroes in the lower band, since they correspond to CVs with no western neighbour. Thus every row which is an integer multiple of N_x , the coupling to previous matrix elements are lost. This means that we can do N_y instances of the Thomas algorithm simultaneously, by using vectorization.

Similarly for the general banded algorithm (applied to y -component only operators), we notice that every block of N_x rows are completely independent from each other and only rely on the previous block of N_x rows. In a similar fashion, this means that we can do N_x instances of this banded algorithm, each responsible for a sequence of N_x rows. Both of these advantages are outlined in Fig.10.



(a) A tridiagonal x -component only matrix operator. Every $N_x = 4$ rows, there's a 0 in the l -band, meaning each block separated by the green dotted lines can be treated independently.

(b) A banded y -component only matrix operator. Since the forward elimination only relies on the N_x -th previous row, neighbouring rows have no coupling. For example, the black and grey sequences are independent.

Figure 10: Sparsity diagrams of x - and y - component operators on the 4×3 grid shown in Fig.6. Squares correspond to non-zero elements. The green squares represent indices which are zeroed out in the forward elimination. The arrows show which rows are utilized for the forward elimination.

This further illustrates the fact that ADI for each time step takes a 2D problem and splits it into a series of 1D problems. Depending on the direction treated implicitly we get either N_y instances of sequences with length N_x to be solved with the Thomas algorithm or N_x instances of sequences with length N_y to be solved with the banded algorithm.

4.3 Testing the algorithms

4.3.1 A quick matrix solver benchmark

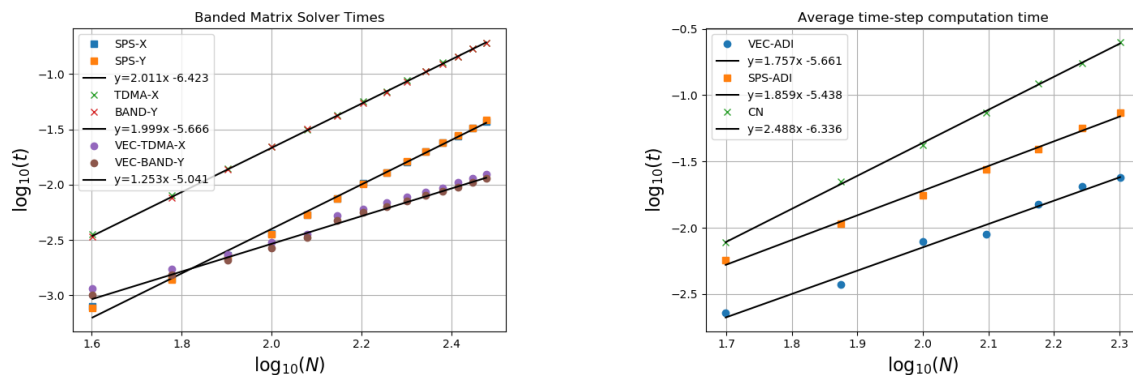
Having outlined how to construct the linear operators, and proposed a few methods for solving them, their computational efficiencies are tested, by averaging the time which it takes to solve similar problems with the different approaches. Two tests are performed: First, the `scipy.sparse` package is compared to the simple Thomas and banded algorithms and their vectorized versions. Secondly, the iteration time of a Crank-Nicholson (CN) and Alternating Direction Implicit (ADI) time step is compared. The CN approach utilizes a 5 point stencil and is thus solved using the `scipy.sparse` direct solver, whereas the ADI approach is solved using both the sparse solver and our vectorized algorithms described earlier. In the second comparison, solving for the flow field and constructing the advection operator is also included.

By doing these simplified tests, one can see how the calculation time scales with amount of CVs. For simplicity, let $N_x = N_y = N$. Each time value is averaged over 1000 executions, to combat clock granularity. Keep in mind that the size of the matrices involved are N^2 , but due to

the sparse nature of the matrices, the amount of non-zero elements are only proportional to N^2 . The Thomas algorithm is known to be linear in time with respect to the matrix size (so $t \propto N^2$), which is indeed the finding on Fig.11a. Interestingly enough however, the built in sparse algorithm has a near identical scaling, but is almost one order of magnitude faster. The `scipy.sparse` direct solver is likely performing a similar solution algorithm, but may be programmed more efficiently, with less overhead. At $N \geq 100$, the vectorized versions however beats both, and exhibits scaling like $t \propto N^{1.25}$, which had been linear in a perfectly parallelized case. Still the speed up is well noticeable, especially for larger systems.

When comparing the calculation times of actual advection-diffusion time steps, we learn even more about the pros and cons of the algorithms. Consider $\log_{10}(N) = 2$. Here we found earlier that the sparse and vectorized banded solvers were very similar in calculation time, or about $\log_{10}(t) \simeq -2.5$. Still the vectorized algorithm is about $10^{0.4} \simeq 2.5$ times faster. Also, we notice that the scaling relations of full time steps are quite different from those showing only the 1D operator solver times. This indicates that a big part of the VEC-ADI speed up comes not only from the efficient solver, but from the fact that we never construct matrices, but rather store only the non-zero element bands in 1D vectors and use vectorized algorithms for working with those.

Another take away lesson from this test is that the VEC-ADI approach is fast enough, so that in the time it takes for it to time-step a system with $N = 200$, the Crank-Nicholson approach can only handle a system with $N = 75$ in the same calculation time. Another way to think about this is that in this case VEC-ADI is able to produce $(200/75)^2 \simeq 7$ times the amount of data in the same calculation time. Due to the slow scaling of CN, this difference is only going to be bigger if larger systems are compared.



(a) Averaged solution times for implicit ADI steps. The x - and y - component problems are overlaid, but are very similar in execution time.

(b) Averaged solution times for a full time step. Here Crank-Nicholson and ADI with two different matrix solvers are compared.

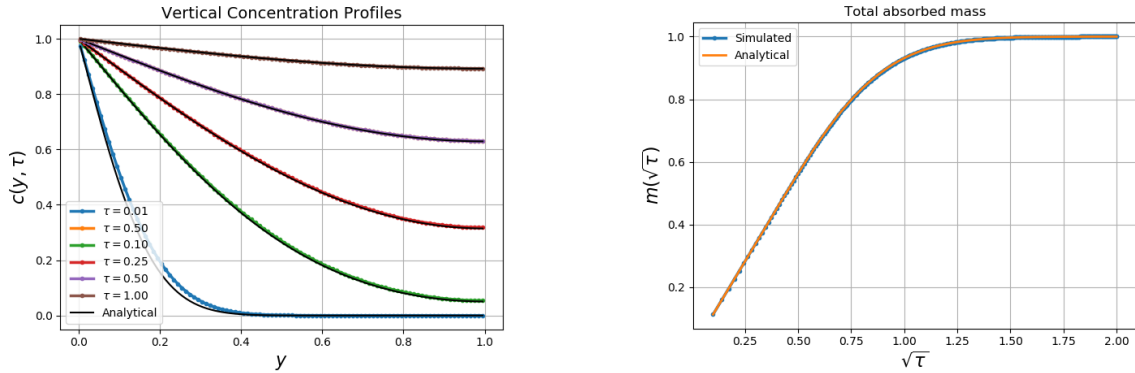
Figure 11: Averaged calculation times on $N_x = N_y = N$ sized grids. SPS refers to the inbuilt sparse solver in `scipy.sparse`, TDMA and BAND refer to the tri-banded algorithms described earlier and the VEC prefix refers to the more efficient vectorized versions of those.

4.3.2 Sanity checks

Due to the non-linear mechanics involved in Eq.14 and Eq.15, it is difficult to find analytical solutions to compare numerical results to. However, there are still some sanity checks one can perform to at least gain more confidence in the numerical model and program. These include:

- Check that the 1D diffusion results of Eq.17 and Eq.18 are recovered, if $Ra=0$.
- Check that chemical mass conservation is fulfilled, if the system is fully isolated.
- Check that an altered permeability field does in fact result in an altered stream function and velocity field.

First the 1D case is tested, by setting $Ra = 0$. A uniform box of width and height equal to 1 is discretized into a 100×100 grid of CVs. The system is time-stepped 200000 times with a time step of $\Delta\tau = 2 \cdot 10^{-5}$. The comparison with Eq.17 and Eq.18 is shown on Fig.12.



(a) A handful of concentration profiles, where the simulated data comes from taking a vertical slice of the 2D field. The analytical values for each time step presented is found using Eq.17.

(b) Total chemical mass as a function of the square root of elapsed time. For the simulated results, this is found simply by summing the concentration value of each CV times the area of the CVs. The analytical values are found using Eq.18.

Figure 12: Comparing the $Ra = 0$ case to a 1 dimensional diffusion only problem. The simulations agree well with analytical results.

The mass conservation condition is checked by imposing no-flux boundary conditions on all edges and initializing the C -field with a Gaussian blob of the form:

$$C_0 = \exp(-50(x - 0.5)^2 - 50(y - 0.5)^2)$$

We then place the blob in a 100×100 box with side length 1. The Rayleigh number is set to 2000 and the time step $2.5 \cdot 10^{-7}$. The system is time stepped 1200 times, and the resulting fields and their total chemical masses are found and displayed on Fig.13. Over the length of the short simulation, the fractional mass deviation was of the order of magnitude 10^{-14} , which is low enough to be attributed to round off error.




Figure 13: A visual representation of the C - fields used check for mass conservation. As expected the heavier blob sinks to the bottom and diffuses.

Finally, one can experiment with applying non-uniform permeability fields. For this we initialize a C -field with the same Gaussian shape as before, and set a permeability fluctuation of the form:

$$f = a (1 - 0.1(x - y)^2)$$

Where a is a normalization constant, such that f averages to 1 over the domain. Notice that this function has a maximum on the line $y = x$, and is an inverted parabola in the direction away from this line. Our intuition would tell us that the blob would be more inclined to travel down the $y = x$ line than before. We then use the `StreamSolver` class to solve for the stream function associated stream function. On Fig.14, we can clearly see that the stream lines are more biased towards the $y = x$ line, indicating that the permeability field does in fact alter

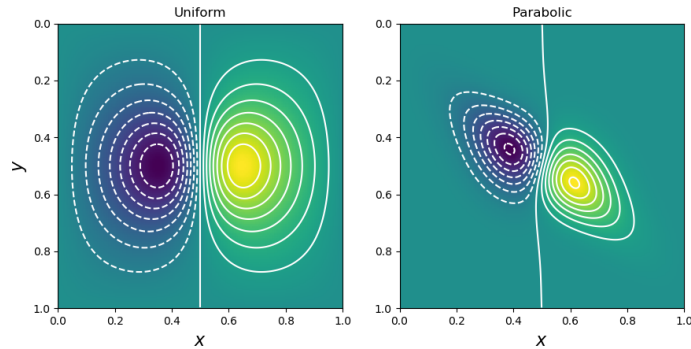


Figure 14: The stream functions of centered Gaussian concentration fields at $t = 0$ compared for a uniform and parabolic permeability field. The white lines are contours of ψ , which correspond to stream lines of the flow field.

4.3.3 The need for perturbation

Since the system starts in a pseudo-equilibrium, the onset of density driven convection requires some small perturbation. If no perturbation is enforced, the inherit round-off error involved in floating point operations will effectively act as a small perturbation. To combat this, for each timestep, the top-boundary condition ($c(x, y = 0) = 1$) is slightly altered to include a source of randomness, such that:

$$c(x, y = 0) = 1 - \delta(x)$$

Where δ is confined to $[0, 10^{-6}]$ and sampled from the uniform distribution. For each time step, a new random number is sampled for each of the top CV boundaries, which is independent of its neighbours both in time and space. Farajzadeh et al solved this by adding a small sinusoidal perturbation to the first time steps, and found that the convective behaviour was largely independent of the imposed sinusoid [20], but to eliminate all bias to specific wavelengths, the uniform distribution with a minuscule amplitude is applied here, imitating for example thermal fluctuations.

5 Numerical results

5.1 Understanding the base case

As an example of the general mechanics, consider a system in which the width is twice the height. Initially, a Rayleigh number of 3000 is tested. The spatial domain is a rectangle of width 2 and height 1, which is discretized on 250x125 points, with a time step of $5 \cdot 10^{-9}$ as per the dimensionless diffusive time scale. Note that all presented results are as per the rescaling of Eq.11.

5.1.1 A qualitative analysis of the observed patterns

Initially, the system exhibits only diffusive behaviour. After some time, while the concentration field is still very uniform in x , as expected by pure diffusion, the flow field slowly starts to become oscillatory, albeit with a low amplitude. This amplitude then grows and causes the initial onset of the convective dynamics. Around the onset, the observed characteristic wavelength is clearly well defined. At later times, the non-linear dynamics commence, and the merging and splitting of convection plumes become increasingly apparent. These dynamics are shown on Fig.15-19.

Note that to properly display the dynamics of interest, the individual figures are cropped slightly differently, as the length scales of these dynamics vary. Also, to reduce cluttering, each arrow represents the flow field averaged over CVs in a square around that vector arrow. When referring to the velocity magnitude, note the large variations between the early snapshots and the later ones.

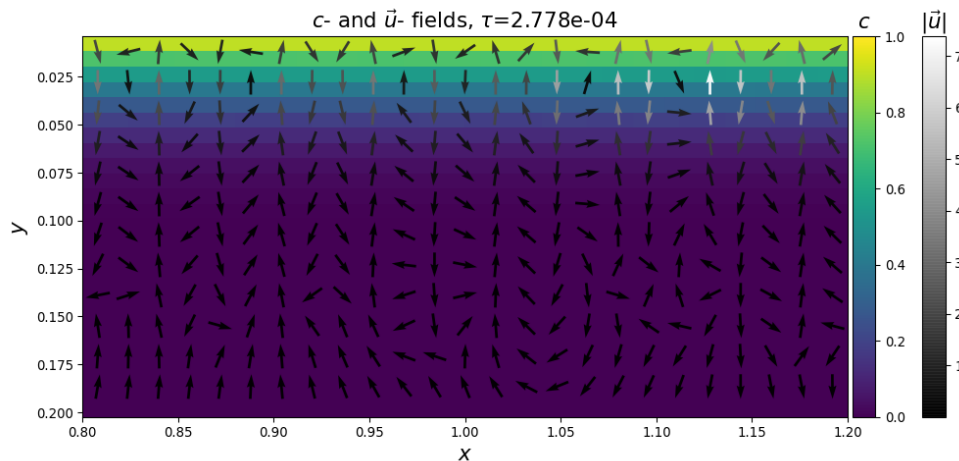


Figure 15: Initially, the concentration field is very uniform. Within the concentration gradient, the flow field is starting to become oscillatory, but with a very low magnitude, with respect to the natural scales of the system.

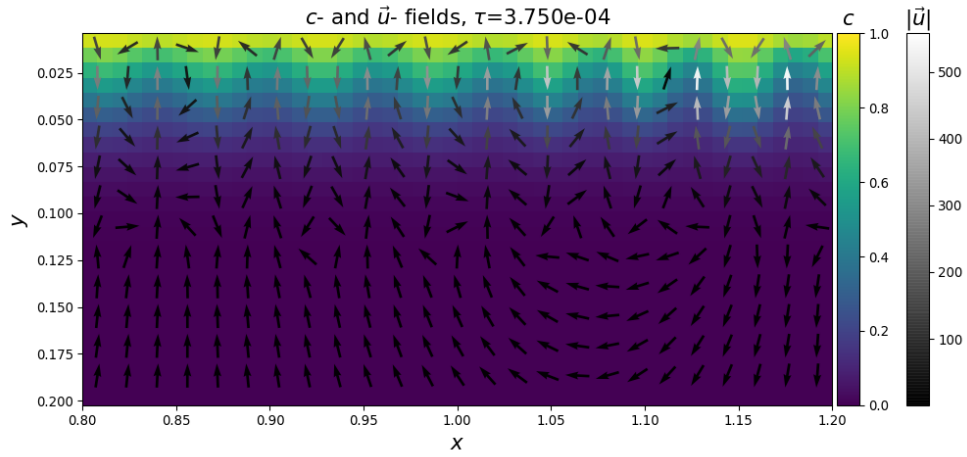


Figure 16: Around the onset time, the oscillatory behaviour is now clearly visible, both in the concentration field and the flow field. Notice the increased magnitude in the flow field compared to Fig.15. At this time, one might expect a characteristic wavelength to be well defined, both in the concentration field and the vertical component of the flow field.




Figure 17: After the onset of initial convection plumes, upward streams of fresh unaffected fluid heavily influence the concentration field. These strong upward flowing regions are spaced with significantly wider intervals than the initial plumes, thus when these ascending plumes reach the surface they push the top of the descending plumes laterally, merging them together.




Figure 18: Another mechanism, which becomes noticeable after the first merging events is plume splitting. As can be seen in this time slice, upward currents are observed to split some previously merged plumes in two. In this time slice, one can notice a clear difference in characteristic wavelength between the horizontal slice near the surface, and the descending front, which has been split.




Figure 19: At significantly later times than the initial onset, the previously mentioned plume formation- merging- and splitting dynamics all intertwine to produce the convective pattern.

5.1.2 A closer look at the top boundary

Given that the mass enters the system from above, understanding how the concentration field near the top boundary varies with time can learn more about the DDC dynamics. To do this, the top 5 rows of CVs (corresponding to $y \in [0, 0.04]$) are averaged vertically for each time step. This data

is then used to produce a so called space-time map of the vertically averaged top portion of the concentration field, to track the regions which emit carbon to the sinking concentration plumes. As shown on Fig.20, one can clearly identify how the system initially exhibits only diffusive behaviour. A short time later, the front spontaneously becomes unstable and becomes oscillatory. This is when DDC sets in, and one can see upward streams diverge and force the yellow plumes to shift laterally, which later merges them together. After the first series of merging events, we notice smaller plumes appear spontaneously, only to merge with the already established one shortly there after. Finally, the dynamics all coexist, and balance each other out, to the point where changes in characteristic length scale between the plumes varies much slower.

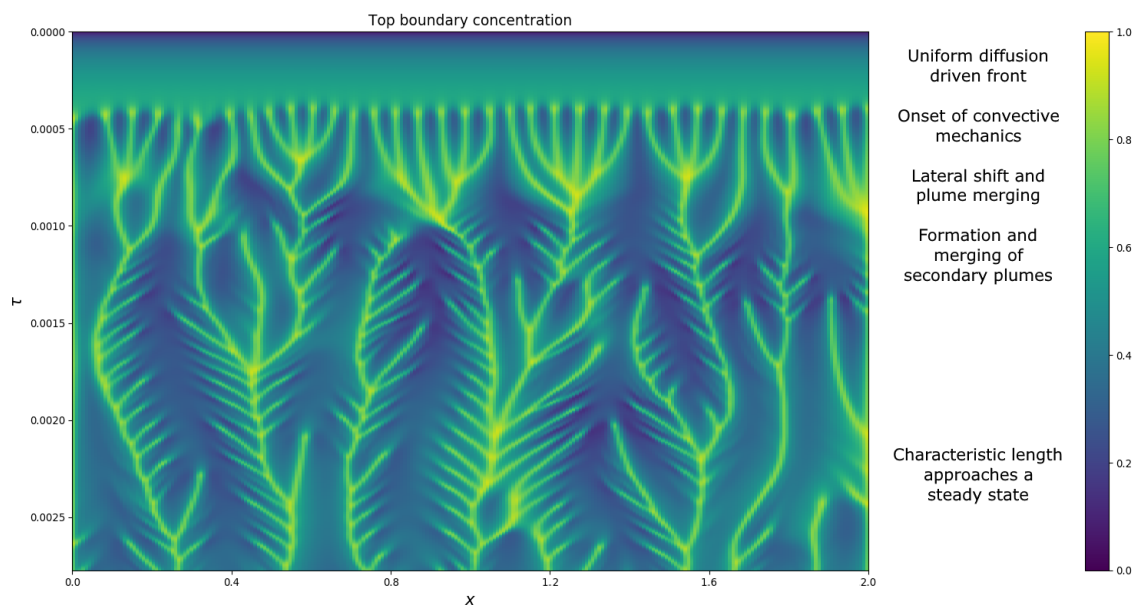


Figure 20: A space-time map of the top 4% of the concentration field for $Ra=3000$. The vertical axis represents time, and the horizontal represents the x coordinate. The ongoing mechanics are annotated roughly at the times at which they start occurring. These dynamics come quickly after one another in the transient time domain from initially uniform to the general DDC behaviour.

5.1.3 Estimating the characteristic wavenumber

By visually interpreting Fig.20, one can clearly see that the system goes through a range of dynamics which change the characteristic length scale of the system. To better quantify this, a power spectrum approach can be used to identify the distribution of different modes, and how that distribution undergoes change over time. To utilize this, a one dimensional function needs to be defined, which Fourier analysis can then be applied on. Since the absorbed mass of the system is continuously growing, we introduce the fluctuation in the vertically averaged concentration field:

$$\zeta(x) = \int_0^1 c(x, y) dy - \int_0^2 \int_0^1 c(x, y) dx dy \quad (48)$$

This is a one dimensional signal which averages over x to 0, so any oscillations found in this signal represent fluctuations with respect to the horizontal coordinate in the concentration field. We then take the fast Fourier transform (FFT) of the signal $\zeta(x)$ using the built-in `numpy.fft` package. The sample count was set to 4096 and a Hanning window was applied to reduce the weight of the boundaries. Let $F(\kappa)$ denote the obtained complex frequency spectrum with κ being the wavenumber. The expectation value and standard deviation of the spectrum of wave numbers can then determined via the following expressions:

$$E[\kappa] = \frac{\sum_{\kappa} \kappa F(\kappa) F^*(\kappa)}{\sum_{\kappa} F(\kappa) F^*(\kappa)} \quad \sigma_{\kappa} = \sqrt{\frac{\sum_{\kappa} \kappa^2 F(\kappa) F^*(\kappa)}{\sum_{\kappa} F(\kappa) F^*(\kappa)} - E^2[\kappa]} \quad (49)$$

Note that when working in the dimensionless x units, κ is also dimensionless, but represents reciprocal wavelength as is convention. As can be seen in Fig.21, the characteristic wavenumber is initially poorly defined, with a very high standard deviation. Around the onset of DDC, the wavenumber suddenly becomes well defined, and in this case evaluates to $E[\kappa] \simeq 17$, where the standard deviation has a local minimum at $\tau \simeq 0.0004$. This agrees well with the visual observation on Fig.20, which shows that there are 34 concentration peaks around the onset of DDC, when one takes into account that the system has a width of 2 dimensionless units. The time dependent statistical values also clearly show that the system undergoes rapid change shortly after the onset of convection, where the characteristic wavenumber plummets, and then stays roughly constant with only a very slow decay.





Figure 21: The time variation in characteristic wavenumber and width of the power spectrum for horizontal fluctuations in the $Ra = 3000$ system.

5.2 Varying the Rayleigh number


As previously discussed, the Rayleigh number balances the rate of advection and diffusion. Therefore, one can expect significantly different patterns, with regard to plume morphology when the Rayleigh number is varied. This is well within the capabilities of the written code, but comes with the added complication that the velocity is also strongly dependant on the Rayleigh number. This implies that one should reduce the time step for high Ra simulations, to increase the precision and stability of the simulations. For this reason, a time step of $\Delta\tau = 5 \cdot 10^{-6} Ra^{-1}$ was settled on. A snapshot of the system is saved every 500 time steps, which proved sufficient to catch all events from start until the fingers reach the bottom of the cell within 1000 snapshots.

5.2.1 The variation in plume patterns

Two visual examples of pattern variations from the previous $Ra = 3000$ case are shown on Fig.22-23. These examples display the cases of $Ra = 1000$ and $Ra = 5000$. Upon investigation it becomes very apparent that the length and time scales vary significantly with the Rayleigh number. This tendency holds also for intermittent Rayleigh numbers, not displayed here to reduce cluttering.

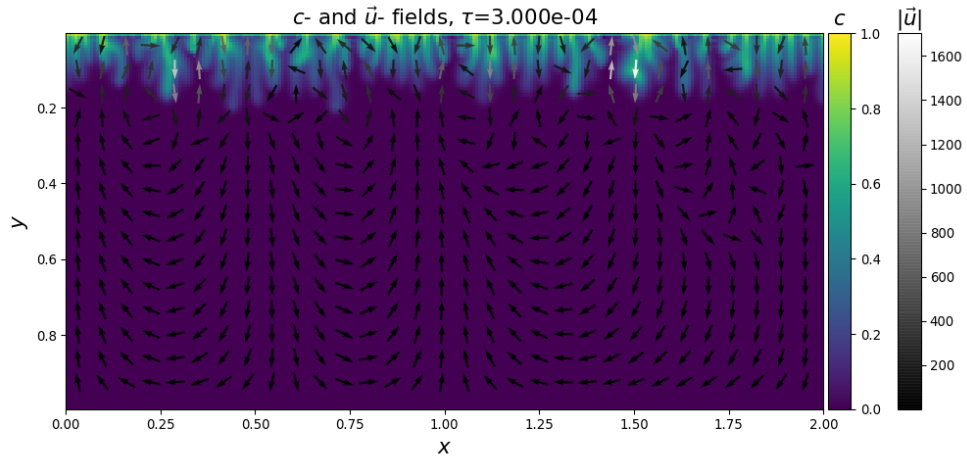


(a) A lower Ra case, shortly after the onset of convective dynamics.




(b) A lower Ra case, shortly before the plumes contact the bottom of the cell.

Figure 22: Two snapshots of the concentration- and velocity fields for $Ra = 1000$. Note how the relatively few plumes merge into a much wider plumes than in the $Ra = 3000$ case shown on Fig.19. The velocity magnitude is also significantly lower.



(a) A higher Ra case, shortly after the onset of convective dynamics.



(b) A higher Ra case, shortly before the plumes contact the bottom of the cell.

Figure 23: Two snapshots of the concentration- and velocity fields for $Ra = 5000$. Note how the plume count is much higher than in the $Ra = 3000$ case shown on Fig.19. In this case, the merging and splitting dynamics become apparent much earlier, and at lower depths.

5.2.2 Determining the onset time of DDC

As shown on the space-time map on Fig.20, the onset of DDC happens when the top of the cell goes from being a uniform front to taking on an oscillatory form. Thus to determine the time at which DDC starts to become apparent, it should suffice to determine a condition on top concentration field at the top of the cell. For each instance of different Ra , a space-time map is created, similar to that on Fig.20, and then for each horizontal slice, the mean value and standard deviation of the top of the concentration field is calculated. These properties of the top of the c -field are then compared, such that when the ratio of standard deviation to the mean value grows beyond 10^{-6} , we conclude that the instability has in fact started. The first time value when this happens is logged for each simulation.

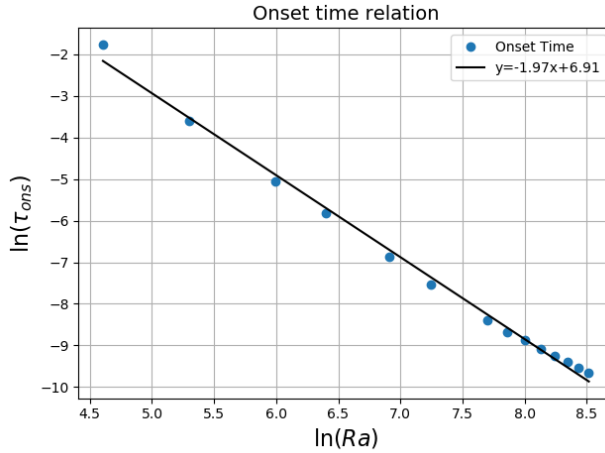


Figure 24: A logarithmic plot of the onset time of DDC as a function of Rayleigh number. The linear fit indicates a relation of $\tau_{ons} \propto Ra^{-1.97}$

As can be seen on Fig.24, when plotted logarithmically, a power law trend is found which suggests that the onset time of convection scales roughly with $\tau_{ons} \propto Ra^{-2}$. The Rayleigh numbers shown here range from 100 to 5000. Simulations for Rayleigh numbers 25, 50 and 75 were also performed, but they never fulfilled the onset condition.

5.2.3 Determining the initial wavenumber

As earlier mentioned, the morphology also depends greatly on the Rayleigh number. Having found the onset times of DDC, the same power spectrum method (Eq.49) is used to identify the characteristic wavenumber of the emerging instability.

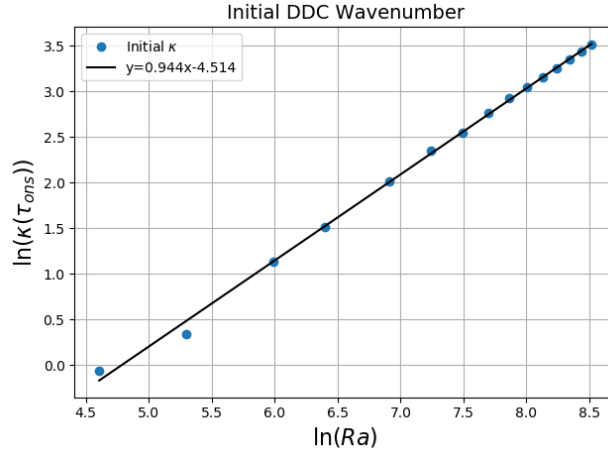


Figure 25: A logarithmic plot of the wavenumbers initially dominant in the instability with respect to Rayleigh number. The linear fit indicates a relation of $\kappa \propto Ra^{0.94}$

As can be seen on Fig.25, the wavelengths which initially dominate the instability appear to scale almost linearly with Rayleigh number. This in turn means that the characteristic length scales of initial plumes should approximately scale with the reciprocal of the Rayleigh number. This along with the findings of onset times agree well with results presented in linear stability analysis studies [8].

5.2.4 The time dependence of characteristic wavenumbers

Using the same power spectrum method eq Eq.49 one can define the time dependant characteristic wavenumber of the concentration fields in simulations with varying Ra . As has been shown in Fig.24 and Fig.25, both the onset time and initial wavenumber have a power law behaviour with Ra . Therefore the presented time dependent characteristic wavenumbers in Fig.26 have been scaled both in the time and wavenumber axes. As one can see, when rescaled in this manor, the dynamics are very similar at or around the onset of DDC. The transient region around $\tau \simeq 4\tau_{ons}$ is quite dependent on the individual Rayleigh number, and thus not a direct consequence of the onset values, but after the traansient region, the characteristic wavenumber appears to settle at roughly a quarter of wavenumber of the initial instability.




Figure 26: The time dependent characteristic wavenumbers of simulations with varying Ra . The time and wavenumber axes have been rescaled in terms of the onset time and initial wavenumber.

5.2.5 Total adsorbed mass

After the diffusive front breaks up and the convective mechanics begin, the rate of mass transfer from gas to liquid increases significantly. To quantify this, the total amount of dissolved carbon species is estimated by integrating the concentration fields over the area domain for each snapshot in time. The results of this are shown in Fig.27 for a handful of selected Ra values. The results are compared to the pure diffusion case of Eq.18. As can be seen, each of the runs start similarly, but diverge from the straight corresponding to their onset times. After the onset, the slopes increase significantly, resulting in a much higher effective macroscopic diffusion constant between the gas and the liquid.

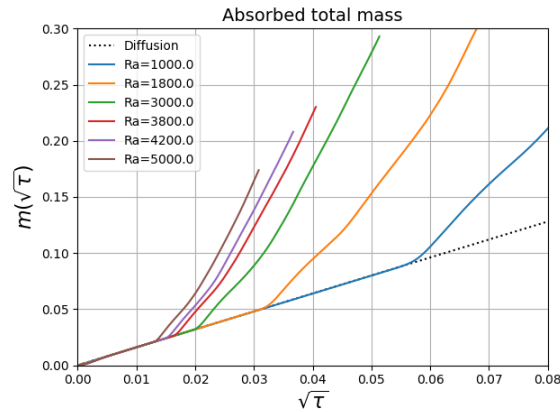


Figure 27: The time dependent absorbed chemical mass as a function of square root of time for various values of Ra .

6 Experimental methods

6.1 Designing the experiment

6.1.1 The experiment holder

For the design of the experimental apparatus itself, the following criteria were imposed:

- The main purpose of the experiment holder is to provide a sturdy base for attaching flow cells illuminated by a back light.
- The plane in which the experiment takes place must be able to tilt from vertical to near horizontal.
- An emphasis is laid on modularity - that is being able to re-purpose the experiment holder for experiments of various sizes. Also if parts need to be replaced for unforeseeable reasons, this should be as easy as possible.

To meet these criteria, a frame was built around a consumer grade LED light box (IKEA FLOALT 60x60cm). The frame consists of machined aluminum blocks and 20x5mm profiles, to which the light box is attached. This frame is hinged to a sturdy base at the bottom, and connected in a triangular fashion to a pair of cylindrical rails. This telescoping action along with locking wing nuts allows for setting the light box at any angle with respect to the ground. The front-most profiles are connected to the rest of the assembly with M6 bolts, which can be attached in multiple places, such that experiments from 30cm to 54cm can be placed on the holder without visual obstruction from the rails.




Figure 28: A CAD drawing of the experiment holder, with its most distinctive features pointed out.

6.1.2 Constructing the flow cell


For the design and construction of the flow cell, the following criteria was imposed:

- As a starting point, the flow cell should be a conventional Hele-Shaw cell, thus have a well defined uniform thickness.
- The flow cell must allow for gas flow in and out. The aim is to maintain a constant atmospheric pressure of CO₂. Other gases should be expelled from the cell during experiments.
- The cell must be easy to clean and rinse, so no leftover gas or dissolved carbon species remain in the cell between experiments, potentially contaminating consecutive experiments.
- The flow cell should be modular, in the sense that one can perform experiments with different plate spacings and with or without glass beads, which can impose a non-uniform permeability.


The requirements listed above were met by constructing a flow cell, consisting of two plates, sealed by a custom molded silicone gasket. The plate dimensions are 50x50cm. The outer plate is conventional glass, and the inner plate which is placed against the model holder is made of clear acrylic. The acrylic plate has machined inlets and outlets that allow for gas and liquid flow in and out of the cell. Depending on the task at hand, the various holes can be blocked or left open.

The gaskets are made of an addition-curing silicone moldmaking compound (Köraform). The benefits of this compound for creating a custom gasket is low pouring viscosity before hardening, fast curing, neglectable shrinkage and high mechanical strength. After mixing the two components, a uniform slab is made by letting the silicone compound cure while being clamped between the two plates. Metal spacers are inserted between the plates at the points of clamping to discourage warping of the plates, to achieve a thickness which is as uniform as possible. After curing, slab is demolded and then trimmed to the desired dimensions. Since the holes in the acrylic backplate are spaced over a region of roughly 30cm, the gasket is trimmed to similar dimensions. During the entire molding and trimming process one has to be mindful of any eventual air bubbles which may be present in the silicone compound. Most of these can be expelled before molding, and any remaining pockets of air can be removed by a suitable trim layout. The entire process is visualized in Fig.29. The silicone gasket used in the majority of the experiments has a thickness of 2mm and is trimmed to contain experiments confined within a 32x32cm region.


When clamping the cell together for experiments, a small amount of silicone grease is spread over the gasket, where it interfaces the two plates. The purpose of this is to ensure a good liquid seal around the entire gasket, without having to resort to high clamping pressures, which might warp the plates and cause a non-uniform permeability within the cell. With the silicone grease in place, one can easily see when a good seal is achieved while clamping the cell together, before adding any liquid.



(a) The liquid silicone compound is distributed around the region of interest, and is clamped between the two plates.



(b) After curing, a pre-cut template of the desired dimensions is used to find the most suitable trim.




(c) The finished gasket after trimming, ready for experiments.

Figure 29: The molding process. Note the holes left behind by the metal spacers used to assure a constant thickness.

6.2 Visualizing acidity

The raw data in the experiments conducted come in the form of images of the liquid under consideration. As previously mentioned, the carbonic acid equilibrium is a system of various dissolved inorganic carbon species which interact with each other, the solvent and the gas phase. Thus a relatively straight forward method of visualizing changes in acidity is to utilize pH indicators. This method is rather common in previously conducted work, but comes with some limitations [14].

A pH indicator is for all intents and purposes a weak acid, which has the added property that the protonated and dissociated forms of the molecule have different absorption spectra in the visible range. As per the definition of weak acids, this implies that each pH indicator has an associated equilibrium constant, according to the reaction:



This implies that there is a similar concentration of the protonated and dissociated forms of the indicator around $\text{pH} \simeq \text{p}K_a$. Due to the logarithmic nature of the pH scale, the relative concentrations of these two forms grow and shrink exponentially around this acidity value. Therefore a color change is only observed when $\text{pH} \simeq \text{p}K_a$. Typically a deviation in pH of about 1 logarithmic unit makes either form completely dominate, and no color change is observed past this point. For the application in question, this means that any single pH indicator is only useful to visualize a limited range of carbonic acid concentration. Therefore, the choice of pH indicator inherently sets an upper and lower bound on the concentrations one can deduce.

6.2.1 Making a suitable indicator liquid

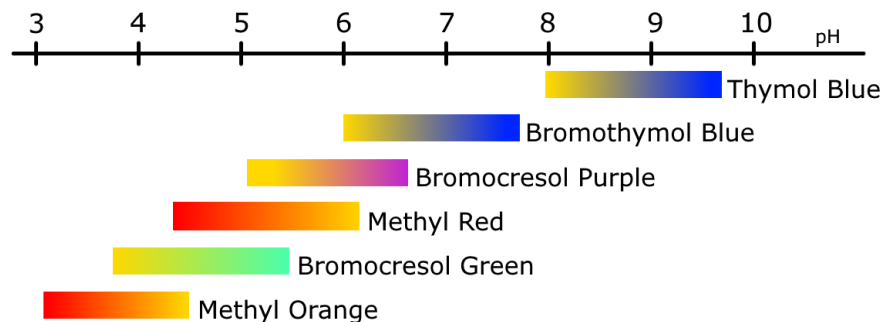


Figure 30: A visual representation of the various pH indicators investigated. The colored regions show at which acidity values the indicators transition, and between what colors. Outside of the shown ranges the indicators hold a solid color. For instance Bromothymol Blue is active in a pH range from 6 to 7.6 In more acidic conditions it stays yellow, and in more basic conditions it stays blue.

To combat this limitation, several pH indicators are combined in an attempt to form a more continuous color spectrum, such that more information about the pH fluctuations can be extracted from the image data. Fig.30 shows the active ranges and colors of the pH indicators considered for this purpose. The goal is to produce a solution of these indicators that is as active as possible down to the equilibrium pH of water in contact with atmospheric pressure of pure CO_2 . Thus the ideal solution goes through significant and distinctive color variations from neutral conditions, to a pH of about 4.

One inconvenient fact about pH indicators, is that they are generally not very water soluble (although this varies greatly between the different compounds). To work around this, the color mixtures under consideration are first dissolved in a small amount of 2-propanol, in which they are all highly soluble. Since the solid indicator powder obtained for these experiments contain the acid form of these compounds, the solution is then diluted with Type II water and neutralized with 0.01M NaOH base.

The mixtures under consideration are then titrated to various acidity levels using 0.01M NaOH and HCl along with a pH electrode. The mixtures are checked for color change in the desired pH ranges and components are added or dropped where appropriate to achieve better color resolution. After successive tests, a satisfactory candidate was found. A concentrated solution was mixed and used as a stock for later dilution. The details of this solution is shown in Tab.2 This indicator solution is then titrated to various pH levels of interest, injected to the flow cell and photographed. This is done to estimate an RGB value of the different values for calibration purposes. This is then interpolated to further visualize how this wide-range indicator fluctuates over a much wider range, compared to the more binary nature of its individual counterparts.

Table 2: The wide-range pH indicator developed for CO₂ DDC experiments. The first column is a stock solution, and the second column is the 1:50 dilution of the stock, which is used for experiments. In addition to the color content, the solvent additives are shown below.

	Stock	1:50 Dilution
Thymol blue [mg/L]	1250	25
Bromothymol blue [mg/L]	2500	50
Methyl red [mg/L]	1600	32
2-pronanol %vol	12.5	0.25
NaOH [mol/L]	$2.5 \cdot 10^{-3}$	$5 \cdot 10^{-5}$

Note: The proposed indicator has distinctive color fluctuations over most of the pH levels of interest. Due to the logarithmic pH scale, definition is most important closer on the acidic end of the scale, i.e. around pH 4. That begs the question why an addition of bromocresol green and/or methyl orange was not applied to this mixture, as they are active in this range. The addition of bromocresol green was found to have neglectable effect, and although methyl orange makes the yellow-orange-red transition more pronounced - the small gain in resolution comes with the drawback that methyl orange is known to be acutely toxic and was not deemed safe considering the fact that the flow cell is a prototype, which will be disassembled and possibly repurposed multiple times. Even with proper personal protective equipment, the risk of exposure during experimental setup or cleaning after completed experiments is deemed too great to warrant it. For future work with a more robust flow cell and experimental routine, it might be worth reconsidering.

A side note on the use of pH indicators

As previously explained, the fundamental working principle of pH indicators is that they themselves are weak acids. This by definition implies that their respective acid-base equilibria couple with the carbonic acid equilibrium system under consideration. Upon taking molar weights into account, one can check from Tab.2 that the concentration of the color components are of the order of magnitude $10^{-5}M$. Literature suggests that the total concentration of CO₂ derivatives in water in equilibrium with atmospheric pressure of pure CO₂ is of the order of magnitude $10^{-2}M$, which is 3 orders of magnitude greater than that of the indicators [16]. That being said, the concentration of the individual carbonic species can be significantly lower as demonstrated on Fig.4b.

Thus the use of color indicators in general could in theory interfere with the very thing under observation. This potential issue is simply postulated here, and the later analysis will neglect any chemical equilibrium interference which might arise from these interference effects. Previous literature suggests this should be a valid assumption [14].

6.2.2 Image analysis

All image analysis is performed in Python 3.7, using the image analysis and computer vision library `OpenCV 4.2`. This library interfaces with `NumPy` for numerical calculations and transformations and `Matplotlib` for plotting and visualization. The applied image analysis consists of taking the raw image, applying a simple geometrical transformation to it, identifying and masking the liquid region and determining the pH field from the color of the liquid.

Before applying the color interpolating method, every picture frame needs to be pre-processed geometrically and masked to display only the liquid region. The steps taken are displayed in Fig.31. Since the camera is never perfectly aligned with respect to the flow cell, a bounding box for the

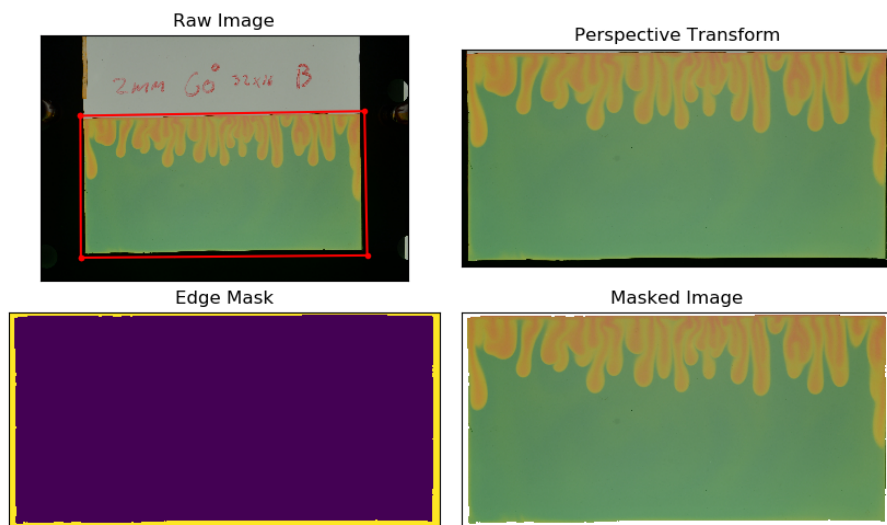


Figure 31: Going from raw image to an image of only the liquid region with the correct perspective.

liquid region is found and a perspective transformation is applied to transform this into a rectangle. By applying this perspective transformation, the plane of the experiment now matches the plane of the transformed image. Afterwards, a mask is applied to identify the non-liquid regions. This is a relatively simple procedure considering three of the edges are black and one of the edges are white. A simple threshold of one of the color channels with a little mask dilation is enough to sufficiently determine the mask. The masked pixels are then removed, and do not take part in the color interpolation algorithm.

Interpolating color


Now that the fluid region has been isolated, the task at hand is finding a meaningful way of correlating a color to a pH level. From the raw image file, the color of each pixel is represented by three 8-bit integers, for respectively the red, green and blue channel. For all intents and purposes, this can be thought of as a 3D vector space, and each pixel can be treated as an element of said vector space.

By interpreting the three channels as coordinates of a vector space, we can produce a scatter plot that visualizes the color change undergone by the liquid throughout the experiment. Fig.32 shows three ways of visualizing said change. Each point of the scatter plots are colored in with the associated value. First, we see that the observed color values are essentially scattered around a path which traverses through the color space. The goal is then to parameterize this path in terms of pH, to correlate a location in color space with an acidity value


Consider a region cropped from the center of the final frame of Fig.31. The region is chosen to be representative of the color spectrum the interpolation method must deal with. Before doing any interpolation, it's worth considering the general properties of the color change under study. By interpreting the three channels as coordinates of a vector space, we can produce a scatter plot that show the placement of each individual pixel in the color space. Upon inspection of Fig.32, one can see that representative collection of pixels are scattered around a path in the three dimensional color

space.

If the coordinates of the color space are transformed from typical Cartesian (R,G,B) to spherical, one obtains a radial coordinate (analogous to brightness) and two angular coordinates, which represent the color hue. Since the brightness of a pixel has no effect on which pH value should be assigned to it, this component can be discarded. The valuable information is in the polar and azimuthal angles, which can then be used to correlate a pixel color to a pH.



(a) The unaltered RGB space. The acidification process can be viewed as a path through this space.



(b) The normalized RGB space represented by a polar and azimuthal angle. The black path is calibrated with known pH values.

Figure 32: Two coordinate systems for the RGB color space. The points represent pixels from the center region of the pre-processed frame in Fig.31. In each case the pixel is placed according to the color space representation, and colored in for clarification.

By titrating the pH indicator solution to various known acidity levels, one can construct a calibration path by injecting the mixtures with known pH levels into the flow cells and imaging them. For each calibration image, a uniform region, free of air bubbles and other artifacts is found, and used to produce an averaged color value associated with that pH value. This is repeated multiple times, to produce a piecewise linear calibration path. To deduce the pH of any unknown pixel, one then finds the shortest distance from it to this calibrated path in the angular representation of the RGB color space. The shortest distance from any point to the calibration path can be found analytically, and said point is then assigned to the line segment to which the distance is shortest. On this line segment, the distance between the two neighbouring calibration colors can be used to interpolate the pH value, which will then be assigned to the pixel in question. This method is less sensitive to image noise, as it relies on both coordinates, rather than for in-





Figure 33: A visual representation of the pH determination procedure. Here a pH value is determined for 4 example colors, by finding their closest point on the calibration path.

stance interpolating pH as a function of one value, such as the color hue of the HSV color space. Despite this, a small Gaussian filter is still applied to the resulting pH field, to filter out pixel-scale fluctuations, and to combat the fact that the raw image data is discrete. The end result of the pH determination algorithm is shown in Fig.34, where the input image is a cropped section of Fig.31, for clarification purposes.

Due to the nature of the data being imported, the quickest way to perform the pH determination is to produce a look-up table of all the possible 256^3 combinations of RGB values. By doing this, every analyzed pixel essentially accesses a pre-calculated array of pH values, which is much faster than performing the coordinate transforms and collapsing them to the calibration path for every pixel of every analyzed image. Additionally, during image analysis, the frame can often be resized severely without loss of pH definition, to reduce the computational requirement and file storage.



(a) A cropped region of the pre processed frame of Fig.31

(b) The result of the pH determination algorithm.

Figure 34: An example input and output of the complete pH determination method. With this pH indicator mixture and color analysis method, a continuous pH field is determined, yielding much more information than a conventional thresholding and/or edge detection method.

6.3 Refining the experiments

The majority of experiments were performed by placing the molded silicone gasket, such that the row of liquid inlet holes were exactly on the lower boundary of the flow cell. The idea behind this was to set up the experiments in such a way that the flow cell can be undisturbed, and simply flushed out and replenished for consecutive experiments. This could be achieved by clamping the gasket between a glass plate and a plexiglass plate with suitable connectivity in the form of tubing inlets, as shown on Fig.35.

Initially, this method appeared to work well, but strong artifacts were noticed. Considerable amounts of time was spent attempting to eliminate these. Among the strongest unknown effects as well as their remedies are:

- Rinsing out the old fluid before a new experiment proved more challenging than first assumed. Despite repeated drainage and refilling, streaks of low pH regions were frequently observed. Swapping around the inlet and outlet ports in order to promote flow across the entire cell helped this issue, but never got rid of it completely.

- Evacuating remaining CO_2 posed an issue, since even small amounts of left over gas would immediately produce a color change when fresh indicator liquid was injected into the cell. This was solved by injecting more fluid after each experiment, to expel gas by reducing the gaseous volume between consecutive experiments.
- After long experiments, some components developed a slight red tint. This is likely due to some discoloration or possibly slight precipitation of methyl red. Because of this and the aforementioned issues, a cleaning solution of dilute NaOH and 2-propanol was mixed, to flush through the flow cell between experiments. This mixture is a good solvent with respect to methyl red, and also encourages dissolution of CO_2 gas due to its relatively high pH.




Figure 35: A diagram explaining the available fluid connectivity for the experiments. For most runs, all the liquid inlets and outlets were connected together with t-joins, and through a set of valves, these could either be isolated, connected to a reservoir of indicator liquid or drain to the sink. This was done by placing the reservoir of fresh fluid high enough and the drain low enough as to utilize gravity for these operations. Three of the gas inlets were connected to CO_2 , while the remaining two served as vents to assure steady atmospheric pressure.


While the initial issues were dealt with in a suitable manor, by the remedies listed previously, another artifact still remained. After the initial onset of convective dynamics, the fingers would rapidly translate towards the edges of the cell and sink there. Further details of this is shown in the results section, but it had the effect of making vast amounts of carbonic acid flow down the sides, to the point where it would clearly dominate the dynamics in the rest of the cell. Troubleshooting this proved a lengthy process. A few plausible reasons for this artifact were considered. Among these were whether or not chemical residue on the gasket could be the cause of the density growing faster there, due to for example dissolution of any possible residue. The concern for this was in part due to some slight discoloration mentioned above, and so proper rinsing should solve that. Also, the tubing inlets used for indicator liquid could potentially be at fault, allowing for some added permeability

between the various holes (promoting flow between the opposing edges of the cell). This was dealt with by clamping each separate liquid inlet, instead of only closing valves to and from the sink or source.


This, while sensible, didn't seem to solve the tendency towards flow down along the cell edges. Therefore, the approach of having a refillable cell was discarded, and it was instead dismantled, cleaned and reassembled between each experiment, with proper cleaning and drying of all components. While the artifact still persisted, it was only significantly later that it was discovered that the plexiglass plate with the tubing inlets and outlets was in fact slightly warped. While near impossible to spot with the bare eye, when placed against a known flat surface, the warp was detectable. It was measured with calipers that the height discrepancy between the middle of the plexiglass plate and the four corners was around 2mm. It was also measured that the distortion was very similar in both directions. Thus one can assume that the plexiglass plate is parabolically shaped with curvature $\simeq 1\text{mm}/250^2\text{mm}^2 = 1.6 \cdot 10^{-5}\text{mm}^{-1}$ in both the X and Y directions, granted the curvature is similar in both directions and half the plate width is 250mm.

Considering how the molding process was performed, this in fact means that the edge of the gasket has a well defined 2mm thickness, but the effective thickness of the flow cell varies within the liquid region. With this previously mentioned curvature, after taking the thickness and placement of the gasket into account, a parabolic approximation estimates that the effective plate spacing in the center is less than 1.6mm. That's a roughly 25% thickness increase, and although the permeability $k = \frac{b^2}{12}$ only applies for uniform thickness, one can assume the permeability discrepancy to be closer to 55%, which is clearly not insignificant.

While this wasn't noticed until late in the experimental phase, another system was constructed without said plexiglass plate. Instead the new system would consist of two glass plates, neither of which had any tubing in- or outlets. The glass plates measured 35cm x 35cm with a thickness of 12mm and 16mm respectively, and are thus much stiffer than the previous setup. Due to the lack of connectivity, the filling of fluid and injection of gas would need to be handled in another manor. A similar gasket was molded, with a flow domain of 30x30cm. Only the bottom half of this region is used, and two slits are made in the top, to allow for Luer lock syringe needles to be inserted into the cell. The layout of this is shown on Fig.36. The needles are flattened to allow them to fit into the 1.0mm gap. Since the previous system was expected to be 2.0mm, but the reality being that the center of it was significantly thinner, it was decided to make this gasket thinner as well.



(a) A schematic of the construction of the revised flow cell. The liquid inlet needle is only used for filling up the cell, and the opening serves as a vent for excess gas during experiments.



(b) Photo of the revised flow cell after clamping the plates together. Outside the liquid region, an experiment identifier is written on the glass with whiteboard marker, and a ruler is attached to scale the images correctly.

Figure 36: The simpler, yet improved flow cell, constructed of two glass plates with a similar silicone gasket as the previous system. The liquid region is 30cm wide and 15cm tall with a thickness of 1mm.

Artifacting due to non-uniform permeability was definitely less in the revised cell, and the edge-seeking trend of the previous flow cell configuration was not appearing any more. However, either due to clamping pressure or some other uncontrollable variable, the newer system appears to exhibit the opposite trend, where most of the carbonic acid would now sink down towards the center. This will be shown in the results section. Therefore, from an experimental design perspective, we come to the undeniable conclusion that the flatness of plates and in turn the uniformity of the permeability field can be incredibly difficult to achieve.

7 Experimental results


7.1 Patterns in the 2mm cell

As discussed in the previous section, the first flow cell had a warped plate and thus a very uneven permeability field. Two concentration fields from an experiment in the $\theta = 90^\circ$ angle are shown in Fig.37. The two snapshots in Fig.37 clearly show how plumes that sink down along the edges cause convection on a cell scale, which in turn washes out the other plumes so they become very distorted. This effect was noticeable in all experiments in this flow cell, and as was discussed in the previous section, the warping of one of the plates was deemed the reason for this artifact.


To demonstrate the importance of the permeability field in this situation, a simulation is performed to see qualitatively how a small plate separation difference can have major consequences for the dynamics and patterns. The developed numerical model is used to simulate a case of $Ra = 3000$ with a permeability fluctuation function:

$$f(x, y) = A \left(1 + a \left((x - 1)^2 + (y - 0.2)^2 \right)^2 \right)$$

Note that this permeability field as well as the results shown for the simulation are in the non-dimensional units as per Eq.11. This function imitates a roughly constant thickness with a small additional parabolic element in both directions, centered at $x, y = (1, 0.2)$ such that the thickness grows away from this center. A modest height difference from center to edge of 10% is applied by setting $a = 0.1$ and A is determined such that the function is correctly normalized. Two snapshots of the corresponding concentration- and flow fields are shown on Fig.38. These results qualitatively show how even a small imbalance in permeability can majorly affect the observed concentration fields and how they develop.




(a) Once the experiment is underway, the plumes translate towards the edges and sink rapidly there.

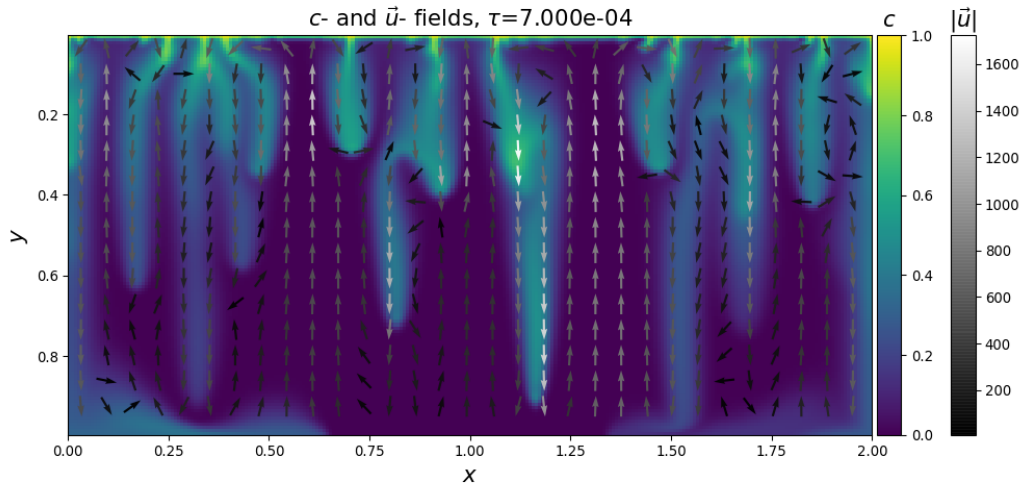


(b) A cell scale convection takes place, since fluid is being forced down along the edges, which completely dominate the dynamics in the rest of the cell.

Figure 37: The edge artifact which persisted in all experiments in the 2mm flow cell with a warped plexiglass plate. Already after the initial onset, plumes shift laterally towards the edges and sink rapidly there, causing cell scale convection which dominates the DDC dynamics in the rest of the cell.



(a) After the DDC dynamics begin, plumes shift towards the edges where they sink rapidly.



(b) Shortly afterwards, the sinking liquid at the cell edges cause cell scale convection. Strong upwards currents are being driven by this on either side of the middle.

Figure 38: A simulation meant to imitate the effect of a plate separation which varies parabolically. The expression for the permeability field is thus assumed to be a fourth order polynomial, in this case centered at $x, y = (1, 0.2)$. The strength of this fluctuation imitates a 10% difference in plate separation from the center to the edges. This can give rise to an experimental artifact which is qualitatively similar to the one observed in Fig.37 Note the dimensionless units as per Eq.11.

7.2 Patterns in the 1mm cell

Images obtained of flow patterns in the refined 1mm cell were used along with the proprietary pH color analysis method and Eq.29 to determine the concentration fields. Numerous experiments were performed with cell tilt angles ranging from $\theta = 30^\circ$ to $\theta = 75^\circ$. Fig.40-Fig.44 show a few snapshots from one of the runs at each angle. Although this cell was refined specifically to improve on the unfavourable permeability fluctuations of the previous one, the results clearly indicate that there is a strong tendency for very asymmetric flow. This is also shown in the space-time maps on Fig.45-49, in which we clearly see the plumes gravitate towards the right. This is most likely due to some remaining permeability fluctuations, which could originate from the experimental design itself (i.e. materials used, such as the plates and/or molded gasket), or simply by the act of opening cleaning and reassembling and potential inaccuracies accompanied with that. One of these possibilities is that the clamping pressure could be warping the plates slightly, which could be enough to cause minor fluctuations in effective plate spacing. There is inherently no way to guarantee that the clamping pressure is similar between the experiments, and as such, this could give rise to some warping which is not necessarily consistent between different experiments.

The time dependent characteristic wavenumbers as per Eq.49 is shown on Fig.39. The results clearly indicate the characteristic wavenumber decreasing with time, but upon further inspection one finds that the height of the wavenumber peaks do not follow the pattern one would have expected. For instance, in the case of 45° and 60° , the latter appears to have a lower peak wavenumber than the former. This goes against the expectation, that for $Ra \propto \sin\theta$ one would expect the more vertical angle to have a noticeably higher finger count and thus higher characteristic wavenumber.

Despite this inconsistency, the obtained concentration field data does demonstrate that the developed pH indicator method can indeed account for concentration variations within and around convection plumes, and give an estimate of the total amount of dissolved carbon. This is data which can not be well defined with single indicator methods or schlieren imaging. Despite the issues with asymmetrical flow, the results indicate do demonstrate how the indicator, color analysis and chemical equilibrium model can help investigate the continuous nature of the concentration fields.




Figure 39: Characteristic wavenumbers for experiments in the 1mm cell.




Figure 40: Snapshots from one run with tilt angle $\theta = 75^\circ$. The middle of the cell clearly dominates the rest.




Figure 41: Snapshots from one run with tilt angle $\theta = 60^\circ$. In this case, the plumes appear to shift laterally to the right.




Figure 42: Snapshots from one run with tilt angle $\theta = 45^\circ$. In this case, the plumes appear to shift laterally to the right.




Figure 43: Snapshots from one run with tilt angle $\theta = 37^\circ$. In this case, the plumes appear to shift laterally to the right.




Figure 44: Snapshots from one run with tilt angle $\theta = 30^\circ$. In this case, the plumes appear to shift laterally to the right.




Figure 45: A space-time map of one of the runs with tilt angle $\theta = 75^\circ$. This clearly demonstrates how the plumes gravitate towards the middle or slightly to the right.




Figure 46: A space-time map of one of the runs with tilt angle $\theta = 60^\circ$. In this case, the plumes move laterally to the right.




Figure 47: A space-time map of one of the runs with tilt angle $\theta = 45^\circ$. In this case, the plumes move laterally to the right.




Figure 48: A space-time map of one of the runs with tilt angle $\theta = 37^\circ$. In this case, the plumes move laterally to the right.




Figure 49: A space-time map of one of the runs with tilt angle $\theta = 30^\circ$. In this case, the plumes move laterally to the right.

8 Conclusions

The inspiration for this project came mainly from previous work on plume patterns in DDC [12][13], and especially an article by Thomas et al. which demonstrated the importance of proper use of pH color indicators for visualization of the acidic plume formations [14]. This was an inspiration to attempt to derive a more rigorous method by using the principles of analytical chemistry, a proprietary color indicator mixture and image analysis to shed light on the continuous nature of the pH and concentration fields that is often neglected in similar experiments. Most comparable experiments rely on a single indicator only, and thus puts an inherent threshold on which concentrations one can detect, and most concentration values beyond that point remain undetectable due to the binary nature of pH indicators.

As a necessary pre-requisite for this, a DDC experiment was designed from scratch, and an emphasis was put on the prototype being easy to re-purpose or reconfigure. A working setup was made, which was able to perform DDC experiments at various angles and which was also easy to reconfigure for other system sizes and/or thicknesses. However while the desire for easy re-configuration was welcomed for prototyping purposes, the conclusion is drawn that this may in turn also be what made the experiments so unreliable.

The first flow cell suffered from a warped plexiglass plate, which meant that the permeability was far from uniform within the liquid region. Various potential issues were tackled before the warping of the plate was later discovered to be the root of the problem. Simulations qualitatively back this up, and demonstrate how a parabolic change in cell aperture of only $\simeq 10\%$ can strongly affect the obtained concentration field patterns. Even after building another flow cell from scratch, this problem was never truly avoided. Therefore the conclusion is drawn that the single biggest improvement one could make to the presented work would be to design a more rigorous flow cell using a completely different build method. On a side note, both Faisal et al. and Kneafsey and

Pruess reported similar issues with non-uniform plate separation, but not the same extent as here [12] [13]. It is quite plausible that the act of clamping together the model is a source of inconsistency. Given that the model needs to be cleaned between each experiment, one needs to be certain that when reassembled, the flow cell geometry is sufficiently consistent in order to investigate the effect of external parameters such as the reduction in gravitational force.

One of the ideas behind the modular design was to simplify adding for instance glass beads or other sources of varying permeability fields, with only little modification to the experimental setup. It was decided to first focus on experiments without beads, to fully develop the color indicator method and leave such experiments for later. This was done with the idea in mind that the uniform case should be easier to start with and develop the concept. However, one might argue that the addition of packed beads could actually help stabilize the permeability field and might reduce some of the experimental artifacts seen here.

While this does mean that the experimentally obtained plume morphologies are heavily affected by experimental artifacts, the developed proprietary color analysis method was clearly able to identify the varying pH field, and coupled with the assumed pseudoequilibrium model, was able to estimate the carbon concentration fields. Therefore one can conclude that using a universal pH indicator, or a mixture of multiple individual color indicators can in fact prove useful to visually investigate patterns of fluids in which chemical reactions take place. A recommendation for further study within this topic would be to more rigorously validate at which indicator-concentrations does the presence of an indicator significantly alter the system dynamics. For this system it has been argued that the concentrations used do not affect the system in a meaningful way [14], but that may not be the case for other systems one might want to apply a similar approach to.

Due to unforeseen circumstances, a significant portion of the project was spent outside of the laboratory. During this time, a numerical scheme was developed, which is able to simulate DDC with various conditions including geometries, boundary conditions and non-uniform permeabilities. The developed code utilizes a Finite Volume Method for the concentration field, a Finite Difference Method for the flow field and an operator splitting ADI method for time propagation. An optimized direct solver algorithm was written which scales incredibly well with growing system size while the error is maintained at second order in time increments and with guaranteed stability. Therefore the presented flow patterns are significantly denser than most other results in literature on similar systems [20]. The results with regards to onset of DDC match those of linear stability analysis to a reasonable degree of accuracy [8]. The code proved especially useful to investigate the general trends in plume morphology, and to give better insight into why the experimental artifacts appeared how they did. The developed code is openly available on GitHub for all to use and/or improve upon [21].

Appendix: Mathematical expressions

An analytical solution of the 1D diffusion equation

Consider Eq.16, which needs to be solved for a concentration field as a function of position and time. By using separation of variables, introduce the basis function:

$$\Phi(y, \tau) = f(\tau)g(y)$$

When acted on with the differential operator of Eq.16, one finds a condition for fractions of derivatives of f and g , which must hold true for all y and τ , thus they must be constant.

$$\frac{f'(\tau)}{f(\tau)} = \frac{g''(y)}{g(y)} = \text{const.} = \alpha$$

Now one can find different types of solutions for different α . Note that if α is positive, that implies that f becomes a diverging exponential, which can not be the case, since we know that as $\tau \rightarrow \infty$ the concentration field remains finite. Additionally, one can check that the corresponding $g(\tau)$ would not fit the BCs. Therefore we focus on the remaining two cases.

If $\alpha = 0$, one finds that $f(\tau) = \text{const.}$ and $g(y) = ay + b$. So in the steady state, the concentration profile is linear. This along with the zero-flux boundary condition at $y = 1$ tells us that there must be a contribution of $\Phi_0 = 1$, in the $\tau \rightarrow \infty$ limit.

Last but not least consider $\alpha < 0$. In this case, the argument in the exponential function for f will be real, but using the characteristic polynomial one finds a complex argument for g . This gives rise to oscillatory solutions:

$$f(\tau) = e^{-|\alpha|\tau} \quad g(y) = A \cos(\sqrt{|\alpha|}y) + B \sin(\sqrt{|\alpha|}y)$$

The cosine function interferes with the boundary condition at $y = 0$, so only the sine function is valid for this solution. Furthermore, given the zero-flux boundary condition at $y = 1$, we must have:

$$g'(y) \Big|_{y=1} = B\sqrt{|\alpha|} \sin(\sqrt{|\alpha|}y) = 0 \quad \sqrt{|\alpha|} = k_n = \frac{\pi(1+2n)}{2}$$

Having identified the possible basis solutions, we construct a linear combination of the form $c(y, \tau) = \sum_n c_n \Phi_n$, as follows:

$$c(y, \tau) = 1 + \sum_{n=1}^{\infty} c_n e^{-k_n^2 \tau} \sin(k_n y)$$

In order to determine the c_n constants, we apply the condition that at $\tau = 0$ the concentration field is 0 everywhere on the domain of the solution. Therefore:

$$\sum_{n=1}^{\infty} c_n \sin(k_n y) = -1$$

And due to the orthogonality of the sine functions, by taking an inner product with a sinusoid of the form $\sin(k_m y)$ on both sides, one finds that $c_n = -2/k_n$. Thus we conclude that:

$$c(y, \tau) = 1 - 2 \sum_{n=1}^{\infty} k_n^{-1} e^{-k_n^2 \tau} \sin(k_n y)$$

A more detailed derivation of Eq.25 and Eq.29

From the charge conservation equation Eq.26, we find that:

$$[\text{HCO}_3^-] + 2[\text{CO}_3^{2-}] = [\text{H}^+] + [\text{Na}^+] - [\text{OH}^-]$$

From the expression for K_{a2} in Tab.1, we find that:

$$[\text{CO}_3^{2-}] = [\text{HCO}_3^-] \frac{K_{a2}}{[\text{H}^+]}$$

Insert this into the charge equation to find that:

$$[\text{HCO}_3^-] = \frac{1}{1 + 2K_{a2}/[\text{H}^+]} \left([\text{H}^+] + [\text{Na}^+] - \frac{k_w}{[\text{H}^+]} \right) \quad (\star)$$

Having found an expression for the bicarbonate ion as a function independent of the other carbon species, we now derive similar relations for the remaining two species, in terms of the bicarbonate concentration:

$$\begin{aligned} [\text{H}_2\text{CO}_3] &= [\text{HCO}_3^-] \frac{[\text{H}^+]}{K_{a1}} \\ [\text{CO}_2] &= [\text{H}_2\text{CO}_3] \frac{1}{K_H} = [\text{HCO}_3^-] \frac{[\text{H}^+]}{K_{a1}K_H} \end{aligned}$$

Insert these in the expression for total dissolved carbon as per Eq.28 to find that:

$$C_T = [\text{HCO}_3^-] \left(\frac{K_{a2}}{[\text{H}^+]} + 1 + \frac{[\text{H}^+]}{K_{a1}} + \frac{[\text{H}^+]}{K_{a1}K_H} \right)$$

Finally insert the independent expression for $[\text{HCO}_3^-]$, (labeled \star) derived from the charge conservation condition to arrive at Eq.29. As discussed in the relevant section, this gives an estimate for the total carbon content of a fluid element in which these four carbon species are in equilibrium with each other, but not with the gaseous CO_2 .

The derivation of Eq.25 is done in a very similar manor. Assuming that $pH \ll pK_{a2}$, we can neglect the fraction in (\star) , to find that:

$$[\text{HCO}_3^-] = [\text{H}^+] + [\text{Na}^+] - \frac{k_w}{[\text{H}^+]}$$

Now in Eq.25, the system is a body of water in total equilibrium with a known partial pressure of CO_2 , and as such we apply all the equilibrium constants of Tab.1 to find that:

$$\frac{K_{a1}K_HK_h}{[\text{H}^+]} p_{\text{CO}_2} = [\text{H}^+] + [\text{Na}^+] - \frac{k_w}{[\text{H}^+]}$$

By multiplying both sides with $[\text{H}^+]$ a parabolic expression is obtained which can be solved for $[\text{H}^+]$. Only the positive solution lies within the domain of possible solutions.

Taylor approximations, and special cases in the FVM

Throughout the computational section about the Finite Volume Method and Finite Difference Method, centered difference quotients were used, as their error is second order in grid spacing [19]. However, one needs a different numerical scheme to handle the coupling between the top row of CVs and the top boundary, since there is no cell on the opposite side of the boundary to access. As a more detailed explanation of the derivations, consider without loss of generality the equations governing the top left corner CV. This cell has all cases of interactions. On its left side, there is no neighbour, and thus a zero-flux BC is introduced. In its top there is a Dirichlet BC, with constant concentration, and on the remaining two sides are other CVs to which advection and diffusion can happen. Recall that this cell is indexed 0, and its neighbours are 1 to the East and N_x to the South. Due to the non-zero BC in the West direction, this flux is omitted, and we need to evaluate the following fluxes:

$$\frac{\partial c}{\partial t} = F_n + F_s + F_e$$

As mentioned in the relevant section, when the path integral over the east facing line segment is evaluated, as per Eq. 32, once we take into account that the stream function is zero on the boundary, we find that:

$$F_e = \frac{1}{h_x} \left(\frac{c_E - c_P}{h_x} - \frac{\psi_{se}}{h_y} \frac{c_P + c_E}{2} \right)$$

Similarly, the southern flux is evaluated

$$F_s = \frac{1}{h_y} \left(\frac{c_S - c_P}{h_x} + \frac{\psi_{se}}{h_y} \frac{c_P + c_S}{2} \right)$$

Thus only the Northern flux where the Dirichlet BC condition is located remains. We start by applying the definition of the flux to be evaluated. Note that the unit vector is now in the negative direction, due to the coordinate system being defined as increasing downwards. The velocity term dies out on the top border because by definition the stream function is constant there

$$F_n = \frac{1}{A} \int_n (\nabla c - \vec{u}c) \cdot (-\hat{e}_y) dx = \frac{1}{h_y} \left(-\frac{\partial c}{\partial y} \right) \Big|_n$$

Since the partial derivative needs to be evaluated at the domain boundary $y = 0$, it becomes helpful to express the values c_P and c_S in terms of a Taylor series around $y = 0$. In general:

$$c(x, y) \simeq c(x, 0) + y \left(\frac{\partial c}{\partial y} \right) \Big|_{y=0} + \frac{y^2}{2} \left(\frac{\partial^2 c}{\partial y^2} \right) \Big|_{y=0} + \mathcal{O}(y^3)$$

This can then be evaluated at c_P and c_S , recall that the y -coordinate of these are $h_y/2$ and $3h_y/2$

$$c_P = c(x, 0) + \frac{h_y}{2} \left(\frac{\partial c}{\partial y} \right) \Big|_{y=0} + \frac{h_y^2}{8} \left(\frac{\partial^2 c}{\partial y^2} \right) \Big|_{y=0} + \mathcal{O}(y^3)$$

$$c_S = c(x, 0) + \frac{3h_y}{2} \left(\frac{\partial c}{\partial y} \right) \Big|_{y=0} + \frac{9h_y^2}{8} \left(\frac{\partial^2 c}{\partial y^2} \right) \Big|_{y=0} + \mathcal{O}(y^3)$$

Ultimately, we need to obtain an expression for the first derivative, so we weight these expressions and subtract them as to cancel the second derivative entirely:

$$c_S - 9c_P = -8c(x, 0) - 3h_y \left(\frac{\partial c}{\partial y} \right) \Big|_{y=0}$$

Solving this for the derivative yields an expression which is now second order in error with respect to h_y .

$$\left(\frac{\partial c}{\partial y}\right)\Big|_{y=0} = \frac{9c_P - c_S}{3h_y} - \frac{8}{3h_y}c(x,0) + \mathcal{O}(h_y^2)$$

Here $c(x,0)$ represents the concentration value given by the boundary condition. Naturally, the x -coordinate matches that of the CV in question, and can in theory vary between neighbouring CVs if one wants to introduce non-uniform behaviour. Inserting this expression for the Northern flux yields:

$$F_n = -\frac{1}{h_y} \left(\frac{9c_P - c_S}{3h_y} - \frac{8}{3h_y}c(x,0) \right) = \frac{c_S - 9c_P}{3h_y^2} + \frac{8}{3h_y^2}c(x,0)$$

One can now evaluate the sum of fluxes to determine the time derivative:

$$\frac{\partial c}{\partial t} = F_n + F_s + F_e = \frac{c_E - c_P}{h_x^2} - \frac{\psi_{se}}{h_x h_y} \frac{c_P + c_E}{2} + \frac{c_S - c_P}{h_y^2} + \frac{\psi_{se}}{h_x h_y} \frac{c_P + c_S}{2} + \frac{c_S - 9c_P}{3h_y^2} + \frac{8}{3h_y^2}c(x,0)$$

When the coefficient for each CV is isolated, and their numeric indexes are inserted, the matrix form of this equation becomes:

$$\left(-\frac{4}{h_x^2} - \frac{1}{h_y^2}\right)c_0 + \left(\frac{1}{h_x^2} - \frac{\psi_{se}}{h_x h_y}\right)c_1 + \left(\frac{1}{h_y^2} + \frac{\psi_{se}}{h_x h_y}\right)c_{N_x} = -\frac{8}{3h_y^2}c(x_0,0)$$

Where $c(x_0,0)$ is the value dictated by the Dirichlet BC on the Northern line segment of CV with index 0. This effectively becomes a source term for the top cell, and is the only way new mass can be introduced to the system. Within the rest of the system (that is rows 1- N_y) mass is only transferred between CVs and can not be created or destroyed. This is in turn how the right-hand-side vector \vec{b} appears in section 4.1.4, but from this derivation one can see that it is only non-zero in the first N_x elements (corresponding to the top row), and zero everywhere else.

9 Bibliography

- [1] Michael L. Szulczewski et al. “Lifetime of carbon capture and storage as a climate-change mitigation technology”. In: *Proceedings of the National Academy of Sciences* 109.14 (2012), pp. 5185–5189. DOI: [10.1073/pnas.1115347109](https://doi.org/10.1073/pnas.1115347109). URL: <https://www.pnas.org/content/109/14/5185>.
- [2] S.Ó. Snæbjörnsdóttir, B. Sigfússon, and C. et al. Marieni. “Carbon dioxide storage through mineral carbonation”. In: *Nat Rev Earth Environ* 1 (2020), pp. 90–102. DOI: <https://doi-org.ezproxy.uio.no/10.1038/s43017-019-0011-8>.
- [3] IPCC. *Carbon Dioxide Capture and Storage*. Ed. by B Metz et al. 2005.
- [4] Jonathan P. Ennis-King and Lincoln Paterson. “Role of Convective Mixing in the Long-Term Storage of Carbon Dioxide in Deep Saline Formations”. In: *SPE Journal* 10.03 (2005), pp. 349–356. DOI: [10.2118/84344-pa](https://doi.org/10.2118/84344-pa).
- [5] Erik Lindeberg and Dag Wessel-Berg. “Vertical convection in an aquifer column under a gas cap of CO₂”. In: *Energy Conversion and Management* 38 (1997). Proceedings of the Third International Conference on Carbon Dioxide Removal, S229–S234. ISSN: 0196-8904. DOI: [https://doi.org/10.1016/S0196-8904\(96\)00274-9](https://doi.org/10.1016/S0196-8904(96)00274-9). URL: <http://www.sciencedirect.com/science/article/pii/S0196890496002749>.
- [6] S. Chandrasekhar. *Hydrodynamic and hydromagnetic stability*. Dover, 1981.
- [7] S. Bories. “Natural Convection in Porous Media”. In: *Advances in Transport Phenomena in Porous Media*. Ed. by J. Bear and M. Y. Corapcioglu. Dordrecht: Martinus Nijhoff Publishers, 1987. Chap. 1, pp. 77–143.
- [8] H. Hassanzadeh, M. Pooladi-Darvish, and D.W. Keith. “Stability of a fluid in a horizontal saturated porous layer: effect of non-linear concentration profile, initial, and boundary conditions.” In: *Transp Porous Med* 65 (2006), pp. 193–211. DOI: <https://doi-org.ezproxy.uio.no/10.1007/s11242-005-6088-1>.
- [9] R. Farajzadeh et al. “Mass Transfer of CO₂ Into Water and Surfactant Solutions”. In: *Petroleum Science and Technology* 25 (Dec. 2007), pp. 1493–1511. DOI: [10.1080/10916460701429498](https://doi.org/10.1080/10916460701429498).
- [10] H. Karimaie and E. Lindeberg. “Experimental Verification of CO₂ Dissolution Rate Due to Diffusion Induced Convection”. In: *Energy Procedia* 114 (2017). 13th International Conference on Greenhouse Gas Control Technologies, GHGT-13, 14-18 November 2016, Lausanne, Switzerland, pp. 4917–4925. ISSN: 1876-6102. DOI: <https://doi.org/10.1016/j.egypro.2017.03.1633>. URL: <http://www.sciencedirect.com/science/article/pii/S1876610217318271>.
- [11] Yongqiang Tang et al. “Experimental Study on the Density-Driven Carbon Dioxide Convective Diffusion in Formation Water at Reservoir Conditions”. In: *ACS Omega* 4.6 (2019), pp. 11082–11092. DOI: [10.1021/acsomega.9b00627](https://doi.org/10.1021/acsomega.9b00627). eprint: <https://doi.org/10.1021/acsomega.9b00627>. URL: <https://doi.org/10.1021/acsomega.9b00627>.
- [12] Titly Farhana Faisal et al. “Quantitative and qualitative study of density driven CO₂ mass transfer in a vertical Hele-Shaw cell”. In: *International Journal of Heat and Mass Transfer* 81 (2015), pp. 901–914. ISSN: 0017-9310. DOI: <https://doi.org/10.1016/j.ijheatmasstransfer.2014.11.017>. URL: <http://www.sciencedirect.com/science/article/pii/S0017931014009892>.
- [13] T.J. Kneafsey and K Pruess. “Laboratory Flow Experiments for Visualizing Carbon Dioxide-Induced, Density-Driven Brine Convection”. In: *Transp Porous Med* 82 (2010), pp. 123–139. DOI: <https://doi-org.ezproxy.uio.no/10.1007/s11242-009-9482-2>.

- [14] C. Thomas et al. “Experimental study of CO₂ convective dissolution: The effect of color indicators”. In: *International Journal of Greenhouse Gas Control* 42 (2015), pp. 525–533. ISSN: 1750-5836. DOI: <https://doi.org/10.1016/j.ijggc.2015.09.002>. URL: <http://www.sciencedirect.com/science/article/pii/S1750583615300700>.
- [15] Ralph H. Petrucci et al. *General chemistry: principles and modern applications*. Prentice Hall, 2016.
- [16] N. N. Greenwood and A. Earnshaw. *Chemistry of the elements*. Elsevier Butterworth-Heinemann, 1997.
- [17] R. Sander. “Compilation of Henry’s law constants (version 4.0) for water as solvent”. In: *Atmospheric Chemistry and Physics* 15.8 (2015), pp. 4399–4981. DOI: [10.5194/acp-15-4399-2015](https://doi.org/10.5194/acp-15-4399-2015). URL: <https://www.atmos-chem-phys.net/15/4399/2015/>.
- [18] Joel H. Ferziger and Perić Milovan. *Computational methods for fluid dynamics*. Springer, 2002.
- [19] Brian Bradie. *A Friendly Introduction to Numerical Analysis*. Pearson, 2007. ISBN: 8131709426.
- [20] Rouhollah Farajzadeh et al. “Numerical simulation of density-driven natural convection in porous media with application for CO₂ injection projects”. In: *International Journal of Heat and Mass Transfer* 50.25 (2007), pp. 5054–5064. ISSN: 0017-9310. DOI: <https://doi.org/10.1016/j.ijheatmasstransfer.2007.08.019>.
- [21] *FYS4150 Github Repository*. URL: <https://github.com/hilmaryb/FYS4150>.