



PAPER

OPEN ACCESS

RECEIVED
17 May 2024REVISED
3 February 2026ACCEPTED FOR PUBLICATION
17 February 2026PUBLISHED
13 March 2026

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Beyond linear decomposition: a nonlinear eigenspace decomposition for a moist atmosphere with clouds

Antoine Remond-Tiedrez^{1,*}, Leslie M Smith² and Samuel N Stechmann² ¹ Department of Mathematics, University of Wisconsin–Madison, Madison, WI, United States of America² Department of Mathematics and Department of Atmospheric and Oceanic Sciences, University of Wisconsin–Madison, Madison, WI, United States of America

* Author to whom any correspondence should be addressed.

E-mail: a.remond-tiedrez@gmail.com**Keywords:** potential vorticity inversion, clouds, wave-vortical decomposition, inertia-gravity waves, water vapour, moisture, phase interface**Mathematics Subject Classification numbers:** Primary: 86A10, 76U60, Secondary: 37N10, 53Z05, 65N99

Recommended by Professor Beatrice Pelloni

Abstract

A linear decomposition of states underpins many classical systems. This is the case of the Helmholtz decomposition, used to split vector fields into divergence-free and potential components, and of the dry Boussinesq system in atmospheric dynamics, where identifying the slow and fast components of the flow can be viewed as a decomposition. The dry Boussinesq system incorporates two leading ingredients of mid-latitude atmospheric motion: rotation and stratification. In both cases the leading order dynamics are linear so we can rely on an eigendecomposition to decompose states. Here we study the extension of dry Boussinesq to incorporate another important ingredient in the atmosphere: moisture and clouds. The key challenge with this system is that nonlinearities are present at leading order due to phase boundaries at cloud edge. Therefore standard tools of linear algebra, relying on eigenvalues and eigenvectors, are not applicable. The question we address in this paper is this: in spite of the nonlinearities, can we find a decomposition for this moist Boussinesq system? We identify such a decomposition adapted to the nonlinear balances arising from water phase boundaries. This decomposition combines perspectives from partial differential equations (PDEs), the geometry, and the conserved energy. Moreover it sheds light on two aspects of previous work. First, this decomposition shows that the nonlinear elliptic PDE used for potential vorticity and moisture inversion can be used outside the limiting system where it was first derived. Second, we are able to rigorously justify, and interpret geometrically, an existing numerical method for this elliptic PDE. This decomposition may be important in applications because, like its linear counterparts, it may be used to analyse observational data. Moreover, by contrast with previous decompositions, its formulation includes the nonlinearity from the presence of clouds and latent heating.

Contents

1. Introduction	2
2. Background and motivation	3
2.1. Helmholtz decomposition	3
2.2. Dry Boussinesq and its decomposition	4
2.3. Moist Boussinesq	5
2.4. Literature on atmospheric dynamics with moisture and clouds, and mathematical theory	8
2.5. Simplified notation	9
3. Decomposition	9
4. Decomposition and slow–fast dynamics	18
5. Decomposition and energy	22
5.1. PDE-centric approach	24
5.2. Energy-centric approach	30
5.3. Metric-less approach	38
5.4. Summary	41
6. Iterative methods for the decomposition and their geometry	42
7. Summary and comparison with Helmholtz and the dry case	43
7.1. Dry Boussinesq	43
7.2. Moist Boussinesq	45
8. Conclusion	48
8.1. Looking forward	49
Acknowledgments	51
Data availability statement	51
Appendix A. Iterative methods for the decomposition and their geometry—details	51
A.1. Convergence guarantees	51
A.2. Geometry	59
Appendix B. The many facets of the Helmholtz decomposition	70
Appendix C	72
C.1. Helmholtz decomposition	72
C.2. Useful computations, identities, and facts	74
C.3. Tools from convex analysis	79
C.4. Exact nonlinear waves	80
C.5. PQG equations	82
References	82

Note to the reader: in section 1 below we provide a ‘shortest path’ to the result, providing just enough background to state and motivate the result. The main results and conclusions are summarised in sections 7 and 8. For a more detailed discussion of the motivation and context we refer the reader to section 2.

1. Introduction

Decomposing a vector field, or a state vector, into several constitutive components is ubiquitous in fluid mechanics and atmospheric science. For example, the Helmholtz decomposition will split a vector field into its potential and divergence-free components, and it is used throughout fluid mechanics [Cho68, Tem69, MB01]. Another example is the vortical-wave decomposition of a state vector which is used in atmospheric and oceanic sciences to understand the fast rotation and strong stratification regime of the Boussinesq equations, e.g. [Blu72, Bar95, Maj03]. Crucially, as discussed in section 2, both of these decompositions can be viewed from multiple perspectives and share salient features such as reconstruction of the original field via elliptic partial differential equation (PDE), a rich geometric interpretation of the decomposition (involving orthogonality and projections), and a connection to slow and fast components of appropriate dynamical systems.

In this paper we consider a moist Boussinesq system from atmospheric dynamics, which brings additional realism and complexity due to clouds. We ask the following question: does this moist, cloudy Boussinesq system also admit such a decomposition?

While these decompositions have applications in slow–fast dynamical systems, they can be investigated without reference to the dynamical evolution, by considering the associated operator. To see this

from a general perspective, consider a dynamical system written abstractly as

$$\partial_t U + \frac{1}{\epsilon} \mathcal{D}_0 U + \mathcal{D}_1 U = 0,$$

where ϵ is a small parameter, $\mathcal{D}_0$ is the leading-order operator, and $\mathcal{D}_1$ is the next-order contribution to the operator for the dynamical system (e.g. [KM81, Sch94, EM96, SS17]). To leading order with respect to ϵ , it is the operator $\mathcal{D}_0$ that contains the essential information. Understanding the dynamics can be tantamount to understanding the operator $\mathcal{D}_0$. Hence, in the remainder of the paper, we will seldom refer to the actual dynamical evolution, and instead the main object of interest is the operator $\mathcal{D}_0$.

In past cases, such as the Helmholtz decomposition or Boussinesq equations (e.g. [KM81, Sch94, EM96]), the leading-order operator is linear ($\mathcal{D}_0 = \mathcal{L}$), and the decomposition involves a linear operator and linear eigenmodes. On the other hand, for an atmosphere with clouds, additional nonlinearity arises from phase changes of water, and $\mathcal{D}_0 = \mathcal{N}$ is a nonlinear operator (e.g. [SS17]). Hence a substantial challenge in the present paper is from nonlinearity, from the distinction

$$\mathcal{D}_0 = \mathcal{L} \quad \text{versus} \quad \mathcal{D}_0 = \mathcal{N}. \quad (1.1)$$

For a nonlinear operator, it is not clear if eigenmodes can be identified. Consequently, a decomposition

$$U = U_1 + U_2 + \cdots + U_{N_{\text{eig}}} \quad (1.2)$$

into N_{eig} eigenmodes is likely to be impossible.

Despite the challenges of nonlinearity in (1.1), (1.2), here we forge ahead with a decomposition in the spirit of (1.2) as the main aim and main theme of the paper.

In section 2 we describe the background and main questions in more detail. Sections 3–6 describe main results, and sections 7 and 8 present a summary of main results and conclusions, including links to the main theorems for easy reference. We point the reader who may skip the background in section 2 to note section 2.5 where the simplified notation used in the main body of the paper is recorded.

2. Background and motivation

In this section we first expound on known decompositions (namely the Helmholtz decomposition and the vortical-wave decomposition of the dry Boussinesq system) in order to set the stage for the decomposition of moist Boussinesq states studied in this paper. We also motivate in more detail the context in which the dry and moist Boussinesq equations arise.

We start with the Helmholtz decomposition for two reasons. First, we present it here for folks familiar with fluid dynamics but not necessarily familiar with geophysical and atmospheric fluid dynamics. Second, we present it here because, as we will detail later, the Helmholtz decomposition plays a central role in the other decompositions, in both the dry and the moist cases.

2.1. Helmholtz decomposition

The Helmholtz decomposition of a vector field u splits it into two components:

$$u = \nabla p + \sigma$$

where σ is divergence-free. One of the main reasons for the ubiquity of this decomposition, especially in the realm of fluid mechanics, is that this single decomposition can be viewed through a variety of lenses. Here we recall how it can be viewed geometrically, viewed in relation to elliptic PDE, and viewed in relation to slow-fast dynamical systems. We emphasise these three perspectives because these are precisely the perspectives that coincide with each other in the dry case but diverge in the moist case.

We can view the Helmholtz decomposition geometrically: it decomposes the space of three-dimensional vector fields into two orthogonal pieces, namely potential vector fields and divergence-free vector fields, and thus each component of the decomposition can be understood as a projection onto the corresponding subspace. The Helmholtz decomposition is also related to elliptic PDE since the scalar potential p solves

$$\Delta p = \nabla \cdot u$$

while, if we write the solenoidal component as $\sigma = \nabla \times \psi + c$, the vector potential ψ solves

$$\nabla \times (\nabla \times \psi) = \nabla \times u$$

with c given by $c = fu$. Here we use the notation ff to denote the average over the spatial domain of any scalar or vector function f .

Finally we can view the Helmholtz decomposition as a slow-fast decomposition for an appropriate dynamical system. The relevant system is the low Mach number limit of the compressible Euler equations (for an ideal isentropic gas) [KM81, KM82, Maj84]. In this setting the slow dynamics correspond to incompressible flow whereas the fast dynamics correspond to the acoustic wave equation. Crucially: the solenoidal component of the Helmholtz decomposition is precisely the slow piece and the potential component is precisely the fast piece. We postpone a detailed discussion of this slow-fast dynamics to section C.1: this low Mach number limit is well-known, so we only provide details later once it helps us compare and contrast these classical results (pertaining to Helmholtz and dry Boussinesq) with the results we have obtained in this paper.

2.2. Dry Boussinesq and its decomposition

Now we turn our attention to the dry Boussinesq system and its decomposition. As we describe this decomposition we will highlight the properties that it shares with the Helmholtz decomposition discussed above.

In oceanic and atmospheric application problems, one may need a decomposition of observational or model data. The latter is often motivated by, or closely related to, normal-mode analysis of a PDE model such as the dry Boussinesq equations, e.g. [LM92, BCF14, LCBA20, ZNV+23]. Furthermore, a normal-mode decomposition is often the starting point for numerical computations [Lei80, Bar95, SW02, HMRP16, ELS20, VZ21]. For the Boussinesq equations on an f -plane, the normal-mode analysis is also a wave-vortical decomposition, as further explained below. Although many numerical computations focused on constant rotation and buoyancy frequencies in periodic domains [Bar95, SW02, HMRP16], other analyses also address more realistic PDE systems, such as variable-stratification Boussinesq dynamics with rigid surface and bottom boundaries [ELS20], varying rotation rate with associated Rossby waves, and the shallow water equations on a sphere [Kas78, VZ21]. Hence, while we consider the Boussinesq approximation here, it is often straightforward to extend results from the simplified setting of the Boussinesq approximation to the more realistic setting of the anelastic approximation (see, e.g. [WSSM19, WSS+20] and sections 3 and 9 of [SS17]).

Let us stop and discuss this terminology for a second: why is the decomposition of dry Boussinesq states called a vortical-wave decomposition? It is called this because the decomposition has two components. The first component is characterised by the potential vorticity. This is why it is referred to as the ‘vortical mode’. (In the dry case this corresponds to an honest-to-goodness eigenmode. This is no longer true in the nonlinear moist case, which is why we use the term ‘component’ instead of ‘mode’ in that case.) The second component corresponds to the fast piece of the decomposition. It can then be shown that the dynamics obeyed by this fast piece are that of linear oscillations. This is why this component is known as the ‘wave mode’. Together the vortical mode and the wave mode thus constitute the vortical-wave decomposition.

The dry Boussinesq system and its fast rotation and strong stratification limit is one of the success stories of atmospheric science. In the associated parameter regime the dry Boussinesq system may be written in terms of conservation laws:

$$\begin{cases} \partial_t u + u \cdot \nabla u + \frac{1}{\varepsilon} (e_3 \times u + \nabla p - \theta e_3) = 0 \text{ and} \\ \partial_t \theta + u \cdot \nabla \theta + \frac{1}{\varepsilon} u_3 = 0, \end{cases} \quad \begin{matrix} (2.1a) \\ (2.1b) \end{matrix}$$

subject the incompressibility assumption $\nabla \cdot u = 0$. Note that for simplicity physical parameters of order one have to be set to unity. Here the unknowns are the velocity vector field u and the potential temperature θ (a scalar field). The first equation corresponds to conservation of linear momentum while the second equation corresponds to the conservation of potential temperature. Of course, as written above (2.1b) makes it look like θ is *not* conserved. This is because there is a large background potential temperature gradient (this is a manifestation of the strong stratification assumption) and so the quantity θ used above is actually the *anomaly* of the potential temperature about that background gradient.

Here the small parameter indicates that we are particularly interested in the behaviour of this system as the stratification is strong and the rotation is fast. In the limit we formally obtain from (2.1b) that the vertical component of the velocity vanishes and we obtain from the (2.1a) the geostrophic balance $u_h = \nabla_h^\perp p$ and the hydrostatic balance $\theta = \partial_3 p$. Here $v_h = (v_1, v_2, 0)$ denotes the horizontal part of a vector v and $v_h^\perp$ denotes its $\frac{\pi}{2}$ -rotation $v_h^\perp = (-v_2, v_1, 0)$. Note that the velocity is horizontal since its vertical component vanishes, yet it still depends on all three spatial coordinates (in this sense the flow is *not* two-dimensional). Also note that the Coriolis parameter is assumed to be a constant here, for simplicity, as in the f -plane approximation, and consequently Rossby waves are not considered here. Furthermore,

since this slow–fast setting is a setting of small Froude and small Rossby numbers, we are considering, for simplicity, a case where any background flow would be weak compared to the inertia-gravity wave speed. The discussion is formal here, but rigorous theorems are available (see [EM96] and [ME98] for cases with either balanced or general initial data).

In particular, the slow dynamics are characterised by the (quasi-geostrophic (QG)) potential vorticity $PV = \nabla_h^\perp \cdot u_h + \partial_3 \theta$. This follows from two observations which hold in the limit: (1) we may compute from (2.1a)–(2.1b) that PV is purely advected by the flow and (2) the geostrophic and hydrostatic balances tell us that p solves $\Delta p = PV$. This is the crux of why the limiting Boussinesq dynamics on the slow timescale, known as the (dry) QG equations, are a success story: everything comes down to a single conserved quantity, the potential vorticity, from which all other information about the system may be recovered (by way of inverting a Laplacian). The potential vorticity not only helps us understand the limit, it also allows us to decompose states of the original dry Boussinesq system into slow and fast components.

We now discuss this decomposition, also known as the vortical-wave decomposition. This decomposition splits the state-vector (u, θ) into two pieces:

$$\begin{pmatrix} u \\ \theta \end{pmatrix} = \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p \end{pmatrix} + \begin{pmatrix} \sigma_h^\perp + w e_3 \\ \sigma_3 \end{pmatrix}$$

where σ is divergence-free. The first component is known as the *vortical mode* and the second component is known as the *wave mode*. As in the case of the Helmholtz decomposition above, this decomposition has several equivalent interpretations. We can view this decomposition geometrically as projecting into two L^2 -orthogonal subspaces: the first piece is the set of balanced states satisfying both geostrophic and hydrostatic balance and the second piece is the set of states with vanishing potential vorticity. This is similar to the Helmholtz decomposition where the first piece, the potential component, satisfies the ‘balance’ $u = \nabla p$ while the second piece belongs to the set of states with vanishing divergence. Note that L^2 is used here, not by default, but because it is precisely the energy conserved by the dry Boussinesq system (2.1a)–(2.1b). We can view this decomposition in relation to elliptic PDE since p solves

$$\Delta p = PV.$$

Meanwhile, since σ is divergence-free, we can write it in terms of a vector potential as $\sigma = \nabla \times \psi + c$ for some constant c where ψ solves

$$\nabla \times (\nabla \times \psi) = \begin{pmatrix} j \\ \partial_3 w \end{pmatrix}$$

for $j = \partial_3 u_h - \nabla_h^\perp \theta$ denoting the thermal wind imbalance³, $c = -f u_h^\perp + f \theta e_3$, and finally $w = u_3$. Finally we can view this decomposition as a slow-fast decomposition for (2.1a) and (2.1b). Here we describe only cursorily how this may be done, postponing a more detailed discussion to section 4. Indeed, the vortical mode corresponds precisely to solutions of the limiting QG system while the wave mode corresponds to the component of the solution with time derivative proportional to $1/\varepsilon$ (and hence fast). We discuss in more detail why this decomposition works as a slow-fast decomposition in section 4, where we recall how the fast-wave averaging framework allows us to make this idea precise in order to then use this framework to discuss the slow-fast decomposition in the moist case.

2.3. Moist Boussinesq

We have seen in sections 2.1 and 2.2 above that both the Helmholtz decomposition and the vortical-wave decomposition possess several interpretations, be they geometric, related to elliptic PDE, or about the slow-fast nature of appropriate dynamical systems. In this section we introduce the moist Boussinesq system, keeping in mind that our goal is to find a decomposition for that system which possesses properties akin to those discussed above.

First, we motivate the moist Boussinesq system itself. Its *dry* counterpart is potent for analytical, computational, and predictive purposes because it focuses on two physical features which dominate mid-latitude atmospheric dynamics: stratification and rotation (i.e. the Coriolis effect). The *moist* Boussinesq

³ A thermal wind refers to a velocity field which satisfies both geostrophic and hydrostatic balance. This terminology is warranted by the fact that such velocity fields are fully determined by the pressure through geostrophic balance, which in turn is defined by the ‘thermal’ variable, namely the potential temperature, through hydrostatic balance. Here the thermal wind *imbalance* vanishes when both the geostrophic and hydrostatic balances hold. This imbalance can therefore be thought of as measuring how far a given state (u, θ) is from satisfying both of these balances.

system introduces another important physical feature: moisture. Even without incorporating all of the complexities of the water cycle, moisture plays an important role in atmospheric dynamics because of the role it plays in the energy budget: as moist air cools, condensing water releases latent heat energy (and vice-versa as air warms and liquid water evaporates). When incorporating moisture into the dry system from (2.1a) and (2.1b), we obtain

$$\begin{cases} \partial_t u + u \cdot \nabla u + \frac{1}{\epsilon} (e_3 \times u + \nabla p - \theta e_3) = 0 & (2.2a) \\ \partial_t \theta + u \cdot \nabla \theta + \frac{\Gamma_\theta}{\epsilon} u_3 = C, & (2.2b) \\ \partial_t q_v + u \cdot \nabla q_v - \frac{\Gamma_q}{\epsilon} u_3 = -C \text{ and} & (2.2c) \\ \partial_t q_l + u \cdot \nabla q_l = C. & (2.2d) \end{cases}$$

In comparison to the dry Boussinesq system, the moist Boussinesq system includes additional evolution equations for the (anomaly of) water vapour mixing ratio, q_v , in (2.2c) and for the liquid water mixing ratio, q_l , in (2.2d). It also includes a source/sink term C for condensation and evaporation, and associated heating and cooling. Condensation occurs for $C > 0$ and represents a phase transition from vapour to liquid, and evaporation occurs for $C < 0$ and represents a phase transition from liquid to vapour. The constants Γ_θ and Γ_q correspond to parameters encoding the strength of the background vertical gradients in potential temperature and water vapour, respectively. The constant Γ_θ had been set to one in the dry case in (2.1b) for simplicity of exposition (and without loss of generality in the stably stratified regime), and it is re-introduced more generally as Γ_θ in the moist case in (2.2b). Note that we will frequently refer to the anomalous water vapour q_v as simply ‘the water vapour’ with ‘anomaly’ omitted for brevity, and we also note that the precise definition of the water content q_v or q_l as a mixing ratio is measured as kilograms of water (vapour or liquid) per kilogram of dry air [KM06, HDMSS13]. Reversible non-precipitating moist thermodynamics will be considered here. While additional cloud microphysics processes will not be considered here, the moist Boussinesq system above is valuable for more complex scenarios as well, since it provides the starting point for extensions that include rainfall, ice, other precipitation, and other complexities [GS96, KM06, SS17, WSSM19, WSS+20].

Second, we rewrite the moist Boussinesq system in a way that facilitates the main goal here of a slow-fast state decomposition. In particular, note that the source terms C in (2.2b)–(2.2d) appear to potentially cause complications. However, the system can be rewritten so that the source terms do not explicitly appear, through a convenient change of variables. This reformulation has been used in several past studies [Bre87, PS10, HDMSS13, SS17, MSS19, KSS23] and has facilitated theoretical advances such as an energy decomposition [SS17, MSS19] and conservation of potential vorticity [KSS23], even in the presence of clouds and phase changes.

For the convenient way of rewriting (2.2a)–(2.2d), we transform to a different set of thermodynamic variables that is conserved. In particular, define the equivalent potential temperature, θ_e , and total water mixing ratio, q_t , as

$$\begin{cases} \theta_e = \theta + q_v, & (2.3a) \\ q_t = q_v + q_l. & (2.3b) \end{cases}$$

Note that the definition of θ_e used here is the same as in prior work (e.g. [HDMSS13, SS17]) and is a linearised version of more comprehensive definitions from moist thermodynamics (see, e.g. [Ema94] or equation (5) of [Ste05]); such a simplification is useful since it includes the essential physics while also allowing for changes of variables with transparent analytical formulas. The evolution equations of θ_e and q_t can be found by taking appropriate linear combinations of (2.2b)–(2.2d), and they take the form

$$\begin{cases} \partial_t \theta_e + u \cdot \nabla \theta_e + \frac{1}{\epsilon} u_3 = 0, & (2.4a) \\ \partial_t q_t + u \cdot \nabla q_t - \frac{\Gamma_q}{\epsilon} u_3 = 0, & (2.4b) \end{cases}$$

where we have now set the background gradient parameters from (2.2b) and (2.2c) to be $\Gamma_\theta = 2$ and $\Gamma_q = 1$ without loss of generality within the stably stratified regime (with positive values of the dry stability parameter Γ_θ and the moist stability parameter $\Gamma_\theta - \Gamma_q$ [SS17]), and for simplicity. The important feature of (2.4a) and (2.4b) is that the source term of condensation and evaporation, C , has been eliminated. As a result, (2.4a) and (2.4b) show that $\theta_e + (z/\epsilon)$ and $q_t - (z/\epsilon)$ (which combine the anomalies θ_e and q_t with their background profiles z/ϵ and $-z/\epsilon$, respectively) are conserved along fluid parcel trajectories. Physically, the equivalent potential temperature is conserved because losses of water vapour q_v are compensated by gains in heat θ , as indicated by the definition in (2.3a); and the total water mixing ratio, $q_t = q_v + q_l$, is conserved because losses of water vapour q_v are compensated by gains in liquid water q_l .

To complete the rewriting, the buoyancy term θ/ϵ from (2.2a) must also be rewritten in terms of θ_e and q_t . To do so, the following transformation is used:

$$\begin{cases} \theta = \theta_e - \min(q_t, q_{vs}), & (2.5a) \\ q_v = \min(q_t, q_{vs}), & (2.5b) \\ q_l = \max(0, q_t - q_{vs}), & (2.5c) \end{cases}$$

which is the reverse of the transformation in (2.3a)–(2.3b). The saturation water vapour value q_{vs} appears here for distinguishing between the unsaturated phase ($q_v < q_{vs}$), when no cloud is present, and the saturated phase ($q_v = q_{vs}$), when a cloud is present. Note that q_{vs} is not appearing here for the first time. It is already implicit in the condensation/evaporation source term C in (2.2b)–(2.2d). Also note that q_{vs} is the anomaly of the saturation water vapour, and the value of q_{vs} will be constant here for simplicity. For explicit expressions for θ in each phase, one can rewrite (2.5a) as

$$\theta = \begin{cases} \theta_e - q_t & \text{if } q_t < q_{vs}, \\ \theta_e - q_{vs} & \text{if } q_t \geq q_{vs}, \end{cases} \quad (2.6)$$

which can be used for another expression of the buoyancy in terms of the variables θ_e and q_t .

The end result of the rewriting is the moist Boussinesq system in the following form:

$$\begin{cases} \partial_t u + u \cdot \nabla u + \frac{1}{\epsilon} (e_3 \times u + \nabla p - (\theta_e - \min_0 q_t) e_3) = 0, & (2.7a) \\ \partial_t \theta_e + u \cdot \nabla \theta_e = -\frac{1}{\epsilon} u_3 \text{ and} & (2.7b) \\ \partial_t q_t + u \cdot \nabla q_t = \frac{1}{\epsilon} u_3, & (2.7c) \end{cases}$$

subject to $\nabla \cdot u = 0$, and following from (2.2a), (2.4a), (2.4b), and (2.5a). Some additional specifications are as follows. The notation $\min_0 q_t = \min(0, q_t)$ has been introduced in (2.7a), and the saturation water vapour parameter q_{vs} (the anomaly with respect to the background state) has been set to zero without loss of generality, and for simplicity. Throughout this paper we work on the three-dimensional torus (i.e. with triply-periodic boundary conditions). Also recall that we set the background gradient parameters from (2.2b) and (2.2c) to be $\Gamma_\theta = 2$ and $\Gamma_q = 1$ without loss of generality in the stably stratified regime, and for simplicity. One can now see the value in rewriting the moist Boussinesq system: while the (θ, q_v, q_l) formulation in (2.2a)–(2.2d) is helpful for seeing the additional cloud physics that is added to the dry system, the (θ_e, q_t) formulation in (2.7a)–(2.7c) is helpful for writing the moist system in a way that is mathematically similar to the dry system. The main differences between the moist and dry systems are that the moist system has two thermodynamic variables (θ_e and q_t) instead of one (θ), and the buoyancy term in the moist case involves the nonlinear expression $\theta_e - \min_0 q_t$.

What makes the moist Boussinesq system particularly interesting is that moisture introduces *phase boundaries*. In concrete terms: here we have assumed without loss of generality that water saturation occurs at $q_t = 0$ and so different physics are in play depending on the sign of q_t , i.e. depending on whether water is above or below saturation (this is discussed in more detail in section 2). We therefore refer to the set $\{q_t = 0\}$ as the *phase boundary* since it delineates the separation between water in vapour form (outside of a cloud) and water in vapour-plus-liquid form (inside of a cloud).

Finally, to end this subsection, we review some properties of the moist Boussinesq system that are potentially relevant to the aim here—i.e. to the search for a state decomposition.

As $\epsilon \rightarrow 0$, solutions of the moist Boussinesq system (2.7a)–(2.7c) are forced to satisfy two *balances*, namely geostrophic balance $u = u_h = \nabla_h^\perp p$ and hydrostatic balance $\theta_e - \min_0 q_t = \partial_3 p$, which come from the fast rotation and strong stratification, respectively. This warrants the following definition.

Definition 2.1 (state space and balanced set). We denote by $\mathbb{L}_\sigma^2 := (L^2)_\sigma^3 \times L^2 \times L^2$ the state space of the moist Boussinesq system (2.7a)–(2.7c). Here $u \in (L^2)_\sigma^3$ means that $u \in L^2(\mathbb{T}^3; \mathbb{R}^3)$ is divergence-free. We then define the *balanced set*

$$\mathcal{B} := \{(u, \theta_e, q_t) \in \mathbb{L}_\sigma^2 : u_h = \nabla_h^\perp p, u_3 = 0, \text{ and } \theta_e - \min_0 q_t = \partial_3 p \text{ for some } p \in \dot{H}^1\} \quad (2.8)$$

where $\dot{H}^1$ is the set of scalar fields in H^1 with vanishing averages.

We note that a state belongs to the balanced set precisely when it satisfies both geostrophic and hydrostatic balance.

Proceeding once again by analogy with the dry case, one may now identify *two* slow quantities: the *equivalent* potential vorticity $PV_e = \nabla_h^\perp \cdot u_h + \partial_3 \theta_e$, defined now in terms of the *equivalent* potential temperature, and the moist variable $M = \theta_e + q_t$ [SS17]. Similar to the dry case, now p and M fully determine a balanced state by way of the PV_e -and- M inversion

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV_e, \quad (2.9)$$

which is obtained as before by inserting the balances into the definition of PV_e . This is a nonlinear elliptic PDE whose well-posedness was established recently [RTSS24]. Using this PDE for PV_e -and- M inversion, one can then phrase the limiting moist (and potentially *precipitating*) QG system as a nonlinear transport equation for PV_e and M [SS17].

It is important to note that the nonlinear elliptic PDE in (2.9) only has meaning for the *balanced states* from (2.8), at this point in the story. This is in contrast to the cases of the Helmholtz decomposition and dry Boussinesq system from sections 2.1 and 2.2, respectively, for which an elliptic PDE is part of a state decomposition for *general* states. One aim of the present paper is to investigate whether the nonlinear PV_e -and- M inversion in (2.9) may have additional meaning as part of a state decomposition for *general* states.

The question is now this: can we find a decomposition for the moist Boussinesq system which has the same favourable properties as the Helmholtz and vortical-wave decompositions?

2.4. Literature on atmospheric dynamics with moisture and clouds, and mathematical theory

In this section we review some literature on topics related to the main focus of the paper, i.e. state decomposition for a moist atmosphere with clouds. We mention both atmospheric sciences literature and mathematics literature.

Decompositions have been useful for numerous applications in atmospheric science, and they are commonly based on decomposing the dry portions (u, θ) of the atmospheric state. Applications include analysis of observational or model data [HMR85, ŽKT+15], data assimilation [Lor81, Dal93, DB99, ZGK04, Ban21], and nonlinear wave interactions [GG99, MB03, BM04, RDTM08, FKN09]. Other references are also described at the beginning of section 2.2. One aim here is to extend the ideas of dry decompositions to the realm of a moist atmosphere with clouds.

For a moist atmosphere with clouds, decompositions can potentially provide valuable insight into the behaviour of rainfall and precipitation [KMS12]. Past work has commonly used linearisation, and either dry or moist eigenmodes. By adapting or modifying the dry linear eigenmodes, some studies have analysed observational and model data for understanding moisture-coupled phenomena such as convectively coupled equatorial waves [YHS03, GK12, OS16, MC18, KGK+22], tropical intraseasonal oscillations such as the Madden–Julian Oscillation (MJO) [ŽF15], and the Walker circulation [SO14]. Moist eigenmodes, obtained by linearisation of moist dynamics, have also been used for analysing observational and model data for the MJO [OS15, SM15], and for asymptotic analysis of weakly nonlinear wave interactions [CMS15, CMS16]. Balanced and unbalanced moisture decompositions have been investigated with a type of linearisation that applies with clouds present [WSSM19, WSS+20]. Applications to data assimilation have also investigated linear eigenmodes that decompose moisture into its balanced and unbalanced components [OWSS25].

Here we aim for a decomposition that does not rely on linearisation, so that it may potentially include a more faithful representation of the nonlinearity and the physics associated with clouds and latent heating.

The PDE community has recently provided rigorous treatment of moist atmospheric dynamics with phase changes. A model with water vapour and liquid water was analysed in [CZT12, CZFTT13, BCZT14], where the velocity field is prescribed. Then followed results for two-phase flow with dynamical velocity evolving via the primitive equations [CZHK+15, LM20, TW15, TW16]. An additional category of water was included in [CHT+18, TL22, HKLT17, HKLT20], such that liquid water is divided into cloud water (that does not fall) and precipitating water. The work on warm-rain microphysics flow may be separated by the microphysics model determining how water is partitioned between vapour, cloud water and rain water, as well as by whether the velocity is prescribed or governed by the primitive equations. The case of prescribed velocity is considered in [CHT+18] for Grabowski microphysics [Gra98], and in [HKLT17] for Klein–Majda microphysics [KM06]. For primitive-equation velocity, theorems for the Grabowski and Klein–Majda models are found, respectively, in [TL22] and [HKLT20]. Incorporation of ice into water microphysics is discussed in [CJTT21].

We note that the aforementioned rigorous studies of atmospheric flows with phase changes treat the velocity as either prescribed, or as governed by the primitive equations. Following [CT07], the

barotropic-baroclinic decomposition of the velocity is the basis for analysis of the primitive equations and its extensions. Recently, the current authors obtained existence, uniqueness and regularity results [RTSS24] for the nonlinear elliptic PDE (2.9) underlying Boussinesq velocity dynamics in the limit of asymptotically fast rotation rate and buoyancy frequencies [SS17, ZSS22]. Here, we show that the same nonlinear elliptic PDE is the foundation for decomposition of the Boussinesq equations more generally, in the regime of finite parameter values.

Pertaining to the dry Boussinesq evolution equations in the limit of asymptotically fast rotation rate and buoyancy frequency, a series of rigorous results established the decoupling between slow, balanced motions and fast, unbalanced waves [EM96, BANZ97, EM98, ME98, Maj03]. For moist atmospheric dynamics with phase changes, first steps to use the fast-wave averaging framework [EM96, EM98, ME98, Maj03] appear in [ZSS21b, ZSS21a]. The latter references present formal asymptotic analysis [ZSS21b] and supporting numerical analyses [ZSS21a]. The inclusion of boundary layers and the diabatic layer have been investigated in formal asymptotic analysis as well [KSPK22, BHK23]. We note that rigorous fast-wave-averaging analysis for multi-phase atmospheric flows remains open. In section 4, we explain the path forward for rigorous fast-wave-averaging analysis using the decomposition defined in section 3. To date, the evidence suggests that nonlinear oscillations lead to coupling between fast and slow flow components, possibly even in the limiting dynamics when distinct phases co-exist (see figure 13 in the appendix and [ZSS21a]).

2.5. Simplified notation

To simplify the notation used throughout the rest of the paper, we will from now on use θ to denote the equivalent potential temperature θ_e (which is not an issue since the potential temperature, denoted θ in the text above, will not make any appearances in the sequel). Similarly we will write PV for the equivalent potential temperature PV_e and q to mean q_t , the total water content.

3. Decomposition

As discussed in sections 1 and 2, we seek a decomposition adapted to the leading-order operator of the moist Boussinesq system. We can see from (2.7a)–(2.7c) that this operator is

$$\mathcal{N}(u, \theta, q) = \begin{pmatrix} \mathbb{P}_L(u_h^\perp - (\theta - \min_0 q) e_3) \\ u_3 \\ -u_3 \end{pmatrix} \quad (3.1)$$

where we note again that θ_e and q_t are written as θ and q for notational simplicity, as specified in section 2.5. Here $\mathbb{P}_L$ denotes the Leray projector onto the space of divergence-free vector fields. On the three-dimensional torus it takes the simple form $\mathbb{P}_L u = -\nabla \times (\Delta^{-1} \nabla \times u)$. Using the Leray projector, as we do here, introduces a trade-off. The cost is that the operator $\mathcal{N}$ is made more complicated, since it now includes the nonlocal Leray projection term. The benefit is that the spaces we work with are simpler: the Leray projection removes the pressure from immediate consideration (since it only appears as a gradient term, which is L^2 -orthogonal to the space of divergence-free vector fields which the Leray projector maps onto), and so $\mathcal{N}$ in (3.1) can now be viewed as a map from $\mathbb{L}_\sigma^2$ into itself.

Given that the operator $\mathcal{N}$ is nonlinear, it is likely impossible to define and characterise a full set of the nonlinear eigenfunctions and eigenvalues. Hence it is likely impossible to define a full eigenfunction decomposition such as in (1.2). An alternative route to a decomposition must be sought.

Here we propose to use the notions of null set and image as concepts that may allow us to extend from the linear to the nonlinear realm. We propose a decomposition that will split any state in the state space $\mathbb{L}_\sigma^2$ into two pieces: one piece residing in the null set of $\mathcal{N}$ and another piece residing in the image of $\mathcal{N}$. In other words we will perform the decomposition

$$\mathbb{L}_\sigma^2 = \{\mathcal{N} = 0\} + \text{im } \mathcal{N}. \quad (3.2)$$

If the operator $\mathcal{N}$ were *linear*, then this decomposition would be a (coarse) eigendecomposition, into the eigenspace of the zero eigenvalue and the eigenspace of all nonzero eigenvalues. Such eigenspace ideas are helpful as motivation for pursuing the null set and image as fundamental concepts for defining a decomposition here. However, the operator $\mathcal{N}$ is nonlinear, so the full meaning of (3.2) is not immediately clear. We aim to demonstrate in this section that (3.2) is able to characterise the desired decomposition in the nonlinear realm.

Note that in this section we will only present and discuss this decomposition itself. We will not discuss the various interpretations of this decomposition. This will be done in later sections. In section 4

we discuss how this decomposition is related to slow-fast dynamics (i.e. to a version of the QG limit with unbalanced or ill-prepared initial data), in section 5 we discuss how this decomposition relates to the conserved energy and projections, and finally in section 6 we discuss iterative methods used to compute this decomposition in practice and their geometric interpretation.

Since we wish to perform the decomposition foreshadowed in (3.2) we must begin by understanding both the null set and the image of $\mathcal{N}$ as well as possible. We first turn our attention to the null set. As shown in the result below, this null set is precisely the set of states consisting of states in both hydrostatic and geostrophic balance. This is why the null set of $\mathcal{N}$ will be referred to, from now on, as the *balanced set*.

Proposition 3.1 (alternate characterisations of the balanced set). *For any state $\mathcal{X} \in \mathbb{L}_\sigma^2$ the following are equivalent.*

1. $\mathcal{X}$ is balanced, i.e. $\mathcal{N}(u, \theta, q) = 0$.
2. $\mathcal{X} = (u, \theta, q)$ satisfies

$$\partial_3 u_h - \nabla_h^\perp (\theta - \min_0 q) = 0, u_3 = 0, \oint u_h = 0, \text{ and } \oint \theta - \min_0 q = 0. \quad (3.3)$$

3. $\mathcal{X}$ may be parameterised as $\mathcal{X} = \Phi(p, M)$ for some unique $p \in \dot{H}^1$ and $M \in L^2$, where

$$\Phi(p, M) = \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p + \frac{1}{2} \min_0 (M - \partial_3 p) \\ M - \partial_3 p - \frac{1}{2} \min_0 (M - \partial_3 p) \end{pmatrix}.$$

Note that this proposition characterises the balanced set in three ways:

- it is the null set of $\mathcal{N}$,
- it is characterised as the null set of the three *dual quantities* j , a , and w given by

$$j(u, \theta, q) = \partial_3 u_h^\perp - \nabla_h^\perp (\theta - \min_0 q), w(u, \theta, q) = u_3 \\ \text{and } a(u, \theta, q) = -\oint u_h^\perp + \oint (\theta - \min_0 q) e_3,$$

and

- it may be parameterised via the *coordinates* p and M .

The dual quantity j is commonly referred to as the *thermal wind imbalance* since it quantifies how far the state is from obeying both geostrophic and hydrostatic balanced. Further down we will see that the image of $\mathcal{N}$ also admits three comparable descriptions. The dual quantity a depends on some spatial averages of the state (u, θ, q) . Why it depends on these particular averages (e.g. it depends on the horizontal velocity average but not the vertical velocity average) and why these averages are combined in this way to constitute a is explained in remark 3.2. Also note that w is equivalent to u_3 (the vertical component of velocity), but it is convenient to introduce the notation of w for use when the horizontal components are velocity are hidden/combined into other quantities like PV , M , p , etc.

Proof of proposition 3.1. It will be particularly convenient to consider the new unknown

$$v := -u_h^\perp + (\theta - \min_0 q) e_3.$$

Indeed, we see that using this unknown we may rephrase items 1, 2, and 3 in a much simpler way. The first item may now be written as

$$v = \nabla p \text{ and } u_3 = 0 \text{ for some } p \in \dot{H}^1$$

while the second item may be written, using lemma C.5, as

$$\nabla \times v = 0, \oint v = 0, \text{ and } u_3 = 0.$$

It then follows directly from the Helmholtz decomposition (see corollary C.2) that items 1 and 2 are equivalent. Finally we note that, by virtue of lemma C.6, item 3 is equivalent to

$$\begin{cases} u = \nabla_h^\perp p, \\ \theta + q = M, \text{ and} \\ \theta - \min_0 q = \partial_3 p, \end{cases}$$

which implies that $v = \nabla p$ and that $u_3 = 0$, from which item 1 follows. $\square$

Remark 3.2 (the ‘good unknown’). The so-called ‘good-unknown’ $-u_h^\perp + (\theta - \min_0 q)e_3$ will appear throughout this paper. This is because this three-dimensional vector field, which combines the horizontal velocity and the buoyancy, turns out to be the one that needs to be decomposed according to a three-dimensional Helmholtz decomposition. This contrasts with the approach commonly used in the literature where instead of a three-dimensional Helmholtz decomposition of this good unknown (or of its dry counterpart, the good unknown $-u_h^\perp + \theta e_3$), the decomposition carried out is a two-dimensional Helmholtz decomposition of only the horizontal velocity. It is precisely because the typical decomposition is two-dimensional (even though the problem is three-dimensional) that vertically sheared horizontal flows often require special treatment. When an honest-to-goodness three-dimensional decomposition is used, these vertically sheared horizontal flows are naturally accounted for (see remark 3.11 where this is discussed in more detail).

Note that the role played by the good unknown in the good decomposition is discussed in more detail in remark 3.9. For now, in order to illustrate why it is such a convenient unknown to work with, we record a couple of computations from the dry case. In the dry case, the good unknown is

$$-u_h^\perp + \theta e_3.$$

We then observe from lemma C.5 that its divergence is precisely the potential vorticity since

$$\nabla \cdot (-u_h^\perp + \theta e_3) = \nabla_h^\perp \cdot u_h + \partial_3 \theta = PV$$

while its curl is precisely the thermal wind imbalance since

$$\nabla \times (-u_h^\perp + \theta e_3) = \partial_3 u_h - \nabla_h^\perp \theta.$$

The physical relevance of thermal wind imbalance is that it vanishes precisely when the (dry) state (u, θ) is in both hydrostatic and geostrophic balance (provided that both u_h and θ have vanishing averages).

To emphasise once more, in a different way: in the dry case, the good unknown $-u_h^\perp + \theta e_3$ is precisely the vector field which is subject to a Helmholtz decomposition in order to extract the balanced and unbalanced components. Indeed, if we write

$$-u_h^\perp + \theta e_3 = \nabla p + \sigma,$$

where the right-hand side is decomposed into the balanced component ∇p and unbalanced component σ , then we see that the balanced component corresponds to

$$u_h^{(bal)} = \nabla_h^\perp p \text{ and } \theta^{(bal)} = \partial_3 p,$$

which are precisely the geostrophic and hydrostatic balances characterizing balanced flows in the dry Boussinesq system.

Connoisseurs of geophysical systems may have paused when noticing that the notation p is used for both the scalar velocity potential in the Helmholtz decompositions *and* the scalar potential appearing in the geostrophic and hydrostatic balances of the dry Boussinesq system. This discussion right here is the reason why: in *both* cases, p denotes the gradient part of an honest-to-goodness Helmholtz decomposition. The wrinkle is that, in the case of the dry Boussinesq system, we are talking about the Helmholtz decomposition of the ‘good unknown’ $-u_h^\perp + \theta e_3$.

While remark 3.9 discusses in more detail the precise role that the *moist* good unknown $-u_h^\perp + (\theta - \min_0 q)e_3$ plays in the decomposition, the computations above, from the dry case, already indicate that essential quantities may readily be computed from the (dry) good unknown.

Now that we understand the null set quite well we turn our attention to the image of $\mathcal{N}$. This set will actually be called the *wave set*. This is because the dynamics induced on that set by the moist Boussinesq system are oscillatory (in the dry case, the dynamics induced on the dry analogue of the wave set by the dry Boussinesq system are precisely the dynamics of linear oscillations, i.e. waves).

Proposition 3.3 (characterisation of the image of $\mathcal{N}$). *For any state $\mathcal{X} \in \mathbb{L}_\sigma^2$ the following are equivalent.*

1. $\mathcal{X}$ belongs to the image of $\mathcal{N}$.
2. $\mathcal{X} = (u, \theta, q)$ satisfies

$$PV = \nabla_h^\perp \cdot u_h + \partial_3 \theta = 0 \text{ and } M = \theta + q = 0. \quad (3.4)$$

3. $\mathcal{X}$ may be parameterised as $\mathcal{X} = \Psi(\sigma, w)$ for some unique $\sigma \in H_\sigma^1$ and $w \in L^2$ satisfying $\nabla_h^\perp \cdot \sigma_h = \partial_3 w$, where

$$\Psi(\sigma, w) = \begin{pmatrix} \sigma_h^\perp + we_3 \\ \sigma_3 \\ -\sigma_3 \end{pmatrix}.$$

In either case we say that $\mathcal{X} \in \mathcal{W}$ for the wave set $\mathcal{W}$. Why this set warrants the name ‘wave set’ will be discussed in section 4 below when we discuss slow-fast decompositions in more detail.

Note that, as was done for the balanced set above, proposition 3.3 characterises the wave set in three ways:

- it is the image of $\mathcal{N}$,
- it is characterised as the null set of the two dual quantities PV and M , and
- it may be parameterised via the coordinates σ and w .

Proof. First we suppose that item 1 holds, i.e. that

$$\begin{cases} u = \mathbb{P}_L(v_h^\perp + (\phi - \min_0 r)e_3), \\ \theta = v_3, \text{ and} \\ q = -v_3 \end{cases}$$

for some $(v, \phi, r) \in \mathbb{L}_\sigma^2$. This means that $u = v_h^\perp + (\phi - \min_0 q)e_3 + \nabla\pi$ for some $\pi \in \dot{H}^1$ and so

$$\nabla_h^\perp \cdot u_h + \partial_3 \theta = \nabla_h^\perp \cdot (v_h^\perp + \nabla_h \pi) + \partial_3 v_3 = \nabla \cdot v = 0$$

while clearly $\theta + q = 0$. Thus item 2 holds.

Now suppose that item 2 holds. We may introduce the new unknown $v = -u_h^\perp + \theta e_3$ such that now item 2 reads

$$\nabla \cdot v = 0 \text{ and } \theta + q = 0$$

(by virtue of lemma C.5). In other words $v = \sigma$ for some divergence-free σ while $\theta = -q$, and so

$$\begin{cases} u = \sigma_h^\perp + u_3 e_3, \\ \theta = \sigma_3, \text{ and} \\ q = -\sigma_3. \end{cases}$$

Therefore, for $w := u_3$, we see that item 3 holds as desired since $\nabla_h^\perp \cdot \sigma_h - \partial_3 w = -\nabla \cdot u = 0$.

Finally we suppose that item 3 holds. Then item 1 follows immediately from choosing $v = \sigma$, $\phi = w$, and $r = 0$ since $\sigma_h^\perp + we_3$ is divergence-free and so the Leray projection of this vector field is equal to itself. $\square$

At this stage we have a good understanding of both the balanced set (which is the null set of $\mathcal{N}$) and of the wave set (which is the image of $\mathcal{N}$). We now seek to leverage this understanding to produce the decomposition formally foreshadowed in (3.2).

In order to do this in a manner that will be useful later, especially when it comes to dynamics, we need to discuss dynamics briefly. Recall that, at leading order, the dynamics of interest are those governed by $\partial_t + \mathcal{N} = 0$. So how do the null set and image of $\mathcal{N}$ relate to this dynamical system?

Well, note that proposition 3.1 tells us that producing *balanced* states is easy: it suffices to specify p and M since then a balanced state may be reconstructed from these two quantities. If the leading operator $\mathcal{N}$ were linear, then this would be enough since balanced solutions of $\partial_t + \mathcal{N}$ would then be guaranteed to be slow, meaning that they would have vanishing time derivatives. However, due to the presence of moisture and phase boundaries, $\mathcal{N}$ is *not* linear. This means that we must *additionally* impose

that balanced solutions be slow. In more prosaic terms: since any choice of the coordinates p and M produces a balanced state, we must make sure that these coordinates are carefully chosen to also be *slow*.

This is where our understanding of the image of $\mathcal{N}$, i.e. of the wave set, comes in handy. Indeed: we have characterised this wave set as the set of states whose PV and M vanish. This is another way of saying that PV and M are precisely a complete set of *time-dependent* slow dual quantities! (Any constant dual quantity would be trivially slow.)

We have found our solution: to guarantee that our balanced component determined by p and M is also slow, it must be reconstructed solely from the slow dual quantities, PV and M . If we think about this geometrically, this is really a matter of *transversality*: we need the wave set $\mathcal{W}$ to be transverse to the balanced set $\mathcal{B}$ so that balanced states in $\mathcal{B}$ may be fully characterised by the dual quantities annihilating the wave set $\mathcal{W}$. This is taken care of in the result below.

We can think about why the balanced set and the wave set need to be transverse another way. Recall that we are after a decomposition of the form (3.2). This decomposition is only unique if its two components are transverse, otherwise the same state could be decomposed in several different ways.

Lemma 3.4 (PV and M characterise balanced states). *Given $PV \in H^{-1}$ and $M \in L^2$, where H^{-1} denotes the dual of $\dot{H}^1$, there exists a unique $\mathcal{X} \in \mathcal{B}$ for which $PV(\mathcal{X}) = PV$ and $\mathcal{M}(\mathcal{X}) = M$.*

Proof. By virtue of proposition 3.1 it suffices to show that there exist a unique $p \in \dot{H}^1$ and a unique $\tilde{M} \in L^2$ for which $(PV \circ \Phi)(p, \tilde{M}) = PV$ and $(\mathcal{M} \circ \Phi)(p, \tilde{M}) = M$. The latter identity tells us that $\tilde{M} = M$, and since

$$(PV \circ \Phi)(p, M) = \Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p)$$

the result follows from the uniqueness of a solution to nonlinear PV -and- M inversion as proved in [RTSS24]. □

Now that we have verified that the wave set and the balanced set are transverse, we are ready to obtain the decomposition that this entire section is building towards.

Theorem 3.5 (pre-decomposition). *We have the decomposition $\mathbb{L}_\sigma^2 = \mathcal{B} + \mathcal{W}$ in the sense that, for any $\mathcal{X} \in \mathbb{L}_\sigma^2$ there exist unique $\mathcal{X}_B \in \mathcal{B}$ and $\mathcal{X}_W \in \mathcal{W}$, called the balanced component of $\mathcal{X}$ and the wave component of $\mathcal{X}$ respectively, such that*

1. $\mathcal{X} = \mathcal{X}_B + \mathcal{X}_W$ and
2. $PV(\mathcal{X}) = PV(\mathcal{X}_B)$ and $\mathcal{M}(\mathcal{X}) = \mathcal{M}(\mathcal{X}_B)$ for $\mathcal{M}(u, \theta, q) = \theta + q$ and $PV(u, \theta, q) = \nabla_h^\perp \cdot u_h + \partial_3 \theta$.

Moreover $\mathcal{X}_B$ and $\mathcal{X}_W$ may be computed explicitly as follows.

- $\mathcal{X}_B = \Phi(p, M)$ for Φ as in proposition 3.1, where $M := \theta + q$ and where $p \in \dot{H}^1$ is the unique solution of

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV. \quad (3.5)$$

- $\mathcal{X}_W = \mathcal{X} - \mathcal{X}_B$, and, in light of proposition 3.3, we also have $\mathcal{X}_W = \Psi(\sigma, w)$ where $w = u_3$ and

$$\begin{aligned} \sigma &= -(u_h^\perp - u_{B,h}^\perp) + (\theta - \theta_B) e_3 \\ &= -(u_h^\perp - \nabla_h p) + \left(\theta - \partial_3 p - \frac{1}{2} \min_0 (M - \partial_3 p) \right) e_3. \end{aligned}$$

In particular σ satisfies

$$\nabla \times (A_B^{-1} \sigma) = \partial_3 u_h - \nabla_h^\perp (\theta - H_B q) + (\partial_3 w) e_3,$$

where $H_B := \mathbb{1}(q_B < 0) = \mathbb{1}(M < \partial_3 p)$ and $A_B := I - \frac{1}{2} H_B e_3 \otimes e_3$, such that $A_B^{-1} = I + H_B e_3 \otimes e_3$, subject to $\nabla \cdot \sigma = 0$ and $f A_B^{-1} \sigma = -f u_h^\perp + f(\theta - H_B q) e_3$.

Remark 3.6. The presence of A_B above may appear mysterious at first. Its presence is motivated as follows, by contrasting the reconstruction method here with how a vector field is reconstructed via the classical Helmholtz decomposition.

In a classical Helmholtz decomposition, both the conservative and solenoidal components of a vector field may be recovered by solving appropriate elliptic equations. Moreover, either component may be recovered *independently* of the other.

The situation is more complicated here: the balanced and wave components may *not* be recovered independently from one another. Instead, we must recover the balanced component first, and then recover the wave component. Nonetheless, there are still, in this more complicated case, elliptic equations which allow us to recover the balanced and wave components. The difference is that the *coefficients* of the elliptic equation used to recover the wave component are themselves dependent on the balanced component. This dependence comes down precisely to the coefficient matrix A_B .

Proof. The uniqueness of this decomposition follows from lemma 3.4 and the fact that, as per proposition 3.3, elements of the wave set $\mathcal{W}$ have vanishing PV and M . The existence of the decomposition follows from the existence of solutions to nonlinear PV -and- M inversion since finding $\mathcal{X}_B$ is, by virtue of proposition 3.1, equivalent to finding p and M and since p may be chosen to be the solution of nonlinear PV -and- M inversion.

The latter characterisation of σ follows from lemmas C.5 and C.7:

$$\nabla \times (A_B^{-1} \sigma) = \nabla \times [-u_h^\perp + (\theta - H_B q) e_3] + \nabla \times \nabla p = \partial_3 u_h - \nabla_h^\perp (\theta - H_B q) + (\partial_3 w) e_3$$

and

$$\oint A_B^{-1} \sigma = \oint u_h^\perp + \oint (\theta - H_B q) e_3.$$

□

It is interesting to see $\nabla \times (A_B^{-1} \sigma)$ arise in the pre-decomposition above, where matrix A_B^{-1} depends on the balanced moisture, q_B . It would be difficult to guess the form of $\nabla \times (A_B^{-1} \sigma)$, with its coefficient matrix that depends on balanced moisture, q_B . Here, the appearance of A_B^{-1} follows from the decomposition into nullspace and image, and from seeking a curl-like operator for determining the unbalanced coordinate σ . Also note that $\nabla \times (A_B^{-1} \sigma)$ involves a linear differential operator, once q_B is known, but it is a nonlinear transformation from (u, θ, q) to σ .

As an immediate corollary of this decomposition we make the following observation: since both the balanced set and the wave set could be parameterised and since any state can be decomposed into a balanced and a wave piece, this means that we have obtained a global parametrisation of the entire state space.

Corollary 3.7. Any state $\mathcal{X} \in \mathbb{L}_\sigma^2$ can be uniquely parameterised by

$$\begin{cases} u = \nabla_h^\perp p + \sigma_h^\perp + w e_3, \\ \theta = \partial_3 p + \frac{1}{2} \min_0 (M - \partial_3 p) + \sigma_3, \text{ and} \\ q = M - \partial_3 p - \frac{1}{2} \min_0 (M - \partial_3 p) - \sigma_3 \end{cases}$$

for $p \in \dot{H}^1$, $M, w \in L^2$, and $\sigma \in (L^2)_\sigma^3$ satisfying $\nabla_h^\perp \cdot \sigma_h = \partial_3 w$.

Proof. theorem 3.5 tells us that any $\mathcal{X} \in \mathbb{L}_\sigma^2$ may be uniquely decomposed as

$$\mathcal{X} = \mathcal{X}_B + \mathcal{X}_W \in \mathcal{B} + \mathcal{W}.$$

Then propositions 3.1 and 3.3 tell us that $\mathcal{X}_B = \Phi(p, M)$ and $\mathcal{X}_W = \Psi(\sigma, w)$ for some uniquely determined p, M, σ , and w . So indeed

$$\mathcal{X} = \Phi(p, M) + \Psi(\sigma, w)$$

as desired. □

We do not even have to stop there: both the balanced set and the wave set had two *dual* descriptions. They were described using a parametrisation, which produced corollary 3.7 above. But they were also described using the dual quantities that annihilate them. This suggests that using the full set of dual quantities, namely both those characterising the balanced set and those characterising the wave set, would yield a complete description of the state space. We verify this in the result below.

Proposition 3.8 (global nonlinear change of coordinates). Consider the dual quantity space

$$\mathfrak{M} := H^{-1} \times L^2 \times (H^{-1})^2 \times L^2 \times \mathbb{R}^3.$$

The map $\mathcal{M} : \mathbb{L}_\sigma^2 \rightarrow \mathfrak{M}$ given by

$$\mathcal{M}(u, \theta, q) = \begin{pmatrix} \nabla_h^\perp \cdot u_h + \partial_3 \theta \\ \theta + q \\ \partial_3 u_h - \nabla_h^\perp (\theta - \min_0 q) \\ u_3 \\ -f u_h^\perp + f(\theta - \min_0 q) e_3 \end{pmatrix} = \begin{pmatrix} PV \\ M \\ j \\ w \\ a \end{pmatrix}$$

is invertible.

Proof. To verify the invertibility of this change of coordinates we will use the coordinates π, M, ζ , and w defined implicitly by

$$\begin{cases} u = \nabla_h^\perp \pi + \zeta_h^\perp + w e_3, \\ \theta = \partial_3 \pi + \zeta_3 + \frac{1}{2} \min_0 (M - \partial_3 \pi - \zeta_3), \text{ and} \\ q = M - \partial_3 \pi - \zeta_3 - \frac{1}{2} \min_0 (M - \partial_3 \pi - \zeta_3) \end{cases} \quad (3.6)$$

where $\nabla_h^\perp \cdot \zeta_h = \partial_3 w$ and $\nabla \cdot \zeta = 0$. Note that the coordinates (π, M, ζ, w) differ from the coordinates (p, M, σ, w) used in corollary 3.7. The key advantage of the coordinates used here is that

$$\text{sign } q = \text{sign } (M - \partial_3 \pi - \zeta_3), \quad (3.7)$$

which is why they are particularly useful here when it comes to inverting $\mathcal{M}$.

However these coordinates have the less advantageous property that π depends on PV, M , and ζ , as we will see below. This is different from p (from corollary 3.7), which only depends on PV and M . That is the reason why these π and ζ coordinates are *not* used for the decomposition (instead of p and σ , respectively): the resulting balanced component computed from π and M would *not* be slow.

We close this parenthesis regarding the choice of coordinates and go back to the matter of invertibility. In light of (3.7), lemma C.6 and lemma C.7 tell us that (3.6) is equivalent to

$$\begin{cases} -u_h^\perp + (\theta - \min_0 q) e_3 = \nabla \pi + \zeta, \\ u_3 = w, \text{ and} \\ \theta + q = M. \end{cases} \quad (3.8)$$

In particular it follows from the uniqueness of the Helmholtz decomposition that these coordinates are uniquely determined by u, θ , and q .

So now it only remains to show that, for any tuple (PV, M, j, w, a) , corresponding coordinates (π, M, ζ, w) exist. Plugging (3.6) into the definition of the dual quantity map $\mathcal{M}$ we see, using (3.8), that

$$\begin{cases} PV = \Delta \pi + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 \pi - \zeta_3), \\ M = M, \\ j = \partial_3 \zeta_h^\perp - \nabla_h^\perp \zeta_3, \\ w = w, \text{ and} \\ a = f \zeta. \end{cases}$$

Since $\nabla_h^\perp \cdot \zeta_h = \partial_3 w$ we see using lemma C.4 that

$$\nabla \times \zeta = j + (\partial_3 w) e_3.$$

We are thus guaranteed a (unique) solution ζ of

$$\begin{cases} \nabla \times \zeta = j + (\partial_3 w) e_3, \\ \nabla \cdot \zeta = 0, \text{ and} \\ f \zeta = a. \end{cases} \quad (3.9)$$

The solvability of nonlinear PV -and- M inversion [RTSS24] then guarantees a unique solution π in $\dot{H}^1$ of

$$\Delta\pi + \frac{1}{2}\partial_3 \min_0 [(M - \zeta_3) - \partial_3 \pi] = PV. \quad (3.10)$$

In other words, if we define

$$\tilde{\mathcal{E}} := \left\{ (\pi, M, \zeta, w) \in \dot{H}^1 \times L^2 \times (L^2)_\sigma^2 \times L^2 : \nabla_h^\perp \cdot \zeta_h = \partial_3 w \right\}$$

to be the space where the coordinates (π, M, ζ, w) reside and denote by $(u, \theta, q) = \mathcal{C}(\pi, M, \zeta, w)$ the map defined by (3.6), we have shown that for any $(PV, M, j, w, a) \in \mathfrak{M}$ there exists a unique $(\pi, M, \zeta, w) \in \tilde{\mathcal{E}}$ such that $(\mathcal{M} \circ \mathcal{C})(\pi, M, \zeta, w) = (PV, M, j, w, a)$. This means that we can take $(u, \theta, q) = \mathcal{C}(\pi, M, \zeta, w)$ to satisfy $\mathcal{M}(u, \theta, q) = (PV, M, j, w, a)$. Moreover, since we know that this state (u, θ, q) is uniquely determined by its coordinates (π, M, ζ, w) this proves that there is a *unique* state (u, θ, q) satisfying $\mathcal{M}(u, \theta, q) = (PV, M, j, w, a)$. This proves that $\mathcal{M}$ is invertible, as desired. Moreover the inverse may be computed by solving, in order, (3.9) then (3.10) and then plugging the solutions into (3.6). $\square$

Remark 3.9 (the global change of coordinates in the dry case). The proof of proposition 3.8 above is a bit of a mouthful. In order to shed some light on the underlying ideas, without having to worry about the details of the proof itself, we discuss here what an analogue global change of coordinates would look like in the dry case.

In that case we could break down the change of coordinates as a composition of invertible pieces, thus illustrating more clearly what the moving pieces are. Indeed, in the dry case we have the following chain of inversions:

$$\begin{aligned} \begin{pmatrix} u \\ \theta \end{pmatrix} &\leftrightarrow \begin{pmatrix} u_h \\ u_3 \\ \theta \end{pmatrix} \leftrightarrow \begin{pmatrix} -u_h^\perp + \theta e_3 \\ u_3 \end{pmatrix} =: \begin{pmatrix} g \\ w \end{pmatrix} \xleftrightarrow{(1)} \begin{pmatrix} \nabla \cdot g \\ \nabla \times g \\ fg \\ w \end{pmatrix} \xleftrightarrow{(2)} \begin{pmatrix} -\nabla_h \cdot u_h^\perp + \partial_3 \theta \\ \partial_3 u_h - \nabla_h^\perp \theta + (\partial_3 w) e_3 \\ -f u_h^\perp + f \theta e_3 \\ w \end{pmatrix} \\ &=: \begin{pmatrix} PV \\ j + (\partial_3 w) e_3 \\ a \\ w \end{pmatrix} \leftrightarrow \begin{pmatrix} PV \\ j \\ a \\ w \end{pmatrix} \end{aligned}$$

where (1) follows from the Helmholtz decomposition (recorded in theorem C.1) and (2) follows from lemma C.5 and the fact that the velocity field u is divergence-free.

It is therefore tempting to assume that a similar chain of invertible maps produces the global nonlinear change of coordinates for the *moist* case recorded in proposition 3.8. This is not the case. Of course, we would need to perform the following inversions initially in the moist case, by virtue of lemma C.6:

$$\begin{pmatrix} u \\ \theta \\ q \end{pmatrix} \leftrightarrow \begin{pmatrix} u_h \\ u_3 \\ \theta + q \\ \theta - \min_0 q \end{pmatrix} =: \begin{pmatrix} u_h \\ w \\ M \\ b \end{pmatrix}.$$

This is not the source of the issue. The source of the issue is that, in the moist case, the potential vorticity PV and the thermal wind imbalance j come from *different* good unknowns. Indeed:

$$PV = \nabla \cdot (-u_h^\perp + \theta e_3) \text{ while } j = [\nabla \times (-u_h^\perp + (\theta - \min_0 q) e_3)]_h.$$

The good unknowns

$$-u_h^\perp + \theta e_3 \text{ and } -u_h^\perp + (\theta - \min_0 q) e_3$$

are not the same. Note that this is *not* a technical issue. This comes from the dynamics in the moist case being fundamentally different. In particular, the operator in the moist case loses the skew-symmetry present in the dry case, and this is precisely the source of the discrepancy between the good unknowns used for PV and j . Indeed, the (θ, q) -dependence in the equation governing $\partial_t u$ is of the form

$$\partial_t u = \cdots + (\theta - \min_0 q) e_3 = \cdots + (\theta - qH) e_3$$

for $H := \mathbb{1}(q < 0)$ while the u -dependence in the equation governing $\partial_t \theta$ and $\partial_t q$ is of the form

$$\begin{aligned}\partial_t \theta &= \cdots - u_3 \text{ and} \\ \partial_t q &= \cdots + u_3.\end{aligned}$$

In matrix form this reads

$$\partial_t \begin{pmatrix} u \\ \theta \\ q \end{pmatrix} = \begin{pmatrix} \cdots & e_3 & -He_3 \\ -e_3^T & \cdots & \cdots \\ e_3^T & \cdots & \cdots \end{pmatrix} \begin{pmatrix} u \\ \theta \\ q \end{pmatrix}.$$

This means that the good unknown for PV comes from the first column, while the good unknown for j comes from the first row. Said another way: PV comes from the *image* of the operator, whereas j comes from the *null set* of the operator. Since that matrix is *not* skew-symmetric, i.e. since the operator is *not* skew-adjoint, it follows that PV and j come from *different* good unknowns.

Remark 3.10. To conclude this section, we note that we now have *three* descriptions of the state space available to us:

- the ‘primitive’ description in terms of the unknowns u , θ , and q ,
- the description in terms of the coordinates p , M , σ , and w (see corollary 3.7), and
- the description in terms of the dual quantities PV , M , j , a , and w (see proposition 3.8).

Crucially: the latter two descriptions are compatible with the decomposition! We can summarise this with the following diagram, where we give a name to the maps converting between each of these descriptions.

$$\begin{array}{ccc} & \nearrow^S & (u, \theta, q) \in \mathbb{L}_\sigma^2 \\ & & \searrow^{\mathcal{M}} \\ (p, M, \sigma, w) \in \mathfrak{C} & \xleftarrow{\mathcal{C}} & (PV, M, j, a, w) \in \mathfrak{M} \end{array}$$

Recall that $\mathbb{L}_\sigma^2$ is the *state space* introduced in definition 2.1 and that $\mathfrak{M}$ is the *dual quantity space* defined in proposition 3.8. Here $\mathfrak{C}$ is the *coordinate space* defined as follows:

$$\mathfrak{C} = \left\{ (p, M, \sigma, w) \in \dot{H}^1 \times L^2 \times (L^2)_\sigma^2 \times L^2 : \nabla_h^\perp \cdot \sigma_h = \partial_3 w \right\}.$$

The state-to-dual quantity map $\mathcal{M}$ is then as defined in proposition 3.8, where it is shown to be invertible, the coordinate-to-state map $\mathcal{S}$ corresponds to the parametrisation recorded in corollary 3.7, which is invertible since theorem 3.5 tells us how to compute its inverse, and the dual quantity-to-coordinate map $\mathcal{C}$ is well-defined by virtue of the invertibility of the other two maps.

Remark 3.11 (vertically sheared horizontal flows). This remark discusses vertically sheared horizontal flows, which are particular solutions to the Boussinesq system characterised by vanishing thermodynamic variables (namely θ and q in the moist case) and by a purely z -dependent and horizontal velocity $u = u_h(z)$.

The purpose of this remark is two-fold. First it shows what the global change of coordinates of corollary 3.7 and proposition 3.8 looks like in a specific case, namely when it is applied to these vertically sheared horizontal flow. Second it shows one benefit of the decomposition presented here: these flows are accounted for *naturally*.

First, we recall that combining corollary 3.7 and proposition 3.8 gives us *three* descriptions of any given state.

- The description in physical space in terms of (u, θ, q) .
- The description in coordinate space in terms of (p, M, σ, w) .
- The description in dual quantity space in terms of (PV, M, j, w, a) .

Let us now see what vertically sheared horizontal flows look like in each of these three descriptions.

- In physical space a vertically sheared horizontal flow corresponds to any state (u, θ, q) where θ and q vanish and where $u = u_h(z)$, i.e. $u_3 = 0$ and u only depends on z .
- In coordinate space a vertically sheared horizontal flow corresponds to any state described by (p, M, σ, w) where p , M , and w vanish and where $\sigma = \sigma_h(z)$.

- In dual quantity space a vertically sheared horizontal flow corresponds to any state described by (PV, M, j, w, a) where PV, M , and w vanish, where $j = j(z)$, and where $a = (a_1, a_2, 0)$ is any purely horizontal vector in $\mathbb{R}^3$.

In other words, since a vertically sheared horizontal flows is fully characterised by a map $\gamma : \mathbb{T}^1 \rightarrow \mathbb{R}^2$, here are three descriptions of the same such flow.

Physical space	Coordinate space	Dual quantity space
$u = \gamma_h(z)$	$p = 0$	$PV = 0$
$\theta = 0$	$M = 0$	$M = 0$
$q = 0$	$\sigma = -\gamma_h^\perp(z)$	$j = \gamma'(z)$
	$w = 0$	$w = 0$
		$a = -\oint \gamma_h^\perp(z) dz$

Second, we note that as described above our decomposition accounts for these vertically sheared horizontal flows very naturally. This comes from the fact that our decomposition leverages a *three-dimensional* Helmholtz decomposition of the *mixed vector* fields $-u_h^\perp + \theta e_3$ (for PV) and $-u_h^\perp + (\theta - \min_0 q) e_3$ (for j). This is different from what is often done where a *two-dimensional* Helmholtz decomposition is performed *solely* on the horizontal velocity u_h . This two-dimensional Helmholtz decomposition then leaves a two-dimensional average term of the form $u_h(z)$ to be determined, which is precisely the term accounting for these vertically sheared horizontal flows. In other words: our decomposition accounts for these intrinsically, instead of having to treat them as a separate add-on after the fact.

Remark 3.12. It is important to note that the decomposition discussed in this section is *always* valid, even when the state in $\mathbb{L}_\sigma^2$ thus decomposed is *not* a solution of the moist Boussinesq dynamical system. In other words: the decomposition of theorem 3.5 is valid for any state in the state space $\mathbb{L}_\sigma^2$, similarly to how the Helmholtz decomposition is valid for any three-dimensional vector field.

However, as discussed in the next section, the decomposition we introduced in this context does also have additional meaning in the context of the *dynamics* of the moist Boussinesq system (2.7a)–(2.7c).

4. Decomposition and slow–fast dynamics

In section 3 above we have introduced a preliminary decomposition of the leading-order dynamics. In particular we have discussed how to then make sure that the balanced component is slow. The key step in doing so was ensuring that the pressure, which acts as a streamfunction for the balanced component of the velocity, be slow. This differs from previous work [ZSS21b, ZSS21a, WSSM19, WSS+20] where the streamfunction constructed was approximately slow but not precisely slow.

A question remains: ‘did we extract *all* of the slow components of a state in this manner?’. In the dry case, that question may be answered in the affirmative since the fast dynamics are that of linear oscillations which do not feed back into the slow component. In the moist case, however, the fast oscillations are nonlinear and so they *may* feed back into the slow component.

For additional context and comparison, a brief summary of the dry case is as follows. References are, for example, [EM96, ME98, Maj03]. First note that the limit of interest is a limit of small Froude and small Rossby numbers, similar to the standard QG limit that is commonly presented in textbooks [Sal98, Maj03, Val17]. However, the difference from standard QG theory is that inertia-gravity waves are allowed to be present (as the $\Psi(\sigma, w)$ component in proposition 3.3 or the $\mathcal{X}_w$ component in theorem 3.5). The inertia-gravity waves are allowed to have the same amplitude as the balanced component. This scenario is typically referred to as ‘unbalanced initial conditions’ or ‘ill-prepared initial data’ (in contrast to the standard QG theory where the unbalanced inertia-gravity waves are assumed to be small-amplitude or zero). Interestingly, in this scenario with unbalanced initial conditions, in the dry case, the balanced component still obeys the QG equations; the inertia-gravity waves are also present, but their influence on the balanced component is averaged out when considering the longer time scale of the balanced dynamics. ‘Fast-wave averaging’ is the multi-scale technique that allows one to derive the long-time-scale dynamics by averaging over the fast-time-scale dynamics.

In more detail, fast-wave averaging is a technique used to identify the limiting behaviour of certain dynamical systems which admit two timescales. The prototypical case in which two timescales would be

present is the following dynamical system in $\mathbb{R}^{n_1+n_2}$:

$$\partial_t \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} f_1(x_1) \\ \frac{1}{\varepsilon} f_2(x_2) \end{pmatrix}$$

for $(x_1, x_2) \in \mathbb{R}^{n_1+n_2}$, where $\varepsilon > 0$ is a small parameter. In this case we would say that the dynamics of x_1 are slow since they are of order one, whereas the dynamics of x_2 are fast since they are of order $1/\varepsilon$. In general, dynamical systems will not admit an *exact* two timescale decomposition as in the prototypical example above. In general, that decomposition is only approximately valid for finite ε and we are then interested in the limiting behaviour of the slow and fast components as $\varepsilon \rightarrow 0$.

Where does fast-wave averaging fit in all this? Well: possessing two timescales is a feature a dynamical system may or may not manifest, whereas fast-wave averaging is a technique used to identify the limiting behaviour of the slow and fast timescales (for such dynamical systems). Fast-wave averaging happens in four steps.

1. Assume a two-timescale asymptotic expansion of the unknowns (which requires a sub-linear growth condition).
2. Plug this expansion into the equations of motion to identify the resulting equations at leading-order $1/\varepsilon$ and at next-order.
3. Solve the leading-order equations over the fast timescale in terms of a yet-to-be-determined ‘background profile’ at the slow timescale.
4. Determine this background profile at the slow timescale by using a solvability condition for the next-order equations.

We will now run through these four steps for the moist Boussinesq system (2.7a)–(2.7c). This constitutes a version of the QG limit with unbalanced or ill-prepared initial conditions. At each step we will highlight the differences with the dry case. This is because the dry case is well-understood (e.g. [EM96, ME98, Maj03]) however in the moist case we only provide here a sketch of the fast-wave averaging process, leaving a rigorous treatment for future work.

First we assume a two-timescale expansion of the unknowns. This is done exactly as in the dry case and we assume that, for $\mathcal{X} = (u, \theta, q)$,

$$\mathcal{X} = \mathcal{X}^0(t, \tau)|_{\tau=t/\varepsilon} + \varepsilon \mathcal{X}^1(t, \tau)|_{\tau=t/\varepsilon} + O(\varepsilon^2) \quad (4.1)$$

for some $\mathcal{X}^0$ and $\mathcal{X}^1$ to be determined where t is the slow timescale and $\tau = t/\varepsilon$ is the fast timescale, subject to the *sub-linear growth condition*

$$\mathcal{X}^1(t, \tau) = o(\tau) \text{ in } L^2 \text{ uniformly in } t \text{ as } \tau \rightarrow \infty, \quad (4.2)$$

or equivalently for every $\delta > 0$ there exists $\tau_* > 0$ such that if $\tau > \tau_*$ then

$$\frac{\|\mathcal{X}^1(t, \tau)\|_{L^2}}{\tau} < \delta \text{ for all } t.$$

Note that this sub-linearity condition is imposed for the following reason. If the next-order term $\mathcal{X}^1$ is allowed to grow linearly in τ then we would have that

$$\varepsilon \mathcal{X}^1(t, \tau)|_{\tau=t/\varepsilon} \sim \varepsilon \left(\mathcal{X}^1(t, 0) + \frac{t}{\varepsilon} \right) = \varepsilon \mathcal{X}^1(t, 0) + t,$$

i.e. the term t of order one would appear in an expression that ought to be of order ε .

The next step is to insert the two-timescale assumption (4.1) into the moist Boussinesq system (2.7a)–(2.7c) and read off the leading-order and next-order equations that come out. In order to carry this out we record here two preliminary computations. First we note that

$$\partial_t [f(t, \tau)|_{\tau=t/\varepsilon}] = \left[(\partial_t f)(t, \tau) + \frac{1}{\varepsilon} (\partial_\tau f)(t, \tau) \right] \Big|_{\tau=t/\varepsilon}.$$

Second we note that, introducing the notation $H := \mathbf{1}(q^0 + \varepsilon q^1 < 0)$ and $H^0 := \mathbf{1}(q^0 < 0)$,

$$\min_0 (q_0 + \varepsilon q^1) = (q^0 + \varepsilon q^1) H = \underbrace{q^0 H^0}_{\min_0 q^0} + q^0 (H - H^0) + \varepsilon q^1 H$$

where

$$\begin{aligned}
 q^0(H - H^0) &= q^0 [\mathbb{1}(q^0 + \varepsilon q^1 < 0) - \mathbb{1}(q^0 < 0)] \\
 &= q^0 [\mathbb{1}(q^0 + \varepsilon q^1 < 0, q^0 \geq 0) - \mathbb{1}(q^0 < 0, q^0 + \varepsilon q^1 \geq 0)] \\
 &= q^0 [\mathbb{1}(0 \leq q^0 < -\varepsilon q^1) - \mathbb{1}(-\varepsilon q^1 \leq q^0 < 0)] \\
 &= \varepsilon \cdot \frac{q^0}{\varepsilon} \left[\mathbb{1} \left(0 \leq \frac{q^0}{\varepsilon} < -q^1 \right) - \mathbb{1} \left(-q^1 \leq \frac{q^0}{\varepsilon} < 0 \right) \right] \\
 &=: \varepsilon R(q^0; q^1).
 \end{aligned}$$

In particular we note that $R(q^0; q^1) \in O(1)$ provided that q^1 is bounded.

Using these two observations we plug (4.1) into (2.7a)–(2.7c) and read off the leading-order equations, at order $1/\varepsilon$, to be

$$\begin{cases} \partial_\tau u^0 + e_3 \times u^0 + \nabla p^0 - \theta^0 e_3 + \min_0 q^0 e_3 = 0, & (4.3a) \\ \partial_\tau \theta^0 + u_3^0 = 0, \text{ and} & (4.3b) \\ \partial_\tau q^0 - u_3^0 = 0 & (4.3c) \end{cases}$$

while the next-order equations, at order one (i.e. order ε^0) to be

$$\begin{cases} \partial_\tau u^1 + e_3 \times u^1 + \nabla p^1 - \theta^1 e_3 + q^1 H e_3 = -\partial_t u^0 - u^0 \cdot \nabla u^0 + R(q^0; q^1), & (4.4a) \\ \partial_\tau \theta^1 + u_3^1 = -\partial_t \theta^0 - u^0 \cdot \nabla \theta^0, \text{ and} & (4.4b) \\ \partial_\tau q^1 - u_3^1 = -\partial_t q^0 - u^0 \cdot \nabla q^0. & (4.4c) \end{cases}$$

We can also write these equations more compactly as

$$\begin{cases} \partial_\tau \mathcal{X}^0 + \mathcal{N}(\mathcal{X}^0) \text{ and} & (4.5a) \\ \partial_\tau \mathcal{X}^1 + \mathcal{L}_H(\mathcal{X}^1) = -\partial_t \mathcal{X}^0 - u^0 \cdot \nabla \mathcal{X}^0 + \mathcal{R}(\mathcal{X}^0; \mathcal{X}^1), & (4.5b) \end{cases}$$

for $\mathcal{N}$ as defined in (3.1), $\mathcal{L}_H$ its formal linearisation given by

$$\mathcal{L}_H \begin{pmatrix} u \\ \theta \\ q \end{pmatrix} = \begin{pmatrix} \mathbb{P}_L [e_3 \times u - (\theta - qH) e_3] \\ u_3 \\ -u_3 \end{pmatrix}, \quad (4.6)$$

and the remainder term $\mathcal{R}$ given by $\mathcal{R}(u^0, \theta^0, q^0; u^1, \theta^1, q^1) = (R(q^0; q^1), 0, 0)$.

Crucially: we see that the dynamics at leading-order are governed by $\partial_\tau + \mathcal{N} = 0$ for $\mathcal{N}$ as introduced in (3.1). The sets $\mathcal{B}$ and $\mathcal{W}$ characterised in propositions 3.1 and 3.3 thus play an essential role since they come from that same operator $\mathcal{N}$. Indeed, as discussed in proposition C.9, a solution $\mathcal{X}$ of $\partial_\tau \mathcal{X} + \mathcal{N}(\mathcal{X}) = 0$ may be written as

$$\mathcal{X}(\tau) = \mathcal{X}_B(0) + e^{-\tau \mathcal{L}_{W, H_B}} \mathcal{X}_W(0). \quad (4.7)$$

Note that here we abused notation a little bit in order to be reminiscent of what happens in the dry case. In the dry case the leading order dynamics are $\partial_\tau + \mathcal{L}$ for $\mathcal{L}$ given by

$$\mathcal{L}(u, \theta) = \begin{pmatrix} \mathbb{P}_L (u_h^\perp - \theta e_3) \\ u_3 \end{pmatrix}, \quad (4.8)$$

as per (2.1a) and (2.1b). Since the operator $\mathcal{L}$ is linear and constant, it follows that the solution map of $\partial_\tau + \mathcal{L}$ is precisely given by the semigroup $e^{-\tau \mathcal{L}}$. The operator $\mathcal{L}_{W, H_B}$ in the moist case is linear but with variable coefficients (namely through its dependence on H_B), yet we use the same notation to represent the solution map, keeping it mind that it is slightly abusive notation since the solution map is *not* a semigroup.

We are thus ready to carry out the third step of the fast-wave averaging process. We deduce from (4.7) that the solution of the leading-order equations (4.5a) is given by

$$\underline{\mathcal{X}}^0(t, \tau) = \underline{\mathcal{X}}_B^0(t) + e^{-\tau \mathcal{L}_{W, H_B}} \underline{\mathcal{X}}_W^0(t)$$

where the background slow profile $\underline{\mathcal{X}}^0(t) = \underline{\mathcal{X}}_B^0(t) + \underline{\mathcal{X}}_W^0(t)$ is to be determined (this is done in the next step of the fast-wave averaging process). This is where the first essential difference with the dry case manifests itself. In the dry case, the leading-order solution takes the form

$$\underline{\mathcal{X}}^0(t, \tau) = \underline{\mathcal{X}}_B^0(t) + e^{-\tau \mathcal{L}_W} \underline{\mathcal{X}}_W^0(t), \quad (4.9)$$

where $\mathcal{L}_W$ denotes the restriction of $\mathcal{L}$ from (4.8) onto its image. Crucially: since the operator $\mathcal{L}_W$ is skew-adjoint in L^2 the oscillations in (4.9) are linear oscillations with constant frequencies. Here, in the moist case, *the frequencies depend on the phase*.

To see why this matters, suppose for now that we have completed the fourth and last step of the fast-wave averaging process. That is, suppose we have determined the dynamics of the slow part $\mathcal{X}^0(t)$. Since we are interested in the limiting behaviour of solutions, we then wish to study the limiting behaviour of $\mathcal{X}^0(t, \tau)$ as $\varepsilon \rightarrow 0$. In the *dry* case this means studying the limiting behaviour of

$$e^{-\tau \mathcal{L}_W} \underline{\mathcal{X}}_W^0(t).$$

Crucially: since the dry oscillations have constant frequencies, these oscillations have vanishing time averages, and so the weak limit of this term is simply zero!

In the moist case we must by contrast study the limiting behaviour of

$$e^{-\tau \mathcal{L}_{W, H_B}} \underline{\mathcal{X}}_W^0(t).$$

But now, as shown in an example first recorded in [ZSS21a] (see corollary C.22, and the simple example it arises from in proposition C.19), phase-dependent oscillations do *not* have vanishing time averages. This wave term therefore converges weakly to a *non-zero* time average!

This tells us exactly how to augment our ‘pre-decomposition’ above, in theorem 3.5, to now make sure we include *all* of the slow components: we take the slow balanced piece from the previous section and add to it the time-average of the wave piece.

This produces the following slow–fast decomposition, where for any space S , $L^1 S$ is used to denote $L^1([0, \infty); S)$.

Theorem 4.1 (slow–fast decomposition). *We have the decomposition*

$$L^1 \mathbb{L}_\sigma^2 = L^1 \mathcal{B} + \overline{\mathcal{W}} + \widetilde{\mathcal{W}}$$

where

$$\overline{\mathcal{W}} := \{ \mathcal{X} \in L^1 \mathcal{W} : \mathcal{X} \text{ is independent of time } \tau \} \text{ and } \widetilde{\mathcal{W}} := \left\{ \mathcal{X} \in L^1 \mathcal{W} : \lim_{\tau \rightarrow \infty} \int_0^\tau \mathcal{X}(\tau') d\tau' = 0 \right\},$$

in the sense that, for every $\mathcal{X} \in L^1 \mathbb{L}_\sigma^2$ there exist unique $\mathcal{X}_B \in L^1 \mathcal{B}$, $\overline{\mathcal{X}}_W \in \overline{\mathcal{W}}$, and $\tilde{\mathcal{X}}_W \in \widetilde{\mathcal{W}}$ such that

1. $\mathcal{X} = \mathcal{X}_B + \overline{\mathcal{X}}_W + \tilde{\mathcal{X}}_W$ and
2. $PV(\mathcal{X}) = PV(\mathcal{X}_B)$ and $\mathcal{M}(\mathcal{X}) = \mathcal{M}(\mathcal{X}_B)$ for PV and $\mathcal{M}$ as in theorem 3.5.

Moreover $\mathcal{X}_B(\tau)$ and $\mathcal{X}_W(\tau)$ may be computed explicitly at every instant τ as in theorem 3.5 such that then $\overline{\mathcal{X}}_W := \lim_{\tau \rightarrow \infty} \int_0^\tau \mathcal{X}_W(\tau') d\tau'$ and $\tilde{\mathcal{X}} := \mathcal{X} - \overline{\mathcal{X}}_W$.

Proof. This follows immediately from applying theorem 3.5 to $\mathcal{X}(\tau)$ at every instant $\tau \geq 0$ and then constructing $\overline{\mathcal{X}}_W$ and $\tilde{\mathcal{X}}_W$ as indicated. $\square$

Note that the technical aspects of this theorem were already taken care of in the proof of the ‘pre-decomposition’ of theorem 3.5. So there is nothing new to prove here. Nonetheless, as explained in this section in the lead-up to theorem 4.1, this statement now has an *interpretation* that theorem 3.5 did not. Once viewed through the lens of fast-wave averaging, we can indeed interpret theorem 4.1 as a *slow–fast* decomposition:

$$\mathcal{X} = \underbrace{\mathcal{X}_B + \overline{\mathcal{X}}_W}_{\mathcal{X}_{\text{slow}}} + \underbrace{\tilde{\mathcal{X}}_W}_{\mathcal{X}_{\text{fast}}}.$$

In other words: the fast-wave averaging framework now, at least formally, ensures that this decomposition accounts for the entirety of the slow components of the solution!

Finally, for completeness we record the fourth and final step of fast-wave averaging. The idea is that we can solve (4.5b) for $\mathcal{X}^1$, but we are not particularly interested in what that solution looks like. What we care about is making sure that this next-order solution satisfies the sub-linear growth condition (4.2). Ultimately, this comes down to a condition on the right-hand side forcing terms of (4.5b). Since these forcing terms depend on the leading-order solution $\mathcal{X}^0$, enforcing the sub-linear growth condition on $\mathcal{X}^1$ turns into a condition on $\mathcal{X}^0$.

In the dry case, this condition takes the form

$$\int_0^\tau \partial_t \underline{\mathcal{X}}^0 + e^{\tau' \mathcal{L}} \left(e^{-\tau' \mathcal{L}} \bar{u}^0 \cdot \nabla e^{-\tau' \mathcal{L}} \underline{\mathcal{X}}_0 \right) d\tau' = 0$$

since we may use the fact that the semigroup $e^{-\tau \mathcal{L}}$ is norm-preserving. In the moist case however the solution map $e^{-\tau \mathcal{L}_{W, H_B}}$ is *not* norm-preserving (see proposition C.12) and so we cannot perform the same trick. Identifying what the solvability condition looks like in the moist case is therefore left open for future work.

Remark 4.2 (slow-fast interpretation in different limiting regimes). In this section we have then described how, once in the limiting regime where both the Froude and Rossby numbers approach zero (as in the QG approximation with unbalanced or ill-prepared initial data), the decomposition introduced in this paper has an additional interpretation as a slow-fast decomposition. Nonetheless, this same decomposition may still be of use in other limiting regimes, provided the interpretation of each of its components as slow or fast is updated appropriately.

Consider for example the case where the Froude number is still taken to approach zero but the Rossby number is kept constant. In that case the vertically sheared horizontal flows discussed in remark 3.11 are no longer fast, but now become slow!

5. Decomposition and energy

In the dry case, the decomposition of states into their slow and fast components has the additional properties that it is an orthogonal decomposition, where the inner product used is the one induced by the conserved energy.

In the moist case, which is the focus of this paper, we therefore ask the following question: how does our decomposition of theorem 3.5 relate to the energy conserved by the moist Boussinesq system? The energy in question is given by

$$E = \frac{1}{2} \int |u|^2 + \theta^2 + \min_0^2 q + M^2. \quad (5.1)$$

This conserved energy (in piecewise form to account for different phases) was demonstrated by [MSS19]. Note that, in the absence of precipitation, the coefficient of the M^2 term is arbitrary and has been set to 1 for simplicity. The two terms θ^2 and $\min_0^2 q$ are needed to account for buoyant energy transfers, and they are related to the two portions of the moist buoyancy, $\theta - \min_0 q$. For a demonstration that this energy is conserved, see proposition C.10 herein.

In the dry case, since the decomposition comes down to linear projections, and since the energy gives rise to an inner product, the relation between the decomposition and the energy can be phrased as a simple question: are these projections orthogonal with respect to this inner product? In the moist case, both the decomposition and the energy are *nonlinear* and so it is not clear how one would go about even *asking* if the two are related. We offer several options in sections 5.1–5.3, summarizing and contrasting each option, along with their benefits and inconveniences, in section 5.4.

Note that, in each of the three cases discussed below, the strategy is the same. Namely we seek to identify a way to quantify how far apart two states (u_1, θ_1, q_1) and (u_2, θ_2, q_2) are in such a manner that

1. this agrees with the decomposition of theorem 3.5, and gives it further meaning as a projection, in the sense that the closest balanced state to any given state must be its balanced component and
2. this agrees with the energy (when the two are comparable).

Finally, note that the authors believe that the quantification provided in section 5.3 is really the right one to use for describing the decomposition from a projection perspective. However, it may appear a bit odd at first. To better motivate it we therefore discuss, in sections 5.1 and 5.2, alternative approaches which are less conclusive but help us understand why the *good* approach, discussed in section 5.3, must take the form it does. Also, the alternative approaches in sections 5.1 and 5.2 have value in their own right, for other applications beyond the focus here on a nonlinear eigenspace decomposition.

Before we dive into the details of each of these three approaches, we must elucidate the following question: what does it mean for two states to be comparable using the energy?

Remark 5.1 (comparing states using the energy). Recall that the energy of a state $\mathcal{X} = (u, \theta, q)$ is

$$E(\mathcal{X}) = \int |u|^2 + \theta^2 + \min_0^2 q + (\theta + q)^2.$$

The question is: given two states $\mathcal{X}_1$ and $\mathcal{X}_2$, *when* can we use the energy to quantify the discrepancy between $\mathcal{X}_1$ and $\mathcal{X}_2$?

A naive attempt would be to use the energy of their *difference*, i.e.

$$E(\mathcal{X}_1 - \mathcal{X}_2) = \int |u_1 - u_2|^2 + (\theta_1 - \theta_2)^2 + \min_0^2(q_1 - q_2) + ((\theta_1 + q_1) - (\theta_2 + q_2))^2. \quad (5.2)$$

This is not physically meaningful. Why? Because of the nonlinearity. When computing $E(\mathcal{X}_1 - \mathcal{X}_2)$ the non-linear term $\min_0^2(q_1 - q_2)$ only turns on when $q_1 - q_2 > 0$. However this nonlinearity is not supposed to care about the sign of $q_1 - q_2$. Instead the nonlinearity is there to account for the fact that, depending on the signs of q_1 and q_2 , the energetic weights of q_1 and q_2 change.

What is the takeaway? The only time we can directly compare two states $\mathcal{X}_1$ and $\mathcal{X}_2$ using the energy is when the signs of q_1 and q_2 agree everywhere in the spatial domain. The sign of q is a mathematical proxy used to tell us which *phase* the state is in at any point in the spatial domain. We may thus rephrase the observation above as follows.

The only time we can directly compare two states $\mathcal{X}_1$ and $\mathcal{X}_2$
is when their phases coincide everywhere in the spatial domain.

It will therefore be convenient to encode the spatial location of a state's phases. This is done by using the indicator function

$$H := \mathbb{1}(q < 0).$$

Note that since the water content $q = q(t, x)$ is a function of t and x , then the indicator H also depends on t and x . As above, we will often use notation which suppresses the (t, x) -dependence in an effort to improve readability.

So let us consider, momentarily, two states $\mathcal{X}_1$ and $\mathcal{X}_2$ whose phase locations agree, i.e. for which $H_1 = H_2 =: H$. On one hand

$$\begin{aligned} E(\mathcal{X}_1) &= \int |u_1|^2 + \theta_1^2 + \min_0^2 q_1 + (\theta_1 + q_1)^2 \\ &= \int |u_1|^2 + \theta_1^2 + q_1^2 H + (\theta_1 + q_1)^2 \end{aligned}$$

while on the other hand

$$\begin{aligned} E(\mathcal{X}_2) &= \int |u_2|^2 + \theta_2^2 + \min_0^2 q_2 + (\theta_2 + q_2)^2 \\ &= \int |u_2|^2 + \theta_2^2 + q_2^2 H + (\theta_2 + q_2)^2. \end{aligned}$$

A natural way to compare $\mathcal{X}_1$ and $\mathcal{X}_2$ using the energy is then to treat the indicator function H as *fixed* in the expressions above for both $E(\mathcal{X}_1)$ and $E(\mathcal{X}_2)$ and to consider

$$\int |u_1 - u_2|^2 + (\theta_1 - \theta_2)^2 + (q_1 - q_2)^2 H + ((\theta_1 + q_1) - (\theta_2 + q_2))^2. \quad (5.3)$$

This is *not* the same as $E(\mathcal{X}_1 - \mathcal{X}_2)$, as can be seen by comparing (5.2) with (5.3). Nonetheless the quantity appearing in (5.3) is, so far, the only way we have of using the energy to compare two states, as long as their phase locations agree.

So we give this quantity a name, and we generalise beyond an H that is a common phase indicator for two states, and now allow for H to be arbitrarily chosen.

Definition 5.2. For *any* indicator function H and *any* state $\mathcal{X} = (u, \theta, q) \in \mathbb{L}_\sigma^2$ let

$$E_H(\mathcal{X}) := \int |u|^2 + \theta^2 + q^2 H + (\theta + q)^2.$$

Note that, in what follows, we will refer to the arbitrarily chosen H in definition 5.2 as a ‘fixed’ Heaviside H . By ‘fixed’ we mean that H can be arbitrarily assigned and fixed, and we are not imposing (at this point) any particular relationship between the fixed H and other quantities such as q, PV, M , etc. In particular, note that a ‘fixed’ H is not necessarily constant; it may have spatial variations that represent cloudy and non-cloudy regions.

Also note that this notation may then be used even when (especially when!) the sign of q does not agree with H . That is precisely what happens in (5.3): H encodes the signs of q_1 and q_2 , *not* the sign of $q_1 - q_2$ (that was precisely the naive nonphysical approach we sought to avoid!)

This ‘fixed H ’ energy will be useful in various ways below: for instance, it is the quadratic approximation of the conserved energy about a state whose phase is encoded in H . So this energy also appears in situations where linearisations are used, since linearising the dynamical equations is essentially the same as working with the quadratic approximation of the conserved energy. This interpretation of the ‘fixed H ’ energy will be more prevalent in later sections, such as section A.2 where geometrical notions relying on tangent spaces (a.k.a. linearisations) are discussed. Linearisations and the ‘fixed H ’ energy are also relevant for applications such as adjoint models, sensitivity analysis, and variational data assimilation, which rely on tangent linear approximations to the dynamics [LDT86, PD97, Err97, ER99, Mah99, BBP+01, AZD08, DARR14].

In particular we emphasise that this ‘fixed H ’ energy should be viewed as a mathematically convenient quantity, and not as an (immediately) relevant physical quantity. It will be convenient in the discussions below, as we seek the *good* approach. Next, in sections 5.1 and 5.2, we consider alternative approaches which are less conclusive but which help to motivate the form of the *good* approach, discussed in section 5.3.

5.1. PDE-centric approach

We first discuss a PDE-centric approach to relating the decomposition and the energy. Why start here? There are two reasons to begin with a PDE-centric approach.

The first reason to take such an approach is that elliptic PDEs underpin the decomposition in both the case of the Helmholtz decomposition and in the dry case. Since a nonlinear elliptic PDE is one of the few structures that have already been identified to be present in the moist case, this PDE therefore presents a natural starting point.

The second reason to take a PDE-centric approach is that a similar approach was fruitful in identifying the slow–fast decomposition. Indeed, the slow–fast decomposition was obtained by recognizing the important role played by the slow dual quantities PV and M , and then inverting the nonlinear PV -and- M inversion PDE (3.5) in order to obtain p . The variables p and M were then used to characterise the balanced part of a state.

Equipped with propositions 3.1 and 3.3 we can actually go a step further: the slow dual quantities PV and M characterise the balanced component of a state while the fast dual quantities j , w , and a characterise its wave component. Indeed, we have the following.

- On one hand proposition 3.1 tells us that the dual quantities

$$j(u, \theta, q) = \partial_3 u_h - \nabla_h^\perp (\theta - \min_0 q), \quad w(u, \theta, q) = u_3, \\ \text{and } a(u, \theta, q) = -\oint u_h^\perp + \oint (\theta - \min_0 \theta) e_3$$

are transversal to the balanced set $\mathcal{B}$, and so we may expect them to characterise elements of the wave set $\mathcal{W}$.

- On the other hand proposition 3.3 tells us that the dual quantities

$$PV(u, \theta, q) = \nabla_h^\perp \cdot u_h + \partial_3 \theta \text{ and} \\ M(u, \theta, q) = \theta + q$$

are transversal to the wave set $\mathcal{W}$, and so we may similarly expect them to characterise elements of the balanced set $\mathcal{B}$.

Both of these assertions turn out to be true, and thus give rise to the global change of coordinates recorded in proposition 3.8 which allows us to translate between a state (u, θ, q) and the tuple of dual quantities (PV, M, j, w, a) .

We will now use this change of variable from proposition 3.8 to define a distance on the state space $\mathbb{L}_\sigma^2$. We will do this by using the change of variables to pull-back a metric on the dual quantity space $\mathcal{M}$ and thus produce a metric on the state space $\mathbb{L}_\sigma^2$ (see proposition C.13). The execution of this strategy requires a few steps. In order to have a clearer idea of where we are headed, we record below what would happen if we tried doing the same thing in the *dry* case. In the dry case we can write the energy in terms of the dual quantities (namely PV , j , w , and a). In the language that will be used in

this section, we can therefore view the energy, and the L^2 metric it corresponds to, as the pullback of a metric on the (dry) dual quantity space! Since we will ultimately, in this section, define the distance of interest in terms of the *coordinates* (think p , M , σ , and w), we also record in proposition 5.3 below how the energy may be written in a third way, in terms of the coordinates (which are p , σ , and w in the dry case).

Proposition 5.3 (Parseval identity in the dry case). *For any $u \in (L^2)_\sigma^3$ and $\theta \in L^2$ we have that*

$$\underbrace{\int |u|^2 + \theta^2}_E = \int |\nabla \Delta^{-1} PV|^2 + \left| (\nabla \times)^{-1} (j + (\partial_3 w) e_3, a) \right|^2 = \underbrace{\int |\nabla p|^2 + |\sigma|^2 + w^2}_{\sim d^2}$$

where p solves

$$\Delta p = \underbrace{\nabla_h^\perp \cdot u_h + \partial_3 \theta}_{PV}$$

while σ solves

$$\nabla \times \sigma = \underbrace{\partial_3 u_h - \nabla_h^\perp \theta}_j + (\partial_3 w) e_3 \text{ subject to } \oint \sigma = \underbrace{\oint -u_h^\perp + \oint \theta e_3}_a,$$

and where $w := u_3$. Here we write $\sigma = (\nabla \times)^{-1}(\omega, a)$ to mean that σ solves

$$\nabla \times \sigma = \omega \text{ subject to } \nabla \cdot \sigma = 0 \text{ and } \oint \sigma = a.$$

Proof. We recall that the decomposition in the dry case is

$$\begin{pmatrix} u \\ \theta \end{pmatrix} = \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p \end{pmatrix} + \begin{pmatrix} \sigma_h^\perp + w e_3 \\ \sigma_3 \end{pmatrix}$$

where σ is divergence-free and satisfies $\nabla_h^\perp \cdot \sigma_h = \partial_3 w$. We may then compute that

$$PV = \nabla_h^\perp \cdot u_h + \partial_3 \theta = \Delta p$$

while lemma C.4 tells us that

$$\nabla \times \sigma = \partial_3 \sigma_h^\perp - \nabla_h^\perp \sigma_3 + (\nabla_h^\perp \cdot \sigma_h) e_3 = \partial_3 u_h - \nabla_h^\perp \theta + (\partial_3 w) e_3,$$

with moreover

$$\oint \sigma = \oint (-u_h^\perp + \theta e_3),$$

as claimed. Finally we may compute that, since σ is divergence-free and hence $\int \sigma \cdot \nabla p = 0$,

$$\begin{aligned} E &= \frac{1}{2} \int |u|^2 + \theta^2 = \frac{1}{2} \int |u_h|^2 + u_3^2 + \theta^2 \\ &= \frac{1}{2} \int |\nabla_h p + \sigma_h|^2 + w^2 + (\partial_3 p + \sigma_3)^2 \\ &= \frac{1}{2} \int |\nabla p|^2 + |\sigma|^2 + w^2, \end{aligned}$$

as desired. $\square$

The benefit of the approach carried out in this section, where we pull back a metric from dual quantity space onto state space, is that it will immediately produce a metric which agrees with the decomposition. The drawback of this approach is that it does not, inherently, relate to the conserved energy in any way. To see why a metric produced in this way will automatically agree with the decomposition relies on the following observation: if two states have the same PV and M , then they must have the same balanced component. This is proved in lemma 5.4 below. Therefore projecting a state onto the balanced set is the same as keeping its PV and M dual quantities fixed and sending its remaining dual quantities, namely j , w , and a , to zero. This is why a metric pulled back from the dual quantity space $\mathcal{M}$ will necessarily agree with the decomposition. As mentioned above, this relies on the following result.

Lemma 5.4 (PV and M characterise balanced components). *Given any $\mathcal{X} \in \mathbb{L}_\sigma^2$ the unique balanced state in $\mathcal{B}$ whose PV and M agree with the PV and M of $\mathcal{X}$ is precisely the balanced component $\mathcal{X}_B$.*

Proof. Let $\mathcal{X} \in \mathbb{L}_\sigma^2$ and let $\mathcal{Y} \in \mathcal{B}$ be a balanced state whose PV and M agrees with those of $\mathcal{X}$, i.e.

$$PV(\mathcal{X}) = PV(\mathcal{Y}) \text{ and } M(\mathcal{X}) = M(\mathcal{Y}).$$

We know from proposition 3.1 that since $\mathcal{Y}$ is balanced it is uniquely characterised by $p \in \dot{H}^1$ and $M \in L^2$ via $\mathcal{Y} = \Phi(p, M)$. Moreover we may compute that

$$M(\Phi(p, M)) = M \text{ and } PV(\Phi(p, M)) = \Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p).$$

By the well-posedness of nonlinear PV-and-M inversion, we know that there is a unique p solving

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV(\mathcal{Y}) = PV(\mathcal{X}),$$

and proposition 3.1 tells us that the resulting state is precisely $\mathcal{X}_B$. This verifies that indeed $\mathcal{Y} = \mathcal{X}_B$. $\square$

Using lemma 5.4 above we could then immediately deduce that a distance on $\mathbb{L}_\sigma^2$ defined via

$$\begin{aligned} d(\mathcal{X}_1, \mathcal{X}_2)^2 := & \|\mathcal{M}(\mathcal{X}_1) - \mathcal{M}(\mathcal{X}_2)\|_{\mathfrak{M}}^2 = \|PV_1 - PV_2\|_{H^{-1}}^2 + \|M_1 - M_2\|_{L^2}^2 \\ & + \|j_1 - j_2\|_{H^{-1}}^2 + \|w_1 - w_2\|_{L^2}^2 + |a_1 - a_2|^2 \end{aligned} \quad (5.4)$$

would automatically agree with the decomposition. However this distance has no chance of agreeing with the energy! This is because its dimensions, or units, are all wrong. Indeed: this naive distance above treats PV and M on the same footing even though they do not have the same dimensions. In order to produce a metric on the state space $\mathbb{L}_\sigma^2$ pulled back from the dual quantity space $\mathcal{M}$ which has a chance of agreeing with the energy we must therefore *dimensionalise* the naive distance from (5.4).

In order to perform this dimensionalisation we will use the fact that the ‘fixed H ’ energy introduced in definition 5.2 may sometimes be written not in terms of the state (u, θ, q) but in terms of its coordinates (p, M, σ, w) . To see this, we first define the following.

Definition 5.5. Fix an indicator function H . We define, for any $(p, M, \sigma, w) \in \mathcal{C}$,

$$\tilde{E}_H(p, M, \sigma, w) := \frac{1}{2} \int \underbrace{A_H \nabla p \cdot \nabla p}_{=: Q_H^+(p)} + \left(1 + \frac{H}{2}\right) M^2 + \underbrace{A_H^{-1} \sigma \cdot \sigma}_{=: Q_H^-(\sigma)} + w^2,$$

where $A_H = I - \frac{H}{2} e_3 \otimes e_3$ and so $A_H^{-1} = I + H e_3 \otimes e_3$.

Remark 5.6 (the elliptic matrix A_H and its inverse). The matrix A_H introduced in definition 5.5 above will be encountered many times in the sequel. Where does it come from? The simplest explanation is that it comes from the linearisation of nonlinear PV-and-M inversion. Indeed, upon linearising PV-and-M inversion

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV$$

we obtain, upon differentiating,

$$\Delta \partial_\alpha p + \frac{1}{2} \mathbb{1}(M < \partial_3 p) (\partial_\alpha M - \partial_\alpha \partial_3 p) = \partial_\alpha PV.$$

In other words $\partial_\alpha p$ satisfies, for $H := \mathbb{1}(M < \partial_3 p)$,

$$\nabla \cdot (A_H \nabla \partial_\alpha p) = \partial_\alpha PV - \frac{1}{2} H \partial_\alpha M.$$

In particular we may readily verify that the identity

$$(1 + H) \left(1 - \frac{H}{2}\right) = 1$$

holds for any H satisfying $H^2 = H$ (which is the case if H is an indicator function), and so the inverse of A_H is indeed given by

$$A_H^{-1} = I + H e_3 \otimes e_3.$$

In general the form of the fixed Heaviside energy provided in definition 5.5 does not agree with the form of the fixed Heaviside energy provided in definition 5.2. Sometimes, they do.

Proposition 5.7 (rewriting the energy). *Let E_H be as introduced in definition 5.2 and $\tilde{E}_H$ be as in definition 5.5. Let $(u, \theta, q) \in \mathbb{L}_\sigma^2$ be corresponding coordinates (p, M, σ, w) as per corollary 3.7. In general we have that*

$$E_H(u, \theta, q) = \frac{1}{2} \int |u_h|^2 + u_3^2 + \left(1 - \frac{H}{2}\right) (\theta - qH)^2 + \left(1 + \frac{H}{2}\right) (\theta + q)^2.$$

Moreover, if q and its balanced component q_B share the same interface, i.e. if

$$H := \mathbb{1}(q < 0) = \mathbb{1}(q_B < 0),$$

where note that $\mathbb{1}(q_B < 0) = \mathbb{1}(M < \partial_3 p)$, then

$$E_H(u, \theta, q) = \tilde{E}_H(p, M, \sigma, w).$$

Proof. The first identity follows from an immediate computation since $H^2 = H$ and hence

$$\theta^2 + q^2 H = \left(1 - \frac{H}{2}\right) (\theta - qH)^2 + \frac{H}{2} (\theta + q)^2.$$

To deduce the second identity we first note that corollary 3.7 and lemma C.7 tell us that, if $H = H_B := \mathbb{1}(q_B < 0) = \mathbb{1}(M < \partial_3 p)$, then

$$\begin{cases} u_h = \nabla_h^\perp p + \sigma_h^\perp, \\ u_3 = w, \\ \theta - qH = \partial_3 p + (1 + H)\sigma_3, \text{ and} \\ \theta + q = M. \end{cases}$$

Therefore

$$\begin{aligned} & \int |u_h|^2 + u_3^2 + \left(1 - \frac{H}{2}\right) (\theta - qH)^2 + \left(1 + \frac{H}{2}\right) (\theta + q)^2 \\ &= \int |\nabla_h p + \sigma_h|^2 + w^2 + \left(1 - \frac{H}{2}\right) (\partial_3 p + (1 + H)\sigma_3)^2 + \left(1 + \frac{H}{2}\right) M^2. \end{aligned}$$

Finally, for A_H and A_H^{-1} as in definition 5.5 we compute that, since A_H and A_H^{-1} are symmetric and since σ is divergence-free,

$$\begin{aligned} \int |\nabla_h p + \sigma_h|^2 + \left(1 - \frac{H}{2}\right) (\partial_3 p + (1 + H)\sigma)^2 &= \int A_H (\nabla p + A_H^{-1} \sigma) \cdot (\nabla p + A_H^{-1} \sigma) \\ &= \int A_H \nabla p \cdot \nabla p + 2 \underbrace{\int \nabla p \cdot \sigma}_{=0} + \int A_H^{-1} \sigma \cdot \sigma. \end{aligned}$$

The second identity follows. □

This is a good point to take a stock. Recall that we seek to define a metric pulled back from the dual quantity space $\mathcal{M}$ onto the state space $\mathbb{L}_\sigma^2$. Such a metric will automatically agree with the decomposition but, in order to make it agree with the energy, we must dimensionalise it appropriately. The identity recorded in proposition 5.7 above tells us how to do this. The last remaining piece required to define our pullback metric is thus to identify how the dual quantities $(PV, M, j, w, \text{ and } a)$ will determine the coordinates $(p, M, \sigma, \text{ and } w)$. Note that such a map is already implicit in corollary 3.7, however we will do something slightly different: since the ‘fixed H ’ energy already uses a fixed Heaviside, we will leverage this to define a map taking dual quantities to coordinates in a *linear* way. This linearity will make the ensuing pullback distance easier to work with.

Lemma 5.8 (definition and invertibility of the fixed Heaviside parametrisation C_H). *Fix an indicator function H and define $A_H := I - \frac{H}{2} e_3 \otimes e_3$, such that $A_H^{-1} = I + H e_3 \otimes e_3$ (since $(1 - H/2)(1 + H) \equiv 1$). We define $C_H : \mathfrak{M} \rightarrow \mathfrak{C}$, where $\mathfrak{C}$ is the coordinate space defined in (3.10). Given $(PV, M, j, w, a) \in \mathfrak{M}$, let $M := M$, $w := w$, and let p and σ solve*

$$\nabla \cdot (A_H \nabla p) = PV - \frac{1}{2} \partial_3 (HM) \tag{5.5}$$

and

$$\nabla \times (A_H^{-1} \sigma) = j + (\partial_3 w) e_3 \text{ subject to } \nabla \cdot \sigma = 0 \text{ and } \oint A_H^{-1} \sigma = a. \quad (5.6)$$

The map C_H is well-defined and invertible.

Regarding the fixed Heaviside H , recall that this H is not necessarily constant; it may have spatial variations that represent cloudy and non-cloudy regions. The H is fixed in the sense that it can be arbitrarily assigned and fixed, and we are not imposing (at this point) any particular relationship between the fixed H and other quantities such as q, PV, M , etc.

Note that if $H = H_B := \mathbb{1}(q_B < 0) = \mathbb{1}(M < \partial_3 p)$ then C_H is precisely the map implicit in corollary 3.7 which assigns the coordinates p, M, σ , and w to a state (u, θ, q) through its dual quantities PV, M, j, w , and a .

Proof of lemma 5.8. It suffices to show that both (5.5) and (5.6) have a unique solution once PV, M, j, w , and a are fixed. The well-posedness of (5.5) follows from the fact that A_H is uniformly elliptic since $A_H(x) \geq \frac{1}{2}I$ for all x . The well-posedness of (5.6) follows from the fact that A_H^{-1} is uniformly elliptic, since $A_H^{-1}(x) \geq I$ for all x , and so proposition C.8 applies to (5.6). $\square$

Remark 5.9. In the definition of C_H we are linearising the coordinates, but *not* the dual quantities. This is because it is the full nonlinear dual quantities which are well-adapted to the balanced $\mathcal{B}$ and to the wave set $\mathcal{W}$, so changing those would mean that the resulting distance (in definition 5.10 below) would fail to be compatible with the energy. It should not yet be clear why using linearised versions of the coordinates is particularly helpful. This will be discussed in section 5.3 below.

We are now ready to define a dimensionalised version of the naive pullback metric of (5.4).

Definition 5.10 (dimensionalised Parseval distance). Fix an indicator function H . Define $d_H : \mathbb{L}_\sigma^2 \times \mathbb{L}_\sigma^2 \rightarrow [0, \infty)$ as follows:

$$\frac{1}{2} d_H(\mathcal{X}_1, \mathcal{X}_2)^2 = \tilde{E}_H(C_H[\mathcal{M}(\mathcal{X}_1) - \mathcal{M}(\mathcal{X}_2)]).$$

In other words, given $\mathcal{X}_1, \mathcal{X}_2 \in \mathbb{L}_\sigma^2$, let

$$M := M(\mathcal{X}_1) - M(\mathcal{X}_2), w := w(\mathcal{X}_1) - w(\mathcal{X}_2), \text{ and } a := a(\mathcal{X}_1) - a(\mathcal{X}_2),$$

let p solve, for $PV := PV(\mathcal{X}_1) - PV(\mathcal{X}_2)$,

$$\nabla \cdot (A_H \nabla p) = PV - \frac{1}{2} \partial_3 (HM)$$

and let σ solve, for $j := j(\mathcal{X}_1) - j(\mathcal{X}_2)$,

$$\nabla \times (A_H^{-1} \sigma) = j + (\partial_3 w) e_3 \text{ subject to } \oint A_H^{-1} \sigma = a \text{ and } \nabla \cdot \sigma = 0.$$

Then

$$\frac{1}{2} d_H(\mathcal{X}_1, \mathcal{X}_2)^2 = \int Q_H^+(p) + \left(1 + \frac{H}{2}\right) M^2 + Q_H^-(\sigma) + w^2$$

for $Q_H^\pm$ as introduced in definition 5.5.

Proposition 5.11. For any indicator function H the function d_H defined in definition 5.10 is a metric which agrees with the extraction of a balanced component. Moreover this distance is compatible with the energy when $H = H_1 = H_{B,1} = H_2 = H_{B,2}$, i.e. in that case

$$\begin{aligned} \frac{1}{2} d_H(\mathcal{X}_1, \mathcal{X}_2)^2 &= E_H(\mathcal{X}_1 - \mathcal{X}_2) \\ &= \int |u_1 - u_2|^2 + (\theta_1 - \theta_2)^2 + (q_1 - q_2)^2 H + (M_1 - M_2)^2 \end{aligned}$$

where $H = \mathbb{1}(q_1 < 0) = \mathbb{1}(q_2 < 0)$.

Proof. First we show that d_H is a metric. Since $\tilde{E}_H$ is positive-definite and viewed here as taking arguments in $\mathfrak{C}$ it is a quadratic form which induces an inner product, and hence a metric which we denote d_E , on $\mathfrak{C}$. Since $\mathcal{C}_H$ is linear and since both $\mathcal{C}_H$ and $\mathcal{M}$ are invertible (see proposition 3.8 and lemma 5.8) and hence injective it follows from proposition C.13 that d_H is the pullback metric of d_E under $\mathcal{C}_H \circ \mathcal{M}$. So d_H is indeed a metric on $\mathbb{L}_\sigma^2$.

Now we show that d_H conditionally agrees with the energy. Suppose that $H = H_1 = H_{B,1} = H_2 = H_{B,2}$. The argument is easier to follow if we rewrite d_H as follows, using the fact that $\mathcal{C}_H$ is linear:

$$\frac{1}{2}d_H(\mathcal{X}_1, \mathcal{X}_2)^2 = \tilde{E}_H[\mathcal{C}_H(\mathcal{M}(\mathcal{X}_1)) - \mathcal{C}_H(\mathcal{M}(\mathcal{X}_2))].$$

Let $(p_1, M_1, \sigma_1, w_1) := \mathcal{C}_H(\mathcal{M}(\mathcal{X}_1))$. Since $H = H_{B,1} = \mathbb{1}(M_1 < \partial_3 p_1)$ we see that

$$\begin{aligned} \nabla \cdot (A_H \nabla p_1) &= PV_1 - \frac{1}{2} \partial_3 (HM_1) \\ \iff \Delta p_1 + \frac{1}{2} \partial_3 \min_0 (M_1 - \partial_3 p_1) &= PV_1. \end{aligned} \quad (5.7)$$

Similarly, since $H_1 = H_{B,1}$ we have that

$$\begin{aligned} j_1 &= \partial_3 u_{h,1} - \nabla_h^\perp (\theta_1 - \min_0 q_1) \\ &= \partial_3 u_{h,1} - \nabla_h^\perp (\theta_1 - H_{B,1} q_1) \end{aligned}$$

and so σ_1 solves

$$\nabla \times (A_H^{-1} \sigma_1) = \partial_3 u_{h,1} - \nabla_h^\perp (\theta_1 - H_{B,1} q_1) + (\partial_3 w_1) e_3. \quad (5.8)$$

We deduce from (5.7) and (5.8) that p_1 and σ_1 are precisely characterised as in theorem 3.5 by

$$\mathcal{X}_1 = \Phi(p_1, M_1) + \Psi(\sigma_1, w_1).$$

In other words: the approximate inversion $\mathcal{C}_H \circ \mathcal{M}$ is *exact* when $H = H_1 = H_{B,1}$. Moreover proposition 5.7 tells us that $\tilde{E}_H$ agrees with E_H and with the nonlinear energy when $H = H_1 = H_{B,1}$ and so we conclude that indeed if $H = H_1 = H_{B,1} = H_2 = H_{B,2}$ then

$$\begin{aligned} \frac{1}{2}d_H(\mathcal{X}_1, \mathcal{X}_2)^2 &= E_H(\mathcal{X}_1 - \mathcal{X}_2) \\ &= \int |u_1 - u_2|^2 + (\theta_1 - \theta_2)^2 + (q_1 - q_2)^2 H + (M_1 - M_2)^2, \end{aligned}$$

where $H = \mathbb{1}(q_1 < 0) = \mathbb{1}(q_2 < 0)$.

Finally we show that d_H agrees with the extraction of a balanced component. We fix $\mathcal{X} \in \mathbb{L}_\sigma^2$ and let $\mathcal{Y} \in \mathcal{B}$. Since $\mathcal{Y} \in \mathcal{B}$ we know from proposition 3.1 that

$$j(\mathcal{Y}) = 0, w(\mathcal{Y}) = 0, \text{ and } a(\mathcal{Y}) = 0.$$

Let us define

$$M := M(\mathcal{X}) - M(\mathcal{Y}) \text{ and } PV := PV(\mathcal{X}) - PV(\mathcal{Y}), \quad (5.9)$$

and let p be the solution of

$$\nabla \cdot (A_H \nabla p) = PV - \frac{1}{2} \partial_3 (HM) \quad (5.10)$$

while σ is the solution of

$$\nabla \times (A_H^{-1} \sigma) = j(\mathcal{X}) + [\partial_3 w(\mathcal{X})] e_3 \text{ subject to } \oint A_H^{-1} \sigma = a(\mathcal{X}) \text{ and } \nabla \cdot \sigma = 0.$$

Then we may write the distance between $\mathcal{X}$ and $\mathcal{Y}$ via

$$\frac{1}{2}d_H(\mathcal{X}, \mathcal{Y})^2 = \int Q_H^+(p) + \left(1 + \frac{H}{2}\right) M^2 + Q_H^-(\sigma) + w^2$$

for $Q_H^\pm$ as in definition 5.5. Crucially: σ and w only depend on $\mathcal{X}$! So we may simplify the minimisation problem as follows:

$$\arg \min_{\mathcal{Y} \in \mathcal{B}} d_H(\mathcal{X}, \mathcal{Y}) = \arg \min_{\mathcal{Y} \in \mathcal{B}} \int Q_H^+(p) + \left(1 + \frac{H}{2}\right) M^2.$$

In particular, since $Q_H^+(v) = A_H v \cdot v \geq \frac{1}{2} |v|^2$, i.e. since Q_H^+ is positive-definite, the unique minimiser corresponds to

$$p = 0 \text{ and } M = 0.$$

Plugging this into (5.9), and then into (5.10), this means that $M(\mathcal{X}) = M(\mathcal{Y})$, and hence that $PV(\mathcal{X}) = PV(\mathcal{Y})$. In other words the unique minimiser is the balanced state $\mathcal{Y}$ whose PV and M agree with $\mathcal{X}$. By lemma 5.4, this means precisely that $\mathcal{Y} = \mathcal{X}_B$, as desired. $\square$

This concludes our first attempt at quantifying how far apart two states are in a way that agrees with

1. the extraction of a balanced component and
2. the conserved energy.

We see that the first of these two requirements was met by the distance introduced in definition 5.10 above. However the second requirement was only met in the restricted setting where the fixed Heaviside agreed with the Heavisides arising from both states being compared.

In some sense, that only the first requirement is satisfied is not surprising. Indeed, the two distances introduced in this section are built on top of the dual quantities from the global change of coordinates of proposition 3.8 and these dual quantities have been chosen precisely to be compatible with the decomposition. However, these dual quantities have, *a priori*, little to do with the conserved energy and so it is actually quite positive that the distance built in this way is in any way comparable to the energy.

This motivates an alternate approach: what if, instead of using decomposition-amenable dual quantities as a starting point, we used the energy itself as a starting point to define a distance? This is the approach we take in the following section.

5.2. Energy-centric approach

Recall that we are after a way to measure how different two states are. Moreover we wish to do so in a manner which is compatible with both the balanced-unbalanced decomposition underpinning theorem 3.5 and the conserved energy. In the previous section we approached this problem with the first requirement in mind, and eventually discovered that this approach failed to (generically) satisfy the second requirement. In this section we proceed the other way around: we enforce the second requirement by putting the conserved energy front and centre in the construction of a new distance. The question will then be to identify whether or not this metric agrees with the decomposition.

Also, recall the difficulty that arises in trying to define a distance or metric or inner product that agrees with the energy in (5.1): the nonlinear switch from the minimum function or Heaviside function. This nonlinear switch is activated based on the total water q , and it is not clear how to define the nonlinear switch when two different total water variables, q_1 and q_2 , are involved in quantifying the distance between two different states. This difficulty and others are discussed in more detail in the text preceding section 5.1.

A key idea is helpful for overcoming the difficulties of the nonlinear switch: in defining a metric, begin by working *pointwise in space*. Then, depending on which phase that point belongs to, the conserved energy (5.1) determines a natural inner product to use.

More precisely, this section proceeds as follows. We begin by defining a metric *pointwise* which measures the discrepancy between two states $\mathcal{X}_1(x)$ and $\mathcal{X}_2(x)$ when evaluated at a specific point x of the domain $\mathbb{T}^3$. This pointwise metric is defined to agree with the conserved energy density and so it is defined piecewise, taking different forms depending on the sign of $q(x)$. We then record the properties of this pointwise metric: it is well-defined, a metric, and agrees with the energy unconditionally. We also record how to compute this distance explicitly in a special case. We then make the key observation of this section: this pointwise metric does *not* agree with the extraction of a balanced component, i.e. it does not agree with the decomposition. Finally we conclude this section by defining the induced metric on the state space $\mathbb{L}_\sigma^2$ obtained by integrating the pointwise metric over the domain and explain why we do not expect this metric to agree with the decomposition either.

As mentioned above, a key idea in the definition of this metric is to work *pointwise in space*. In particular, depending on which phase that point belongs to, the conserved energy (5.1) determines a natural inner product to use. We introduce appropriate notation with this in mind.

Definition 5.12. For any $S_1 = (u_1, \theta_1, q_1) \in \mathbb{R}^5$ and $S_2 = (u_2, \theta_2, q_2) \in \mathbb{R}^5$, where $u_1, u_2 \in \mathbb{R}^3$, we define the inner products

$$\langle S_1, S_2 \rangle_u := u_1 \cdot u_2 + \theta_1 \theta_2 + q_1 q_2 + (\theta_1 + q_1)(\theta_2 + q_2)$$

and

$$\langle S_1, S_2 \rangle_s := u_1 \cdot u_2 + \theta_1 \theta_2 + (\theta_1 + q_1)(\theta_2 + q_2).$$

We denote by $\|\cdot\|_u$ and $\|\cdot\|_s$ the corresponding norms.

The inner products introduced in definition 5.12 are chosen precisely so that the ensuing norms $\|\cdot\|_u$ and $\|\cdot\|_s$ agree with the conserved energy density. Indeed, when $q < 0$ we have that

$$\|(u, \theta, q)\|_u^2 = |u|^2 + \theta^2 + q^2 + (\theta + q)^2$$

and when $q \geq 0$ we have that

$$\|(u, \theta, q)\|_s^2 = |u|^2 + \theta^2 + (\theta + q)^2,$$

such that indeed they agree with the conserved energy density as recorded in (5.1).

Moreover these inner products admit orthonormal bases which relate naturally to the moist variable M and the buoyancy b . This is depicted and explained in figure 1.

Next we wish to define the *pointwise* metric. As alluded to at the start of this section, this metric will be defined *piecewise* (depending on the sign of q) and so we introduce notation which describes each of these pieces.

Definition 5.13. We consider the following subsets of $\mathbb{R}^5$.

- $\mathcal{U} := \{(u, \theta, q) \in \mathbb{R}^5 : q < 0\}$ is the set of unsaturated states,
- $\mathcal{S} := \{(u, \theta, q) \in \mathbb{R}^5 : q \geq 0\}$ is the set of saturated states, and
- $\mathcal{I} := \{(u, \theta, q) \in \mathbb{R}^5 : q = 0\}$ is the state-space interface between unsaturated states and saturated states.

We have the necessary notation in hand to define the following *Snell-type* metric. We refer to this metric as being of Snell-type since shortest paths between states in different phases exhibit the piecewise affine behaviour observed by light passing across a boundary (see also figure 2).

Definition 5.14 (snell-type metric, pointwise). We define $d_{\text{Snell}} : \mathbb{R}^5 \rightarrow \mathbb{R}^5 \rightarrow [0, \infty)$ as follows.

- If $S_1, S_2 \in \mathcal{U}$ then $d_{\text{Snell}}(S_1, S_2) := \|S_1 - S_2\|_u$.
- If $S_1, S_2 \in \mathcal{S}$ then $d_{\text{Snell}}(S_1, S_2) := \|S_1 - S_2\|_s$.
- If $S_1 \in \mathcal{S}$ and $S_2 \in \mathcal{U}$ then

$$d_{\text{Snell}} := \min_{I \in \mathcal{I}} \|S_1 - I\|_s + \|S_2 - I\|_u.$$

- If $S_1 \in \mathcal{U}$ and $S_2 \in \mathcal{S}$ then $d_{\text{Snell}}(S_1, S_2) := d_{\text{Snell}}(S_2, S_1)$.



First we verify that this distance function is well-defined.

Lemma 5.15 (the Snell-type distance is well-defined). For any $S_1 \in \mathcal{S}$ and any $S_2 \in \mathcal{U}$ there is a unique minimiser of

$$\min_{I \in \mathcal{I}} \|S_1 - I\|_s + \|S_2 - I\|_u.$$

Proof. To verify that d_{Snell} is well-defined it suffices to verify that, for any $S_1 \in \mathcal{S}$ and any $S_2 \in \mathcal{U}$, there exists a unique minimiser of

$$\min_{I \in \mathcal{I}} \|S_1 - I\|_s + \|S_2 - I\|_u. \quad (5.11)$$



To do so, fix $S \in \mathbb{R}^5 \setminus \mathcal{I}$ and consider $f: \mathbb{R}^4 \rightarrow \mathbb{R}^5$ defined by, for every $I = (u_h, w, \theta, 0) \in \mathcal{I}$,

$$f(u_h, w, \theta) = \|S - I\|$$

where $\|\cdot\|$ denotes $\|\cdot\|_s$ or $\|\cdot\|_u$. To verify that (5.11) has a unique minimiser it suffices to show that f is strictly convex. Since f is smooth (as S is not in $\mathcal{I}$, and f and its derivatives would blow up precisely when $S = I$, but $I \in \mathcal{I}$) it suffices to verify that

$$\nabla^2 f > 0.$$

So we compute:

$$\partial_i f = \partial_i \|S - I\| = \frac{\langle I - S, \partial_i I \rangle}{\|S - I\|}$$

and hence

$$\partial_i \partial_j f = \partial_j \left(\frac{\langle I - S, \partial_i I \rangle}{\|S - I\|} \right) = \frac{\langle \partial_i I, \partial_j I \rangle \|S - I\|^2 - \langle \partial_i I, S - I \rangle \langle \partial_j I, S - I \rangle}{\|S - I\|^3}.$$

To make the remaining computation easier to follow, let us write

$$x^i := \partial_i I \text{ and } z := S - I.$$

Then

$$\|z\|^3 \nabla^2 f = \langle x^i, x^j \rangle \|z\|^2 - \langle x^i, z \rangle \langle x^j, z \rangle =: A$$

and so it suffices to show that $A > 0$ (since $\|z\| \neq 0$ as long as $S \notin \mathcal{I}$). To simplify further, let us introduce

$$D := \text{diag}(\|x^i\|)_{i=1}^4.$$

Then

$$A > 0 \iff D^T A D > 0$$

where, since $(x^i)_{i=1}^4$ forms an orthogonal set with respect to both $\|\cdot\|_s$ and $\|\cdot\|_u$, we have that $D^T A D = I\|z\|^2 - B$ for

$$B_{ij} := \langle \hat{x}^i, z \rangle \langle \hat{x}^j, z \rangle.$$

To conclude the proof it thus suffices to prove that $B < \|z\|^2 I$. First we show that $B \leq \|z\|^2 I$. Well, for any $y \in \mathbb{R}^4$,

$$By \cdot y = y_i \langle \hat{x}^i, z \rangle y_j \langle \hat{x}^j, z \rangle = y \cdot (\langle \hat{x}^i, z \rangle)_{i=1}^4 \leq |y|^2 \left(\sum_{i=1}^4 \langle \hat{x}^i, z \rangle^2 \right).$$

Since $(\hat{x}^i)_{i=1}^4$ is an orthonormal set, we have that

$$\sum_{i=1}^4 \langle \hat{x}^i, z \rangle^2 \leq \|z\|^2,$$

which verifies that $B \leq \|z\|^2 I$. Finally we show that $B < \|z\|^2 I$. This follows from the fact that $(\hat{x}^i)_{i=1}^4$ span $\mathcal{I}$, but not $\mathbb{R}^5$, and so

$$\sum_{i=1}^4 \langle \hat{x}^i, z \rangle^2 = \|\text{proj}_{\mathcal{I}} z\|^2 < \|z\|^2$$

since $z = S - I \notin \mathcal{I}$ since $I \in \mathcal{I}$ but $S \notin \mathcal{I}$. In other words: since $S \notin \mathcal{I}$ we have that $B < \|z\|^2 I$, and so $\nabla^2 f > 0$. □


Now we verify that the distance function introduced in definition 5.14 is indeed a metric. We will also show that this Snell-type distance also agrees with the energy, when they are comparable.

Lemma 5.16 (the snell-type distance is a metric). *As introduced in definition 5.14, d_{Snell} is a metric on $\mathbb{R}^5$. Moreover it agrees with the energy density in the sense that, for any $S_1, S_2 \in \mathbb{R}^5$, if the signs of q_1 and q_2 agree (i.e. if the states S_1 and S_2 are in the same phase) then*

$$d_{\text{Snell}}(S_1, S_2)^2 = |u_1 - u_2|^2 + (\theta_1 - \theta_2)^2 + (q_1 - q_2)^2 h + ((\theta_1 + q_1) - (\theta_2 + q_2))^2$$

where

$$h = \begin{cases} 1 & \text{if } q_1, q_2 < 0 \text{ and} \\ 0 & \text{if } q_1, q_2 \geq 0. \end{cases}$$



Proof. Positive-definiteness and symmetry are immediate so we only need to prove that the triangle inequality holds. Indeed: the agreement with the energy follows immediately from the definition of $\|\cdot\|_u$ and $\|\cdot\|_s$ in definition 5.12.

For $S_1, S_2, S_3 \in \mathbb{R}^5$ we therefore wish to show that

$$d_{\text{Snell}}(S_1, S_2) \leq d_{\text{Snell}}(S_1, S_3) + d_{\text{Snell}}(S_3, S_2).$$

Clearly the triangle inequality holds when all three states belong to the same phase (either $\mathcal{U}$ or $\mathcal{S}$). So we need only consider the case where two of the states S_1, S_2 , and S_3 are in one phase and the last state is in the other. Since all that matters is whether or not the intermediate point S_3 is alone in its phase, there are without loss of generality two cases to consider.

- Case 1: $S_1, S_2 \in \mathcal{S}$ and $S_3 \in \mathcal{U}$. Let us denote by I_{13} the point intersecting the geodesic between S_1 and S_3 and the state-space interface $\mathcal{I}$, and similarly I_{23} for the point arising from the geodesic between S_2 and S_3 . We use the triangle inequality in $\mathcal{U}$, then the triangle inequality in $\mathcal{S}$:

$$\begin{aligned} d_{\text{Snell}}(S_1, S_3) + d_{\text{Snell}}(S_3, S_2) &= d_{\text{Snell}}(S_1, I_{13}) + d_{\text{Snell}}(I_{13}, S_3) + d_{\text{Snell}}(S_3, I_{23}) + d_{\text{Snell}}(I_{23}, S_2) \\ &\geq d_{\text{Snell}}(S_1, I_{13}) + d_{\text{Snell}}(I_{13}, I_{23}) + d_{\text{Snell}}(I_{23}, S_2) \\ &\geq d_{\text{Snell}}(S_1, S_2). \end{aligned}$$

This chain of inequalities is depicted pictorially in figure 3.

- Case 2: $S_1, S_3 \in \mathcal{S}$ and $S_2 \in \mathcal{U}$. Let I_{12} and I_{23} be defined similarly to in case 1. We use the triangle inequality in $\mathcal{S}$ and the definition of d_{Snell} as a minimisation problem when crossing the interface to deduce that

$$\begin{aligned} d_{\text{Snell}}(S_1, S_3) + d_{\text{Snell}}(S_3, S_2) &= d_{\text{Snell}}(S_1, S_3) + d_{\text{Snell}}(S_3, I_{23}) + d_{\text{Snell}}(I_{23}, S_2) \\ &\geq d_{\text{Snell}}(S_1, I_{23}) + d_{\text{Snell}}(I_{23}, S_2) \\ &\geq d_{\text{Snell}}(S_1, S_2). \end{aligned}$$

This chain of inequalities is depicted pictorially in figure 4. □

We now seek to determine whether or not this metric, once ported over to measure the distance between spatially-varying states, agrees with the balanced-unbalanced decomposition of theorem 3.5.

To do so we will need to understand this Snell-type metric in the pointwise case (which we have discussed so far) in more detail. In particular we will rely on the explicit computation of geodesic trajectories in the case where $u_1 = u_2 = 0$. In that case, the states live in the (θ, q) -plane. We consider two states in that plane living on either side of the state-space interface $\mathcal{I}$, i.e. where one state is saturated and the other is not. The geodesic between these two points is then piecewise affine, changing direction when it reaches the interface. Crucially: in this simplified setting where $u_1 = u_2 = 0$ we can compute *explicitly* where the geodesic intersects the interface.




Figure 4. A pictorial depiction of the two steps used in case 2 of the proof of lemma 5.16.

Lemma 5.17 (explicit computation of the intersection point). For any $S_1 \in \mathcal{S}$ and any $S_2 \in \mathcal{U}$, if $u_1 = u_2 = 0$ then, for $I(\theta) := (0, 0, 0, \theta, 0) \in \mathcal{I} \subseteq \mathbb{R}^5$,

$$\arg \min_{\theta \in \mathbb{R}} \|S_1 - I(\theta)\|_s + \|S_2 - I(\theta)\|_u \in \left\{ \frac{\alpha - \sqrt{\delta}}{\beta}, \frac{\alpha + \sqrt{\delta}}{\beta} \right\}$$

for

$$\begin{aligned} \alpha &= 4q_1^2(2\theta_2 + q_2) - 6q_2^2(2\theta_1 + q_1), \\ \beta &= 8q_1^2 - 24q_2^2, \text{ and} \\ \delta &= 12q_1^2q_2^2(2\theta_1 + q_1 + 4\theta_2 + 2q_2)^2. \end{aligned}$$

Figure 5 depicts an example where the two possible intersection points are computed explicitly.

Proof. The key idea is that the first derivative condition for minimisers may be squared to produce a fourth-order polynomial equation. Solving this equation using symbolic computation software then produces the claim. More precisely, let us define $f: \mathbb{R} \rightarrow \mathbb{R}$ via

$$f(\theta) := \|S_1 - I(\theta)\|_s + \|S_2 - I(\theta)\|_u.$$

Then, proceeding as in the proof of lemma 5.15 we see that

$$f'(\theta) = \frac{\langle I - S_1, \partial_\theta I \rangle_s}{\|S_1 - I\|_s} + \frac{\langle I - S_2, \partial_\theta I \rangle_u}{\|S_2 - I\|_u}$$

and so, since $\partial_\theta I = e_\theta$, we see that if $f'(\theta) = 0$ then

$$\langle I - S_2, e_\theta \rangle_u^2 \|S_1 - I\|_s^2 = \langle I - S_1, e_\theta \rangle_s^2 \|S_2 - I\|_u^2. \quad (5.12)$$

Since

$$\langle I - S_i, e_\theta \rangle_u = \langle I - S_i, e_\theta \rangle_s = 2\theta - 2\theta_i - q_i$$

while

$$\begin{aligned} \|S_1 - I\|_s^2 &= (\theta - \theta_1)^2 + (\theta - \theta_1 - q_1)^2 \text{ and} \\ \|S_2 - I\|_u^2 &= \frac{1}{2}(\theta - \theta_2 + q_2)^2 + \frac{3}{2}(\theta - \theta_2 - q_2)^2, \end{aligned}$$

where we have used the identity

$$\theta^2 + q^2 h = (1 - h/2)(\theta - qh)^2 + (h/2)(\theta + q)^2,$$

the equation (5.12) comes down to a polynomial equation of order four for θ . We may solve this equation explicitly using symbolic computation software, thus producing the claim. $\square$




Figure 5. This shows an example where the two possible intersection points are computed from lemma 5.17. In practice, it is very straightforward to identify the correct intersection point: the lengths of both paths may readily be computed, and the intersection point corresponding to the shortest path is then the correct one.

We now discuss how to extract balanced components in this pointwise setting where the state space is $\mathbb{R}^5$ (and not $\mathbb{L}_\sigma^2$). The system of interest is

$$\begin{cases} \frac{du_h}{dt} + \frac{1}{\varepsilon} u_h^\perp = 0, & (5.13a) \\ \frac{dw}{dt} + \frac{1}{\varepsilon} (\theta - \min_0 q) = 0, & (5.13b) \\ \frac{d\theta}{dt} + \frac{1}{\varepsilon} w = 0, \text{ and} & (5.13c) \\ \frac{dq}{dt} - \frac{1}{\varepsilon} w = 0, & (5.13d) \end{cases}$$

which fits into our generic framework as it can be written in the form, for $S = (u_h, w, \theta, q)$, as

$$\frac{dS}{dt} + \frac{1}{\varepsilon} N(S) = 0 \text{ for } N(S) = \begin{pmatrix} u_h^\perp \\ \theta - \min_0 q \\ w \\ -w \end{pmatrix}.$$

As detailed in section 3, the first step in identifying an appropriate decomposition for this system is to study the kernel / zero set of N and the image of N . We see that the kernel of N may be characterised (akin to proposition 3.1) as a zero set of appropriate dual quantities or may be explicitly parameterised:

$$\begin{aligned} N(S) = 0 &\iff S \in \ker N =: B \\ &\iff u_h = 0, w = 0, \text{ and } b(S) := \theta - \min_0 q = 0 \\ &\iff S = \begin{pmatrix} 0 \\ 0 \\ \frac{1}{2} \min_0 M \\ M - \frac{1}{2} \min_0 M \end{pmatrix} \end{aligned}$$

while, akin to proposition 3.3, the image of N may be characterised as follows:

$$S \in \text{im } N =: W \iff M(S) := \theta + q = 0 \iff S = \begin{pmatrix} u_h \\ w \\ \eta \\ -\eta \end{pmatrix}.$$

This leads to the following decomposition of $\mathbb{R}^5$:

$$S = \begin{pmatrix} u_h \\ w \\ \theta \\ q \end{pmatrix} = \underbrace{\begin{pmatrix} 0 \\ 0 \\ \frac{1}{2} \min_0 M \\ M - \frac{1}{2} \min_0 M \end{pmatrix}}_{\in \ker N = B} + \underbrace{\begin{pmatrix} u_h \\ w \\ \eta \\ -\eta \end{pmatrix}}_{\in \text{im } N = W}. \quad (5.14)$$

Now let us restrict our attention once again to the (θ, q) -plane. We seek to determine whether or not the Snell-type metric agrees with the decomposition (5.14). In the (θ, q) -plane, the balanced component takes the form

$$\begin{pmatrix} \theta_B \\ q_B \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \min_0 M \\ M - \frac{1}{2} \min_0 M \end{pmatrix}.$$

It turns out that this decomposition *fails to agree* with the Snell-type energy.




Figure 6. A graphical depiction of proposition 5.18. Note that the location of the points is not exact: this is an artistic rendering of the setup of the proof of proposition 5.18 used because it makes it easier to see what is going on.

Proposition 5.18 (the snell-type distance does not agree with the pointwise decomposition). As introduced in definition 5.14, d_{Snell} does not agree with the extraction of a balanced component, in the sense that there exists $S \in \mathbb{R}^5$ and $T_* \in B$ such that

$$d_{\text{Snell}}(S, T_*) < d_{\text{Snell}}(S, S_B).$$

In other words:

$$S_B \neq \arg \min_{T \in B} d_{\text{Snell}}(S, T).$$

A simple scenario where this disagreement occurs is depicted in figure 6.

Proof. Choose $S = (-3, 1)$ such that $M = -3 + 1 = -2$ and so its balanced component is $S_B = (-1, -1)$ and choose $T_* = (-1.1, -1.1)$. Note that S is saturated whereas S_B and T_* are unsaturated. And, of course, both S_B and T_* belong to the balanced set $B = \ker N$ since they have vanishing M .

On the one hand, equipped with lemma 5.17 we see that

$$d_{\text{Snell}}(S, S_B) = \|S - I_1\|_s + \|S_B - I_1\|_u$$

where

$$I_1 := I(\theta_1) \text{ for } \theta_1 = -3 + \frac{\sqrt{3}}{2}$$

and hence

$$d_{\text{Snell}}(S, S_B) = \sqrt{\frac{5}{2} - \sqrt{3}} (1 + \sqrt{3}) > 2.394.$$

On the other hand, for $I_2 := I(\theta_2)$ where $\theta_2 = \frac{11}{526}(-135 + 17\sqrt{3})$,

$$d_{\text{Snell}}(S, T_*) = \|S - I_2\|_s + \|T_* - I_2\|_u = \left(\sqrt{101488 - 31790\sqrt{3}} (10 + 11\sqrt{3}) \right) / 2630 < 2.381.$$

Since

$$d_{\text{Snell}}(S, T_*) < d_{\text{Snell}}(S, S_B)$$

this proves the claim. $\square$

We conclude this section by extending the Snell-type distance discussed so far *pointwise* to a distance between spatially-varying states.

Definition 5.19 (Snell-type metric, spatially-varying states). We define $d_S : \mathbb{L}_\sigma^2 \times \mathbb{L}_\sigma^2 \rightarrow [0, \infty)$ as follows. For any $\mathcal{X}, \mathcal{Y} \in \mathbb{L}_\sigma^2$,

$$d_S(\mathcal{X}, \mathcal{Y}) := \|d_{\text{Snell}}(\mathcal{X}, \mathcal{Y})\|_{L^2} = \left(\int d_{\text{Snell}}^2(\mathcal{X}(x), \mathcal{Y}(x)) dx \right)^{1/2}.$$

This distance is indeed a metric, and it agrees with the conserved energy when comparable to it.

Proposition 5.20. As introduced in definition 5.19, d_S is a metric on $\mathbb{L}_\sigma^2$ which agrees with the energy in the sense that, if $\mathcal{X}_1, \mathcal{X}_2 \in \mathbb{L}_\sigma^2$ have the same phases, meaning that the signs of q_1 and q_2 agree everywhere, then

$$\begin{aligned} (d_S(\mathcal{X}_1, \mathcal{X}_2))^2 &= E_H(\mathcal{X}_1 - \mathcal{X}_2) \\ &= \int |u_1 - u_2|^2 + (\theta_1 - \theta_2)^2 + (q_1 - q_2)^2 H + ((\theta_1 + q_1) - (\theta_2 + q_2))^2 \end{aligned}$$

for $H := \mathbb{1}(q_1 < 0) = \mathbb{1}(q_2 < 0)$.

Proof. First we verify that d_S is a metric. The positive-definiteness and symmetry follow immediately from the fact that d_{Snell} is a metric, since the L^2 norm is also positive-definite.

The triangle inequality of d_S then follows from the triangle inequalities of d_{Snell} and of L^2 and from the fact that $s \mapsto \sqrt{s}$ and $s \mapsto s^2$ are increasing:

$$\begin{aligned} d_S(\mathcal{X}_1, \mathcal{X}_2) &= \left(\int d_{\text{Snell}}^2(\mathcal{X}_1, \mathcal{X}_2) \right)^{1/2} \\ &\leq \left(\int (d_{\text{Snell}}(\mathcal{X}_1, \mathcal{X}_3) + d_{\text{Snell}}(\mathcal{X}_3, \mathcal{X}_2))^2 \right)^{1/2} \\ &= \|d_{\text{Snell}}(\mathcal{X}_1, \mathcal{X}_3) + d_{\text{Snell}}(\mathcal{X}_3, \mathcal{X}_2)\|_{L^2} \\ &\leq \|d_{\text{Snell}}(\mathcal{X}_1, \mathcal{X}_3)\|_{L^2} + \|d_{\text{Snell}}(\mathcal{X}_3, \mathcal{X}_2)\|_{L^2} \\ &= d_S(\mathcal{X}_1, \mathcal{X}_3) + d_S(\mathcal{X}_3, \mathcal{X}_2). \end{aligned}$$

Finally we prove that d_S agrees with the energy, when comparable. So suppose that $q_1(x)$ and $q_2(x)$ have the same phases, i.e. their signs agree everywhere in the spatial domain. Then, by virtue of lemma 5.16, we have that

$$\begin{aligned} d_S^2(\mathcal{X}_1, \mathcal{X}_2) &= \int d_{\text{Snell}}^2(\mathcal{X}_1(x), \mathcal{X}_2(x)) dx \\ &= \int |u_1 - u_2|^2 + (\theta_1 - \theta_2)^2 + (q_1 - q_2)^2 H + ((\theta_1 + q_1) - (\theta_2 + q_2))^2 \\ &= E_H(\mathcal{X}_1 - \mathcal{X}_2) \end{aligned}$$

for $H := \mathbb{1}(q_1 < 0) = \mathbb{1}(q_2 < 0)$. □

Remark 5.21. Finally we note that, although this distance agrees with the energy, the evidence from the restricted case of spatially constant cases strongly suggests that this Snell-type distance does *not* agree with the extraction of a balanced component, especially in light of lemma C.23 which tells us that spatially constant solutions are honest-to-goodness solutions of the moist Boussinesq system considered here.

While this Snell-type distance may not agree with the balanced–unbalanced decomposition, it could be useful for other applications such as data-driven methods that identify nonlinear features in dynamical systems and utilise a metric, distance function, and/or kernel function [CL06, GM12, BH16, DG19, Gia19].

5.3. Metric-less approach

To make the approach of this section easy to follow we must be clear about what the last two sections, sections 5.1 and 5.2, achieved.

- In section 5.1 we devised a metric which agreed with the decomposition, but which failed to agree with the energy (it agreed only *conditionally*).

- In section 5.2 we devised a metric which agreed with the energy, but which failed to agree with the decomposition.

In this section we will see that it is possible to agree with both the decomposition and the energy, at a price. That price is that we lose the triangle inequality. Instead of a metric, we will therefore be dealing with a statistical divergence.

Definition 5.22 (statistical divergence). Let $\mathcal{H}$ be a Hilbert space. A *statistical divergence* on $\mathcal{H}$ is a map $D : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$ satisfying the following properties.

1. Non-negative: $D(x, y) \geq 0$ for every $x, y \in \mathcal{H}$.
2. Definite: $D(x, y) = 0$ if and only if $x = y$.
3. Limiting quadratic: for every $x \in \mathcal{H}$ there is a positive-definite quadratic form $\mathcal{Q}_x$ such that, for every $y \in \mathcal{H}$,

$$\frac{1}{\varepsilon^2} D(x, x + \varepsilon y) \rightarrow \mathcal{Q}_x(y)$$

as $\varepsilon \downarrow 0$. Since this quadratic form gives rise to a positive-definite inner product, $\mathcal{Q}_x$ is often referred to as the *Riemann metric induced by the divergence*.

The prototypical example of a statistical divergence is the relative entropy, or Kullback–Leibler divergence, from information theory [CT12, BM12, CMT14]. Another reminiscent notion from the PDE literature is the *relative energy* [Daf79, FJN12, FKNZ16, BFH17].

The statistical divergence here is built in a simple way: we use the metric d_H introduced in section 5.1, but now let H vary depending on the states whose divergence is being measured.

Definition 5.23. For any $\mathcal{X}_1, \mathcal{X}_2 \in \mathbb{L}_\sigma^2$ we define

$$D(\mathcal{X}_1; \mathcal{X}_2) := \frac{1}{2} d_{H_{B_1}}(\mathcal{X}_1, \mathcal{X}_2)^2$$

where d_H is as in definition 5.10, where $H_{B_1} := \mathbb{1}(q_{B,1} < 0)$ denotes the *balance* interface of the first state $\mathcal{X}_1$.

Remark 5.24. We detail how to compute the statistical divergence D introduced in definition 5.23 (similarly to how we detailed the computation of the metric d_H in definition 5.10). So let $\mathcal{X}_1 = (u_1, \theta_1, q_1)$ and $\mathcal{X}_2 = (u_2, \theta_2, q_2)$ be states in $\mathbb{L}_\sigma^2$.

First we compute $\mathcal{X}_{B,1}$ in order to determine the indicator H_{B_1} . This is done by setting

$$M_1 := M(\mathcal{X}_1) = \theta_1 + q_1 \text{ and } PV_1 = PV(\mathcal{X}_1) = \nabla_h^\perp \cdot u_{h,1} + \partial_3 \theta_1$$

and finding the solution $p_1 \in \mathring{H}^1$ of

$$\Delta p_1 + \frac{1}{2} \partial_3 \min_0 (M_1 - \partial_3 p_1) = PV_1.$$

The phases of the balanced component of the first state are then encoded in

$$H_{B_1} := \mathbb{1}(q_{B,1} < 0) = \mathbb{1}(\partial_3 p_1 < M_1).$$

Second we proceed as in definition 5.10 and compute $d_{H_{B_1}}(\mathcal{X}_1, \mathcal{X}_2)$. This means that we let

$$PV := PV(\mathcal{X}_1) - PV(\mathcal{X}_2), M := M(\mathcal{X}_1) - M(\mathcal{X}_2), \\ j := j(\mathcal{X}_1) - j(\mathcal{X}_2), w := w(\mathcal{X}_1) - w(\mathcal{X}_2), \text{ and } a := a(\mathcal{X}_1) - a(\mathcal{X}_2),$$

and that we let $p \in \mathring{H}^1$ and $\sigma \in (L^2)_\sigma^3$ solve

$$\nabla \cdot (A_{H_{B_1}} \nabla p) = PV - \frac{1}{2} \partial_3 (HM)$$

and

$$\nabla \times (A_{H_{B_1}}^{-1} \sigma) = j + (\partial_3 w) e_3 \text{ subject to } \oint A_{H_{B_1}}^{-1} \sigma = a \text{ and } \nabla \cdot \sigma = 0$$

for A and A^{-1} as in definition 5.5.

So finally we have that

$$D(\mathcal{X}_1, \mathcal{X}_2) = \frac{1}{2} d_{H_{B_1}}(\mathcal{X}_1, \mathcal{X}_2)^2 = \int Q_{H_{B_1}}^+(p) + \left(1 + \frac{H_{B_1}}{2}\right) M^2 + Q_{H_{B_1}}^-(\sigma) + w^2$$

for $Q^\pm$ as introduced in definition 5.5. We note that, in the more concise notation of definition 5.10 we may write this as

$$D(\mathcal{X}_1, \mathcal{X}_2) = \tilde{E}_{H_{B_1}}(\mathcal{C}_{H_{B_1}}[\mathcal{M}(\mathcal{X}_1) - \mathcal{M}(\mathcal{X}_2)]).$$

Remark 5.25. We saw in remark 5.24 that we may write the statistical divergence in terms of $\mathcal{C}_H$. This map is an analogue of $\mathcal{C}$, which is the map from dual quantity space to coordinate space which associates to a state's dual quantities its corresponding coordinates (see remark 3.10). The analogous map $\mathcal{C}_H$ proceeds similarly, except that now the three transformations (between state space, dual quantity space, and coordinate space) are *linear*. The linearity comes from treating a particular indicator function H as fixed.

More precisely, given an indicator function H , we have the following.

- The ‘fixed H ’ coordinate-to-state map $\mathcal{S}_H : \mathfrak{C} \rightarrow \mathbb{L}_\sigma^2$ is given by

$$\mathcal{S}_H(p, M, \sigma, w) = (u, \theta, q) \iff \begin{cases} u = \nabla_h^\perp p + \sigma_h^\perp \\ \theta = \partial_3 p + \frac{1}{2} H(M - \partial_3 p) + \sigma_3 \\ q = M - \partial_3 p - \frac{1}{2} H(M - \partial_3 p) - \sigma_3, \end{cases}$$

where recall that $\mathfrak{C}$ defined in remark 3.10 is the coordinate space.

- The ‘fixed H ’ state-to-dual quantity map $\mathcal{M}_H : \mathbb{L}_\sigma^2 \rightarrow \mathfrak{M}$, where $\mathfrak{M}$ denotes the dual quantity space as in proposition 3.8, is given by

$$\mathcal{M}_H(u, \theta, q) = (PV, M, j_H, w, a_H)$$

where

$$j_H(u, \theta, q) = \partial_3 u_h - \nabla_h^\perp(\theta - qH) \text{ and } a_H = -\oint u_h^\perp + \oint(\theta - qH)e_3.$$

This means that only j and a are modified into j_H and a_H , respectively, but the remaining dual quantities are unchanged. This is because they are the only nonlinear dual quantities in the first place since they are the only dual quantities involving the buoyancy.

- The ‘fixed H ’ dual quantity-to-coordinate map $\mathcal{C}_H : \mathfrak{M} \rightarrow \mathfrak{C}$ is as defined in lemma 5.8.

These maps will all be used below to compute the Riemann metric induced by the statistical divergence introduced in definition 5.23.

Proposition 5.26. *The map $D : \mathbb{L}_\sigma^2 \times \mathbb{L}_\sigma^2 \rightarrow [0, \infty)$ introduced in definition 5.23 is a statistical divergence in the sense of definition 5.22. In particular, for any $\mathcal{X}, \mathcal{Y} \in \mathbb{L}_\sigma^2$,*

$$\frac{1}{\varepsilon^2} D(\mathcal{X}; \mathcal{X} + \varepsilon \mathcal{Y}) \rightarrow E_B([\mathcal{Y}]_{\mathcal{X}})$$

where $[\mathcal{Y}]_{\mathcal{X}} := (\mathcal{S}_B \circ \mathcal{C}_B \circ \mathcal{M}_H)(\mathcal{Y})$ is known as the $\mathcal{X}$ -representative of $\mathcal{Y}$ (such that if $H_{\mathcal{X}} = H_{\mathcal{Y}}$ and $H_{B, \mathcal{X}} = H_{B, \mathcal{Y}}$ then $[\mathcal{Y}]_{\mathcal{X}} = \mathcal{Y}$). Moreover

1. this statistical divergence agrees with the conserved energy in the sense that, if $H_1 = H_2 = H_{B,1} = H_{B,2}$ then

$$D(\mathcal{X}_1; \mathcal{X}_2) = E_H(\mathcal{X}_1 - \mathcal{X}_2)$$

where $H = \mathbb{1}(q_1 < 0) = \mathbb{1}(q_2 < 0)$ and

2. this statistical divergence agrees with the extraction of a balanced component in the sense that

$$\mathcal{X}_B = \arg \min_{\mathcal{Y} \in \mathcal{B}} D(\mathcal{X}; \mathcal{Y}).$$

Proof. The fact that D is non-negative and definite follows immediately from the positive-definiteness of d_H , which is established in proposition 5.11. The fact that D agrees with the energy and is compatible with the extraction of a balanced component also follows from the proof of proposition 5.11. So finally we turn our attention to the Riemann metric induced by this statistical divergence. Let $\mathcal{X}, \mathcal{Y} \in \mathbb{L}_\sigma^2$, which we write as $\mathcal{X} = (u, \theta, q)$ and $\mathcal{Y} = (v, \phi, r)$, and let $\varepsilon > 0$.

We note that

$$\begin{aligned} j(\mathcal{X} + \varepsilon\mathcal{Y}) - j(\mathcal{X}) &= \partial_3(u_h + \varepsilon v_h) - \nabla_h^\perp(\theta + \varepsilon\phi) + \nabla_h^\perp(\min_0(q + \varepsilon r)) \\ &\quad - \partial_3 u_h + \nabla_h^\perp \theta - \nabla_h^\perp \min_0 q \\ &= \varepsilon \partial_3 v_h - \varepsilon \nabla_h^\perp \phi + \nabla_h^\perp [\min_0(q + \varepsilon r) - \min_0 q] \\ &= \varepsilon \partial_3 v_h - \varepsilon \nabla_h^\perp \phi + \nabla_h^\perp [\varepsilon r \mathbb{1}(q + \varepsilon r < 0) + R_\varepsilon(q; r)] \\ &= \varepsilon j_{H(\mathcal{X} + \varepsilon\mathcal{Y})}(\mathcal{Y}) + \nabla_h^\perp R_\varepsilon(q; r) \end{aligned}$$

for

$$\begin{aligned} R_\varepsilon(q; r) &:= q[\mathbb{1}(q + \varepsilon r < 0) - \mathbb{1}(q < 0)] \\ &= \varepsilon \cdot \underbrace{\frac{q}{\varepsilon} \left[\mathbb{1}\left(\frac{q}{\varepsilon} + r < 0\right) - \mathbb{1}\left(\frac{q}{\varepsilon} < 0\right) \right]}_{\rightarrow q(\delta_0 \circ q)r} \end{aligned}$$

where δ_0 denotes the Dirac delta distribution at zero such that, since $q(\delta \circ q) = 0$, $\frac{1}{\varepsilon} R_\varepsilon(q; r) \rightarrow 0$ as $\varepsilon \downarrow 0$.

Therefore, since all other dual quantities in $\mathcal{M}$ besides j are linear,

$$\begin{aligned} \mathcal{M}(\mathcal{X} + \varepsilon\mathcal{Y}) - \mathcal{M}(\mathcal{X}) &= \begin{pmatrix} \varepsilon PV(\mathcal{Y}) \\ \varepsilon M(\mathcal{Y}) \\ \varepsilon j_{H(\mathcal{X} + \varepsilon\mathcal{Y})}(\mathcal{Y}) + \nabla_h^\perp R_\varepsilon(q; r) \\ \varepsilon w(\mathcal{Y}) \\ \varepsilon a(\mathcal{Y}) \end{pmatrix} = \varepsilon \begin{pmatrix} PV(\mathcal{Y}) \\ M(\mathcal{Y}) \\ j_{H(\mathcal{X} + \varepsilon\mathcal{Y})}(\mathcal{Y}) \\ w(\mathcal{Y}) \\ a(\mathcal{Y}) \end{pmatrix} + \varepsilon \begin{pmatrix} 0 \\ 0 \\ \nabla_h^\perp \frac{1}{\varepsilon} R \\ 0 \\ 0 \end{pmatrix} \\ &=: \varepsilon \mathcal{M}_{H_\varepsilon}(\mathcal{Y}) + \varepsilon \mathcal{R}, \end{aligned}$$

where as noted above $\mathcal{R} \rightarrow 0$ when $\varepsilon \downarrow 0$.

So finally we have that

$$\begin{aligned} D(\mathcal{X}; \mathcal{X} + \varepsilon\mathcal{Y}) &= E_B(\mathcal{C}_B[\mathcal{M}(\mathcal{X} + \varepsilon\mathcal{Y}) - \mathcal{M}(\mathcal{X})]) \\ &= E_B(\mathcal{C}_B[\varepsilon \mathcal{M}_{H_\varepsilon}(\mathcal{Y}) + \varepsilon \mathcal{R}]) \\ &= \varepsilon^2 E_B([Y]_{\mathcal{X} + \varepsilon\mathcal{Y}} + \mathcal{C}_B \mathcal{R}) \end{aligned}$$

such that indeed, since $H(\mathcal{X} + \varepsilon\mathcal{Y}) \rightarrow H(\mathcal{X}) = H$,

$$\frac{1}{\varepsilon^2} D(\mathcal{X}; \mathcal{X} + \varepsilon\mathcal{Y}) \rightarrow E_B([Y]_{\mathcal{X}}),$$

as claimed. □

5.4. Summary

The takeaway message from this section is the following: if we wish to quantify the discrepancy between two states in $\mathbb{L}_\sigma^2$ in a manner which is both consistent with the decomposition and with the conserved energy then we must give up on the triangle inequality. As sections 5.1 and 5.2 show, we are unable to obtain a metric which has both of these properties. We therefore identify, in section 5.3, a *statistical divergence* which fails to satisfy the triangle inequality but which is consistent with both the decomposition and the conserved energy. This is summarised in the table below.

Approach	Agrees with the decomposition	Agrees with the energy	Satisfies the triangle inequality
Parseval-type metric	Yes	No	Yes
Snell-type metric	No	Yes	Yes
Statistical divergence	Yes	Yes	No

6. Iterative methods for the decomposition and their geometry

In this section we record the key results that enable us to compute the decomposition of a state into its slow and fast components in *practice*. The details are relegated to section A.

The only computationally non-trivial step in that process is to extract the balanced component, since this requires the inversion of the nonlinear elliptic PDE

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV$$

where PV and M are given and we solve for p .

Our ability to solve this nonlinear elliptic PDE comes from identifying a *variational form* of this equation. That is the observation underpinning [RTSS24]. We recall this variational formulation here.

Theorem 6.1 (variational formulation of PV-and-M inversion). *For any $M \in L^2$, $PV \in H^{-1}$, and $p \in \dot{H}^1$, p is a global minimiser of E_{var} , defined as*

$$E_{\text{var}} := \int_{\mathbb{T}^3} \frac{1}{2} |\nabla p|^2 - \frac{1}{4} \min_0 (M - \partial_3 p)^2 + \langle PV, p \rangle_{H^{-1} \times \dot{H}^1},$$

if and only if p is an H^1 -weak solution of

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV$$

in the sense that

$$\int_{\mathbb{T}^3} \nabla p \cdot \nabla \phi + \frac{1}{2} \min_0 (M - \partial_3 p) \partial_3 \phi = -\langle PV, \phi \rangle_{H^{-1} \times \dot{H}^1} \text{ for every } \phi \in \dot{H}^1. \quad (6.1)$$

Proof. See lemma 4.9 in [RTSS24]. □

With this variational formulation of PV-and-M inversion in hand, we can now leverage iterative methods from optimisation in order to find computational methods.

These methods are discussed in detail in section A: in section A.1 we discuss convergence guarantees for these methods while in section A.2 we relate one of these iterative methods (arising from Newton's methods) to the geometry of the balanced set $\mathcal{B}_0$, obtaining a geometric interpretation of this descent method. We leave the details for the interested reader in section A. Here we only record the key results and observations supported by section A.

The first key result is a convergence guarantee. We provide an informal statement here. The full result, including a precise statement, is available in section A.1 as theorem A.19.

Theorem 6.2 (convergence rates for our variational energy). *Let E_{var} be the variational energy recorded in theorem 6.1 and let p^* denote the unique minimiser of E_{var} in $\dot{H}^1$. For every $p \in \dot{H}^1$ there exists a descent direction $\pi \in \dot{H}^1$ such that*

$$E_{\text{var}}(p + \pi) - E_{\text{var}}(p^*) \leq (1 - \gamma) (E_{\text{var}}(p) - E_{\text{var}}(p^*))$$

for some $\gamma \in (0, 1)$ independent of p .

We actually devise *several* iterative methods that rely on distinct descent direction, and thus produce different convergence rates (this is recorded precisely in Theorem A.19).

One such descent direction is known as the *Newton descent method*. It is defined discussed in more detail in section A.1, but for our purposes here it suffices to know that the Newton descent method π is chosen to be $\pi \sim (\nabla^2 E_{\text{var}})^{-1} \nabla E_{\text{var}}$, which is done precisely as follows.

Proposition 6.3. Let E_{var} be the variational energy introduced in theorem 6.1. For any $p \in \dot{H}^1$, $PV \in H^{-1}$, and $M \in L^2$, the Newton descent direction π is the $\dot{H}^1$ -weak solution of

$$-\nabla \cdot \left[\left(I - \frac{1}{2} \mathbf{1}(M < \partial_3 p) e_3 \otimes e_3 \right) \nabla \pi \right] = \Delta p + \frac{1}{2} \partial_3 (\min(M - \partial_3 p, 0)) - PV. \quad (6.2)$$

In other words, for every $\phi \in \dot{H}^1$,

$$\int_{\mathbb{T}^3} \nabla \pi \cdot \nabla \phi - \frac{1}{2} \mathbf{1}(M < \partial_3 p) (\partial_3 \pi) \partial_3 \phi \quad (6.3)$$

$$= \int_{\mathbb{T}^3} -\nabla p \cdot \nabla \phi - \frac{1}{2} \min(M - \partial_3 p, 0) \partial_3 \phi - \langle PV, \phi \rangle. \quad (6.4)$$

In particular: the Newton descent direction introduced in proposition 6.3 is one of the descent directions to which theorem 6.2 applies. This means that we can find minima of the variational energy by simply iteratively following the Newton descent direction. This iterative method is discussed in greater detail in section A and it stands out for two reasons.

1. This method actually coincides with the numerical method used in previous work. This means that the convergence rates presented here are applicable to previous work (where convergence guarantees where not present). This is discussed in more detail in remark A.8.
2. This method may be interpreted as a geometric flow on the set of balanced states (once that set is equipped with the structure of a Hilbert manifold). More precisely: PV and M may be used as dual coordinates on that set, and this method is then a discretisation of the geometric flow which keeps the value of M fixed while varying PV to reach the desired target. This is discussed in more detail, once the relevant geometric terminology and machinery has been introduced, in remark A.39.

7. Summary and comparison with Helmholtz and the dry case

We are approaching the conclusion of this paper. Before we conclude, we recall in this section that the decomposition we introduced for the moist Boussinesq system under consideration has many different facets. We summarise these many facets here. In order to make it as easy as possible for the reader to follow this discussion, we first recall the many facets of the vortical-wave decomposition of the dry Boussinesq system. We also include a discussion of the many facets of the Helmholtz decomposition in appendix B. The reader familiar with the dry Boussinesq system may readily skip the Helmholtz discussion. Conversely, the reader less familiar with geophysical systems may find the Helmholtz discussion helpful to contextualise the many facets discussed here.

This is not merely a recapitulation of the discussion of these two classical systems done in the introduction. Throughout the paper we have introduced a variety of concepts and terminology for the sake of the moist decomposition, but we see that many of these have counterparts in these classical systems. It is these counterparts that we summarise first in 7.1 below before summarizing the various facets of the moist decomposition in section 7.2.

We could more prosaically call this section: ‘the twelve aspects of a decomposition found in both dry and moist Boussinesq’.

7.1. Dry Boussinesq

We begin with the various facets of the vortical-wave decomposition used for the dry Boussinesq system.

1. For the dry Boussinesq system discussed in this paper, the decomposition of a state consisting of a vector velocity field u and a scalar potential temperature field θ takes the form

$$\begin{pmatrix} u \\ \theta \end{pmatrix} = \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p \end{pmatrix} + \begin{pmatrix} \sigma_h^\perp + w e_3 \\ \sigma_3 \end{pmatrix},$$

where σ is divergence-free.

2. In the dry case, the balanced component satisfies the geostrophic and hydrostatic balances:

$$u_B = \nabla_h^\perp p \text{ and}$$

$$\theta_B = \partial_3 p.$$

3. The dual quantity characterising the wave set is, in the dry case, the potential vorticity. Indeed in this case the wave set is precisely the set of states whose potential vorticity vanishes:

$$\begin{aligned} PV(\sigma_h^\perp + we_3, \sigma_3) &= \nabla_h^\perp \cdot (\sigma_h^\perp + we_3)_h + \partial_3 \sigma_3 \\ &= \nabla_h \cdot \sigma_h + \partial_3 \sigma_3 = 0 \end{aligned}$$

since σ is divergence-free.

4. In the dry case the good unknown is the following three-dimensional vector field combining both the horizontal velocity and the potential temperature:

$$-u_h^\perp + \theta e_3.$$

5. The elliptic equation used to recover the pressure p , and hence the balanced component, is the Laplace/Poisson equation

$$\Delta p = PV.$$

6. The wave component is recovered by setting $w = u_3$ and finding σ to be the solution of the following div-curl system:

$$\nabla \times \sigma = \partial_3 u_h - \nabla_h^\perp \theta + (\partial_3 w) e_3 = j + (\partial_3 w) e_3 \text{ subject to } \nabla \cdot \sigma = 0 \text{ and } \oint \sigma = -\oint u_h^\perp + \oint \theta e_3 = a.$$

7. We may describe states of the dry Boussinesq system in three equivalent ways.
 (a) Describe them using u and θ .
 (b) Describe them using the dual quantities PV , j , w , and a .
 (c) Describe them using the coordinates p , σ , and w .
 8. The leading-order operator of the dynamics of the dry Boussinesq system considered in this paper is given by

$$\mathcal{L} \begin{pmatrix} u \\ \theta \end{pmatrix} = \begin{pmatrix} \mathbb{P}_L(u_h^\perp - \theta e_3) \\ u_3 \end{pmatrix}.$$

The decomposition is adapted to this operator since the balanced set is its kernel and the wave set is its image.

9. We may compute immediately that the wave component of a state undergoes linear oscillations under the leading-order operator since, when restricted to the wave set, the square of the operator behaves like multiplication by -1 :

$$\mathcal{L}^2 \begin{pmatrix} \sigma_h^\perp + we_3 \\ \sigma_3 \end{pmatrix} = - \begin{pmatrix} \sigma_h^\perp + we_3 \\ \sigma_3 \end{pmatrix}.$$

10. The conserved energy in the dry case is simply the L^2 norm given by

$$E = \int |u|^2 + \theta^2.$$

11. The Parseval-type identity in the dry case lets us rewrite the energy in terms of the coordinates or the dual quantities as follows:

$$\int |u|^2 + \theta^2 = \int |\nabla p|^2 + |\sigma|^2 + w^2 = \int |\nabla \Delta^{-1} PV|^2 + |\nabla \times^{-1}(j + (\partial_3 w) e_3, a)|^2.$$

12. For the dry Boussinesq case the balanced set takes the form

$$\mathcal{B}_{dry} := \left\{ (u, \theta) \in (L^2)_{\sigma}^3 \times L^2 : u = \nabla_h^{\perp} p, u_3 = 0, \text{ and } \theta = \partial_3 p \text{ for some } p \in \dot{H}^1 \right\}.$$

Therefore the extraction of a balanced component once again has a projection perspective:

$$(u_B, \theta_B) = \arg \min_{\mathcal{Y} \in \mathcal{B}_{dry}} \|(u, \theta) - \mathcal{Y}\|_{L^2}.$$

7.2. Moist Boussinesq

Finally we perform the same exercise as in appendix B and 7.1: we list the twelve aspects of the moist decomposition presented in this paper, in the same order as the above sections to make the comparison easy to follow.

Recall that here, in a manner consistent with the rest of this paper (barring the introduction) we write θ to denote the *equivalent* potential temperature (denoted θ_e in the introduction).

1. The decomposition of a state (u, θ, q) is given by

$$\begin{pmatrix} u \\ \theta \\ q \end{pmatrix} = \begin{pmatrix} \nabla_h^{\perp} p \\ \partial_3 p + \frac{1}{2} \min_0 (M - \partial_3 p) \\ M - \partial_3 p - \frac{1}{2} \min_0 (M - \partial_3 p) \end{pmatrix} + \begin{pmatrix} \sigma_h^{\perp} + w e_3 \\ \sigma_3 \\ -\sigma_3 \end{pmatrix},$$

where p has vanishing average and σ is divergence-free.

2. In the moist case the balanced component satisfies the geostrophic and hydrostatic balances, meaning that

$$\begin{aligned} u_B &= \nabla_h^{\perp} p \text{ and} \\ \theta_B - \min_0 q_B &= \partial_3 p. \end{aligned}$$

To fully characterise the balanced component we must also include the moist variable M and the following ‘auxiliary’ balance:

$$\theta_B + q_B = M.$$

3. The slow dual quantities characterising the wave set are now PV and M , and indeed

$$PV(\sigma_h^{\perp} + w e_3, \sigma_3, -\sigma_3) = \nabla \cdot \sigma = 0 \text{ and } M(\sigma_h^{\perp} + w e_3, \sigma_3, -\sigma_3) = \sigma_3 - \sigma_3 = 0.$$

4. There are *two* good unknowns in the moist case, owing to the fact that the leading-order operator is *not* skew-symmetric (this is discussed in more detail in remark 3.9). One good unknown is

$$-u_h^{\perp} + (\theta - \min_0 q) e_3,$$

which is related to j and the null set of the leading-order operator, while its ‘adjoint’ good unknown is

$$-u_h^{\perp} + \theta e_3,$$

which is related to PV and the image of the leading-order operator.

5. In order to reconstruct the balanced component we set $M = \theta + q$ and then let p be the solution of nonlinear PV -and- M inversion

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV.$$

6. To reconstruct the wave component we let $w = u_3$ and let σ solve

$$\begin{aligned} \nabla \times (A_B^{-1} \sigma) &= \partial_3 u_h - \nabla_h^\perp (\theta - \theta H_B) + (\partial_3 w) e_3 =: j_B + (\partial_3 w) e_3 \\ \text{subject to } \nabla \cdot \sigma &= 0 \text{ and } \oint A_B^{-1} \sigma = - \oint u_h^\perp + \oint (\theta - q H_B) e_3, \end{aligned}$$

where $A_B^{-1} = I + H_B e_3 \otimes e_3$ for $H_B := \mathbb{1}(M < \partial_3 p)$.

7. We can describe moist Boussinesq states in three ways.
- Describe them using u , θ , and q .
 - Describe them using the dual quantities PV , M , j , w , and a .
 - Describe them using the coordinates p , M , σ , and w .
8. The leading-order operator of the dynamics in the moist case is

$$\mathcal{N} \begin{pmatrix} u \\ \theta \\ q \end{pmatrix} = \begin{pmatrix} \mathbb{P}_L [u_h^\perp - (\theta - \min_0 q) e_3] \\ u_3 \\ -u_3 \end{pmatrix}.$$

9. As detailed in section 4 the operator of interest acting on the wave set, denoted by $\mathcal{L}_{W, H_B}$, is the linearisation of $\mathcal{N}$ about a balanced state with phases characterised by H_B . We then observe that the restriction of this operator to the wave set produces oscillatory dynamics. By contrast with the case of the Helmholtz decomposition and of the dry Boussinesq system, the frequencies of these oscillations are *phase-dependent*:

$$\mathcal{L}_{W, H_B}^2 \begin{pmatrix} \sigma_h \\ 0 \\ 0 \end{pmatrix} = - \begin{pmatrix} \sigma_h \\ 0 \\ 0 \end{pmatrix} \text{ and } \mathcal{L}_{W, H_B}^2 \begin{pmatrix} w e_3 \\ \sigma_3 \\ -\sigma_3 \end{pmatrix} = - (1 + H_B) \begin{pmatrix} w e_3 \\ \sigma_3 \\ -\sigma_3 \end{pmatrix}.$$

10. In the moist case the conserved energy is no longer quadratic. It is now nonlinear and given by

$$E = \int |u|^2 + \theta^2 + \min_0^2 q + (\theta + q)^2.$$

11. In the moist case there is no simple Parseval-type identity which holds: the entirety of section 5 is devoted to finding an alternative! Ultimately we identify in section 5.3 a *statistical divergence* given by

$$D(\mathcal{X}_1; \mathcal{X}_2) = \frac{1}{2} \int A_{H_1} \nabla p \cdot \nabla p + \left(1 + \frac{H_1}{2}\right) (M_1 - M_2)^2 + A_{H_1}^{-1} \sigma \cdot \sigma + (w_1 - w_2)^2$$

where, for $H_1 := \mathbb{1}(q_{B,1} < 0)$,

- p solves $\nabla \cdot (A_{H_1} \nabla p) = PV_1 - PV_2 - \frac{1}{2} \partial_3 [H_1 (M_1 - M_2)]$ and
- σ solves $\nabla \times (A_{H_1}^{-1} \sigma) = j_1 - j_2 + \partial_3 (w_1 - w_2) e_3$ subject to $\nabla \cdot \sigma = 0$ and $\oint A_{H_1}^{-1} \sigma = a_1 - a_2$.

12. The projection perspective in the moist case takes the following form:

$$\mathcal{X}_B = \arg \min_{\mathcal{Y} \in \mathcal{B}} D(\mathcal{X}; \mathcal{Y}).$$

Above we have summarised in detail how the many aspects of the dry and moist Boussinesq decompositions are analogous to one another, highlighting their differences. Below we provide a table summarizing sections 7.1 and 7.2 and which may serve either as a reference or as a more visual way to compare and contrast the various aspects of each decomposition.

	Dry Boussinesq	Moist Boussinesq
Decomposition	$\begin{pmatrix} u \\ \theta \end{pmatrix} = \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p \end{pmatrix} + \begin{pmatrix} \sigma_h^\perp + w e_3 \\ \sigma_3 \end{pmatrix}$	$\begin{pmatrix} u \\ \theta \\ q \end{pmatrix} = \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p + \frac{1}{2} \min_0 (M - \partial_3 p) \\ M - \partial_3 p + \frac{1}{2} \min_0 (M - \partial_3 p) \end{pmatrix} + \begin{pmatrix} \sigma_h^\perp + w e_3 \\ \sigma_3 \\ -\sigma_3 \end{pmatrix}$
Balanced set	$\mathcal{B}_{dry} = \left\{ u = \nabla_h^\perp p, u_3 = 0, \text{ and } \theta = \partial_3 p \text{ for some } p \in \dot{H}^1 \right\}$	$\mathcal{B}_{dry} = \left\{ u = \nabla_h^\perp p, u_3 = 0, \text{ and } \theta - \min_0 q = \partial_3 p \text{ for some } p \in \dot{H}^1 \right\}$
Slow dual quantities, which characterise balanced components	$PV = \nabla_h^\perp \cdot u_h + \partial_3 \theta$	$PV = \nabla_h^\perp \cdot u_h + \partial_3 \theta \text{ and } M = \theta + q$
Fast dual quantities, which characterise wave components	$\begin{aligned} j &= \partial_3 u_h - \nabla_h^\perp \theta, \\ a &= -f u_h^\perp + f \theta e_3, \\ \text{and } w &= u_3 \end{aligned}$	$\begin{aligned} j &= \partial_3 u_h - \nabla_h^\perp (\theta - \min_0 q) e_3, \\ a &= -f u_h^\perp + f (\theta - \min_0 q) e_3, \\ \text{and } w &= u_3 \end{aligned}$
Elliptic PDE for the balanced component	$\Delta p = PV$	$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = PV$
Elliptic PDE for the wave component	$\begin{aligned} \nabla \times \sigma &= j + (\partial_3 w) e_3 \\ \text{subject to } \nabla \cdot \sigma &= 0 \text{ and} \\ f \sigma &= -f u_h^\perp + f \theta e_3 \end{aligned}$	$\begin{aligned} \nabla \times (A_B^{-1} \sigma) &= j_B + (\partial_3 w) e_3 \\ \text{subject to } \nabla \cdot \sigma &= 0 \text{ and} \\ f A_B^{-1} \sigma &= -f u_h^\perp + f (\theta - q H_B) e_3 \end{aligned}$
Projection perspective	$(u_B, \theta_B) = \arg \min_{\mathcal{Y} \in \mathcal{B}_{dry}} \ (u, \theta) - \mathcal{Y}\ _{L^2}$	$\mathcal{X}_B = \arg \min_{\mathcal{Y} \in \mathcal{B}} D(\mathcal{X}; \mathcal{Y})$
Global changes of coordinates	State: u and θ dual quantities: PV, j, w , and a Coordinates: p, σ , and w	State: u, θ , and q dual quantities: PV, M, j, w , and a Coordinates: p, M, σ , and w
Leading-order operator	$\mathcal{L} \begin{pmatrix} u \\ \theta \end{pmatrix} = \begin{pmatrix} \mathbb{P}_L(u_h^\perp - \theta e_3) \\ u_3 \end{pmatrix}$	$\mathcal{N} \begin{pmatrix} u \\ \theta \\ q \end{pmatrix} = \begin{pmatrix} \mathbb{P}_L[u_h^\perp - (\theta - \min_0 q) e_3] \\ u_3 \\ -u_3 \end{pmatrix}$
Conserved energy	$E = \int u ^2 + \theta^2$	$E = \int u ^2 + \theta^2 + \min_0^2 q + (\theta + q)^2$

8. Conclusion

In this section we discuss the main contributions of this paper and the main takeaways for the reader.

The core contribution of this paper is the following. We provided a blueprint for decomposing the solutions of $\partial_t + \mathcal{N}$, specifically when $\mathcal{N}$ is nonlinear. Recall that otherwise, if $\mathcal{N}$ is linear, then we can proceed via spectral methods and identify eigenvalues and eigenvectors of that operator.

This blueprint requires a good understanding of the null set and the image of $\mathcal{N}$, which we refer to as the balanced set and the wave set in the context of this paper.

We actually provide *two* descriptions of each of these sets: (1) an explicit parametrisation and (2) a characterisation as the joint zero set of some dual quantities. Note that the dual quantities annihilating the wave set are known as slow dual quantities while those annihilating the balanced set are known as fast dual quantities.

A crucial condition to check in order to follow this general blueprint is that the slow dual quantities be characteristic of balanced states. In other words: there must be only one balanced set with any given PV and M . Geometrically this is asking that the wave set (i.e. the joint zero set of PV and M) be *transversal* to the balanced set.

This allows us to define a global decomposition of state space:

$$\mathbb{L}_G^2 = \mathcal{B} + \mathcal{W}.$$

In particular this decomposition gives broader meaning to nonlinear PV -and- M inversion. Previously, this inversion was understood in the context of balanced states: given PV and M , the inversion was used to produce the corresponding balanced state. Now this inversion has relevance for *any* state: it is used to produce its balanced *component*.

We then discuss three aspects of this decomposition.

1. The first aspect we discuss is its interpretation as a slow-fast decomposition. We obtain the three-part decomposition

$$\mathcal{X} = \underbrace{\overbrace{\mathcal{X}_B}^{\text{balanced}} + \overbrace{\bar{\mathcal{X}}_W}^{\text{wave}}}_{\text{slow}} + \underbrace{(\mathcal{X}_W - \bar{\mathcal{X}}_W)}_{\text{fast}} \quad (8.1)$$

where $\bar{\mathcal{X}}_W$ denotes the time average of the wave component.

Contrast this with the dry case: in that case the wave component undergoes linear oscillations, their time averages thus vanish, and so they do *not* contribute to the slow component at lowest order (they do contribute, a little, when the small parameter ε is finite). Here, in the moist case, the wave component undergoes nonlinear oscillations and so its time average contributes to the slow component.

2. The second aspect of the decomposition we discuss is its connection to the conserved energy. (Recall: when the operator $\mathcal{N}$ is linear, the eigendecomposition obtained often has favourable orthogonality properties, e.g. with respect to the inner product generated by the conserved energy as is the case in the case for the dry Boussinesq system.) We obtain a statistical divergence which quantifies the discrepancy between states in a manner consistent with both the conserved energy and the decomposition.

This statistical divergence is inspired by the Parseval-type identities which hold for the Helmholtz decomposition and the dry Boussinesq decomposition. For example in the dry case we can write the energy either in terms of the state (u, θ) or in terms of the dual quantities (PV, j, w, a) . We proceed similarly in the moist case: we use the nonlinear change of variables from the state (u, θ, q) to the dual quantities (PV, M, j, w, a) to define the statistical divergence in terms of these dual quantities.

3. The third aspect of the decomposition we discuss is how to compute it in practice. We introduce iterative methods that may be used to compute the decomposition. (These methods produce the balanced component $\mathcal{X}_B$ and the wave component may then be obtained as a residual $\mathcal{X}_W := X - \mathcal{X}_B$.) One of these methods, the Newton descent method, has the added benefit of shedding light on past work on this moist Boussinesq system.

This Newton descent iteration corresponds to a previously used numerical method. Our analysis here thus provides the first rigorous proof that this numerical method converges.

Table 1. The main results of this paper and some important auxiliary results.

Result 1: Decomposition	Theorem 3.5
The balanced set	Proposition 3.1
The wave set	Proposition 3.3
Nonlinear change of variables	Proposition 3.8
Result 2: Slow-fast interpretation of the decomposition	Theorem 4.1
Result 3: Statistical divergence agreeing with the decomposition and the conserved energy	Proposition 5.26
Definition of the statistical divergence	Definition 5.23 (and Definition 5.10)
Result 4: Convergence rate for the Newton descent iteration	Theorem A.19
The Newton descent search direction	Proposition A.7
Justification of a numerical method used in the literature	Remark A.8
Geometric interpretation	Remark A.39

This Newton descent iteration may also be interpreted geometrically: it is precisely the discretisation of the flow along the balanced set where the moist variable M is kept constant and where we follow the potential vorticity coordinate direction until the desired potential vorticity is attained.

Table 1 below serves as an easy reference for the main results of this paper. This table also lists a few important auxiliary results, remarks, and definitions which either highlight the key ideas underpinning the main results or provide essential context for these results.

We now pause to highlight the main difficulty in constructing a decomposition with the favourable properties discussed above. Naively, the difficulty may appear to be the nonlinearities or the lack of skew-symmetry. Both of these challenges come from the phase boundaries, and so the real difficulty is more specifically this: the phase boundaries are *dynamic*, meaning that they evolve over time.

To explain this we consider what would happen if this difficulty were removed. That is, suppose that (u, θ, q) is a solution of the moist Boussinesq system (2.7a)–(2.7c) whose phase boundaries do *not* evolve over time. In other words, if $H := \mathbb{1}(q < 0)$, suppose that $\partial_t H = 0$.

Then we may treat H as a fixed indicator function and observe that $\mathcal{N}(u, \theta, q) = \mathcal{L}_H(u, \theta, q)$. Crucially: $\mathcal{L}_H$ is L_H^2 -skew-symmetric (see lemma C.11). This means that, when the phase boundaries are *static*, the leading-order operator becomes both linear *and* skew-symmetric. Both symptomatic difficulties (nonlinearities and lack of skew-symmetry) are removed when the root difficulty (dynamic phase boundaries) is removed.

8.1. Looking forward

In this section we record some thoughts on future research direction stemming from the topics discussed here.

- Can we prove that the precipitating QG system (PQG; see [SS17], or appendix C.5 for a self-contained introduction), or a variant thereof is the rigorous fast-wave averaging limit of the moist Boussinesq system, in the limit of fast rotation and strong stratification (i.e. as both the Froude and Rossby numbers approach zero, at the same rate)?

In the case of balanced, or well-prepared initial data we expect to be able to obtain the PQG system as a rigorous limit of the moist Boussinesq system discussed here.

However, in the case of unbalanced, or ill-prepared initial data, there is mounting evidence that the appropriate limiting system is not the PQG system. The crux of the matter are the nonlinear oscillations of the fast waves: for ill-prepared data these waves are present in the system and, in the moist case, they feed back into the slow components through their average. We thus expect that, at the very least, a modification needs to be made to the PQG system if we hope to obtain it as a rigorous limit of the moist Boussinesq system for generic initial data.

Since this is a challenging problem, and here are two more restricted questions to look at first and whose resolution will shed light on the general case.

- Treat the ODE case first. The ODE problem is recorded in lemma C.23, which also shows that ODE solutions are PDE solutions. Crucially: this ODE system possesses one of the key difficulties found in the PDE case: its wave solutions, i.e. its fast components, are nonlinear oscillations (which are piecewise linear oscillations, meaning that depending on the sign of q they have fixed frequencies, but their frequencies change as the sign of q changes) with *non-zero* averages.
- Treat the case of steady interfaces. In other words, look at solutions to the full moist Boussinesq system whose interface (i.e. the zero level set of q) is fixed in time and does not evolve dynamically, even though q itself evolves dynamically. Such solutions are readily obtained by considering velocity fields whose support is localised away from the interface. In this case, the Heaviside $H = \mathbb{1}(q < 0)$ which appears in the moist Boussinesq system can be considered to be fixed, and so the leading-order operator is linear! This removes a key difficulty from the general, nonlinear, moist Boussinesq system due to dynamically-evolving phase boundaries.

It is important to note what we expect to glean from studying this ‘fixed H ’ case. Indeed, if we consider this example from a physical perspective it is severely contrived and hence there is little hope of extracting meaningful physical intuition from this test case. Nonetheless, studying this ‘fixed H ’ case may prove helpful in isolating the mathematical difficulties one at a time, and this is where we may benefit from studying this simpler case first before handling the general case.

Note that carrying out the fast-wave averaging process will require a precise identification of the solvability condition for the next-order equations, something which we have not carried out, even formally, in this paper. Recall that this next-order solvability condition is used to determine the behaviour of the ‘background profile’ at the slow timescale which arises from the solution of the leading-order equation.

- Can we find an explicit example of a PDE state for which its Snell-type distance projection does not agree with our decomposition? We strongly suspect that such examples exist since they are easy to find in the ODE case.
- In section A.2 we discussed how the iterative method used to solve nonlinear PV -and- M inversion numerically could be interpreted geometrically. This interpretation viewed the iterative method as the discretisation of a flow *along* the balanced set. It may be possible to re-interpret this same iterative method in a different way as follows. First, the detailed geometric understanding of the balanced set, in terms of a chart, coordinate vectors, and dual coordinate vectors, would need to be extended to the entire state space, thus describing both the balanced set *and* the wave set. We may then be able to re-interpret this same iterative method geometrically as a flow in state space which starts at some (possibly unbalanced) initial state and flows *towards* the balanced set, only reaching it when the iterative method terminates. The appeal of such an interpretation would be that it matches the following physical intuition: a balanced state is obtained by starting from *any* state and then progressively removing wave components (i.e. imbalances) until we reach the balanced set.
- The two metrics/distance functions defined in sections 5.1 and 5.2 may potentially be used for other applications, such as adjoint models or data-driven models. The ‘fixed H ’ energy from section 5.1 appears in situations where linearisation is used, and hence it could find application to adjoint models, sensitivity analysis, and variational data assimilation, which rely on tangent linear approximations to the dynamics [LDT86, PD97, Err97, ER99, Mah99, BBP+01, AZD08, DARR14]. The ‘fixed H ’ energy also appears here in the numerical/iterative methods, in section A.2 where geometrical notions relying on tangent spaces (a.k.a. linearisations) are discussed. The Snell-like distance in section 5.2 here could be useful for applications such as data-driven methods that identify nonlinear features in dynamical systems and utilise a metric, distance function, and/or kernel function [CL06, GM12, BH16, DG19, Gia19].
- Further complexity could potentially be explored in the treatment of cloud microphysics. For the topic of slow-fast decompositions, a first treatment was considered here for the case of reversible non-precipitating moist thermodynamics. For the related topic of the QG limit as an asymptotic limit, further complexity in cloud microphysics has been considered in some prior work that has considered rainfall and precipitation processes [SS17, WSSM19, WSS+20, BHK23]. Some other treatments of clouds do not use the types of cloud microphysics that are used in cloud-resolving models; instead they use a large-scale parameterisation of convection and/or latent heating. For such large-scale parameterisations, it is common to add the parameterisations to a dry QG framework [LH04, LLR11, DVMFH10] without considering a moist QG limit as a moist asymptotic limit, and the consideration of such QG frameworks as the result of a moist asymptotic limit could be a topic for future research. From this brief discussion of other treatments of cloud microphysics or large-scale parameterisations, one can see several different topics that could be interesting avenues for future research.

- The main scaling scenario considered here was a QG scaling. Semi-geostrophic scaling is another paradigm [Hos82, TE85, Cul18, Cul21] which has played a fundamental role in understanding frontogenesis and other phenomena. A decomposition based on semi-geostrophic theory has been used for data analysis, including the use of a modified static stability to represent latent heat release [Cul18]. It could also be interesting to consider a nonlinear eigenspace decomposition in the spirit of section 3 herein, in the different setting of semi-geostrophic theory.
- Can the nonlinear eigenspace decomposition be extended to cases with a background flow that is as strong as the inertia-gravity wave speed? Such a situation can arise, for instance, in the presence of a strong jet flow. In such cases, the leading-order operator $\mathcal{D}_0$ can include linear advection by the background flow, which can cause a more complex decomposition into balanced and unbalanced components.
- Can the ocean (or other geophysical or astrophysical fluids) have a nonlinear eigenspace decomposition, similar to the atmospheric decomposition presented here? The nonlinearity in the buoyancy here was a nonlinear switch due to clouds, but the blueprint of the techniques here may also be applicable for other types of nonlinearity. For instance, in oceanography, a nonlinear equation of state describes the nonlinear dependence of oceanic buoyancy on temperature and salinity, and plays an important role in a variety of processes [KM10, HN13, NHR15, RMBN15, Kor21]. Also, PV inversion is used in oceanography as well [Mül95, OSB+01], and it would be interesting to investigate whether the nonlinear eigenspace decomposition here may be applicable for incorporating the effects of a nonlinear equation of state for seawater.

Acknowledgments

The authors thank two anonymous reviewers for helpful comments that improved the presentation. The authors are grateful for the support of NSF DMS-1907667. A R-T gratefully acknowledges support from the Simons Foundation, the Isaac Newton Institute for Mathematical Sciences, and the Department of Applied Mathematics and Theoretical Physics of the University of Cambridge where part of this work was carried out.

Data availability statement

No new data were created or analysed in this study.

Appendix A. Iterative methods for the decomposition and their geometry—details

A.1. Convergence guarantees

In this section we introduce and compare several gradient descent methods that can be employed to minimise the variational energy introduced in theorem 6.1, and thus find solutions of PV - and M -inversion. We will discuss three descent methods: $\dot{H}^1$ -gradient descent, L^2 -gradient descent, and Newton descent (which, as we will discuss, can be interpreted locally as a gradient descent), recording corresponding convergence rates for each method. Ultimately, the convergence rates of interest are those obtained for *backtracking line search* (see definition A.17) which are recorded in theorem A.19 but we also record analogous convergence rates in theorem A.15 for *exact* line search. This is done to make the rates obtained in the main result of theorem A.19 easier to contextualise, by comparing them with what can be obtained in the idealised scenario of exact line search.

We then conclude this section with a consistency result that guarantees that minimisers of the variational energy restricted to an (appropriate) finite-dimensional subspace of $\dot{H}^1$ do converge to the true minimiser of the energy. Put together, the convergence rates of theorem A.19 and the consistency result of theorem A.21 tell us that we have a ready-to-implement method for finding approximate minimisers which are guaranteed to converge to the true minimiser.

Before we get started we recall results from [RTSS24] that will be used in this section.

Theorem A.1 (properties of the variational energy). *The variational energy E_{var} defined in theorem 6.1 has the following properties.*

1. E_{var} is $\frac{1}{2}$ -strongly convex, weakly lower semi-continuous over $\dot{H}^1$, and coercive over $\dot{H}^1$ in the sense that $E_{\text{var}}(p) \geq \frac{1}{16} \|p\|_{\dot{H}^1}^2 - \frac{3}{4} \|M\|_{L^2}^2 - 4 \|PV\|_{\dot{H}^{-1}}^2$ for every $p \in \dot{H}^1$.

2. We may write

$$E_{\text{var}}(p) = \int_{\mathbb{T}^3} e_M(x, \nabla p(x)) dx + \langle PV, p \rangle \quad (\text{A.1})$$

for $e_M(x, u) := \frac{1}{2}|u|^2 - \frac{1}{4}\min(M(x) - u_3, 0)^2$.

3. E_{var} is Gâteaux differentiable on $\dot{H}^1$ with Gâteaux derivative at any $p \in \dot{H}^1$, denoted by $DE_{\text{var}}(p)$, given by

$$DE_{\text{var}}(p)\phi = \int_{\mathbb{T}^3} \nabla p \cdot \nabla \phi + \frac{1}{2} \min(M - \partial_3 p, 0) \partial_3 \phi + \langle PV, \phi \rangle \text{ for every } \phi \in \dot{H}^1. \quad (\text{A.2})$$

4. The Gâteaux derivative E_{var} is Lipschitz, with specifically the estimate $\|DE_{\text{var}}(p_1) - DE_{\text{var}}(p_2)\|_{H^{-1}} \leq \frac{3}{2}\|p_1 - p_2\|_{\dot{H}^1}$ for every $p_1, p_2 \in \dot{H}^1$, and so E_{var} is Fréchet differentiable.

5. Suppose that E_{var} has a minimiser $p^* \in \dot{H}^1$. For any $p \in \dot{H}^1$ the following estimate holds: $\|p - p^*\|_{\dot{H}^1}^2 \leq 4(E_{\text{var}}(p) - E_{\text{var}}(p^*)) = 4(E_{\text{var}}(p) - \min E_{\text{var}})$.

Proof. See proposition 4.4, lemma 4.6, corollary 4.7, and lemma 4.8 of [RTSS24]. $\square$

In particular [RTSS24] establishes the well-posedness of nonlinear PV-and-M inversion.

Theorem A.2 (existence and uniqueness). For every $M \in L^2$, and $PV \in H^{-1}$ the variational energy E_{var} introduced in theorem 6.1 has a unique global minimiser in $\dot{H}^1$ which is also the unique weak solution of (6.1).

Moreover we have the following H^2 regularity result ([RTSS24] proves higher-order regularity, but this H^2 is all that is required below). Let $M \in H^1$ and $PV \in L^2$. Any $\dot{H}^1$ -weak solution p of (6.1) belongs to H^2 and satisfies $\|p\|_{\dot{H}^2} \leq \|M\|_{\dot{H}^1} + 2\|PV\|_{L^2}$.

Proof. This is theorem 5.1 and lemma 5.2 in [RTSS24]. $\square$

Now to kick things off in earnest we define the gradient of a functional, taking particular care to highlight the role played by the choice of the underlying inner product.

Definition A.3 (gradient). Consider a functional $\mathcal{F} : B \rightarrow \mathbb{R}$ on a Banach space B equipped with an inner product $(\cdot, \cdot)$ which is not necessarily complete and suppose that there exists a map $D : B \rightarrow B$ such that $D\mathcal{F}(u)\phi = (D(u), \phi)$ for every $u, \phi \in B$. We call D the gradient of $\mathcal{F}$ with respect to the inner product $(\cdot, \cdot)$. Moreover, if $(\cdot, \cdot)$ is known as the ‘X-inner product’ then we write $D = \nabla_X \mathcal{F}$ and call it the X-gradient of $\mathcal{F}$.

With definition A.3 in hand we may now turn our attention towards two gradients of the variational energy that we will use throughout: the $\dot{H}^1$ and L^2 -gradients. We record expressions for both of these below.

Lemma A.4 ($\dot{H}^1$ and L^2 -gradients for our variational energy). Consider the variational energy E_{var} introduced in theorem 6.1. The $\dot{H}^1$ -gradient of E_{var} is given by

$$(\nabla_{\dot{H}^1} E_{\text{var}})(p) = (-\Delta)^{-1} \left[-\Delta p - \frac{1}{2} \partial_3 (\min(M - \partial_3 p, 0)) + PV \right],$$

meaning that $u = (\nabla_{\dot{H}^1} E_{\text{var}})(p)$ is the unique weak solution of

$$-\Delta u = -\Delta p - \frac{1}{2} \partial_3 (\min(M - \partial_3 p, 0)) + PV \quad (\text{A.3})$$

satisfying $\int_{\mathbb{T}^3} u = 0$, i.e. $u \in \dot{H}^1$ such that, for every $\phi \in \dot{H}^1$,

$$\int_{\mathbb{T}^3} \nabla u \cdot \nabla \phi = \int_{\mathbb{T}^3} \nabla p \cdot \nabla \phi + \frac{1}{2} \min(M - \partial_3 p, 0) \partial_3 \phi + \langle PV, \phi \rangle. \quad (\text{A.4})$$

Moreover, if $p \in H^2$, $M \in H^1$, and $PV \in L^2$ then the L^2 -gradient of E_{var} is given by

$$(\nabla_{L^2} E_{\text{var}})(p) = -\Delta p - \frac{1}{2} \partial_3 (\min(M - \partial_3 p, 0)) + PV. \quad (\text{A.5})$$

Proof. Item 3 of theorem A.1 tells us that $u = \nabla_{\dot{H}^1} E_{\text{var}}(p)$ satisfies

$$\int_{\mathbb{T}^3} \nabla u \cdot \nabla \phi = (u, \phi)_{\dot{H}^1} = (\nabla_{\dot{H}^1} E_{\text{var}}(p), \phi)_{\dot{H}^1} = DE_{\text{var}}(p)\phi$$

$$= \int_{\mathbb{T}^3} \nabla p \cdot \nabla \phi + \frac{1}{2} \min(M - \partial_3 p, 0) \partial_3 \phi + \langle PV, \phi \rangle,$$

which is precisely (A.4), the weak formulation of (A.3). In order to establish (A.5) it suffices to integrate by parts the expression for the Gâteaux derivative of E_{var} obtained in item 3 of theorem A.1. This yields

$$(\nabla_{L^2} E_{var}(p), \phi)_{L^2} = DE_{var}(p) \phi = \int_{\mathbb{T}^3} \left[-\Delta p - \frac{1}{2} \partial_3 (\min(M - \partial_3 p, 0)) + PV \right] \phi,$$

from which the claim follows. $\square$

Remark A.5 (comparing the L^2 and $\dot{H}^1$ gradients). It is not a coincidence that in lemma A.4 above the L^2 and $\dot{H}^1$ gradients are related via $\nabla_{\dot{H}^1} E_{var} = (-\Delta)^{-1} (\nabla_{L^2} E_{var})$. This identity actually follows from the following fact. For any Hilbert space $H \subseteq L^2$ and any Fréchet differentiable $F: L^2 \rightarrow \mathbb{R}$, if we denote by $R_H: R \rightarrow R^*$ the Riesz representation map on H then $\nabla_H F = R_H(\nabla_{L^2} F)$. Indeed, for any $u, \phi \in H$: $DF(u)\phi = (\nabla_{L^2} F(u), \phi)_{L^2} = \langle \nabla_{L^2} F(u), \phi \rangle_{H^* \times H} = (R_H(\nabla_{L^2} F)(u), \phi)_H$.

In our case we have that $R_{\dot{H}^1} = (-\Delta)^{-1}$ in the sense that, for any $f \in H^{-1}$, $u = R_{\dot{H}^1} F$ if and only if u is the $\dot{H}^1$ -weak solution of $-\Delta u = f$ in $\mathbb{T}^3$. The general fact above thus tells us that $\nabla_{\dot{H}^1} E_{var} = R_{\dot{H}^1}(\nabla_{L^2} E_{var}) = (-\Delta)^{-1}(\nabla_{L^2} E_{var})$, as observed in lemma A.4.

Besides gradient descent methods using the $\dot{H}^1$ and L^2 gradients, the third descent method we will discuss is *Newton descent*. This comes from using, locally, the inner product generated by the *Hessian* of the variational energy to define a gradient direction. This idea is now made precise.

Definition A.6 (Newton descent). Let $\mathcal{F}: H \rightarrow \mathbb{R}$ be a twice Fréchet differentiable map over a Hilbert space H . Suppose there exists $C > 0$ such that, for every $u \in H$, $D_u^2 \mathcal{F}(\cdot, \cdot)$ induces an inner product, written $(\cdot, \cdot)_{D_u^2 \mathcal{F}}$ and called the *Hessian inner product*, such that $\frac{1}{C} \|\phi\|_H \leq \|\phi\|_{D_u^2 \mathcal{F}} \leq C \|\phi\|_H$ for every $\phi \in H$ (i.e. $\|\cdot\|_{D_u^2 \mathcal{F}}$ is uniformly comparable to $\|\cdot\|_H$). The negative of the gradient of $\mathcal{F}$ with respect to the Hessian inner product is called the *Newton descent direction*.

Having defined Newton descent we now record its form for our variational energy of interest.

Proposition A.7 (Newton descent for our variational energy). Let E_{var} be the variational energy introduced in theorem 6.1. The Hessian inner product induced by E_{var} is, for any $p \in \dot{H}^1$ and $M \in L^2$,

$$(\phi, \psi)_{D_p^2 E_{var}} = \int_{\mathbb{T}^3} \nabla \phi \cdot \nabla \psi - \frac{1}{2} \mathbb{1}(M < \partial_3 p) (\partial_3 \phi) (\partial_3 \psi) \quad (\text{A.6})$$

and so the Newton descent direction π is the $\dot{H}^1$ -weak solution of

$$-\nabla \cdot \left[\left(I - \frac{1}{2} \mathbb{1}(M < \partial_3 p) e_3 \otimes e_3 \right) \nabla \pi \right] = \Delta p + \frac{1}{2} \partial_3 (\min(M - \partial_3 p, 0)) - PV, \quad (\text{A.7})$$

for $PV \in H^{-1}$, i.e. for all $\phi \in \dot{H}^1$,

$$\begin{aligned} & \int_{\mathbb{T}^3} \nabla \pi \cdot \nabla \phi - \frac{1}{2} \mathbb{1}(M < \partial_3 p) (\partial_3 \pi) \partial_3 \phi \\ &= \int_{\mathbb{T}^3} -\nabla p \cdot \nabla \phi - \frac{1}{2} \min(M - \partial_3 p, 0) \partial_3 \phi - \langle PV, \phi \rangle. \end{aligned} \quad (\text{A.8})$$

Proof. Note that $(\phi, \psi)_{D_p^2 E_{var}} = D^2 E_{var}(\phi, \psi)$ and so (A.6) follows from differentiating the expression for DE_{var} recorded in item 1 of theorem A.1, since $\min(\cdot, 0)$ is Lipschitz with derivative $\min(\cdot, 0)' = \mathbb{1}(\cdot < 0)$. Since the Newton descent direction π is characterised by the fact that $(\pi, \phi)_{D_p^2 E_{var}} = -DE_{var}(p)\phi$ for every $\phi \in \dot{H}^1$, (A.8) then follows immediately from (A.6) and the expression for DE_{var} of item 3 of theorem A.1. To see that (A.7) is indeed the strong form of (A.8) we integrate by parts. $\square$

Remark A.8 (interpretation of Newton descent). Section A.2 below details how the Newton descent iteration, recorded in proposition A.7 above, may be interpreted geometrically.

Another interpretation is already available to us now. First: note that once Newton descent enters the so-called ‘quadratic regime’, backtracking line search will systematically select steps of unit size. Given a current iterate p_k , the next iterate p_{k+1} will thus be given by $p_{k+1} = p_k + \pi$ where π solves (A.7), which we may rewrite as

$$-\nabla \cdot (A_{H_k} \nabla \pi) = \nabla \cdot (A_{H_k} \nabla p_k) + \frac{1}{2} \partial_3 (H_k M) - PV.$$

We may rearrange this equation and deduce that the new iterate $p_{k+1} = p_k + \pi$ solves

$$\nabla \cdot (A_{H_k} \nabla p_{k+1}) = PV - \frac{1}{2} \partial_3 (H_k M).$$

In other words: the next iterate p_{k+1} is obtained by solving the linearised PV -and- M inversion where the Heaviside characterising the linearisation comes from the previous step. Crucially: this is precisely the method that was used in previous works [HESS21] to numerically solve the nonlinear PV -and- M inversion. This means that the convergence rates presented here for Newton descent are applicable to previous works. Moreover this means that the convergence *guarantees* for Newton descent are also the first convergence guarantees of this numerical method which is already used in practice.

The remainder of this section is devoted to the obtention of convergence rates for each of the three descent methods discussed above. There are two key assumptions that must be verified about our variational energy in order for such convergence rates to be obtained. The first comes from convexity and takes the form of a quadratic lower bound on the variational energy. The second, which takes the form of a quadratic *upper* bound on the variational energy, is that to which we turn our attention now.

The basic ideas are that (1) we will introduce a simpler energy density for which, due to its piecewise quadratic nature, we can easily find such a quadratic upper bound, and (2) we will show that this simpler energy density can be turned into our true variational energy density up to some affine transformations. The result will then follow from the fact that such affine transformations will only interact with these quadratic upper bounds in predictable ways. In order to carry out this argument it is particularly convenient to introduce the notion of a *quadratic remainder*, which we define below.

Definition A.9 (quadratic remainder). Consider a continuously differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. For any $x \in \mathbb{R}^d$ we call function $T_x^f : \mathbb{R}^d \rightarrow \mathbb{R}$ defined as $T_x^f(y) := f(y) - f(x) - Df(x)(y - x)$ for every $y \in \mathbb{R}^d$ the *quadratic remainder of f at x* .

We now show that the aforementioned ‘simpler’ energy density—introduced in the result below—does indeed have a quadratic upper bound.

Lemma A.10 (quadratic upper bound for a piecewise quadratic function). Consider $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by $f(x) = \frac{1}{2}|x|^2 - \frac{1}{4} \min(x_3, 0)^2$ for all $x \in \mathbb{R}^3$. For every $x \in \mathbb{R}^3$, $f(y) \leq f(x) + Df(x)(y - x) + \frac{1}{2}|y - x|^2$.

Proof. To make the computations easier we introduce the following function: $A(x) := I - \frac{1}{2} \mathbb{1}(x_3 < 0) e_3 \otimes e_3$, such that $f(x) = \frac{1}{2} A(x) x \cdot x$, and hence $\nabla f(x) = A(x)x$. In particular note that, for every x , the largest eigenvalue of $A(x)$ is equal to 1, and so $A(x)\xi \cdot \xi \leq |\xi|^2$ for all $\xi \in \mathbb{R}^3$.

In terms of the quadratic remainder T_x^f , we may write the target inequality as $T_x^f(y) \leq \frac{1}{2}|y - x|^2$, so we begin by using the fact that A is pointwise symmetric to rewrite T_x^f in a more convenient way. We obtain

$$T_x^f(y) = \frac{1}{2} [A(y) - A(x)] y \cdot y + \frac{1}{2} A(x) (y - x) \cdot (y - x).$$

Since $A(y) - A(x) = \frac{1}{2} [\mathbb{1}(x_3 < 0 \leq y_3) - \mathbb{1}(y_3 < 0 \leq x_3)] e_3 \otimes e_3$ we compute that

$$T_x^f(y) = \begin{cases} -\frac{y_3^2}{4} + \frac{1}{2} A(x) (y - x) \cdot (y - x) & \text{if } y_3 < 0 \leq x_3, \\ \frac{x_3^2}{4} + \frac{1}{2} A(x) (y - x) \cdot (y - x) & \text{if } x_3 < 0 \leq y_3, \text{ and} \\ \frac{1}{2} A(x) (y - x) \cdot (y - x) & \text{otherwise.} \end{cases}$$

Only one of these three cases is troublesome since, if $y_3 < 0 \leq x_3$ or if x_3 and y_3 have the same sign then $T_x^f(y) \leq \frac{1}{2} A(x) (y - x) \cdot (y - x) \leq \frac{1}{2} |y - x|^2$. In the troublesome case, i.e. if $x_3 < 0 \leq y_3$, then we need to be more careful. In this case writing $z_h = (z_1, z_2)$ for any $x \in \mathbb{R}^3$ tells us that

$$T_x^f(y) = \frac{1}{2} A(x) (y - x) \cdot (y - x) + \frac{1}{4} x_3^2 = \frac{1}{2} |x_h - y_h|^2 + \frac{1}{4} (x_3 - y_3)^2 + \frac{1}{4} x_3^2.$$

Crucially: if $x_3 < 0 \leq y_3$ then $x_3 y_3 \geq 0$ and hence $\frac{1}{4} (x_3 - y_3)^2 + \frac{1}{4} x_3^2 \leq \frac{1}{2} (x_3 - y_3)^2$, which proves that $T_x^f(y) \leq \frac{1}{2} |x - y|^2$ even in this case, concluding the proof. $\square$

With lemma A.10 in hand we may now use the fact that quadratic upper bounds on quadratic remainders behave nicely under transformations (see lemmas C.16–C.18) to deduce that the true variational energy density of interest also has a quadratic upper bound.

Lemma A.11 (quadratic upper bound for our variational energy density). Consider $e : \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by $e(x) := \frac{1}{2}|x|^2 - \frac{1}{4}\min(r - x_3, 0)^2$, where r is fixed. For every $x, y \in \mathbb{R}^3$, $e(y) \leq e(x) + De(x)(y - x) + \frac{1}{2}|x - y|^2$.

Proof. We note that, for f as in lemma A.10, $f(ru_3 - x) = e(x) + a(x)$ for an affine function $a(x) = -rx_3 + \frac{1}{2}r^2$. Introducing the affine transformation $\Phi(x) = me_3 - x$ this means that $e = f \circ \Phi - a$. Therefore we may use lemmas C.16–C.18 to see that the quadratic remainder of e is given by $T_x^e = T_x^{f \circ \Phi} - T_x^a = T_{\Phi(x)}^f \circ \Phi$. In particular we deduce from lemma A.10 that $T_x^e(y) \leq \frac{1}{2}|\Phi(y) - \Phi(x)|^2 = \frac{1}{2}|y - x|^2$, which is precisely the desired inequality. $\square$

Our search for a quadratic upper bound on the variational energy now concludes as we leverage the quadratic upper bound on the variational energy density obtained in lemma A.11 to deduce that the variational energy itself does indeed have a comparable quadratic upper bound.

Lemma A.12 (quadratic upper model for our variational energy). Let E_{var} be the variational energy introduced in theorem 6.1. For every $p, r \in \dot{H}^1$, the following inequality holds:

$$E_{var}(r) \leq E_{var}(p) + DE_{var}(p)(r - p) + \frac{1}{2}\|r - p\|_{\dot{H}^1}^2.$$

Proof. We write E_{var} as in (A.1), abusively writing $e_M(\nabla p)$ to mean $e_M(x, \nabla p(x))$ in the sequel, and note that item 3 of theorem A.1 tells us that

$$DE_{var}(p)\phi = \int_{\mathbb{T}^3} \nabla_u e_M(\nabla p) \cdot \nabla \phi + \langle PV, \phi \rangle. \quad (\text{A.9})$$

Similarly we may reinterpret lemma A.11 as telling us that

$$e_M(\nabla r) \leq e_M(\nabla p) + \nabla_u e_M(\nabla p) \cdot \nabla(r - p) + \frac{1}{2}|\nabla(r - p)|^2. \quad (\text{A.10})$$

Putting (A.1), (A.9), and (A.10) together then yields the claim. $\square$

We now have all of the ingredients we need to prove the convergence rates for our three descent methods in the context of *exact* line search. First we state a generic version of that result in proposition A.13 before specialising it to our energy afterwards in theorem A.15. Note that the result below is classical—see for example [BV04]. We record it here, along with its proof, in order to ensure that good care is taken of tracking the constants arising from comparable norms involved in the descent method. This is particularly important for us since ultimately it will lead to significant differences between the convergence rate of the L^2 -gradient descent compared with the other descent methods.

Proposition A.13 (key inequality for the convergence rate of gradient descent with exact line search in Hilbert spaces). Let $(\cdot, \cdot)_X$ be a complete inner product on a Hilbert space H for which there exists a constant $C > 0$ such that

$$\|u\|_H \leq C\|u\|_X \text{ for every } u \in H. \quad (\text{A.11})$$

Let $\mathcal{F} : H \rightarrow \mathbb{R}$ be a continuously Fréchet differentiable functional satisfying, for some constants $m, M > 0$,

$$\frac{m}{2}\|v - u\|_H^2 \leq \mathcal{F}(v) - \mathcal{F}(u) - D\mathcal{F}(u)(v - u) \leq \frac{M}{2}\|v - u\|_H^2 \quad (\text{A.12})$$

for every $u, v \in H$, and let u^* denote the unique minimiser of $\mathcal{F}$. Then

$$\mathcal{F}\left(u - \frac{1}{C^2M}\nabla_X \mathcal{F}(u)\right) \leq \mathcal{F}(u) - \frac{1}{2C^2M}\|\nabla_X \mathcal{F}(u)\|_X^2 \quad (\text{A.13})$$

and

$$\mathcal{F}(u) - \mathcal{F}(u^*) \leq \frac{1}{2m}\|\nabla_H \mathcal{F}(u)\|_H^2 \quad (\text{A.14})$$

such that

$$\mathcal{F}\left(u - \frac{1}{C^2M}\nabla_X \mathcal{F}(u)\right) - \mathcal{F}(u^*) \leq \left(1 - \frac{m}{C^4M}\right)(\mathcal{F}(u) - \mathcal{F}(u^*)). \quad (\text{A.15})$$

Remark A.14 (on the importance of the key inequality). The so-called ‘key inequality’ (A.15) warrants such a lofty name since the actual convergence rate follows from that inequality immediately. Indeed, if we denote let $x_0 \in \mathbb{T}^3$ and denote by x_k the iterates obtained via gradient descent with exact line search, i.e.

$$x_k := \arg \min_{t \in \mathbb{R}} \mathcal{F}(x_{k-1} - t \nabla \mathcal{F}(x_{k-1})),$$

then inducting on the key inequality tells us that

$$\mathcal{F}(x_{k-1}) - \mathcal{F}(x^*) \leq \left(1 - \frac{m}{C^4 M}\right)^k (\mathcal{F}(x_0) - \mathcal{F}(x^*)),$$

which guarantees that indeed $\mathcal{F}(x_k)$ converges to $\mathcal{F}(x_0)$.

Proof of proposition A.13. Plugging $v = u - t \nabla_X \mathcal{F}(u)$ into the second inequality in (A.12) tells us that

$$\begin{aligned} \mathcal{F}(u - t \nabla_X \mathcal{F}(u)) &\leq \mathcal{F}(u) - t \|\nabla_X \mathcal{F}(u)\|_X^2 + \frac{M t^2}{2} \|\nabla_X \mathcal{F}(u)\|_H^2 \\ &\leq \mathcal{F}(u) + \left(\frac{C^2 M t^2}{2} - t \right) \|\nabla_X \mathcal{F}(u)\|_X^2. \end{aligned}$$

Since the quadratic $q(t) := \frac{C^2 M t^2}{2} - t$ is minimised at $t^* = \frac{1}{C^2 M}$, where it attains the value $q(t^*) = -\frac{1}{2C^2 M}$, we deduce (A.13) from the inequality above.

To establish (A.14) we use the first inequality in (A.12). We begin by writing $D\mathcal{F}(u)(v - u) = (\nabla_H \mathcal{F}(u), v - u)_H$, which tells us that with respect to v the right-hand side of that inequality is minimised at $v^* = u - \frac{1}{m} \nabla_H \mathcal{F}(u)$. Therefore, for any $u, v \in H$,

$$F(v) \geq \left[\mathcal{F}(u) + (\nabla_H \mathcal{F}(u), v - u)_H + \frac{m}{2} \|v - u\|_H^2 \right]_{v=v^*} = \mathcal{F}(u) - \frac{1}{2m} \|\nabla_H \mathcal{F}(u)\|_H^2.$$

Setting $v = u^*$ produces (A.14).

Finally we turn our attention to the key inequality (A.15). (A.11) and (A.14) tell us that $\|\nabla_X \mathcal{F}(u)\|_X^2 \geq \frac{2m}{C^2} (\mathcal{F}(u) - \mathcal{F}(u^*))$ and so combining this inequality with (A.13) yields (A.15). □

We now specialise the generic result proved above to the case of our variational energy.

Theorem A.15 (convergence rates for our variational energy with exact line search). Let E_{var} be the variational energy introduced in theorem 6.1, let V be a non-empty subspace of $\dot{H}^1$ (possibly equal to $\dot{H}^1$), and let p^* be the unique minimiser of E_{var} in V .

1. Let $p \in V$ and let $\pi = -\nabla_{\dot{H}^1} E_{var}|_V(p)$ denote the $\dot{H}^1$ -gradient descent direction for $E_{var}|_V$ at p . Then $E_{var}(p + \pi) - E_{var}(p^*) \leq \frac{1}{2} (E_{var}(p) - E_{var}(p^*))$.
2. Let $p \in V$ and let π denote the Newton descent direction for $E_{var}|_V$ at p . Then $E_{var}(p + \frac{1}{4}\pi) - E_{var}(p^*) \leq \frac{31}{32} (E_{var}(p) - E_{var}(p^*))$.
3. Suppose that $V \subseteq H^2$, suppose that there exist a constant $C_n > 0$ such that $\|p\|_{\dot{H}^1} \leq C_n \|p\|_{L^2}$ for every $p \in V$, and suppose that $M \in H^1$ and $PV \in L^2$. Let $p \in V$ and let $\pi = -\nabla_{L^2} E_{var}|_V(p)$ denote the L^2 -gradient descent direction for $E_{var}|_V$ at p . Then $E_{var}(p + \frac{1}{C_n^2} \pi) - E_{var}(p^*) \leq \left(1 - \frac{1}{2C_n^4}\right) (E_{var}(p) - E_{var}(p^*))$.

Proof. Since all of these descent directions are gradient descents, we need to verify that the hypotheses of proposition A.13 are verified. Item 4 of theorem A.1 guarantees that E_{var} is continuously Fréchet differentiable, item 1 of theorem A.1, lemma A.12, and lemma C.15 show that (A.12) is verified with $m = \frac{1}{2}$ and $M = 1$, and theorem A.2 verifies that E_{var} does indeed have a unique minimiser. The only hypothesis left to check is (A.11), and the value of the constant C varies from case to case since the inner product $(\cdot, \cdot)_X$ used to define the gradient differs in each case.

1. In this case $(\cdot, \cdot)_X = (\cdot, \cdot)_{\dot{H}^1}$ and so $C = 1$.
 2. In this case we may use the Hessian inner product recorded in proposition A.7 to see that $\|\phi\|_{D_p^2 E_{var}}^2 = \int_{\mathbb{T}^3} |\nabla \phi|^2 - \frac{1}{2} \mathbb{1}(M < \partial_3 p) \geq \frac{1}{2} \|\phi\|_{\dot{H}^1}^2$, i.e. $C = 2$.
 3. In this case $C = C_n$, by assumption.
-

Remark A.16 (on the convergence rate of L^2 -gradient descent). We comment on the constant C_n which appears in the third item of theorem A.15 above. Note that in practice such a constant exists since we work with V a finite-dimensional subspace of $\dot{H}^1$, in which case all norms on V are equivalent. Moreover, say that in such an instance the dimension of V is given by n (typically commensurate with the number of elements in a finite-element scheme). Then, as n approaches infinity, the constant C_n would necessarily grow unboundedly, which would in turn mean that $1 - \frac{1}{2C_n^2} \uparrow 1$. This means that in the case of L^2 -gradient descent the convergence rate remains linear but the constant involved worsens (i.e. approaches one) as the dimension of the subspace we work with increases (e.g. as the mesh is refined).

Note that this remark also apply when backtracking is used instead of exact line search—c.f. theorem A.19 below.

We now turn our attention towards the second, and arguably more important, set of convergence rates recorded in this section: those arising from *backtracking* line search. For we define backtracking line search below. Recall that, by contrast with exact line search discussed so far, backtracking line search can be implemented immediately in practical situations (which is typically not the case for exact line search) while retaining enough favourable structure such that provable convergence rates are still in play with backtracking line search.

Definition A.17 (backtracking line search). Consider a functional $\mathcal{F} : B \rightarrow \mathbb{R}$ on a Banach space B , let $x \in B$ be a current iterate, let $\Delta x \in B$ be a search direction, and consider two parameters $\alpha \in (0, 1/2)$ and $\beta \in (0, 1)$. Backtracking line search proceeds as follows: initiate the step size $t_0 = 1$ and, for $i \geq 0$, if

$$f(x + t_i \Delta x) > f(x) + \alpha t_i Df(x) \Delta x$$

then set $t_{i+1} := \beta t_i$, while if

$$f(x + t_i \Delta x) \leq f(x) + \alpha t_i Df(x) \Delta x \quad (\text{A.16})$$

then set the step size $t_* := t_i$ and terminate.

As we did above in the context of exact line search when specialising from the generic result of proposition A.13 to theorem A.15 which applies to the variational energy of interest, we first record a generic result for the convergence rate of gradient descent methods with *backtracking* line search. As with proposition A.13 above, such a result is classical (see [BV04]) and its proof is provided here in order to carefully track the constants that appear due to the presence of comparable norms.

Proposition A.18 (key inequality for the convergence rate of gradient descent with exact line search in Hilbert spaces). Consider the same notation and assumptions as proposition A.13 and let $\alpha \in (0, 1/2)$ and $\beta \in (0, 1)$. For any $u \in H$ and any $t \in [0, \frac{1}{C^2 M}]$,

$$\mathcal{F}(u - t \nabla_X \mathcal{F}(u)) \leq \mathcal{F}(u) - \alpha t \|\nabla_X \mathcal{F}(u)\|_X^2, \quad (\text{A.17})$$

i.e. the backtracking exit condition (A.16) holds. Moreover, for any $u \in H$ if t_* denotes the step size obtained via backtracking line search with parameters α and β and search direction $-\nabla_X \mathcal{F}(u)$ then

$$t_* \geq \min \left(1, \frac{\beta}{C^2 M} \right) \quad (\text{A.18})$$

and

$$\mathcal{F}(u - t_* \nabla_X \mathcal{F}(u)) - \mathcal{F}(u^*) \leq \left(1 - \frac{2\alpha m}{C^2} \min \left(1, \frac{\beta}{C^2 M} \right) \right) (\mathcal{F}(u) - \mathcal{F}(u^*)).$$

Proof. Proceeding as in the proof of proposition A.13 we deduce from (A.12) that

$$\mathcal{F}(u - t \nabla_X \mathcal{F}(u)) \leq \mathcal{F}(u) + \left(\frac{C^2 M t^2}{2} - t \right) \|\nabla_X \mathcal{F}(u)\|_X^2.$$

In particular we note that if $t \in [0, \frac{1}{C^2 M}]$ then $\frac{C^2 M t^2}{2} - t \leq -\frac{t}{2} < -\alpha t$, and so indeed (A.17) holds. As a consequence this means that either $t_* = t_0 = 1$ or $t_* = \beta t_i$ for some t_i which did *not* satisfy the backtracking exit condition, and so which must satisfy $t_i > \frac{1}{C^2 M}$. Putting these two options together establishes (A.18). Finally we note that t_* must satisfy the backtracking exit condition (A.17) and so combining that inequality with (A.12), (A.14) and (A.18) allows us to obtain the main inequality. $\square$

We now specialise proposition A.18 above to our variational energy.

Theorem A.19 (convergence rates for our variational energy with backtracking line search). Let E_{var} be the variational energy introduced in theorem 6.1, let V be a non-empty subspace of $\dot{H}^1$ (possibly equal to $\dot{H}^1$), let p^* denote the unique minimiser of E_{var} in V , and let $\alpha \in (0, 1/2)$ and $\beta \in (0, 1)$. In each case below let t_* denote the step size obtained via backtracking line search with direction π (chosen differently in each case) and parameters α and β .

1. Let $p \in V$ and let $\pi = -\nabla_{\dot{H}^1} E_{\text{var}}|_V(p)$ denote the $\dot{H}^1$ -gradient descent direction for $E_{\text{var}}|_V$ at p . Then

$$t_* \geq \beta \text{ and } E_{\text{var}}(p + t_*\pi) - E_{\text{var}}(p^*) \leq (1 - \alpha\beta)(E_{\text{var}}(p) - E_{\text{var}}(p^*)).$$

2. Let $p \in V$ and let π denote the Newton descent direction for $E_{\text{var}}|_V$ at p . Then

$$t_* \geq \frac{\beta}{4} \text{ and } E_{\text{var}}(p + t_*\pi) - E_{\text{var}}(p^*) \leq \left(1 - \frac{\alpha\beta}{16}\right)(E_{\text{var}}(p) - E_{\text{var}}(p^*)).$$

3. Suppose that $V \subseteq H^2$, suppose that there exist a constant $C_n > 0$ such that $\|p\|_{\dot{H}^1} \leq C_n \|p\|_{L^2}$ for every $p \in V$, and suppose that $M \in H^1$ and $PV \in L^2$. Let $p \in V$ and let $\pi = -\nabla_{L^2} E_{\text{var}}|_V(p)$ denote the L^2 -gradient descent direction for $E_{\text{var}}|_V$ at p . Then $t_* \geq \min\left(1, \frac{\beta}{C_n^2}\right)$ and

$$E_{\text{var}}(p + t_*\pi) - E_{\text{var}}(p^*) \leq \left(1 - \frac{\alpha}{C_n^2} \min\left(1, \frac{\beta}{C_n^2}\right)\right)(E_{\text{var}}(p) - E_{\text{var}}(p^*)).$$

Proof. This follows from proposition A.18 just as theorem A.15 follows from proposition A.13, noting that the appropriate values of the constant C in each case were identified in the proof of theorem A.15. $\square$

With the convergence rates in hand we change gear and turn our attention towards a consistency result. We seek to prove that minimisers of the variational energy restricted to (appropriate) finite-dimensional subspaces of $\dot{H}^1$ do converge to the true minimiser. In order to do so we first prove the following elementary estimate.

Lemma A.20 (explicit local modulus of continuity of the variational energy). Let E_{var} be the variational energy introduced in theorem 6.1. For any $p_1, p_2 \in \dot{H}^1$ the following estimate holds.

$$E_{\text{var}}(p_1) - E_{\text{var}}(p_2) \leq \left(\frac{1}{2}\|M\|_{L^2} + \frac{3}{4}\|\nabla p_1\|_{L^2} + \frac{3}{4}\|\nabla p_2\|_{L^2}\right)\|\nabla(p_1 - p_2)\|_{L^2}. \quad (\text{A.19})$$

Proof. We write $g_0(s) := \min(s, 0)$ such that

$$E_{\text{var}}(p_1) - E_{\text{var}}(p_2) = \int_{\mathbb{T}^3} \frac{1}{2} \left(|\nabla p_1|^2 - |\nabla p_2|^2\right) - \frac{1}{4} [g_0(M - \partial_3 p_1) - g_0(M - \partial_3 p_2)].$$

We then estimate

$$\left||\nabla p_1|^2 - |\nabla p_2|^2\right| = |\nabla(p_1 + p_2)| |\nabla(p_1 - p_2)| \leq (|\nabla p_1| + |\nabla p_2|) |\nabla(p_1 - p_2)|$$

and, since g_0 is 1-Lipschitz, we may proceed in the same way to estimate

$$\left|g_0(M - \partial_3 p_1)^2 - g_0(M - \partial_3 p_2)^2\right| \leq (2|M| + |\partial_3 p_1| + |\partial_3 p_2|) |\partial_3(p_1 - p_2)|.$$

Putting it all together produces the desired inequality. $\square$

We are now equipped to prove the last main result of this section, which pertains to numerical consistency.

Theorem A.21 (convergence of the minimisers of the approximate variational energy to the minimiser of the true variational energy). Let $V_n \subseteq \dot{H}^1$ be a closed subspace of $\dot{H}^1$, define $E_n := E_{\text{var}}|_{V_n}$, and let $p^* := \arg \min E_{\text{var}}$ and $p_n^* := \arg \min E_n$. For any $p_n^\varepsilon, p_n \in V_n$,

$$\begin{aligned} & \|p_n^\varepsilon - p^*\|_{\dot{H}^1} \\ & \leq 4[E_n(p_n^\varepsilon) - \min E_n] + (8\|M\|_{L^2} + 12\|PV\|_{H^{-1}} + 3\|p_n - p^*\|_{\dot{H}^1})\|p_n - p^*\|_{\dot{H}^1}. \end{aligned} \quad (\text{A.20})$$

In particular, if p_n^ε and p_n are chosen to satisfy, for some $\alpha, C > 0$,

$$E_n(p_n^\varepsilon) - \min E_n \leq \varepsilon \text{ and } \|p_n - p^*\|_{\dot{H}^1} \leq Cn^{-\alpha} \text{ for all } n \quad (\text{A.21})$$

then, for n sufficiently large to ensure $Cn^{-\alpha} \leq 1$,

$$\|p_n^\varepsilon - p^*\|_{\dot{H}^1} \lesssim \varepsilon + (1 + \|M\|_{L^2} + \|PV\|_{H^{-1}}) n^{-\alpha}, \quad (\text{A.22})$$

i.e. ε -minimisers of the discretised variational energy E_n can be made arbitrarily close to the minimiser of the true variational energy E_{var} .

Proof. First note that the p_n^* 's are well-defined since the properties of E_{var} parlay over to properties of E_n and the proof of theorem A.2 can be reproduced for E_n .

Since (A.22) follows from (A.20) under the additional assumptions of (A.21) it suffices to establish (A.20). First we note that, since p^* and p_n^* are minimisers of E_{var} and E_n , respectively, item 5 of theorem A.1 (which holds *mutatis mutandis* for E_n) tells us that

$$\begin{aligned} \|p_n^\varepsilon - p^*\|_{\dot{H}^1} &\leq \|p_n^\varepsilon - p_n^*\|_{\dot{H}^1} + \|p_n^* - p^*\|_{\dot{H}^1} \\ &\leq 4[E_n(p_n^\varepsilon) - E_n(p_n^*)] + 4[E_{\text{var}}(p_n^*) - E_{\text{var}}(p^*)]. \end{aligned} \quad (\text{A.23})$$

By virtue of theorem A.2 and A.20 we may estimate the second term above via

$$\begin{aligned} E_{\text{var}}(p_n^*) - E_{\text{var}}(p^*) &= E_n(p_n^*) - E_{\text{var}}(p^*) = \min E_n - E_{\text{var}}(p^*) \leq E_{\text{var}}(p_n) - E_{\text{var}}(p^*) \\ &\leq \left(\frac{1}{2} \|M\|_{L^2} + \frac{3}{4} [\|p_n\|_{\dot{H}^1} + \|p^*\|_{\dot{H}^1}] \right) \|p_n - p^*\|_{\dot{H}^1} \\ &\leq \left(2\|M\|_{L^2} + 3\|PV\|_{H^{-1}} + \frac{3}{4} \|p_n - p^*\|_{\dot{H}^1} \right) \|p_n - p^*\|_{\dot{H}^1}. \end{aligned} \quad (\text{A.24})$$

Putting (A.23) and (A.24) together then produces (A.20). $\square$

A.2. Geometry

In this section we do two things.

- Foreshadowed by proposition 3.1 we see that p and M may be viewed as *coordinates* along the balanced set $\mathcal{B}$. Moreover PV and M then arise naturally as *dual coordinates* (with respect to some natural inner product arising from the energy).
- Equipped with this geometric language we interpret the Newton descent iteration of section A.1 above in geometric terms. We explain how the Newton iteration can be viewed as a flow along the balanced set $\mathcal{B}$. This flow follows the coordinate direction along M while driving PV towards its true value.

We begin by recalling a definition from section 3.

Definition A.22 (global chart for the balanced set). For $\mathbb{H} := \dot{H}^1 \times L^2$ we define $\Phi : \mathbb{H} \rightarrow \mathcal{B}$ via

$$\Phi(p, M) = \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p + \frac{1}{2} \min_0 (M - \partial_3 p) \\ M - \partial_3 p - \frac{1}{2} \min_0 (M - \partial_3 p) \end{pmatrix}.$$

We now collect some more properties of this map.

Proposition A.23. Φ is bi-Lipschitz, i.e. Lipschitz with Lipschitz inverse.

In particular, since the parametrisation Φ of $\mathcal{B}$ is bi-Lipschitz it follows that the balanced set is an honest-to-goodness manifold. This is depicted in figure 7.

Corollary A.24. The balanced set $\mathcal{B}$ is a Lipschitz-regular Hilbert manifold, modeled after the Hilbert space $\mathbb{H}$, with global chart Φ .

Proof of corollary A.24. This follows from immediately from proposition A.23. $\square$

Proof of proposition A.23. First we note that the inverse of Φ is given by

$$\Phi^{-1}(u, \theta, q) = \begin{pmatrix} \Delta^{-1} [\nabla_h^\perp \cdot u_h + \partial_3 (\theta - \min_0 q)] \\ \theta + q \end{pmatrix}.$$




Figure 7. A depiction of the global chart for the balanced set introduced in definition A.22.

Indeed, suppose that $(u, \theta, q) = \Phi(p, M)$. We see that

$$\text{sign } q = \text{sign}(M - \partial_3 p)$$

and so lemmas C.6 and C.5 tell us that

$$\begin{aligned} (u, \theta, q) = \Phi(p, M) &\iff \begin{cases} u = \nabla_h^\perp p \\ \theta = \partial_3 p + \frac{1}{2} \min_0(M - \partial_3 p) \\ q = M - \partial_3 p - \frac{1}{2} \min_0(M - \partial_3 p) \end{cases} \\ &\iff \begin{cases} u = \nabla_h^\perp p \\ \theta + q = M \\ \theta - \min_0 q = \partial_3 p \end{cases} \\ &\iff \begin{cases} -u_h^\perp + (\theta - \min_0 q) e_3 = \nabla p \\ u_3 = 0 \\ \theta + q = M \end{cases} \\ &\iff \begin{cases} p = \Delta^{-1} \nabla \cdot [-u_h^\perp + (\theta - \min_0 q) e_3] \\ u_3 = 0 \\ M = \theta + q \end{cases} \\ &\iff \begin{cases} p = \Delta^{-1} [\nabla_h^\perp \cdot u_h + \partial_3 (\theta - \min_0 q)] \\ M = \theta + q \\ u_3 = 0 \end{cases} \\ &\iff (p, M) = \Phi^{-1}(u, \theta, q). \end{aligned}$$

The fact that both Φ and Φ^{-1} are Lipschitz then follows from the fact that $s \mapsto \min_0 s$ is Lipschitz, that $\Delta : \dot{H}^1 \rightarrow H^{-1}$ is an isometry, and from lemma C.14. Indeed, for Φ we have that

$$\begin{aligned} \|\Phi(p_1, M_1) - \Phi(p_2, M_2)\|_{L^2}^2 &\lesssim \|\nabla p_1 - \nabla p_2\|_{L^2}^2 + \|M_1 - M_2\|_{L^2}^2 \\ &= \|(p_1, M_1) - (p_2, M_2)\|_{\dot{H}^1 \times L^2}^2 \end{aligned}$$

while for Φ^{-1} we have that

$$\begin{aligned} &\|\Phi^{-1}(u_1, \theta_1, q_1) - \Phi^{-1}(u_2, \theta_2, q_2)\|_{\dot{H}^1 \times L^2}^2 \\ &= \|\Delta^{-1} [\nabla_h^\perp \cdot u_{h,1} + \partial_3 (\theta_1 - \min_0 q_1) - \nabla_h^\perp \cdot u_{h,2} - \partial_3 (\theta_2 - \min_0 q_2)]\|_{\dot{H}^1}^2 \\ &\quad + \|(\theta_1 + q_1) - (\theta_2 + q_2)\|_{L^2}^2 \\ &= \|\nabla_h^\perp \cdot u_{h,1} + \partial_3 (\theta_1 - \min_0 q_1) - \nabla_h^\perp \cdot u_{h,2} - \partial_3 (\theta_2 - \min_0 q_2)\|_{H^{-1}}^2 \\ &\quad + \|(\theta_1 + q_1) - (\theta_2 + q_2)\|_{L^2}^2 \\ &\lesssim \|u_{h,1} - u_{h,2}\|_{L^2}^2 + \|(\theta_1 - \min_0 q_1) - (\theta_2 - \min_0 q_2)\|_{L^2}^2 + \|(\theta_1 + q_1) - (\theta_2 + q_2)\|_{L^2}^2 \\ &\lesssim \|(u_1, \theta_1, q_1) - (u_2, \theta_2, q_2)\|_{L^2}^2. \end{aligned}$$

□

Remark A.25 (alternate form of Φ^{-1}). We typically write Φ^{-1} in terms of the nonlinear PV -and- M inversion, and not in terms of the inversion of a Laplacian as is done in the proof of proposition A.23 above. This is because the Laplacian inversion does not use PV , but instead uses a potential vorticity based on the *buoyancy* (by contrast with PV which is based on the equivalent potential temperature θ) given by

$$PV_b = \nabla_h^\perp \cdot u_h + \partial_3 (\theta - \min_0 q).$$

Typically we do *not* want to use PV_b since it is not slow. That is why PV is usually preferable. However, for the purpose of the proof of proposition A.23, PV_b is preferable to PV since it lets us prove more easily that Φ^{-1} is Lipschitz.

So far we have been able to endow the balanced set $\mathcal{B}$ with the structure of a Lipschitz-regular manifold. Unfortunately, that is as regular as the balanced set gets due to the presence of $\min_0 q$ in the buoyancy term, which is Lipschitz but not differentiable.

Nonetheless we can formally differentiate $\min_0 q$ and thus endow the balanced set $\mathcal{B}$ with a formal differentiable structure. In order to do so, we introduce notation for its formal tangent space, for the formal derivative of the chart Φ , and for the inner product we will use (which comes from the conserved energy).

Definition A.26. Fix an indicator function H .

1. We define

$$T_H \mathcal{B} := \left\{ (u, \theta, q) \in \mathbb{L}_\sigma^2 : \partial_3 u_h - \nabla_h^\perp (\theta - qH) = 0, u_3 = 0, \oint u_h = 0, \text{ and } \oint (\theta - qH) = 0 \right\}.$$

2. We define $D_H \Phi : \mathbb{H} \rightarrow T_H \mathcal{B}$ via

$$D_H \Phi (p, M) := \begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p + \frac{H}{2} (M - \partial_3 p) \\ M - \partial_3 p - \frac{H}{2} (M - \partial_3 p) \end{pmatrix}.$$

3. For any $(u, \theta, q) \in \mathbb{L}_\sigma^2$ we define

$$\|(u, \theta, q)\|_{L_H^2}^2 := \int |u|^2 + \theta^2 + q^2 H + (\theta + q)^2$$

and let $\langle \cdot, \cdot \rangle_{L_H^2}$ denote the corresponding inner product.

As expected, the formal derivative of Φ completely characterises the formal tangent space. (This would be trivially true if the balanced set $\mathcal{B}$ was endowed with an honest-to-goodness differentiable structure by its chart Φ .)

Lemma A.27. $T_H \mathcal{B} = \text{im } D_H \Phi$.

Proof. Proceeding similarly to the proof of proposition 3.1 we introduce the ‘good unknown’ (see remark 3.2)

$$v_H := -u_h^\perp + (\theta - qH) e_3.$$

For $\mathcal{X} = (u, \theta, q)$ we have on one hand that, as per lemma C.5,

$$\mathcal{X} \in T_H \mathcal{B} \iff \nabla \times v_H, \oint v_H = 0, \text{ and } u_3 = 0.$$

On the other hand, proceeding as in the proof of proposition A.23 and using lemma C.6 tells us that

$$\mathcal{X} = D_H \Phi (p, M) \iff \begin{cases} -u_h^\perp + (\theta - qH) e_3 = \nabla p \\ u_3 = 0 \\ \theta + q = 0 \end{cases}$$




Figure 8. The formal tangent space to the balanced set introduced in definition A.26 is the image of the formal derivative of the global chart.

$$\iff \begin{cases} v_H = \nabla p \\ u_3 = 0 \\ \theta + q = M \end{cases}.$$

The Helmholtz decomposition (corollary C.2) allows us to conclude:

$$\begin{aligned} \mathcal{X} \in T_H \mathcal{B} &\iff \nabla \times v_H = 0, \oint v_H = 0, \text{ and } u_3 = 0 \\ &\iff \nabla \times v_H = 0, \oint v_H = 0, u_3 = 0, \text{ and } \theta + q =: M \\ &\iff v_H = \nabla p \text{ for some } p \in \dot{H}^1, u_3 = 0, \text{ and } \theta + q = M \\ &\iff \mathcal{X} = D_H \Phi(p, M) \text{ for some } (p, M) \in \mathbb{H} \\ &\iff \mathcal{X} \in \text{im } D_H \Phi. \end{aligned}$$

□

We now remark that, as suggested by the notation we chose to employ, $D_H \Phi$ may be viewed formally as the derivative of Φ . Since the balanced set $\mathcal{B}$ is the image of the chart Φ , it follows that the image of its formal derivative $D_H \Phi$ may be interpreted formally as the tangent space to that same balanced set. Since lemma A.27 above tells us that the image of $D_H \Phi$ is precisely $T_H \mathcal{B}$, we see that once again the notation was chosen to be suggestive: $T_H \mathcal{B}$ may be viewed formally as the tangent space to the balanced set $\mathcal{B}$. This is made precise in remark A.28 below and is depicted in figure 8.

Remark A.28. Formally, we may compute that, for $H := \mathbb{1}(q < 0)$ and for any derivative ∂ ,

$$\partial(\min_0 q) = (\partial q) \mathbb{1}(q < 0) = (\partial q) H.$$

This shows that $D_H \Phi$ is, formally, the derivative of Φ . Since Φ is a global chart for the Hilbert manifold $\mathcal{B}$ this means that we may formally view $T_H \mathcal{B}$, which is the image of $D_H \Phi$, as the tangent space of $\mathcal{B}$.

An alternative justification, more direct albeit still formal, is as follows. Let $\mathcal{X} = (u, \theta, q) : (-1, 1) \rightarrow \mathcal{B}$ be a curve of balanced states and let $\mathcal{Z}$ denote its derivative at zero, i.e.

$$\mathcal{Z} = (v, \phi, r) := \frac{d}{ds} \mathcal{X}(s) \Big|_{s=0}.$$

Since $\mathcal{X} \in \mathcal{B}$ it follows that

$$j(\mathcal{X}) = 0, w(\mathcal{X}) = 0, \text{ and } a(\mathcal{X}) = 0.$$

Both w and a are linear, so we deduce immediately that

$$w(\mathcal{Z}) = 0 \text{ and } a(\mathcal{Z}) = 0.$$

What about j ? Using the formal identity $\partial(\min_0 q) = (\partial q)H$ discussed above we compute that

$$0 = \frac{d}{ds} j(\mathcal{X}) \Big|_{s=0} = \frac{d}{ds} [\partial_3 u_h - \nabla_h^\perp (\theta - \min_0 q)] \Big|_{s=0} = \partial_3 v_h - \nabla_h^\perp (\phi - rH) =: j_H(\mathcal{Z}).$$

This means that $\mathcal{Z}$ belongs to $T_H \mathcal{B}$. In other words: if $\mathcal{X}$ is a curve of balanced states then its formal tangent vector $\mathcal{Z}$ lies in $T_H \mathcal{B}$. This warrants treating formally $T_H \mathcal{B}$ as the tangent space of the balanced set $\mathcal{B}$.

Note that since the chart Φ is invertible and since $D_H \Phi$ is its formal derivative, we may expect $D_H \Phi$ to inherit this invertibility. This is indeed the case. We record this result here since it will be of use later.

Lemma A.29. $D_H \Phi$ is invertible.

Proof. This follows as in the proof that Φ is invertible (proposition A.23) since $(1+H)(1-H/2) = 1$ and so

$$\begin{aligned} (u, \theta, q) = D_H \Phi(p, M) &\iff \begin{cases} u = \nabla_h^\perp p \\ \theta = \partial_3 p + \frac{H}{2}(M - \partial_3 p) \\ q = M - \partial_3 p - \frac{H}{2}(M - \partial_3 p) \end{cases} \\ &\iff \begin{cases} -u_h^\perp + (\theta - qH)e_3 = \nabla p \\ u_3 = 0 \\ \theta + q = M \end{cases} \\ &\iff \begin{cases} p = \Delta^{-1} \nabla \cdot [-u_h^\perp + (\theta - qH)e_3] \\ u_3 = 0 \\ M = \theta + q \end{cases} \\ &\iff \begin{cases} p = \Delta^{-1} [\nabla_h^\perp \cdot u_h + \partial_3 (\theta - qH)] \\ u_3 = 0 \\ M = \theta + q \end{cases} \end{aligned}$$

and so the inverse is given by

$$(D_H \Phi)^{-1}(u, \theta, q) = \begin{pmatrix} \Delta^{-1} [\nabla_h^\perp \cdot u_h + \partial_3 (\theta - qH)] \\ \theta + q \end{pmatrix}.$$

□

Note that in particular this allows us to view $D_H \Phi$ as recording the *coordinate vectors* arising from the chart Φ . This is made precise in remark A.30 below and depicted in figure 9.

Remark A.30 (coordinate vectors). Since $D_H \Phi$ is the formal derivative of Φ and since $D_H \Phi$ is an invertible linear map whose image is precisely the formal tangent space $T_H \mathcal{B}$, we may view $D_H \Phi$ as recording the coordinate vectors corresponding to Φ . Indeed, for any $\mathcal{Z} = (u, \theta, q) \in T_H \mathcal{B}$ there is a unique $(\pi, \mu) \in \mathbb{H}$ such that

$$\mathcal{Z} = (D_H \Phi)(\pi, \mu) = \underbrace{\begin{pmatrix} \nabla_h^\perp \pi \\ (1 - \frac{H}{2}) \partial_3 \pi \\ -(1 - \frac{H}{2}) \partial_3 \pi \end{pmatrix}}_{=: e_p(\pi)} + \underbrace{\begin{pmatrix} 0 \\ \frac{H}{2} \mu \\ (1 - \frac{H}{2}) \mu \end{pmatrix}}_{=: e_M(\mu)}$$

where e_p and e_M denote the coordinate ‘vectors’ along the p and M coordinates, respectively. Note that these are not actually vectors, but rather coordinate *operators*, namely $e_p : \dot{H}^1 \rightarrow T_H \mathcal{B}$ and $e_M : L^2 \rightarrow T_H \mathcal{B}$ such that $T_H \mathcal{B} = \text{im } e_p + \text{im } e_M$. This is depicted in figure 9.




Figure 9. The ‘columns’ of the formal derivative $D_H\Phi$, denoted by e_p and e_M , may be viewed as coordinate vectors on the formal tangent space.

A minor miracle now occurs: while p and M are coordinates along the balanced set $\mathcal{B}$, their dual coordinates turn out to be, essentially, PV and M ! We see this by computing the L_H^2 -adjoint of $D_H\Phi$.

Lemma A.31. *Note that the dual of $\mathbb{H}$ is $\mathbb{H}^* = H^{-1} \times L^2$. Equipping $T_H\mathcal{B}$ with the inner product L_H^2 , the adjoint $(D_H\Phi)^* : T_H\mathcal{B} \rightarrow \mathbb{H}^*$ is given by*

$$(D_H\Phi)^*(u, \theta, q) = \begin{pmatrix} -\nabla_h^\perp \cdot u_h - \partial_3 \theta + \frac{1}{2} \partial_3 [H(\theta + q)] \\ (1 + \frac{H}{2})(\theta + q) \end{pmatrix} = \begin{pmatrix} -PV + \frac{1}{2} \partial_3 (HM) \\ (1 + \frac{H}{2})M \end{pmatrix}.$$

Proof. For any $(p, M) \in \mathbb{H}$ and any $(u, \theta, q) \in T_H\mathcal{B}$ we compute by integrating by parts that

$$\begin{aligned} \langle D_H\Phi(p, M), (u, \theta, q) \rangle_{L_H^2} &= \int u \cdot \nabla_h^\perp p + \theta \left[\partial_3 p + \frac{H}{2}(M - \partial_3 p) \right] \\ &\quad + qH \left[M - \partial_3 p - \frac{H}{2}(M - \partial_3 p) \right] + (\theta + q)M \\ &= \int u \cdot \nabla_h^\perp p + \theta \partial_3 p - \frac{H}{2}(\theta + q) \partial_3 p + \left(1 + \frac{H}{2}\right)(\theta + q)M \\ &= \int -(\nabla_h^\perp \cdot u_h) p - (\partial_3 \theta) p + \frac{1}{2} \partial_3 [H(\theta + q)] + \left(1 + \frac{H}{2}\right)(\theta + q)M \\ &= \left\langle -\nabla_h^\perp \cdot u_h - \partial_3 \theta + \frac{1}{2} \partial_3 [M(\theta + q)], p \right\rangle_{H^{-1} \times \dot{H}^1} \\ &\quad + \left\langle \left(1 + \frac{H}{2}\right)(\theta + q), M \right\rangle_{L^2} \\ &= \left\langle \begin{pmatrix} -\nabla_h^\perp \cdot u_h - \partial_3 \theta + \frac{1}{2} \partial_3 [H(\theta + q)] \\ (1 + \frac{H}{2})M \end{pmatrix}, \begin{pmatrix} p \\ M \end{pmatrix} \right\rangle_{H^* \times H}. \end{aligned}$$

□

Remark A.32 (dual coordinate vectors). Just as in remark A.30 we viewed $D_H\Phi$ as recording the coordinate vectors corresponding to the chart Φ , we can view $(D_H\Phi)^*$ as recording their dual coordinate one-forms. Indeed, if we define, for any $(u, \theta, q) \in T_H\mathcal{B}$,

$$\begin{aligned} e_p^b(u, \theta, q) &:= -PV + \frac{1}{2} \partial_3 (HM) \text{ and} \\ e_M^b(u, \theta, q) &:= \left(1 + \frac{H}{2}\right)(\theta + q) \end{aligned}$$

then we see that, for any $(u, \theta, q) \in T_H\mathcal{B}$ and any $(\pi, \mu) \in \mathbb{H}$,

$$\begin{aligned} \langle (u, \theta, q), e_p(\pi) \rangle_{L_H^2} &= \langle e_p^b(u, \theta, q), p \rangle_{\dot{H}^1} \text{ and} \\ \langle (u, \theta, q), e_M(\mu) \rangle_{L_H^2} &= \langle e_M^b(u, \theta, q), M \rangle_{L^2}. \end{aligned}$$

This is depicted in figure 10.

We can see that these coordinates p and M are actually orthogonal. This is recorded below, in a result that will also come in handy when interpreting the Newton descent iteration geometrically.




Figure 10. A depiction of the dual coordinates $e_M^b = (1 + \frac{H}{2})M$ and $e_p^b = -PV + \frac{1}{2}\partial_3(HM)$ on the formal tangent space $T_H\mathcal{B}$.

Lemma A.33. For any $(p, M) \in \mathbb{H}$,

$$(D_H\Phi)^*(D_H\Phi)(p, M) = \begin{pmatrix} -\nabla \cdot (A_H \nabla p) \\ (1 + \frac{H}{2})M \end{pmatrix}$$

for $A_H := I - \frac{H}{2}e_3 \otimes e_3$ (as usual).

Remark A.34. Note that $(D_H\Phi)^*(D_H\Phi) = (-\Delta) \oplus 1$ when $H = 0$, which is the canonical isometry from $\mathbb{H} = \dot{H}^1 \times L^2$ to its dual $H^* = H^{-1} \times L^2$. This is because when $H = 0$ the dry case may be viewed as a special case of the moist Boussinesq system which corresponds to neglecting q (i.e. setting $q = 0$).

Proof of lemma A.33. Let $(p, M) \in \mathbb{H}$ and let us write $(u, \theta, q) := D_H\Phi(p, M)$. Then $M(u, \theta, q) := \theta + q = M$ while

$$PV(u, \theta, q) := \nabla_h^\perp \cdot u_h + \partial_3 \theta = \nabla \cdot (A_H \nabla p) + \frac{1}{2} \partial_3(HM).$$

Therefore

$$\begin{aligned} (D_H\Phi)^*(D_H\Phi)(p, M) &= \begin{pmatrix} -PV + \frac{1}{2}\partial_3(HM) \\ (1 + \frac{H}{2})M \end{pmatrix} \\ &= \begin{pmatrix} -\nabla \cdot (A_H \nabla p) - \frac{1}{2}\partial_3(HM) + \frac{1}{2}\partial_3(HM) \\ (1 + \frac{H}{2})M \end{pmatrix} \\ &= \begin{pmatrix} -\nabla \cdot (A_H \nabla p) \\ 1 + H/2 \end{pmatrix} \end{aligned}$$

as claimed. □

Remark A.35. Written in block form,

$$(D_H\Phi)^*(D_H\Phi) = \begin{pmatrix} -\nabla \cdot (A_H \nabla) & 0 \\ 0 & 1 + H/2 \end{pmatrix}$$

which is block-diagonal. This means that the coordinates p and M of the Hilbert manifold $\mathcal{B}$, which arise from the global chart Φ , may formally be viewed as L_H^2 -orthogonal (but not orthonormal).

In other words, for any $\mathcal{Z} \in T_H\mathcal{B}$, which may be written uniquely as

$$\mathcal{Z} = e_p(\pi) + e_M(\mu)$$

for $(\pi, \mu) \in \mathbb{H}$ and e_p and e_M as in remark A.30, we have that

$$\begin{aligned} \langle e_p(\pi), e_M(\mu) \rangle_{L_H^2} &= \int 0 + \left(1 - \frac{H}{2}\right) (\partial_3 \pi) \cdot \frac{H}{2} \mu - \left(1 - \frac{H}{2}\right) (\partial_3 \pi) \cdot \left(1 - \frac{H}{2}\right) \mu \cdot H + 0 \\ &= \int \frac{H}{4} (\partial_3 \pi) \mu - \frac{H}{4} (\partial_3 \pi) \mu = 0. \end{aligned}$$

This allows us to refine our comment of remark A.30: the decomposition $T_H\mathcal{B} = \text{im } e_p + \text{im } e_M$ is actually L_H^2 -orthogonal.

Moreover, for the dual one-forms $e_p^\flat$ and $e_M^\flat$ as defined in remark A.32 we have that

$$\mu = \theta + q \text{ and } \pi \text{ solves } \nabla \cdot (A_H \nabla \pi) = \nabla_h^\perp \cdot u_h + \partial_3 \theta. \quad (\text{A.25})$$

Indeed:

$$\left(1 + \frac{H}{2}\right) (\theta + q) = e_M^\flat(\mathcal{Z}) = e_M^\flat(e_p(\pi)) + e_M^\flat(e_M(\mu)) = 0 + \left(1 + \frac{H}{2}\right) \mu.$$

Similarly: on one hand

$$e_p^\flat(\mathcal{Z}) = -\nabla_h^\perp \cdot u_h - \partial_3 \theta + \frac{1}{2} \partial_3 [H(\theta + q)] = -\nabla_h^\perp \cdot u_h - \partial_3 \theta + \frac{1}{2} \partial_3 (HM) \quad (\text{A.26})$$

while on the other hand

$$e_p^\flat(e_p(\pi)) + e_p^\flat(e_M(\mu)) = -\nabla \cdot (A_H \nabla \pi) + \frac{1}{2} \partial_3 (H\mu) = -\nabla \cdot (A_H \nabla \pi) + \frac{1}{2} \partial_3 (HM). \quad (\text{A.27})$$

Equating (A.26) and (A.27) yields the equation for π recorded in (A.25), as desired.

Remark A.36 (comparison of the L_H^2 -orthogonal decomposition with well-known orthogonal decompositions). In remark A.32 above we have explained how any formal tangent vector may be written in a unique way in terms of the ‘coordinate vectors’ e_p and e_M .

More precisely, we know that for any formal tangent vector $\mathcal{Z} = (u, \theta, q) \in T_H\mathcal{B}$ there exists a unique pair $(\pi, \mu) \in \mathbb{H}$ such that

$$\begin{aligned} \mathcal{Z} = e_p(\pi) + e_M(\mu) \text{ where } & \begin{cases} e_M^\flat(e_M(\mu)) = e_M^\flat(\mathcal{Z}) \text{ and} \\ e_p^\flat(e_p(\pi)) = e_p^\flat(\mathcal{Z}) \end{cases} \\ \iff & \begin{cases} \mu = \theta + q \text{ and} \\ \pi \text{ solves } \nabla \cdot (A_H \nabla \pi) = PV. \end{cases} \end{aligned} \quad (\text{A.28})$$

We will now help the reader make better sense of this by showing how their favourite orthogonal decomposition, namely that of Fourier series, may be written in this way.

Fourier series tell us that every $f \in L^2(\mathbb{T}^3)$ has uniquely determined coordinates $(a_k)_k \in \ell^2(\mathbb{Z}^3)$ such that

$$f(x) = \sum_k a_k e^{2\pi i k \cdot x}. \quad (\text{A.29})$$

Given a function f , the process producing its Fourier series can be split into two steps.

1. We compute the amplitudes a_k .
2. We reconstruct f from its amplitudes.

Let us rewrite each of these two steps in a manner which will make it easier to directly compare with the L_H^2 -orthogonal decomposition discussed above.

- Step 1: The computation of each amplitude corresponds to a map $a_k : L^2(\mathbb{T}^3) \rightarrow \mathbb{C}$ defined by

$$a_k(f) = \int_{\mathbb{T}^3} f(x) e^{-2\pi i k \cdot x} dx.$$

- Step 2: We write the reconstruction step in an admittedly unusual manner. This will make it easier to compare Fourier series to the L_H^2 -orthogonal decomposition of our formal tangent vectors. If we view each a_k as corresponding to a coordinate direction in $L^2(\mathbb{T}^3)$ we may then encode this coordinate direction as a map $e_k : \mathbb{C} \rightarrow L^2(\mathbb{T}^3)$ given by

$$[e_k(a)](x) = a e^{2\pi i k \cdot x}.$$

Combining these formulations of the two steps used to produce Fourier series we may write (A.29) in the form

$$f = \sum_k e_k(a_k) \text{ where } a_k = \langle f, e^{2\pi i k \cdot x} \rangle_{L^2}. \quad (\text{A.30})$$

In particular, in the language we used to describe the formal tangent space $T_H\mathcal{B}$, each ‘coordinate vector’ e_k has a ‘dual coordinate vector’ $e_k^\flat$, which is really just its adjoint and is a map $e_k^\flat: L^2(\mathbb{T}^3) \rightarrow \mathbb{C}$. Using the L^2 inner product on $L^2(\mathbb{T}^3)$ and the standard complex inner product on $\mathbb{C}$ defined by $\langle x, y \rangle_{\mathbb{C}} = x\bar{y}$ for every $x, y \in \mathbb{C}$ we may compute these dual coordinates vectors: for any $k \in \mathbb{Z}^3$, any $f \in L^2(\mathbb{T}^3)$, and any $a \in \mathbb{C}$,

$$\langle e_k^\flat(f), a \rangle_{\mathbb{C}} = \langle f, e_k(a) \rangle_{L^2} = \int_{\mathbb{T}^3} f(x) \overline{ae^{2\pi i k \cdot x}} dx = \bar{a} \int_{\mathbb{T}^3} f(x) e^{-2\pi i k \cdot x} dx = \bar{a} a_k(f) = \langle a_k(f), a \rangle_{\mathbb{C}}.$$

In other words: $e_k^\flat = a_k$.

We may thus write (A.30), already an alternate formulation of (A.29), in a manner even more reminiscent of our treatment of coordinate vectors and dual coordinate vectors in $T_H\mathcal{B}$, namely

$$f = \sum_k e_k(a_k) \text{ where } a_k = e_k^\flat(f). \quad (\text{A.31})$$

Finally, to turn this expression into an exact parallel of our treatment of $T_H\mathcal{B}$ we must use the orthonormality of the Fourier basis. That is, we use the fact that

$$e_k^\flat(e_k(a)) = a_k(e_k(a)) = \int_{\mathbb{T}^3} ae^{2\pi i k \cdot x} e^{-2\pi i k \cdot x} dx = a.$$

We may therefore rewrite (A.31) in its final form as

$$f = \sum_k e_k(a_k) \text{ where } e_k^\flat(e_k(a_k)) = e_k^\flat(f). \quad (\text{A.32})$$

This is exactly the same form as (A.28)!

This section is concerned with the geometry of the balanced set $\mathcal{B}$ on its own, and not with how it is embedded in the broader state space $\mathbb{L}_\sigma^2$. Nonetheless, at this stage it is worth commenting on how the L_H^2 inner product endowed on the formal tangent space $T_H\mathcal{B}$ of $\mathcal{B}$ interacts with the wave set. The result below is depicted in figure 11.

Lemma A.37. *The wave set $\mathcal{W}$ is L_H^2 -orthogonal to the formal tangent space $T_H\mathcal{B}$ in the sense that, for any $\mathcal{Z} \in T_H\mathcal{B}$ and any $\mathcal{Y} \in \mathcal{W}$, $\langle \mathcal{Z}, \mathcal{Y} \rangle_{L_H^2} = 0$.*

Proof. lemma A.27 tells us that $\mathcal{Z} = D_H\Phi(\pi, \mu)$ for some $(\pi, \mu) \in \mathbb{H}$ while proposition 3.3 tells us that $\mathcal{Y} = \Psi(\sigma, w)$ for some $\sigma \in \mathring{H}_\sigma^1$ and $w \in L^2$. Then

$$\begin{aligned} \langle \mathcal{Z}, \mathcal{Y} \rangle_{L_H^2} &= \langle D_H\Phi(\pi, M), \Psi(\sigma, w) \rangle_{L_H^2} \\ &= \int \nabla_h^\perp \pi \cdot \sigma_h^\perp + \left(1 - \frac{H}{2}\right) (\partial_3 \pi) \cdot \sigma_3 + \left(1 - \frac{H}{2}\right) (\partial_3 \pi) \cdot \sigma_3 \cdot H + 0 \\ &= \int \nabla_h \pi \cdot \sigma_h + (\partial_3 \pi) \sigma_3 = \int \nabla \pi \cdot \sigma = 0 \end{aligned}$$

since σ is divergence-free. □

We are now ready to use the geometric language introduced in this section in order to provide a geometric interpretation of the Newton descent iteration presented in section A.1.

First, we note that we can rewrite the defining equation of the Newton descent search direction, namely (A.7), as follows.

Remark A.38. The Newton descent direction π solves, for $H := \mathbb{1}(M < \partial_3 p)$,

$$-\nabla \cdot (A_H \nabla \pi) = PV(\Phi(p, M)) - PV.$$

This follows from (A.7) upon observing that $\min_0(M - \partial_3 p) = H(M - \partial_3 p)$.

In other words, and this will play a key role in the geometric interpretation of Newton descent: π is determined by the potential vorticity discrepancy, i.e. by how far the potential vorticity at the current guess is from the true potential vorticity.




Figure 11. Depiction of lemma A.37: the wave set is L_H^2 -orthogonal to the formal tangent space to the balanced set.

We are now ready to describe the geometric interpretation of Newton descent.

Remark A.39 (geometric interpretation of Newton descent). Consider the following formal flow on the balanced set $\mathcal{B}$ and its formal tangent space $T_H \mathcal{B}$.

$$\begin{cases} \dot{\mathcal{X}}(s) = (D_H \Phi)^{-*}(PV - PV(\mathcal{X}), 0) \text{ and} \\ \mathcal{X}(0) = \Phi(p_0, M) \end{cases} \quad (\text{A.33})$$

where PV and M are given and where the starting point $p_0 \in \dot{H}^1$ may be freely chosen.

We claim that, formally, the Newton descent iteration may be viewed as a discretisation of this flow. First: what does this flow even do?

Well, since

$$e_p^b = -PV + \frac{1}{2} \partial_3(HM) \text{ and } e_M^b = \left(1 + \frac{H}{2}\right) M$$

we may view PV and M as dual coordinates along the balanced set $\mathcal{B}$. The map $(D_H \Phi)^{-*} : \mathbb{H}^* \rightarrow T_H \mathcal{B}$ thus associates to $(PV, M) \in \mathbb{H}^*$ the corresponding tangent vector in $T_H \mathcal{B}$. This means that the flow follows the PV -direction, keeping the value of M fixed, so as to approach the given target value of PV . This is depicted in figure 12. (An analogue of this flow is presented in remark A.40 below in a much simpler setting in order to make the intuition easier to follow.)

To see that the Newton iteration is a discretisation of this flow we rewrite the Newton iteration in a different way.

The Newton iteration proceed as follows.

- PV and M are given.
- Initialise the iteration by choosing an arbitrary $p_0 \in \dot{H}^1$.
- Iterate, for $k \geq 0$,
 1. Let $H_k := \mathbb{1}(M < \partial_3 p_k)$ and $PV_k := \Delta p_k + \frac{1}{2} \partial_3 \min_0(M - \partial_3 p_k)$.
 2. π solves $\nabla \cdot (A_{H_k} \nabla \pi) = PV - PV_k$.
 3. Let $p_{k+1} := p_k + \pi$.

Note that step 3 is written this way because, once the Newton iteration reaches the so-called ‘quadratic regime’ (which is guaranteed to occur once iterates are sufficiently close to the true solution), then exact line search always produces a unit step size. This is why we can update $p_{k+1} = p_k + t\pi$ with $t = 1$.

We are now ready to relate the Newton iteration to the flow above. We may interpret (really, rewrite) the Newton iteration geometrically as follows.




Figure 12. The Newton descent method follows curves of constant M until the appropriate value of PV is reached: at a point $\mathcal{X}$ on the balanced set, the flow corresponding to Newton descent follows the direction of the formal vector $\mathcal{Z}$ whose p -coordinate is given by the Newton descent search direction π .

- PV and M are given.
- Initialise the iteration by choosing an arbitrary $p_0 \in \hat{H}^1$, i.e. choosing an arbitrary $\mathcal{X}_0 = \Phi(p_0, M) \in \mathcal{B}$ with the correct M -coordinate.
- Iterate, for $k \geq 0$.
 1. Let $(p_k, M) := \Phi^{-1}(\mathcal{X}_k)$ and $PV_k := PV(\mathcal{X}_k)$.
 2. π solves $(D_H \Phi)^*(D_H \Phi)(\pi, 0) = (PV - PV_k, 0)$.
 3. Let $\mathcal{X}_{k+1} := \Phi((p_k, M) + (\pi, 0))$.

Note that step 2 is the same as above, merely written differently (actually: every step coincides!) Indeed lemma A.33 tells us that

$$(D_H \Phi)^*(D_H \Phi)(\pi, 0) = (-\nabla \cdot (A_H \nabla \pi), 0).$$

In other words the Newton iteration may be written in one line as

$$\mathcal{X}_{k+1} = \Phi(\Phi^{-1}(\mathcal{X}_k) + (\pi, 0)) \text{ where } \pi \text{ solves } (D_H \Phi)^*(D_H \Phi)(\pi, 0) = (PV - PV_k, 0).$$

When $\mathcal{X}_{k+1}$ is close to $\mathcal{X}_k$, which equivalently means that π is small, we may thus formally approximate

$$\mathcal{X}_{k+1} \sim \mathcal{X}_k + (D_H \Phi)(\pi, 0).$$

We compute that

$$\begin{aligned} (D_H \Phi)(\pi, 0) &= (D_H \Phi) \left[(D_H \Phi)^*(D_H \Phi) \right]^{-1} (PV - PV_k, 0) \\ &= (D_H \Phi)^{-*} (PV - PV(\mathcal{X}_k), 0). \end{aligned}$$

In other words:

$$\mathcal{X}_{k+1} \sim \mathcal{X}_k + (D_H \Phi)^{-*} (PV - PV(\mathcal{X}_k), 0)$$

which verifies that indeed the Newton iteration is a formal discretisation of the flow in (A.33) above.

Remark A.40 (a simple flow). We discuss here a very simple analogue of the flow presented in remark A.39. Consider the following flow on the two-dimensional plane:

$$\begin{aligned}\dot{x}(s) &= x_* - x(s) \\ \dot{y}(s) &= 0\end{aligned}$$

with initial conditions

$$\begin{aligned}x(0) &= x_0 \\ y(0) &= y_*\end{aligned}$$

where x_* and y_* are given (they play the role of the given target values of PV and M) and where x_0 (just like p_0) may be freely chosen. This means that the x -coordinate plays the role of the PV -coordinate whereas the y -coordinate plays the role of the M -coordinate.

Then the new unknown

$$\tilde{x}(s) := x_* - x(s)$$

satisfies

$$\dot{\tilde{x}}(s) = -\dot{x}(s) = \tilde{x}(s)$$

and so we deduce that the solution is given by

$$\begin{aligned}\tilde{x}(s) &= \tilde{x}(0) e^{-s} = (x_* - x_0) e^{-s} \\ y(s) &= y_*.\end{aligned}$$

In terms of the original unknowns x and y this yields

$$\begin{aligned}x(s) &= x_* - \tilde{x}(s) = (1 - e^{-s}) x_* + e^{-s} x_0 \\ y(s) &= y_*.\end{aligned}$$

Therefore the solution $(x, y)(s)$ approaches the point with coordinates (x_*, y_*) exponentially fast, along a trajectory of constant y .

Appendix B. The many facets of the Helmholtz decomposition

In this section we summarise the various facets of the Helmholtz decomposition. This is in addition to similar summaries being carried out in sections 7.1 and 7.2 where the decompositions underpinning the dry and moist Boussinesq systems, respectively, are summarised in a similar manner.

Note that the numbering below is the same as the numbering in the discussion of the dry and moist decompositions in sections 7.1 and 7.2 above. That is not an accident: it serves to highlight the parallels between these decompositions. Then again we caution the reader that the roles played by the potential and solenoidal pieces are now flipped: ∇p appeared in the slow balanced components in section 7.1, it now appears in the fast waves; σ appeared in the fast waves, it now appears in the slow balanced components.

1. For any three-dimensional vector field u the Helmholtz decomposition is

$$u = \nabla p + \sigma$$

where σ is a divergence-free vector field.

2. In the context of the Helmholtz decomposition, if we wish to use terminology from atmospheric science and talk about *balanced states*, a state u will be balanced precisely when it is a solenoidal vector field. In other words the balance in question is

$$u_B = \sigma,$$

where the subscript B denotes the balanced component of u .

While introducing this terminology in the context of the Helmholtz decomposition may appear overzealous, it does serve to make the parallels with the dry Boussinesq system, and ultimately with the *moist* Boussinesq system, more evident.

3. Since the second component of the Helmholtz decomposition is the balanced component, by analogy with the terminology of this paper we refer to the first component as the *wave component*⁴. By analogy with proposition 3.3 we may identify the *slow measurement* in this setting as the dual quantity which annihilates the wave set. Since the wave set is the set of potential vector fields this means that the slow dual quantities are precisely the *curl* and the *spatial average*: ∇p is the wave component of u and

$$\nabla \times \nabla p = 0 \text{ and } \oint \nabla p = 0.$$

4. In this case there is no subtlety to identifying the ‘good unknown’ since it is the vector field u itself. (See remark 3.2 for a discussion of what is meant by the term ‘good unknown’ and what role they play.)
5. The elliptic PDE used to reconstruct the balanced component can be viewed as a div-curl system. Indeed σ solves

$$\nabla \times \sigma = \nabla \times u \text{ subject to } \nabla \cdot \sigma = 0 \text{ and } \oint \sigma = \oint u.$$

Note that the forcing is precisely the slow dual quantities: here the curl and the spatial average are the analogs of the potential vorticity in the dry case since they are slow dual quantities which characterise the balanced component.

Alternatively we can view this div-curl system as a curl-squared system for the vector potential of σ . Writing $\sigma = c + \nabla \times \psi$ we see that $c = \oint u$ while ψ solves $\nabla \times (\nabla \times \psi) = \nabla \times u$. (This curl-squared system may be written as a Laplacian equation since we may without loss of generality impose that the vector potential ψ be divergence-free, and in that case the curl-squared operator acting on ψ reduces to a Laplacian.)

6. The elliptic PDE used to reconstruct the wave component is Laplace’s equation:

$$\Delta p = \nabla \cdot u.$$

Note that it is not the fast dual quantity, the divergence, which plays the role of a forcing term for the wave component.

7. Using the decomposition we can obtain a couple of global changes of coordinates. We can fully describe the vector field u in three ways.
 - (a) Describe it using u itself.
 - (b) Describe it using the dual quantities $\nabla \cdot u$, $\nabla \times u$, and $\oint u$.
 - (c) Describe it using the coordinates p and σ .
8. Recall that the Helmholtz decomposition may be viewed as a slow-fast decomposition for the low Mach number limit of the compressible Euler equations. In that case the leading-order operator of the dynamics is linear, and it is recorded in remark C.3. We bring this up now in order to later compare the leading order operators in the case of dry and moist Boussinesq. For now it suffices to note that, as is discussed in more detail in remark C.3, the Helmholtz decomposition is a decomposition into the kernel and the image for that operator: the balanced set is its kernel and the wave set is its image.
9. As discussed in remark C.3, the wave component of the Helmholtz decomposition can be seen to obey a wave equation. Once again, this is not necessarily worthy of note by itself, but we record it here to compare it with the dry and moist Boussinesq decompositions in sections 7.1 and 7.2 below.
10. When viewed as a slow-fast decomposition, the Helmholtz decomposition inherits the additional feature of being orthogonal with respect to the conserved energy, which happens to simply be the L^2 norm.
11. We note that the Helmholtz decomposition and its induced changes of coordinates discussed above yield Parseval-type identities which allow us to rewrite the energy in three different ways: in terms of the state u , in terms of the coordinates, and in terms of the dual quantities. We obtain that

$$\int |u|^2 = \int |\nabla p|^2 + |\sigma|^2 = \int |\nabla \Delta^{-1} (\nabla \cdot u)|^2 + \left| \nabla \times^{-1} \left(\nabla \times u, \oint u \right) \right|^2.$$

⁴ This terminology may also be justified more directly as follows, without appealing to how the Helmholtz decomposition is similar to the dry and moist Boussinesq decompositions. Indeed, the Helmholtz decomposition may be interpreted as a slow-fast decomposition for the low Mach limit of the compressible Euler equations (see remark C.3) and in this context the first component of the decomposition will precisely obey a wave equation.

Here we write $\sigma = (\nabla \times)^{-1}(\omega, a)$ to mean that σ solves

$$\nabla \times \sigma = \omega \text{ subject to } \nabla \cdot \sigma = 0 \text{ and } \oint \sigma = a.$$

12. Finally we observe that the extraction of the balanced component of a vector field u may be viewed as a projection. Indeed, if we define the balanced set, in the context of the Helmholtz decomposition, to be $\mathcal{B}_{Hel} := (L^2)^3_\sigma$, the space of divergence-free three-dimensional vector fields, then we see that

$$u_B = \arg \min_{v \in \mathcal{B}_{Hel}} \|u - v\|_{L^2}.$$

Appendix C

C.1. Helmholtz decomposition

In this section we record a precise statement of the Helmholtz decomposition. This is a well-known result which plays an important role in this paper and so we record it here, if only to record it in a manner with the notation consistent with the notation in the rest of this paper.

Theorem C.1 (Helmholtz). *For every $u \in L^2$ there exists a unique $p \in \dot{H}^1$ and a unique $\sigma \in L^2_\sigma$ such that $u = \sigma + \nabla p$.*

Proof. We define p to be the solution of $\nabla p = \nabla \cdot u$, ψ to be the solution of $\nabla \times (\nabla \times \psi) = \nabla \times u$, and let $c = \oint u$. Then $u = \nabla p + \sigma$ for $\sigma = \nabla \times \psi + c$ (as can be directly verified using the Fourier transform since here we are working on the three-dimensional torus). $\square$

In particular we will use the following corollary of the Helmholtz decomposition.

Corollary C.2. *For every $u \in L^2$, if $\nabla \times u = 0$ and $\oint u = 0$ then there exists a unique $p \in \dot{H}^1$ such that $u = \nabla p$.*

Proof. It follows from the Helmholtz decomposition that $u = \nabla p + \sigma$ for some uniquely determined p and σ . Since σ is uniquely determined by the curl and the average of u , which both vanish, we see that indeed $\sigma = 0$. $\square$

To conclude this section we discuss, as was alluded to in the introduction, how the Helmholtz decomposition may be viewed as a slow-fast decomposition for the compressible Euler equations.

Remark C.3 (the Helmholtz decomposition is a slow-fast decomposition for the compressible Euler equations). In this remark we discuss how the Helmholtz decomposition may be interpreted as a slow-fast decomposition for compressible Euler equations: the slow piece is governed by the incompressible Euler equations whereas the fast piece is governed by the wave equation of acoustics.

As starting point we take the following isentropic compressible Euler equations.

$$\begin{cases} \partial_t \rho + \nabla \cdot (\rho u) = 0 \text{ and} \\ \rho (\partial_t u + u \cdot \nabla u) + \frac{1}{\varepsilon^2} \nabla p = 0 \end{cases} \quad \begin{matrix} (C.1a) \\ (C.1b) \end{matrix}$$

where ε denotes the Mach number. Here we are considering the ideal gas law $p = \rho^\gamma$ for some $\gamma > 1$.

Then, for $D_t := \partial_t + \cdot \nabla$ denoting the advective derivative,

$$D_t p = \gamma \rho^{\gamma-1} D_t \rho = -\gamma \rho^{\gamma-1} \rho \nabla \cdot u = -\gamma p \nabla \cdot u$$

and so we may rewrite the system in terms of the pressure p and the velocity field u only as

$$\begin{cases} \partial_t p + u \cdot \nabla p + \gamma p \nabla \cdot u = 0 \text{ and} \\ p^{1/\gamma} (\partial_t u + u \cdot \nabla u) + \frac{1}{\varepsilon^2} \nabla p = 0. \end{cases} \quad \begin{matrix} (C.2a) \\ (C.2b) \end{matrix}$$

We may now *symmetrise* the system by changing the pressure coordinate and using instead $q := \frac{p - \bar{p}}{\varepsilon}$ for some constant $\bar{p}$, such that $p = \bar{p} + \varepsilon q$, which yields

$$\begin{cases} \partial_t q + u \cdot \nabla q + \frac{\gamma \bar{p}}{\varepsilon} \nabla \cdot u = 0 \text{ and} \\ p^{1/\gamma} (\partial_t u + u \cdot \nabla u) + \frac{1}{\varepsilon} \nabla q = 0, \end{cases} \quad \begin{matrix} (C.3a) \\ (C.3b) \end{matrix}$$

which is now symmetrised.

We now make the two timescale assumption (see section 4)

$$\mathcal{X}(t) = \mathcal{X}\left(t, \frac{t}{\varepsilon}\right),$$

denoting the fast timescale $\tau := t/\varepsilon$ and we make the expansion assumption

$$\mathcal{X} = \mathcal{X}_0 + \varepsilon \mathcal{X}_1 + O(\varepsilon^2)$$

for $\mathcal{X} = u, q$. In particular we see that

$$p = \bar{p} + \varepsilon q_0 + \varepsilon^2 q_1 + O(\varepsilon^3).$$

This yields

$$\frac{1}{\varepsilon} (\partial_\tau q_0 + \gamma \bar{p} \nabla \cdot u_0) + O(1) = 0$$

and

$$\bar{p}^{1/\gamma} \cdot \frac{1}{\varepsilon} \partial_\tau u_0 + \frac{1}{\varepsilon} \nabla q_0 + O(1) = 0.$$

This leading order this comes down to

$$\begin{cases} \partial_\tau q_0 + \gamma \bar{p} \nabla \cdot u_0 = 0 \text{ and} \\ \bar{p}^{1/\gamma} \partial_\tau u_0 + \nabla q_0 = 0. \end{cases} \quad \begin{matrix} (C.4a) \\ (C.4b) \end{matrix}$$

From now on in this remark we will focus on the leading order system, so we drop the subscript and consider

$$\begin{cases} \partial_t q + \gamma \bar{p} \nabla \cdot u = 0 \text{ and} \\ \bar{p}^{1/\gamma} \partial_t u + \nabla q = 0. \end{cases} \quad \begin{matrix} (C.5a) \\ (C.5b) \end{matrix}$$

Therefore the energy

$$E = \frac{1}{2} \int \gamma \bar{p}^{1+1/\gamma} |u|^2 + q^2 \quad (C.6)$$

is conserved since

$$\begin{aligned} \int \gamma \bar{p}^{1+1/\gamma} u \cdot \partial_t u + q \partial_t q &= \int \gamma \bar{p}^{1+1/\gamma} u \cdot \left(-\frac{1}{\bar{p}^{1/\gamma}} \nabla q \right) + q \cdot (-\gamma \bar{p} \nabla \cdot u) \\ &= -\gamma \bar{p} \int [u \cdot \nabla q + q (\nabla \cdot u)] = 0. \end{aligned}$$

We now set $\gamma = \bar{p} = 1$ for simplicity, reducing the leading order system to

$$\begin{cases} \partial_t q + \nabla \cdot u = 0 \text{ and} \\ \partial_t u + \nabla q = 0 \end{cases} \quad \begin{matrix} (C.7a) \\ (C.7b) \end{matrix}$$

which we may rewrite as

$$\partial_t \begin{pmatrix} q \\ u \end{pmatrix} = \mathcal{L} \begin{pmatrix} q \\ u \end{pmatrix} \text{ where } \mathcal{L} = \begin{pmatrix} 0 & \text{div} \\ \text{grad} & 0 \end{pmatrix}.$$

When $\gamma = \bar{p} = 1$ the conserved energy reduces to the standard L^2 inner product. Here we note that here the leading order operator $\mathcal{L}$ is skew-symmetric with respect to L^2 since the L^2 adjoint of div is $-\text{grad}$.

We then decompose $(L^2)^3 = (L^2)_\sigma^3 + \hat{H}^1$, i.e. writing $u = \sigma + \nabla \phi$ where σ is divergence-free, such that

$$\begin{pmatrix} q \\ u \end{pmatrix} = \begin{pmatrix} q \\ \nabla \phi \end{pmatrix} + \begin{pmatrix} 0 \\ \sigma \end{pmatrix}.$$

The first term in this decomposition is the fast acoustic term. The second term in this decomposition is the slow term obeying incompressible Euler. In particular: the kernel of the leading order operator $\mathcal{L}$ is precisely

the space $\{0\} \times (L^2)_\sigma^3$ to which the second component belongs, and so this is an honest-to-goodness slow-fast decomposition. (We could also readily verify that the image of $\mathcal{L}$ is the space in which the first component lives.)

In other words: the Helmholtz decomposition is precisely the slow-fast decomposition for this compressible Euler system!

Moreover considering $\partial_t(C.7a) - \nabla \cdot (C.7b)$ produces

$$(\partial_t^2 - \Delta)q = 0$$

and so, for S denoting the semigroup associated with the wave equation above,

$$q(t) = S\left(\frac{t}{\varepsilon}\right)\bar{q}(t),$$

where $\bar{q}(t)$ would then be determined by a solvability condition arising from the next-order dynamics.

C.2. Useful computations, identities, and facts

In this section we record a variety of computations, identities, and facts which are of use throughout the paper.

First we record a useful splitting of the curl operator into its horizontal and vertical components.

Lemma C.4 (horizontal–vertical decomposition of the curl). *For any sufficiently regular vector field v , $\nabla \times v = \partial_3 v_h^\perp - \nabla_h^\perp v_3 + (\nabla_h^\perp \cdot v_h)e_3$.*

Proof. Since $\nabla \times v = (\partial_2 v_3 - \partial_3 v_2, \partial_3 v_1 - \partial_1 v_3, \partial_1 v_2 - \partial_2 v_1)$ the result follows from appropriately grouping terms. $\square$

We now record a couple of identities concerning the divergence and curl of a vector field of the form of the ‘good unknowns’ (see remark 3.2). This is a very useful identity since it helps us tie the good unknown to the potential vorticity PV and to the thermal wind imbalance j (see also remark 3.2).

Lemma C.5. *Consider a sufficiently regular vector field v and scalar field ϕ . Then*

$$\nabla \cdot (-v_h^\perp + \phi e_3) = \nabla_h^\perp \cdot v_h + \partial_3 \phi \text{ and } \nabla \times (-v_h^\perp + \phi e_3) = \partial_3 v_h - \nabla_h^\perp \phi - (\nabla_h \cdot v_h)e_3.$$

In particular if v is divergence-free then $\nabla \times (-v_h^\perp + \phi e_3) = \partial_3 v_h - \nabla_h^\perp \phi + (\partial_3 v_3)e_3$.

Proof. This follows immediately from lemma C.4. $\square$

We now record a useful algebraic inversion which tells us that we may equivalently characterise the thermodynamic variables θ and q by instead using the buoyancy $b = \theta - \min_0 q$ and the moist variable $M = \theta + q$.

Lemma C.6. *Given any two scalars b and M the nonlinear algebraic system*

$$\begin{cases} \theta + q = M \text{ and} \\ \theta - \min_0 q = b \end{cases}$$

is invertible, with solution given uniquely by

$$\begin{cases} \theta = b + \frac{1}{2} \min_0 (M - b) \text{ and} \\ q = M - b - \frac{1}{2} \min_0 (M - b). \end{cases}$$

Proof. We can do this easily by introducing $H = \mathbb{1}(q < 0)$. Treating H as fixed, the system then becomes linear (and invertible) and we obtain

$$\begin{cases} \theta = b + \frac{1}{2} H (M - b) \text{ and} \\ q = (1 - H/2) (M - b). \end{cases}$$

In particular we note that $\text{sign } q = \text{sign}(M - b)$ and so $H(M - b) = \min_0 (M - b)$, from which the claim follows. $\square$

Below we relate another identity related to ‘good unknowns’. This time we obtain an identity for the vector field $-u_h^\perp + (\theta - qH)e_3$ which appears in the ‘fixed H ’ version of the dual quantity-to-coordinates map (see lemma 5.8 and remark 5.9).

Lemma C.7. *If, for some indicator function H ,*

$$\begin{aligned} u &= \nabla_h^\perp p + \sigma_h^\perp + we_3, \\ \theta &= \partial_3 p + \frac{H}{2}(M - \partial_3 p) + \sigma_3, \text{ and} \\ q &= M - \partial_3 p - \frac{H}{2}(M - \partial_3 p) - \sigma_3 \end{aligned}$$

for some divergence-free σ satisfying $\nabla_h^\perp \cdot \sigma_h = \partial_3 w$ then

$$-u_h^\perp + (\theta - qH)e_3 = \nabla p + A_H^{-1}\sigma$$

where $A_H = I - \frac{H}{2}e_3 \otimes e_3$ and so $A_H^{-1} = I_H e_3 \otimes e_3$ (since $(1 - H/2)(1 + H) \equiv 1$).

Proof. Since $(v^\perp)^\perp = -v_h$ we see immediately that $-u_h^\perp = \nabla_h p + \sigma_h$. Then we compute that

$$\theta - qH = \left(1 - \frac{H}{2}\right)\partial_3 p + \frac{H}{2}M + \sigma_3 - \frac{1}{2}H(M - \partial_3 p) + H\sigma_3 = \partial_3 p + (1 + H)\sigma_3.$$

The claim follows from combining these two observations. $\square$

We now record a well-posedness result for a div-curl system involving an elliptic matrix. This is used to verify that the ‘fixed H ’ version of the dual quantity-to-coordinate map is invertible.

Proposition C.8 (solvability of a div-curl system). *For any uniformly elliptic $B \in L^\infty$, $f \in (\dot{H}_\sigma^1)^*$, and $g \in \mathbb{R}^3$ there is a unique $\sigma \in (L^2)_\sigma^3$ satisfying*

$$\nabla \times (B\sigma) = f, \nabla \cdot \sigma = 0, \text{ and } \oint \sigma = g. \quad (\text{C.8})$$

Proof. First we show that a solution exists. We let $\psi \in \dot{H}_\sigma^1$ be the unique solution of

$$\nabla \times (\nabla \times \psi) = f$$

and then let $\pi \in \dot{H}^1$ be the unique solution of

$$\nabla \cdot (B^{-1}\nabla \pi) = -\nabla \cdot [B^{-1}(\nabla \times \psi + g)].$$

Then

$$\sigma := B^{-1}(\nabla \pi + \nabla \times \psi + g)$$

is a solution. Indeed

$$\oint B\sigma = g \text{ and } \nabla \times (B\sigma) = \nabla \times (\nabla \times \psi) = f$$

while

$$\nabla \cdot \sigma = \nabla \cdot (B^{-1}\nabla \pi) + \nabla \cdot [B^{-1}(\nabla \times \psi + g)] = 0.$$

The uniqueness of a solution then follows from the uniqueness of the Helmholtz decomposition since we have characterised the solution σ by the Helmholtz decomposition of $B\sigma$. If σ_1 and σ_2 are solutions of (C.8) then their Helmholtz decompositions

$$B\sigma_i = \nabla \pi_i + \nabla \times \psi_i + c_i,$$

for $i = 1, 2$, where $c_i \in \mathbb{R}^3$, satisfy

$$c_1 = \oint B\sigma_1 = g = \oint B\sigma_2 = c_2$$

and

$$\nabla \times (\nabla \times \psi_1) = \nabla \times (B\sigma_1) = f = \nabla \times (B\sigma_2) = \nabla \times (\nabla \times \psi_2),$$

which ensures that $c_1 = c_2$ and $\psi_1 = \psi_2$. This means that

$$B(\sigma_1 - \sigma_2) = \nabla(\pi_1 - \pi_2)$$

and so

$$\nabla \cdot [B^{-1} \nabla(\pi_1 - \pi_2)] = \nabla \cdot (\sigma_1 - \sigma_2) = 0,$$

which also ensures that $\pi_1 = \pi_2$, and hence that $\sigma_1 = \sigma_2$. □

The result below pertains to the leading order dynamics discussed in section 4. It tells us how the balanced and wave sets relate to the dynamics of solutions of $\partial_t \mathcal{X} + \mathcal{N}\mathcal{X} = 0$ for $\mathcal{N}$ as defined in (3.1). Then

Proposition C.9. *Let $\mathcal{X} = (u, \theta, q)$ be a solution of $\partial_t \mathcal{X} + \mathcal{N}\mathcal{X} = 0$ for $\mathcal{N}$ as defined in (3.1). Then*

$$\mathcal{X} = \underbrace{\begin{pmatrix} \nabla_h^\perp p \\ \partial_3 p + \frac{1}{2} \min_0 (M - \partial_3 p) \\ M - \partial_3 p - \frac{1}{2} \min_0 (M - \partial_3 p) \end{pmatrix}}_{\mathcal{X}_B} + \underbrace{\begin{pmatrix} \sigma_h^\perp + w e_3 \\ \sigma_3 \\ -\sigma_3 \end{pmatrix}}_{\mathcal{X}_W}$$

where $\mathcal{X}_B \in \mathcal{B}$ and $\mathcal{X}_W \in \mathcal{W}$ for $\mathcal{B}$ and $\mathcal{W}$ as in propositions 3.1 and 3.3, respectively, for $M := \theta + q$ and p the solution of

$$\Delta p + \frac{1}{2} \partial_3 \min_0 (M - \partial_3 p) = \underbrace{\nabla_h^\perp \cdot u_h + \partial_3 \theta}_{PV},$$

and where $\mathcal{X}_W := \mathcal{X} - \mathcal{X}_B$. Moreover

$$\partial_t \mathcal{X}_B = 0 \text{ and } \partial_t \mathcal{X}_W = \begin{pmatrix} \mathbb{P}_L(-\sigma_h - (1 + H_B)\sigma_3 e_3) \\ w \\ -w \end{pmatrix} =: \mathcal{L}_{W, H_B} \mathcal{X}_W \in \mathcal{W}$$

for $H_B := \mathbb{1}(q_B < 0) = \mathbb{1}(M < \partial_3 p)$. Finally: if $\mathcal{X}_W = (\sigma_h, 0, 0)$ then

$$\mathcal{L}_{W, H_B}^2 \mathcal{X}_W = -\mathcal{X}_W$$

and if $\mathcal{X}_W = (w e_3, \sigma_3, -\sigma_3)$ then

$$\mathcal{L}_{W, H_B}^2 \mathcal{X}_W = -(1 + H_B) \mathcal{X}_W.$$

In other words: the $\mathcal{B}$ -component is constant whereas the dynamics of the $\mathcal{W}$ -component induce an $\mathcal{X}_B$ -dependent dynamical system on $\mathcal{W}$.

Proof. The existence and uniqueness of p is established in [RTSS24]. The fact that the $\mathcal{B}$ -component is constant follows from the fact that $\partial_t PV = \partial_t M = 0$ and the fact that $\partial_t p$ solves

$$\nabla \cdot (A_B \nabla \partial_t p) = 0$$

for $A_B := I - \frac{1}{2} \mathbb{1}(M < \partial_3 p) e_3 \otimes e_3$ uniformly elliptic, such that indeed $\partial_t p$ vanishes. Then $\mathcal{X}_W := \mathcal{X} - \mathcal{X}_B$ belongs to $\mathcal{W}$ and takes the prescribed form by virtue of proposition 3.3. Moreover we note that $\text{im } \mathcal{L}_{W, H_B} \subseteq \mathcal{W}$ since

$$PV \begin{pmatrix} \mathbb{P}_L(-\sigma_h - (1 + H_B)\sigma_3 e_3) \\ w \\ -w \end{pmatrix} = -\nabla_h^\perp \cdot \sigma_h + \partial_3 w = 0$$

and

$$\mathcal{M} \begin{pmatrix} \mathbb{P}_L(-\sigma_h - (1 + H_B)\sigma_3 e_3) \\ w \\ -w \end{pmatrix} = w - w = 0.$$

Finally the identities involving $\mathcal{L}^2$ follow from direct computations. □

We now turn our attention to identities involving the conserved energy. First we record the computation verifying that the energy is indeed conserved.

Proposition C.10 (conserved energy). *Let (u, θ, q) be a solution of the moist Boussinesq system (2.7a)–(2.7c). Then the energy*

$$E = \frac{1}{2} \int |u|^2 + \theta^2 + \min_0^2 q + (\theta + q)^2$$

is conserved, meaning that $\frac{dE}{dt} = 0$.

Proof. We immediately note that, for $M := \theta + q$,

$$(\partial_t + u \cdot \nabla) M = 0.$$

Therefore, since u is divergence-free,

$$\frac{d}{dt} \int \frac{1}{2} (\theta + q)^2 = \int (\partial_t + u \cdot \nabla) \left(\frac{1}{2} M^2 \right) = \int M (\partial_t + u \cdot \nabla) M = 0.$$

Now observe that $\frac{d}{ds} \min_0^2 s = \min_0 s$ and so

$$(\partial_t + u \cdot \nabla) \left(\frac{1}{2} \min_0^2 q \right) = (\min_0 q) (\partial_t + u \cdot \nabla) q.$$

Therefore

$$\begin{aligned} \frac{d}{dt} \int |u|^2 + \frac{1}{2} \theta^2 + \frac{1}{2} \min_0^2 q &= \int u \cdot (\partial_t + u \cdot \nabla) u + \theta (\partial_t + u \cdot \nabla) \theta + (\min_0 q) (\partial_t + u \cdot \nabla) q \\ &= \frac{1}{\varepsilon} \int (-u \cdot e_3 \times u - u \cdot \nabla p + u_3 (\theta - \min_0 q)) - \frac{1}{\varepsilon} \int \theta u_3 + \frac{1}{\varepsilon} \int (\min_0 q) u_3 \\ &= \frac{1}{\varepsilon} \int (\nabla \cdot u) p = 0, \end{aligned}$$

and so $\frac{dE}{dt} = 0$ as claimed. $\square$

We now turn our attention to another variant of the energy, namely the ‘fixed H ’ energy used in section 5. We first show that the (formally) linearised operator is skew-adjoint with respect to that energy, which is also known in the geometric context of section A.2 as the L_H^2 inner product.

Lemma C.11. *Let H be an indicator function. The L_H^2 -adjoint of $\mathcal{L}_H$ is $-\mathcal{L}_H$.*

Proof. For any (u, θ, q) and (v, ϕ, r) in $\mathbb{L}_\sigma^2$ we compute that

$$\begin{aligned} \langle \mathcal{L}_H(u, \theta, q), (v, \phi, r) \rangle_{L_H^2} &= \int \mathbb{P}_L(u_h^\perp - \theta e_3 + q H e_3) \cdot v + u_3 \phi - u_3 r H + 0 \\ &= \int -u \cdot v_h^\perp - \theta v_3 + q v_3 H + u_3 \phi - u_3 r H \\ &= \int u \cdot (-v_h^\perp + \phi e_3 - r H e_3) + \theta (-v_3) + q \cdot v_3 \cdot H \\ &= \langle (u, \theta, q), (\mathbb{P}_L(-v_h^\perp + \phi e_3 - r H e_3), -v_3, v_3) \rangle_{L_H^2} \end{aligned}$$

and so indeed

$$\mathcal{L}_H^* \begin{pmatrix} v \\ \phi \\ r \end{pmatrix} = \begin{pmatrix} \mathbb{P}_L(-v_h^\perp + \phi e_3 - r H e_3) \\ -v_3 \\ v_3 \end{pmatrix} = -\mathcal{L}_H \begin{pmatrix} v \\ \phi \\ r \end{pmatrix}.$$

$\square$

We now record another property of the ‘fixed H ’ energy, showing that if the indicator is time-dependent, this energy *fails* to be conserved.

Proposition C.12. Suppose that H is a given indicator function and that $\mathcal{X} = (u, \theta, q)$ is a solution of $\partial_t \mathcal{X} + \mathcal{L}_H \mathcal{X} = 0$ for $\mathcal{L}_H$ as defined in (4.6). For

$$E_H(\mathcal{X}) := \frac{1}{2} \int |u|^2 + \theta^2 + q^2 H + (\theta + q)^2$$

we have that

$$\frac{d}{dt} E_H = \frac{1}{2} \int q^2 \partial_t H.$$

In particular E_H is constant if and only if the H -interface is contained in the q -interface, i.e. $\partial \{H = 1\} \subseteq \{q = 0\}$.

Proof. This follows from the fact that if $\mathcal{X}$ solves $\partial_t \mathcal{X} + \mathcal{L}_H \mathcal{X}$ then $M := \theta + q$ satisfies $\partial_t M = 0$ and so we may compute that

$$\int u \cdot \partial_t u + \theta \partial_t \theta + q (\partial_t q) H + M (\partial_t M) = \int u_h \cdot u_h^\perp - u_3 (\theta - qH) + \theta u_3 - q H u_3 = 0$$

(in other words: if H is constant in time then the operator $\mathcal{L}_H$ is skew-adjoint with respect to the inner product induced by E_H). □

We now turn our attention to a standard result verifying that pullback metrics are themselves metrics. This plays a crucial role in section 5 in order to eventually define the statistical divergence we introduce in this paper and which agrees with both the decomposition and the conserved energy.

Proposition C.13 (Pullback metric). Let S be a set, let (M, d) be a metric space, and let $F : S \rightarrow M$ be injective. The map $F^* d : S \times S \rightarrow [0, \infty)$ defined by

$$F^* d(x, y) = d(F(x), F(y))$$

for every $x, y \in S$ is a metric on S known as the pullback metric of d under F .

Proof. The pullback metric immediately inherits symmetry and the triangle inequality from d . Its positive-definiteness then follows from the positive-definiteness of d and the injectivity of F : for any $x, y \in S$,

$$F^* d(x, y) = 0 \iff d(F(x), F(y)) = 0 \iff F(x) = F(y) \iff x = y.$$

□

The last result we record in this appendix is a result used in section A.2 to verify that the balanced set is an honest-to-goodness Lipschitz-regular Hilbert manifold. The result in question is an elementary estimate.

Lemma C.14. For any $v \in (L^2)^3$ and any $\phi \in L^2$,

$$\|\nabla_h^\perp \cdot v_h + \partial_3 \phi\|_{H^{-1}} \leq \|v_h\|_{L^2} + \|\phi\|_{L^2}.$$

Proof. This follows from a direct computation since, for any $\psi \in \dot{H}^1$,

$$\begin{aligned} \langle \nabla_h^\perp \cdot v_h + \partial_3 \phi, \psi \rangle_{H^{-1} \times \dot{H}^1} &= \int (\nabla_h^\perp \cdot v_h + \partial_3 \phi) \psi \\ &= - \int v_h \cdot \nabla_h^\perp \psi + \phi \partial_3 \psi \leq (\|v_h\|_{L^2} + \|\phi\|_{L^2}) \|\psi\|_{\dot{H}^1}, \end{aligned}$$

as desired. □

C.3. Tools from convex analysis

In this section we record various results from convex analysis that are of use to us. This begins with a result on various characterisations of strong convexity and concludes with results pertaining to the quadratic remainder (which are used in section A.1 to obtain a quadratic upper bound on our variational energy, one of the two key ingredients used to deduce convergence rates for various descent methods).

First we recall a few first-order characterisations of strong convexity.

Lemma C.15 (equivalent characterisations of strong convexity). *Let $(H, \langle \cdot, \cdot \rangle)$ be a Hilbert space and let $f: H \rightarrow \mathbb{R}$ be Fréchet differentiable. For any $\mu > 0$ the following are equivalent, where each inequality holds for every $x, y \in H$ and $\theta \in [0, 1]$.*

$$f(\theta x + (1 - \theta)y) \leq \theta f(x) + (1 - \theta)f(y) - \theta(1 - \theta) \frac{\mu}{2} \|x - y\|^2 \quad (\text{C.9})$$

$$f(y) \geq f(x) + \langle Df(x), y - x \rangle + \frac{\mu}{2} \|x - y\|^2 \quad (\text{C.10})$$

$$\langle Df(x) - Df(y), x - y \rangle \geq \mu \|x - y\|^2. \quad (\text{C.11})$$

If any of these inequalities hold recall that we say that f is μ -strongly convex.

Proof. See definition 2.1.2 and theorem 2.1.9 in [Nes04]. □

We now turn our attention towards properties of the quadratic remainder. First we note that the map $f \mapsto T^f$ is linear.

Lemma C.16 (linearity of the quadratic remainder map). *For any $x \in \mathbb{R}^d$ the map $T_x: C^1(\mathbb{R}^d; \mathbb{R}) \rightarrow C^0(\mathbb{R}^d; \mathbb{R})$ defined by $f \mapsto T_x^f$ is linear.*

Proof. In light of the definition of the quadratic remainder this follows immediately from the derivative being linear. □

Second we observe that quadratic remainders can be used to characterise affine functions. This is also the first of two results teasing out the relation between affine functions and quadratic remainders

Lemma C.17. Vanishing quadratic remainders characterise affine functions *A continuously differentiable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is affine, meaning that there exist $v \in \mathbb{R}^d$ and $s \in \mathbb{R}$ for which $f(x) = v^T x + s$ for every $x \in \mathbb{R}^d$, if and only if its quadratic remainder T_x^f vanishes everywhere, i.e. $T_x^f(y) = 0$ for every $x, y \in \mathbb{R}^d$.*

Proof. The ‘only if’ direction follows from a direct computation. Conversely suppose that $T_x^f \equiv 0$. In particular this means that $T_0^f(y) = 0$, which may be rearranged into $f(y) = f(0) + \nabla f(0)^T y$, proving that f is affine. □

To conclude this section we note how the quadratic remainder behaves under affine transformations—this is the second of two results on the relation between affine functions and quadratic remainders.

Lemma C.18 (equivariance of the quadratic remainder map under affine transformations). *Consider a continuously differentiable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ and an affine function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}^d$. Then $T_x^{f \circ \Phi}(y) = T_{\Phi(x)}^f(\Phi(y))$ for every $x, y \in \mathbb{R}^d$.*


Proof. We compute that

$$T_x^{f \circ \Phi}(y) = f(\Phi(y)) - f(\Phi(x)) - (D\Phi)^T(x) (\nabla f)(\Phi(x)) \cdot (y - x). \quad (\text{C.12})$$

Since we seek to write $T^{f \circ \Phi}$ in terms of T^f the key step in (C.12) is to write $y - x$ in terms of $\Phi(y) - \Phi(x)$. To do so we note that Φ is affine and so lemma C.17 tells us that T^Φ vanishes, i.e. $\Phi(y) - \Phi(x) = D\Phi(x)(y - x)$ for every $x, y \in \mathbb{R}^d$. Plugging this into (C.12) yields

$$\begin{aligned} T_x^{f \circ \Phi}(y) &= f(\Phi(y)) - f(\Phi(x)) - (\nabla f)(\Phi(x)) \cdot (D\Phi)(x)(y - x) \\ &= f(\Phi(y)) - f(\Phi(x)) - (\nabla f)(\Phi(x)) \cdot (\Phi(y) - \Phi(x)) = T_{\Phi(x)}^f(\Phi(y)), \end{aligned}$$

as claimed. □



C.4. Exact nonlinear waves

In this section we record results pertaining to the nonlinear oscillations with non-vanishing averages which are present in the moist Boussinesq system. The key results are proposition C.19, which shows that a simple ODE leads to nonlinear oscillations with non-vanishing averages, and corollary C.22 which explains how this solutions of this simple ODE give rise to honest-to-goodness solutions of the moist Boussinesq system with oscillations of the same nature.

First we find an exact solution of this simple ODE.

Proposition C.19. *The ODE $q'' + q + \min_0 q = 0$ admits periodic weak solutions with nonzero average over their period.*

Proof. Define

$$Q(t) := \begin{cases} \sqrt{2} \sin t & \text{if } 0 \leq t \leq \pi \text{ and} \\ \sin \sqrt{2}(t - \bar{t}) & \text{if } \pi \leq t < T \end{cases}$$

for $T := \pi + \frac{\pi}{\sqrt{2}}$ and $\bar{t} := \pi - \frac{\pi}{\sqrt{2}}$, and note that Q is differentiable. Then define $q : [0, \infty) \rightarrow \mathbb{R}$ to be the T -periodic extension of Q , noting that q is also differentiable. Moreover: q' is continuous on $\{0, \pi\} + T\mathbb{N}$, i.e. on the set where it changes sign. Therefore q is a weak solution of $q'' + q + \min_0 q = 0$ since

$$q'' = q \text{ when } q > 0 \text{ and } q'' = 2q \text{ when } q < 0.$$

Finally the simple substitution $s = \sqrt{2}(t - \pi) + \pi$ shows that

$$\int_0^T q = \int_0^T Q = \int_0^\pi \sqrt{2} \sin t \, dt + \int_\pi^{2\pi} \frac{1}{\sqrt{2}} \sin s \, ds = \left(\sqrt{2} - \frac{1}{\sqrt{2}} \right) \int_0^\pi \sin t \, dt \neq 0,$$

as desired. □

We now seek to use proposition C.19 above to show that the moist Boussinesq system (2.7a)–(2.7c) admits solutions with the same nonlinear oscillations. As an intermediate step we consider an ODE system contained within the moist Boussinesq system (this notion is made precise in lemma C.21 below).

Corollary C.20. *The ODE system*

$$\begin{cases} \frac{dw}{dt} = \theta - \min_0 q, & \text{(C.13a)} \\ \frac{d\theta}{dt} = -w, \text{ and} & \text{(C.13b)} \\ \frac{dq}{dt} = w & \text{(C.13c)} \end{cases}$$

admits periodic solutions with nonzero averages over their period. Moreover these solutions may be chosen to satisfy $M := \theta + q = 0$.

Proof. Suppose that $M := \theta + q$ vanishes. Then $q = -\theta$ and so the ODE system reduces to

$$\frac{dw}{dt} = -q - \min_0 q \text{ and } \frac{dq}{dt} = w$$

such that q must satisfy

$$\frac{d^2 q}{dt^2} = \frac{dw}{dt} = -q - \min_0 q,$$

or equivalently

$$q'' + q + \min_0 q = 0, \quad (\text{C.14})$$

proposition C.19 then guarantees the existence of a periodic weak solution q_* of (C.14) with nonzero average over its period. Then (w_*, θ_*, q_*) is a solution of the original ODE system (C.13a)–(C.13c) satisfying $\theta_* + q_* = 0$ provided that

$$\theta_* = -q_* \text{ and } w_* = \frac{dq}{dt}.$$

So finally, for the period T and for $\bar{t}$ being defined as in proposition C.19 we have that

$$w_*|_{(0,T)} = \begin{cases} \sqrt{2} \cos t & \text{if } 0 \leq t < \pi \text{ and} \\ \sqrt{2} \cos \sqrt{2}(t - \bar{t}) & \text{if } \pi \leq t < T \end{cases}$$

and so

$$\int_0^T w_* = \int_0^\pi \sqrt{2} \cos t dt + \int_\pi^{2\pi} \cos s ds = (\sqrt{2} - 1) \int_0^\pi \cos t dt \neq 0.$$

So indeed w_* and θ_* also have nonzero averages over their period. $\square$

As alluded to earlier we now show that solutions of the ODE system of corollary C.20 give rise to solutions of the moist Boussinesq system (2.7a)–(2.7c).

Lemma C.21. *Let $\bar{w}, \bar{\theta}, \bar{q} : [0, \infty) \rightarrow \mathbb{R}$ be weak solutions of the ODE system (C.13a)–(C.13c). If we define*

$$u(t, x) = \bar{w}(t) e_3, \theta(t, x) = \bar{\theta}(t), \text{ and } q(t, x) = \bar{q}(t)$$

then (u, θ, q) is a weak solution of the moist Boussinesq system (2.7a)–(2.7c) with $\varepsilon = 1$.

Proof. When u, θ , and q are constant in space the advective terms vanish. Since the pressure necessarily solves $-\Delta p = \nabla \cdot (u \cdot \nabla u)$ it must therefore also vanish. Finally, since $e_3 \times u = u_h^\perp$, which also vanishes when $u = \bar{w}e_3$, the claim follows. $\square$

We are finally ready to show that the moist Boussinesq system (2.7a)–(2.7c) admits solutions with nonlinear oscillations whose time averages do not vanish.

Corollary C.22. *The moist Boussinesq system (2.7a)–(2.7c) admits periodic weak solutions with nonzero averages over their period. Moreover these solutions may be chosen to belong to the wave set $\mathcal{W}$ at every instant (for $\mathcal{W}$ as defined in proposition 3.3).*

Proof. This follows from corollary C.20 and lemma C.21 since together they guarantee the existence of weak solutions of the moist Boussinesq system (2.7a)–(2.7c) with $M := \theta + q = 0$. Moreover we know from lemma C.21 that these solutions take the form $u = \bar{w}e_3$ and $\theta = \bar{\theta}$. Therefore these solutions have vanishing potential vorticity, and they satisfy $M = 0$. proposition 3.3 therefore tells us that, indeed, these solutions belong to the wave set $\mathcal{W}$. $\square$

To conclude this section we record a result analogous to lemma C.21 above: now the full velocity is incorporated into the ODE system, not just its vertical component. This is of independent interest from lemma C.21 since it is used in the study of the Snell-type metric in section 5.2.

Lemma C.23. *Let $\bar{u}, \bar{\theta}, \bar{q} : [0, \infty) \rightarrow \mathbb{R}$ be weak solutions of the ODE system*

$$\begin{cases} \frac{du}{dt} = -u_h^\perp + (\theta - \min_0 q) e_3, & (\text{C.15a}) \\ \frac{d\theta}{dt} = -u_3, \text{ and} & (\text{C.15b}) \\ \frac{dq}{dt} = u_3. & (\text{C.15c}) \end{cases}$$

If we define

$$u(t, x) = \bar{u}(t), \theta(t, x) = \bar{\theta}(t), \text{ and} \\ q(t, x) = \bar{q}(t)$$

then (u, θ, q) is a weak solution of the moist Boussinesq system (2.7a)–(2.7c) with $\varepsilon = 1$.

Proof. This follows exactly as for lemma C.21. $\square$

C.5. PQG equations

Although the PQG equations are not the main object of study of this paper, as mentioned in section 8 they are a natural next step to study. We therefore record this systems here for the reader's convenience.

The dynamic equations of the PQG system are

$$\begin{cases} (\partial_t + u \cdot \nabla) PV = -\partial_3 u \cdot \nabla \theta \text{ and} \\ (\partial_t + u \cdot \nabla) M = \partial_3 \min_0 q. \end{cases}$$

The unknowns are u (with $u_3 = 0$), θ , and q , following the notation that was introduced in section 2.5. They are determined by p and M : u is determined by the geostrophic balance $u_h = \nabla_h^\perp p$, while θ and q can be recovered from inverting the hydrostatic balance $\theta - \min(0, q) = \partial_3 p$ and the definition of the moist variable $M := \theta + q$. Note that M is defined as $M := \theta + q$ based on the evolution equations for θ and q in (2.4a) and (2.4b) including the assumption (without loss of generality in the stably stratified regime) of $\Gamma_\theta = 2$ and $\Gamma_q = 1$; for other values of Γ_θ and Γ_q , an appropriate definition of M is $M := (\Gamma_q/(\Gamma_\theta - \Gamma_q))\theta + q$ [SS17]. Also note that u is purely horizontal since its vertical component vanishes, but still depends on all three spatial coordinates. The pressure p is in turn determined by the potential vorticity PV and M as the unique solution of the PV -and- M inversion

$$\Delta p + \frac{1}{2} \partial_3 (\min(M - \partial_3 p, 0)) = PV.$$

In particular, by the solvability of PV -and- M inversion established in [RTSS24] we see that u , θ , and q are entirely determined by PV and M . We can therefore view the full system as a nonlinear and nonlocal transport equation for PV and M . Alternatively, akin to what is done for dry QG, we can view the entire system as dependent solely on p and M .

The forcing term $\partial_3 \min_0 q$ in the dynamic equation for the moist variable M is the eponymous precipitation term. Indeed, when saturation is taken to occur at $q = 0$ as is done here, $\min_0 q$ describes precisely *liquid rainwater*, which is assumed in this model to fall directly downward at a constant speed. For simplicity all physical constants have been set to unity—the appropriate constants can be found in [SS17].

ORCID iD

Samuel N Stechmann  0000-0001-7188-7946

References

- [AZD08] Amerault C, Zou X and Doyle J 2008 Tests of an adjoint mesoscale model with explicit moist physics on the cloud scale *Mon. Weather Rev.* **136** 2120–32
- [Ban21] Bannister R N 2021 Balance conditions in variational data assimilation for a high-resolution forecast model *Q. J. R. Meteorol. Soc.* **147** 2917–34
- [BANZ97] Babin A, Mahalov A, Nicolaenko B and Zhou Y 1997 On the asymptotic regimes and the strongly stratified limit of rotating Boussinesq equations *Theoret. Computat. Fluid Dyn.* **9** 223–51
- [Bar95] Bartello P 1995 Geostrophic adjustment and inverse cascades in rotating, stratified turbulence *J. Atmos. Sci.* **52** 4410–28
- [BBP+01] Barkmeijer J, Buizza R, Palmer T, Puri K and Mahfouf J-F 2001 Tropical singular vectors computed with linearized diabatic physics *Q. J. R. Meteorol. Soc.* **127** 685–708
- [BCF14] Buhler O, Callies J and Ferrari R 2014 Wave–vortex decomposition of one-dimensional ship-track data *J. Fluid Mech.* **756** 1007–26
- [BCZT14] Bousquet A, Coti Zelati M and Temam R 2014 Phase transition models in atmospheric dynamics *Milan J. Math.* **82** 99–128
- [BFH17] Breit D, Feireisl E and Hofmanová M 2017 Compressible fluids driven by stochastic forcing: the relative energy inequality and applications *Commun. Math. Phys.* **350** 443–73
- [BH16] Berry T and Harlim J 2016 Variable bandwidth diffusion kernels *Appl. Comput. Harmon. Anal.* **40** 68–96
- [BHK23] Bäumer D, Hittmeir S and Klein R 2023 Scaling approaches to quasigeostrophic theory for moist, precipitating air *J. Atmos. Sci.* **80** 1771–86
- [Blu72] Blumen W 1972 Geostrophic adjustment *Rev. Geophys.* **10** 485–528
- [BM04] Biello J A and Majda A J 2004 The effect of meridional and vertical shear on the interaction of equatorial baroclinic and barotropic Rossby waves *Stud. Appl. Math.* **112** 341–90
- [BM12] Branicki M and Majda A J 2012 Quantifying uncertainty for predictions with model error in non-Gaussian systems with intermittency *Nonlinearity* **25** 2543
- [Bre87] Bretherton C S 1987 A theory for nonprecipitating moist convection between two parallel plates. Part I: thermodynamics and “linear” solutions *J. Atmos. Sci.* **44** 1809–27
- [BV04] Boyd S and Vandenberghe L 2004 *Convex Optimization* (Cambridge University Press)
- [Cho68] Chorin A J 1968 Numerical solution of the Navier-Stokes equations *Math. Comput.* **22** 745–62
- [CHT+18] Cao Y, Hamouda M, Temam R, Tribbia J and Wang X 2018 The equations of the multi-phase humid atmosphere expressed as a quasi variational inequality *Nonlinearity* **31** 4692–723

- [CJT21] Cao Y, Jia C, Temam R and Tribbia J 2021 Mathematical analysis of a cloud resolving model including the ice micro-physics *Discrete Contin. Dyn. Syst.* **41** 131–67
- [CL06] Coifman R R and Lafon S 2006 Diffusion maps *Appl. Comput. Harmon. Anal.* **21** 5–30
- [CMS15] Chen S, Majda A J and Stechmann S N 2015 Multiscale asymptotics for the skeleton of the Madden-Julian oscillation and tropical–extratropical interactions *Math. Clim. Weather Forecast.* **1** 43–69
- [CMS16] Chen S, Majda A J and Stechmann S N 2016 Tropical–extratropical interactions with the MJO skeleton and climatological mean flow *J. Atmos. Sci.* **73** 4101–16
- [CMT14] Chen N, Majda A J and Tong X T 2014 Information barriers for noisy Lagrangian tracers in filtering random incompressible flows *Nonlinearity* **27** 2133
- [CT07] Cao C and Titi E S 2007 Global well-posedness of the three-dimensional viscous primitive equations of large scale ocean and atmosphere dynamics *Ann. Math.* **166** 245–67
- [CT12] Cover T M and Thomas J A 2012 *Elements of Information Theory* (Wiley)
- [Cul18] Cullen M 2018 The use of semigeostrophic theory to diagnose the behaviour of an atmospheric GCM *Fluids* **3** 72
- [Cul21] Cullen M 2021 *The Mathematics of Large-Scale Atmosphere and Ocean* (World Scientific)
- [CZFT13] Coti Zelati M, Frémond M, Temam R and Tribbia J 2013 The equations of the atmosphere with humidity and saturation: uniqueness and physical bounds *Physica D* **264** 49–65
- [CZHK+15] Coti Zelati M, Huang A, Kukavica I, Temam R and Ziane M 2015 The primitive equations of the atmosphere in presence of vapour saturation *Nonlinearity* **28** 625–68
- [CZT12] Coti Zelati M and Temam R 2012 The atmospheric equation of water vapor with saturation *Boll. Unione Mat. Italy* **5** 309–36
- [Daf79] Dafermos C M 1979 The second law of thermodynamics and stability *Arch. Ration. Mech. Anal.* **70** 167–79
- [Dal93] Daley R 1993 *Atmospheric Data Analysis* (Cambridge University Press)
- [DARR14] Doyle J D, Amerault C, Reynolds C A and Reinecke P A 2014 Initial condition sensitivity and predictability of a severe extratropical cyclone using a moist adjoint *Mon. Weather Rev.* **142** 320–42
- [DB99] Derber J and Bouttier F 1999 A reformulation of the background error covariance in the ECMWF global data assimilation system *Tellus A* **51** 195–221
- [DG19] Das S and Giannakis D 2019 Delay-coordinate maps and the spectra of Koopman operators *J. Stat. Phys.* **175** 1107–45
- [DVMFH10] De Vries H, Methven J, Frame T H and Hoskins B J 2010 Baroclinic waves with parameterized effects of moisture interpreted using Rossby wave components *J. Atmos. Sci.* **67** 2766–84
- [ELS20] Early J J, Lelong M-P and Sundermeyer M 2020 A generalized wave-vortex decomposition for rotating Boussinesq flows with arbitrary stratification *J. Fluid Mech.* **912** 1–29
- [EM96] Embid P F and Majda A J 1996 Averaging over fast gravity waves for geophysical flows with arbitrary potential vorticity *Commun. PDEs* **21** 619–58
- [EM98] Embid P F and Majda A J 1998 Low Froude number limiting dynamics for stably stratified flow with small or finite Rossby numbers *Geophys. Astrophys. Fluid Dyn.* **87** 1–50
- [Ema94] Emanuel K A 1994 *Atmospheric Convection* (Oxford University Press)
- [ER99] Errico R M and Raeder K D 1999 An examination of the accuracy of the linearization of a mesoscale model with moist physics *Q. J. R. Meteorol. Soc.* **125** 169–95
- [Err97] Errico R M 1997 What is an adjoint model? *Bull. Am. Meteorol. Soc.* **78** 2577–92
- [FJN12] Feireisl E, Jin B J and Novotný A 2012 Relative entropies, suitable weak solutions and weak-strong uniqueness for the compressible Navier–Stokes system *J. Math. Fluid Mech.* **14** 717–30
- [FKN09] Ferguson J, Khouider B and Namazi M 2009 Two-way interactions between equatorially-trapped waves and the barotropic flow *Chinese Ann. Math. B* **30** 539–68
- [FKNZ16] Feireisl E, Klein R, Novotný A and Zatorska E 2016 On singular limits arising in the scale analysis of stratified fluid flows *Math. Models Methods Appl. Sci.* **26** 419–43
- [GG99] Gottwald G and Grimshaw R 1999 The formation of coherent structures in the context of blocking *J. Atmos. Sci.* **56** 3640–62
- [Gia19] Giannakis D 2019 Data-driven spectral decomposition and forecasting of ergodic dynamical systems *Appl. Comput. Harmon. Anal.* **47** 338–96
- [GK12] Gehne M and Kleeman R 2012 Spectral analysis of tropical atmospheric dynamical variables using a linear shallow-water modal decomposition *J. Atmos. Sci.* **69** 2300–16
- [GM12] Giannakis D and Majda A J 2012 Nonlinear Laplacian spectral analysis for time series with intermittency and low-frequency variability *Proc. Natl Acad. Sci.* **109** 2222–7
- [Gra98] Grabowski W W 1998 Toward cloud resolving modeling of large-scale tropical circulations: a simple cloud microphysics parametrisation *J. Atmos. Sci.* **55** 3283–98
- [GS96] Grabowski W W and Smolarkiewicz P K 1996 Two-time-level semi-Lagrangian modeling of precipitating clouds *Mon. Weather Rev.* **124** 487–97
- [HDMSS13] Hernandez-Duenas G, Majda A J, Smith L M and Stechmann S N 2013 Minimal models for precipitating turbulent convection *J. Fluid Mech.* **717** 576–611
- [HESS21] Hu R, Edwards T K, Smith L M and Stechmann S N 2021 Initial investigations of precipitating quasi-geostrophic turbulence with phase changes *Res. Math. Sci.* **8** 25
- [HKLT17] Hittmeir S, Klein R, Li J and Titi E S 2017 Global well-posedness for passively transported nonlinear moisture dynamics with phase changes *Nonlinearity* **30** 3676–718
- [HKLT20] Hittmeir S, Klein R, Li J and Titi E S 2020 Global well-posedness for the primitive equations coupled to nonlinear moisture dynamics with phase changes *Nonlinearity* **33** 3206–36
- [HMR85] Hoskins B J, McIntyre M E and Robertson A W 1985 On the use and significance of isentropic potential vorticity maps *Q. J. R. Meteorol. Soc.* **111** 877–946
- [HMRP16] Herbert C, Marino R, Rosenberg D and Puoquet A 2016 Waves and vortices in the inverse cascade of stratified turbulence with and without rotation *J. Fluid Mech.* **806** 165–204
- [HN13] Hieronymus M and Nycander J 2013 The buoyancy budget with a nonlinear equation of state *J. Phys. Oceanogr.* **43** 176–86
- [Hos82] Hoskins B J 1982 The mathematical theory of frontogenesis *Annu. Rev. Fluid Mech.* **14** 131–51

- [Kas78] Kasahara A 1978 Further studies on a spectral model of global barotropic primitive equations with Hough harmonic expansion *J. Atmos. Sci.* **35** 2043–51
- [KGK+22] Knippertz P et al 2022 The intricacies of identifying equatorial waves *Q. J. R. Meteorol. Soc.* **148** 2814–52
- [KM81] Klainerman S and Majda A 1981 Singular limits of quasilinear hyperbolic systems with large parameters and the incompressible limit of compressible fluids *Commun. Pure Appl. Math.* **34** 481–524
- [KM82] Klainerman S and Majda A 1982 Compressible and incompressible fluids *Commun. Pure Appl. Math.* **35** 629–51
- [KM06] Klein R and Majda A J 2006 Systematic multiscale models for deep convection on mesoscales *Theor. Comput. Fluid Dyn.* **20** 525–51
- [KM10] Klockner A and McDougall T J 2010 Influence of the nonlinear equation of state on global estimates of diapycnal advection and diffusion *J. Phys. Oceanogr.* **40** 1690–709
- [KMS12] Khouider B, Majda A J and Stechmann S N 2012 Climate science in the tropics: waves, vortices and PDEs *Nonlinearity* **26** R1
- [Kor21] Korn P 2021 Global well-posedness of the ocean primitive equations with nonlinear thermodynamics *J. Math. Fluid Mech.* **23** 71
- [KSPK22] Klein R, Schielicke L, Pfahl S and Khouider B 2022 QG–DL–Ekman: dynamics of a diabatic layer in the quasi-geostrophic framework *J. Atmos. Sci.* **79** 887–905
- [KSS23] Kooloth P, Smith L M and Stechmann S N 2023 Hamilton’s principle with phase changes and conservation principles for moist potential vorticity *Q. J. R. Meteorol. Soc.* **149** 1056–72
- [LCBA20] Lelong M-P, Cuyper Y and Bouruet-Aubertot P 2020 Near-inertial energy propagation inside a Mediterranean anti-cyclonic eddy *J. Phys. Oceanogr.* **50** 2271–88
- [LDT86] Dimet F-X L and Talagrand O 1986 Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects *Tellus A* **38** 97–110
- [Lei80] Leith C 1980 Nonlinear normal mode initialization and quasi-geostrophic theory *J. Atmos. Sci.* **37** 958–68
- [LH04] Lapeyre G and Held I 2004 The role of moisture in the dynamics and energetics of turbulent baroclinic eddies *J. Atmos. Sci.* **61** 1693–710
- [LLR11] Lainé A, Lapeyre G and Riviere G 2011 A quasigeostrophic model for moist storm tracks *J. Atmos. Sci.* **68** 1306–22
- [LM92] Lien R-C and Muller P 1992 Normal-mode decomposition of small-scale oceanic motions *J. Phys. Oceanogr.* **22** 1583–95
- [LM20] Lian R and Ma J 2020 Existence of a strong solution to moist atmospheric equations with the effects of topography *Bound. Value Probl.* **103** 34
- [Lor81] Lorenc A 1981 A global three-dimensional multivariate statistical interpolation scheme *Mon. Weather Rev.* **109** 701–21
- [Mah99] Mahfouf J-F 1999 Influence of physical processes on the tangent-linear approximation *Tellus A* **51** 147–66
- [Maj84] Majda A 1984 *Compressible Fluid Flow and Systems of Conservation Laws in Several Space Variables* vol 53 (Springer Science & Business Media)
- [Maj03] Majda A 2003 *Introduction to PDEs and Waves for the Atmosphere and Ocean (Courant Lecture Notes in Mathematics vol 9)* (New York University, Courant Institute of Mathematical Sciences; American Mathematical Society)
- [MB01] Majda A J and Bertozzi A L 2001 *Vorticity and Incompressible Flow* (Cambridge University Press)
- [MB03] Majda A J and Biello J A 2003 The nonlinear interaction of barotropic and equatorial baroclinic Rossby waves *J. Atmos. Sci.* **60** 1809–21
- [MC18] Marques C A and Castanheira J M 2018 Diagnosis of free and convectively coupled equatorial waves *Math. Geosci.* **50** 585–606
- [ME98] Majda A J and Embid P 1998 Averaging over fast gravity waves for geophysical flows with unbalanced initial data *Theor. Comput. Fluid Dyn.* **11** 155–69
- [MSS19] Marsico D H, Smith L M and Stechmann S N 2019 Energy decompositions for moist Boussinesq and anelastic equations with phase changes *J. Atmos. Sci.* **76** 3569–87
- [Mül95] Müller P 1995 Ertel’s potential vorticity theorem in physical oceanography *Rev. Geophys.* **33** 67–97
- [Nes04] Nesterov Y 2004 A basic course *Introductory Lectures on Convex Optimization (Applied Optimization vol 87)* (Kluwer Academic Publishers)
- [NHR15] Nycander J, Hieronymus M and Roquet F 2015 The nonlinear equation of state of sea water and the global water mass distribution *Geophys. Res. Lett.* **42** 7714–21
- [OS15] Ogrosky H R and Stechmann S N 2015 The MJO skeleton model with observation-based background state and forcing *Q. J. R. Meteorol. Soc.* **141** 2654–69
- [OS16] Ogrosky H R and Stechmann S N 2016 Identifying convectively coupled equatorial waves using theoretical wave eigenvectors *Mon. Weather Rev.* **144** 2235–64
- [OSB+01] Ochoa J, Sheinbaum J, Badan A, Candela J and Wilson D 2001 Geostrophy via potential vorticity inversion in the Yucatan Channel *J. Mar. Res.* **59** 725–47
- [OWSS25] Ogrosky H R, Wetzel A N, Smith L M and Stechmann S N 2025 Data assimilation errors due to incomplete accounting of balanced versus unbalanced moisture *Q. J. R. Meteorol. Soc.* **151** e4942
- [PD97] Park S K and Droegemeier K K 1997 Validity of the tangent linear approximation in a moist convective cloud model *Mon. Weather Rev.* **125** 3320–40
- [PS10] Pauluis O and Schumacher J 2010 Idealized moist Rayleigh–Bénard convection with piecewise linear equation of state *Commun. Math. Sci.* **8** 295–319
- [RDTM08] Raupp C F, Dias P L S, Tabak E G and Milewski P 2008 Resonant wave interactions in the equatorial waveguide *J. Atmos. Sci.* **65** 3398–418
- [RMBN15] Roquet F, Madec G, Brodeau L and Nycander J 2015 Defining a simplified yet “realistic” equation of state for seawater *J. Phys. Oceanogr.* **45** 2564–79
- [RTSS24] Remond-Tiedrez A, Smith L M and Stechmann S N 2024 A nonlinear elliptic PDE from atmospheric science: well-posedness and regularity at cloud edge *J. Math. Fluid Mech.* **26** 30
- [Sal98] Salmon R 1998 *Lectures on Geophysical Fluid Dynamics* (Oxford University Press)
- [Sch94] Schochet S 1994 Fast singular limits of hyperbolic PDEs *J. Differ. Equ.* **114** 476–512
- [SM15] Stechmann S N and Majda A J 2015 Identifying the skeleton of the Madden–Julian oscillation in observational data *Mon. Weather Rev.* **143** 395–416

- [SO14] Stechmann S N and Ogrosky H R 2014 The Walker circulation, diabatic heating and outgoing longwave radiation *Geophys. Res. Lett.* **41** 9097–105
- [SS17] Smith L M and Stechmann S N 2017 Precipitating quasigeostrophic equations and potential vorticity inversion with phase changes *J. Atmos. Sci.* **74** 3285–303
- [Ste05] Stevens B 2005 Atmospheric moist convection *Annu. Rev. Earth Planet. Sci.* **33** 605–43
- [SW02] Smith L M and Waleffe F 2002 Generation of slow large scales in forced rotating stratified turbulence *J. Fluid Mech.* **451** 145–68
- [TE85] Thorpe A and Emanuel K 1985 Frontogenesis in the presence of small stability to slantwise convection *J. Atmos. Sci.* **42** 1809–24
- [Tem69] Temam R 1969 Sur l'approximation de la solution des équations de Navier-Stokes par la méthode des pas fractionnaires (II) *Arch. Ration. Mech. Anal.* **33** 377–85
- [TL22] Tan S and Liu W 2022 The strong solutions to the primitive equations coupled with multi-phase moisture atmosphere *Physica D* **440** 133442
- [TW15] Temam R and Wu K J 2015 Formulation of the equations of the humid atmosphere in the context of variational inequalities *J. Funct. Anal.* **269** 2187–221
- [TW16] Temam R and Wang X 2016 Approximation of the equations of the humid atmosphere with saturation *System Modeling and Optimization* ed L Bociu, J-A Désidéri and A Habbal (Springer International Publishing) pp 21–42
- [Val17] Vallis G K 2017 *Atmospheric and Oceanic Fluid Dynamics* (Cambridge University Press)
- [VZ21] Vasylyevych S and Zagar N 2021 A high-accuracy global prognostic model for the simulation of Rossby and gravity wave dynamics *Q. J. R. Meteorol. Soc.* **147** 1989–2007
- [WSS+20] Wetzel A N, Smith L M, Stechmann S N, Martin J E and Zhang Y 2020 Potential vorticity and balanced and unbalanced moisture *J. Atmos. Sci.* **77** 1913–31
- [WSSM19] Wetzel A N, Smith L M, Stechmann S N and Martin J E 2019 Balanced and unbalanced components of moist atmospheric flows with phase changes *Chin. Ann. Math. B* **40** 1005–38
- [YHS03] Yang G-Y, Hoskins B and Slingo J 2003 Convectively coupled equatorial waves: a new methodology for identifying wave structures in observational data *J. Atmos. Sci.* **60** 1637–54
- [ŽF15] Žagar N and Franzke C L 2015 Systematic decomposition of the Madden-Julian Oscillation into balanced and inertio-gravity components *Geophys. Res. Lett.* **42** 6829–35
- [ZGK04] Zagar N, Gustafsson N and Kallen E 2004 Variational data assimilation in the tropics: the impact of a background-error constraint *Q. J. R. Meteorol. Soc.* **130** 103–25
- [ŽKT+15] Žagar N, Kasahara A, Terasaki K, Tribbia J and Tanaka H 2015 Normal-mode function representation of global 3-D data sets: open-access software for the atmospheric research community *Geosci. Model Dev.* **8** 1169–95
- [ZNV+23] Zagar N, Neduhál V, Vasylyevych S, Zaplotnik Z and Tanaka H L 2023 Decomposition of vertical velocity and its zonal wavenumber kinetic energy spectra in the hydrostatic atmosphere *J. Atmos. Sci.* **80** 2747–67
- [ZSS21a] Zhang Y, Smith L M and Stechmann S N 2021 Effects of clouds and phase changes on fast-wave averaging: a numerical assessment *J. Fluid Mech.* **920** A49
- [ZSS21b] Zhang Y, Smith L M and Stechmann S N 2021 Fast-wave averaging with phase changes: asymptotics and application to moist atmospheric dynamics *J. Nonlinear Sci.* **31** 46
- [ZSS22] Zhang Y, Smith L M and Stechmann S N 2022 Convergence to precipitating quasi-geostrophic equations with phase changes: asymptotics and numerical assessment *Phil. Trans. R. Soc. A* **380** 20210030