Halocline mixing due to heavy rainfall in the cenote's of the Yucatan Peninsula.

by

Aaron Coutino

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Applied Mathematics

Waterloo, Ontario, Canada, 2016

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The hydrodynamics of the Yucatan peninsula are extremely important due to the strain on water resources caused by increasing development along the eastern coast. The geology of the peninsula is unique geologically in that it is dominated by porous limestone (karst) which results in seawater permeating underneath a fresh water lens. This geology allows for large fractures or cave networks to form underground and for large sink holes to form on the surface. These sink holes are traditionally called cenotes. The goal of this research was to investigate how heavy rainfall impacts the system and to what extent mixing occurs. This was investigated using sensor data and through numerical simulations. Using salinity and temperature measurements along with rainfall data we were able to strongly match heavy rainfall events (such as the hurricanes in 2011 and 2013) to spikes in salinity and increases in temperature. Through using wavelet analysis we were able to find longer term trends in the data along with the first eight tidal components. Small scale experiments were performed to observe the impact that rainfall has on a stable stratification. Through this observation an analogue of oscillating jets was devised since rainfall is near impossible to simulate numerically. This analogue was then experimentally shown to produce similar results to rainfall. A numerical model (SPINS) was then used to simulate an experiment. A column of tracer was also added to the simulation to model a column of potassium permanganate. The simulation was analysed using spanwise averaged (through y) 2D plots, 2D slices (at a particular y or z value) and 3D volumetric plots. These revealed the generation of a turbulent region which propagated away from the jets and entrained quiescent fluid. In addition to traditional parameters such as horizontal velocity u, density ρ and kinetic energy $\frac{1}{2}(u^2 + v^2 + w^2)$, new parametrizations for mixing and stirring were introduced. These were used to analyse how the dye was impacted by the jets, specifically it's interaction with the turbulent front. This work is the first step to understanding this complicated physical system and has links across many disparate disciplines, from turbulence theory to anthropology.

Acknowledgements

First off, I would like to thank my supervisor Marek Stastna who has been the best mentor any graduate student could ever ask for; I have no idea what I would be doing now if I had not met him. I would also like to thank Ed Reinhardt who has given me an amazing opportunity in Mexico and has opened my eyes to being in the field. I would like to thank Jared Penney for reading, and editing this thesis (it really needed it). I would like to thank Kris Rowe for being an amazing friend to me through my growth in academia. I would also like to thank all my friends in the lab who have helped me along the way. And finally, I must thank my parents for always supporting me and giving me everything they could. I can never repay them for all they have given me.

Dedication

Esta tesis está dedicado a mis papas. Nuca los voy a poder repagar por todo lo que me han dado.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background

The Yucatan peninsula (see Figure 1.1) is a large geological formation primarily comprised of limestone. It encompasses three Mexican states and has a total area of 165,000 km^2. The porosity and permeability of the karst (limestone) results in a lack of surface flows as any surface precipitation quickly enters the extensive aquifer within the bedrock. Over large (geological) time scales, dissolution of the karst has led to a vast network of fractures and caves throughout the region. Within the peninsula, there are two types of large structures comprised of fractures and conduits: a 'Ring of Cenotes' [24], located along the impact crater of the Chicxulub asteroid that resulted in the extinction of the dinosaurs [27], and several zones of long linear fractures (the Sierrita de Ticul fault line, the Holbox fracture zone, the Rio Hondo block fault zone and the La Libertad fault zone [37]). These large structures provide a preferential flow direction for the groundwater within the aquifer that exists in the karst. Furthermore, due to the depth of the karst geology (below sea level), seawater can intrude from the coast below the groundwater within the aquifer. This leads to a system where a freshwater lens (which varies in thickness depending on proximity to the coast) exists above saline seawater below a certain depth [30]. See [2],[15],[28] for a more in depth discussion of the geology of this system.

Throughout human history, the most important resource to sustain human, and more specifically societal, development has been water. It is necessary for consumption, agriculture and residential uses. Specifically humans require fresh, potable water. According to the World Health Organization (WHO) guidelines [23], the upper limits for water potability are 1000 mg/L (or 1 ppt) of salt. Drinking water above these limits results in sodium toxicity as the kidneys cannot remove the salt quickly enough. Due to the lack of natural surface flowing rivers within the Yucatan peninsula, groundwater has been the primary source of fresh water since the area was first inhabited around 10,000 BCE [10],[5]. Specifically, discounting the southern-most portion of the peninsula, people have relied on the freshwater lens of the aquifer, which they have accessed through the cave networks of the region, through small freshwater lakes, and through large sinkholes called 'cenotes' (see Figure 1.2). These cenotes were of particular importance to the Mayan civilization which inhabited the region between 300 CE and 1100 CE. For example, there is evidence that the Maya would create channels for rainwater into these cenotes to resupply their freshwater [36]. The importance of water to their society can be highlighted by the many rituals and festivals dedicated to the water gods [26].

Recent research has posited a link between the 'collapse' of the Mayan civilization and climatic changes. The Mayans were used to temporary changes in rainfall from year to year, channelling rainwater into cenotes, specifically during rainy periods, for storage during

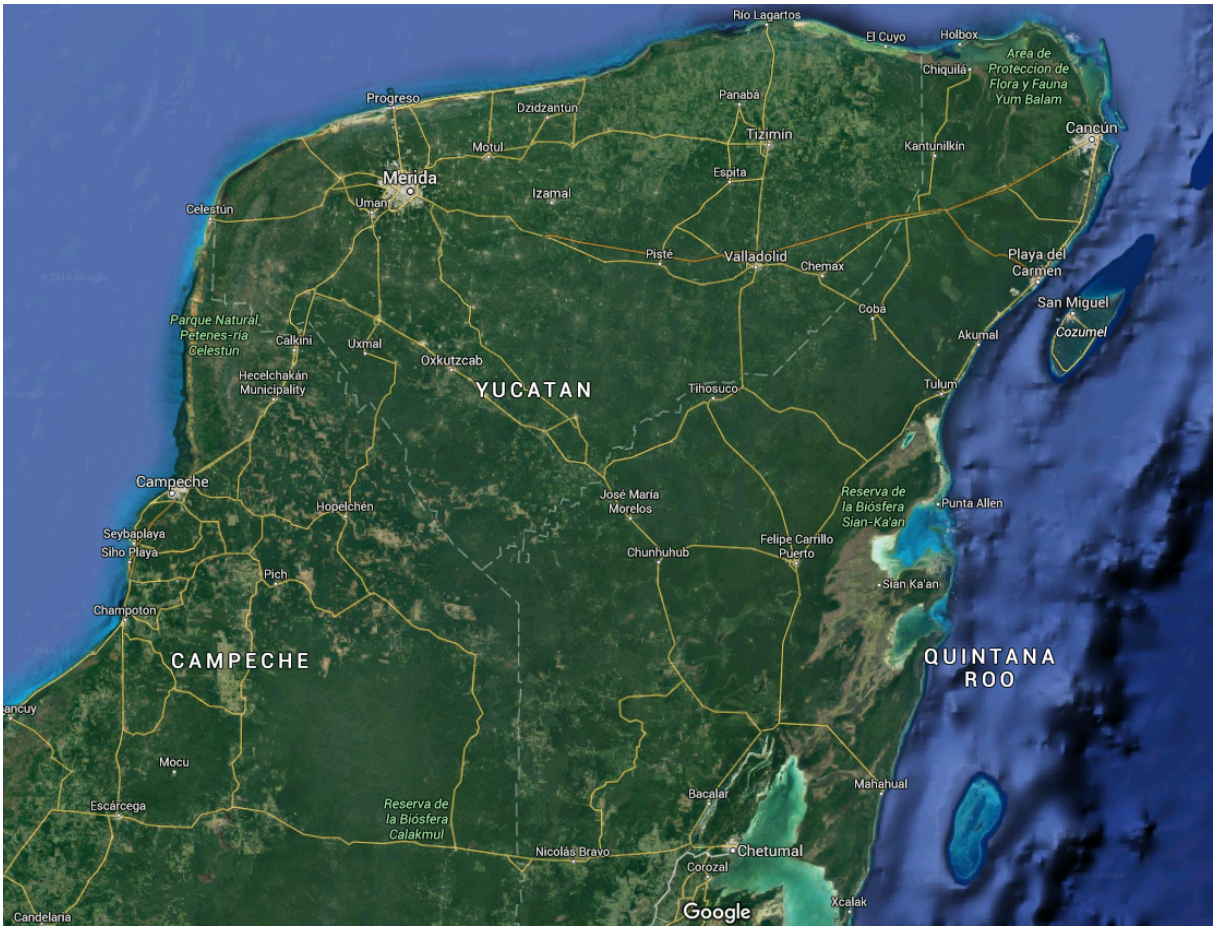

Figure 1.1: A map of the Yucatan peninsula (Google Maps). This thesis will primarily be concerned with the eastern coast, particularly the area surrounding the city of Tulum. (retrieved from https://www.google.ca/maps/@19.9129436,-88.5771341,425145m/data=!3m1!1e3?hl=en)

droughts [19], as mentioned previously. However, sediment records from the Cariaco Basin, Venezuela show that during the Terminal Classic Mayan period (800 CE-1000 CE), there were a series of prolonged multi-year droughts [14]. It is argued that the stress on water resources during this period contributed to the decline of the civilization. It is important to note that the 'collapse' was more of a collapse of the traditional social structures and a dispersion of people away from the previous metropolitan areas [16].

More recently, there has been a large amount of continuous development along the

Figure 1.2: The sacred cenote at Chichén Itzá which was used extensively for religious ceremonies. (retrieved from https://upload.wikimedia.org/wikipedia/commons/9/97/Mexico_Cenotes.jpg)

coast of Quintana Roo, the eastern shore of the Yucatan peninsula. Out of necessity, this development uses groundwater as its primary water supply. Despite this dependence on the groundwater system, there is comparatively little known about the interactions of the potable freshwater lens with the underlying seawater. This is particularly important given the number of large weather events that occur in this region which could impact this balance. In 1992, Moore et al. [21] performed a series of velocity and conductivity measurements along 70 km of the Quinatana Roo coast. They estimated velocities of $2.1 \cdot 10^{-4}$ m/s in the freshwater lens and $8.2 \cdot 10^{-4}$ m/s in the seawater within the dual-porosity karst. Inside large fractures within the freshwater lens, the authors measured a velocity of $1 \cdot 10^{-2}$ m/s 10 km inland and $1.2 \cdot 10^{-1}$ m/s along the coast. They attribute this change in velocity to the decrease in depth of the freshwater lens closer to the coast. From their conductivity measurements, the authors state that the halocline is in a steady-state position due to the rapid flow of fresh and brackish water towards the coast balanced by the convection of the seawater underneath. A more recent study was carried out in 2006 [3], where conductance and temperature measurements were taken in roughly the same section of coast. Their measurements within 1 km of the coast showed a very rapid increase in freshwater temperature and conductance, indicating substantial mixing with the underlying seawater. Inland, the rate of change of temperature and conductivity is much less, indicating less mixing. The authors further found that inland, within fractures, the

morphology and geometry are important factors in controlling the quantity of mixing that occurs. Finally, they showed that the underlying seawater cools as it travels inland, reaching an equilibrium with the freshwater lens 10 km inland. Though studies such as [21] and [3] provide some dynamic information about the aquifer system one key problem is that they only provide information at a specific point in time. These studies cannot provide any information about how the system varies through time and how it is affected by daily, seasonal or yearly cycles in the environment. A sustained measurement campaign carried out over multiple seasons would allow for a link to be made from a system driven by the current climate to the same hydrological system driven by past or future climates. Furthermore, without time variation we cannot observe how specific events (such as hurricanes) affect the amount of mixing that occurs.

To fill this gap in knowledge, a team headed by Dr. Eduard Reinhardt began placing temperature, salinity, and pressure sensors within the Yax Chen cave system (part of the larger Ox Bel Ha cave system, Figure 1.3 shows the entrance to the Yax Chen network), which remained submerged for long durations of time (approximately six months) and managed to capture several large scale weather events. This thesis is an analysis of the data from several of these sensors, along with a set of numerical simulations which will be shown to have forcing which is qualitatively similar effect to rain. This work is organized as follows: the remainder of Chapter 1 will review the background theory of wavelet analysis, provide examples of different signals and compare their wavelet spectrum, and then discuss the hydrodynamic background and theory used in the simulations. Chapter 2 describes the methodology and background information about the sensors and their placement, and is followed by analysis of the sensor data. Chapter 3 begins by showing how the numerical set up was generated from a series of reductions and simplifications of the original problem. We then outline the simulation set up and describe the numerical code used. We continue by analysing the simulation through primary variables, u, $KE = \frac{1}{2}(u^2+v^2+w^2)$ and ρ. We then examine the impact of the forcing on a dye column placed within the simulation. New parameters are introduced to quantify and distinguish stirring and mixing effects. Chapter 4 concludes the thesis by linking the numerical simulation to physical observations and placing the results of the sensor data in a larger scientific context. Finally, we discuss the current and future work that is planed.

Figure 1.3: The entrance to the Yax Chen cave system is located at the bottom and far end of this cenote. At this distance from the coast, there is heavy mangrove cover on the surface. Personal picture.

1.2 Wavelet analysis theory

Time-series analysis is a major part of the study of physical systems. In real-world systems, measurements are often gathered over long periods of time and include many irrelevant fluctuations. Many physical processes are periodic, though the exact period may vary over time or may be on very long time-scales. As a result, one of the primary ways in which time signals are analysed is through Fourier analysis. The Fourier transform of some temporal function $f(t)$ is given by

$$\mathcal{F}[f] \equiv \hat{f}(\omega) = \int_{-\infty}^{\infty} f(t) e^{-i\omega t} dt, \tag{1.1}$$

where ω is the frequency, t is the time and i is the square-root of negative one. The Fourier transform decomposes the function $f(t)$ into an infinite series of sine and cosine waves with differing frequencies. Provided that the function is sufficiently smooth this series will converge exactly to the original function. Two important theorems from Fourier analysis are called the Parseval and Placherel theorems (for proofs and additional information see standard references such as [13] or [12]).

Theorem 1.2.1 (Plancherel's theorem)**.** *If f and h are in $\mathbf{L^1}(\mathbb{R}) \bigcap \mathbf{L^2}(\mathbb{R})$ then,*

$$\int_{-\infty}^{\infty} f(t) h^*(t) dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(\omega) \hat{h}^*(\omega) d\omega, \tag{1.2}$$

where $*$ denotes the complex conjugate, $\wedge$ denotes the Fourier transform, $\mathbf{L^1}(\mathbb{R})$ denotes the space of Lebesgue integrable real functions and $\mathbf{L^2}(\mathbb{R})$ denotes the space of square integrable real functions.

Theorem 1.2.2 (Parseval's theorem)**.** *For $h = f$ it follows that*

$$\int_{-\infty}^{\infty} |f(t)|^2 dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |\hat{f}(\omega)|^2 d\omega. \tag{1.3}$$

These two results are extremely powerful because they apply to any function h that satisfies the conditions of Theorem 1.2.1, which means that we can use Fourier transforms to analyse more exotic transforms (such as the wavelet transform). They also show us that we can relate the 'energy' (norm squared) of any function to its spectral power. These results will become important later when we consider computing the wavelet transform.

While this definition of the Fourier transform is useful in theory, in practice, data is not given by a continuous signal. Rather, it is measured at discrete times with a finite

number of samples. If we assume a uniform sampling rate, and a spacing of Δt between our measurements, we can write a sample of some original function at time t as $f[t] = f[n\Delta t], n = 0...(N-1)$. Using the Dirac delta function we can write our discrete function as

$$f[t] = \sum_{n=0}^{N-1} f(n\Delta t)\delta(t - n\Delta t). \tag{1.4}$$

The Fourier transform of $\delta(t - n\Delta t)$ is $e^{-in\Delta t\omega}$ and so we can write the discrete Fourier transform as

$$\hat{f}[\omega] = \sum_{n=0}^{N-1} f(n\Delta t)e^{\frac{-i2\pi\omega n\Delta t}{N}}. \tag{1.5}$$

This process is very general and can be applied to all the continuous theory we will discuss. For this reason we will continue our discussion of the background theory using continuous notation but it can be easily extended to the discrete case.

The Fourier transform has been extremely influential throughout the development of modern physics, having been applied in areas from quantum theory to signal analysis. However, a major drawback of the analysis is that it provides no information about the location of the frequencies in time. This effect is a result of f and $\hat{f}$ being conjugate variables, and as such must obey an uncertainty principle. Under the Fourier transform each point in the frequency domain requires information about all points in the time domain, and vice-versa. Practically, this means that the more localized our function is in the time domain (the more precise it is) the broader its Fourier transform will be (less precise). This result also means that when we are in the Fourier domain we have zero information about the time localization and when we are in the time domain we have zero information about the frequency localization. For example consider Figure 1.4. In panels (a) and (b) we see two distinct signals that both result in the same spectrum (panels (c) and (d) respectively). This effect is a result of Fourier analysis not providing any time localization information. Since both frequencies of 10 seconds and 30 seconds appear in both signals (and with the same power) their spectra look the same. As a result of this effect, Fourier analysis is not adequate for extracting spectrum information from a signal when we would also like to know at what time a specific frequency is being activated. However, we can modify the traditional Fourier transform to acquire this information. The windowed Fourier transform, also called the Gabor transform (first introduced in [8]), is similar to the traditional transform, only with the addition of a windowing function. Mathematically this is written as

$$\hat{f}_{m,l}(\omega) = \int_{-\infty}^{\infty} f(t)g(t - m)e^{-ilt}dt, \tag{1.6}$$

where $g(t)$ is the windowing function which is real, symmetric, has compact support, and is centred around 0. The windowing function allows the transform to measure the contributions of sine waves of frequency l, locally around the point m. Traditionally, this window function and the sinusoidal components are combined and called the integration kernel, $g_{m,l}(t) \equiv g(t-m)e^{-itl}$. In relation to the uncertainty principle, the windowed Fourier transform can be thought of as losing certainty in the frequency domain but gaining some time localization information. Following the procedure in Chapter 4.1 of [20], in order to examine the uncertainty in each domain, we must look at the properties of the power of the integration kernels, $|g_{m,l}|^2$ and $|\hat{g}_{m,l}|^2$. Our transform no longer brings us to a one-dimensional domain, but we are now in a two-dimensional domain as a result of preserving some information from the time domain. Now consider a point in the time-frequency plane (m, l). Since our windowing function is guaranteed to have compact support, there is a specific 'width' of the function where it is non-negligible. The point (m, l) can only resolve times within this width, the same is true for frequency. To measure the extent to which each domain matters, we can calculate the variance for each domain centred around our point,

$$\sigma_t^2 = \int_{-\infty}^{\infty} (t-m)^2 |g_{m,l}(t)|^2 dt = \int_{-\infty}^{\infty} (t-m)^2 |g(t)|^2 dt, \tag{1.7}$$

and

$$\sigma_\omega^2 = \int_{-\infty}^{\infty} (\omega-l)^2 |\hat{g}_{m,l}(\omega)|^2 dt = \int_{-\infty}^{\infty} (t-l)^2 |\hat{g}(\omega)|^2 dt. \tag{1.8}$$

Using the substitution $v = t - m$ and cancelling the exponential terms, we can see that the variances are independent of m and l. Thus we will have a tiling of similarly-sized boxes all across the time-frequency domain as we vary (m, l). This is the price that we must pay to obtain information about both time and frequency, there is uncertainty as to the exact values. A depiction of these 'Heisenberg Boxes' can be seen in Figure 1.5.

The problem with the windowed Fourier transform is that we are limited by the choice of window, or more specifically by the maximum time-scale of our window (and correspondingly the maximum frequency-scale). Any frequency that is longer than our scales will be impossible to resolve. One solution to the problems with the windowed Fourier transform is the idea of the wavelet. The wavelet transform, first developed in [11], decomposes a signal into a wavelet basis, where a wavelet, ψ, is defined to have compact support, along with the following properties,

$$\int_{-\infty}^{\infty} \psi(t) dt = 0, \ ||\psi||^2 = 1. \tag{1.9}$$

Following the analysis in Chapter 4.3 of [20] and [17], there are two types of wavelets: real and complex. Real wavelets cannot provide amplitude and frequency information, but can

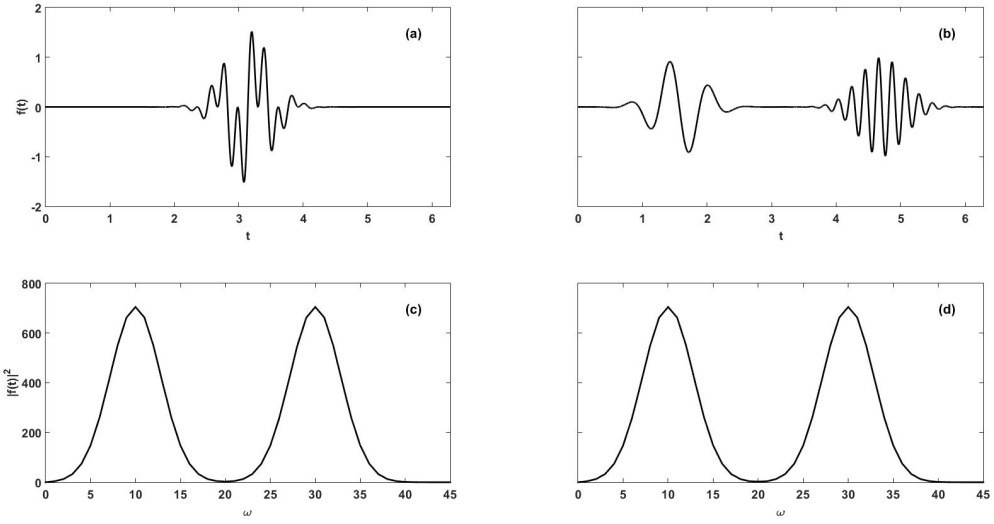

Figure 1.4: The signal in (a) is $\sin(10t)+\sin(30t)$ multiplied by a Gaussian envelope centred at π, while (b) is $\sin(10t)$ multiplied by the same Gaussian centred at $\pi/2$ plus $\sin(30t)$ multiplied by the same Gaussian centred at $3\pi/2$. Their spectra (c) and (d) respectively, appear almost identical. Since Fourier analysis does not provide any information about the time localization of the frequencies within a signal, cases where the same frequencies are active but at different times will result in the same spectra.

easily pick out isolated peaks or discontinuities. These properties are often used to detect sharp transitions and thus predominantly used in image processing. Complex wavelets are similar to the windowed Fourier transform and are used to separate both time and frequency information. For this reason we will be focusing on complex analytic wavelets, since they are the type that will be used for signal processing. Once an original 'mother' wavelet is chosen, it is scaled and translated to form a 'dictionary' (not necessarily a basis) of 'daughter' wavelets. Mathematically, we can write this as

$$\left\{\psi_{u,s}(t) = \frac{1}{\sqrt{s}}\psi\left(\frac{t-u}{s}\right)\right\}_{u\in\mathbb{R},s>0}. \tag{1.10}$$

The continuous wavelet transform is defined as

$$\mathcal{W}[f(u,s)] = \langle f, \psi_{u,s}\rangle = \int_{-\infty}^{\infty} f(t)\frac{1}{\sqrt{s}}\psi^*\left(\frac{t-u}{s}\right)dt. \tag{1.11}$$

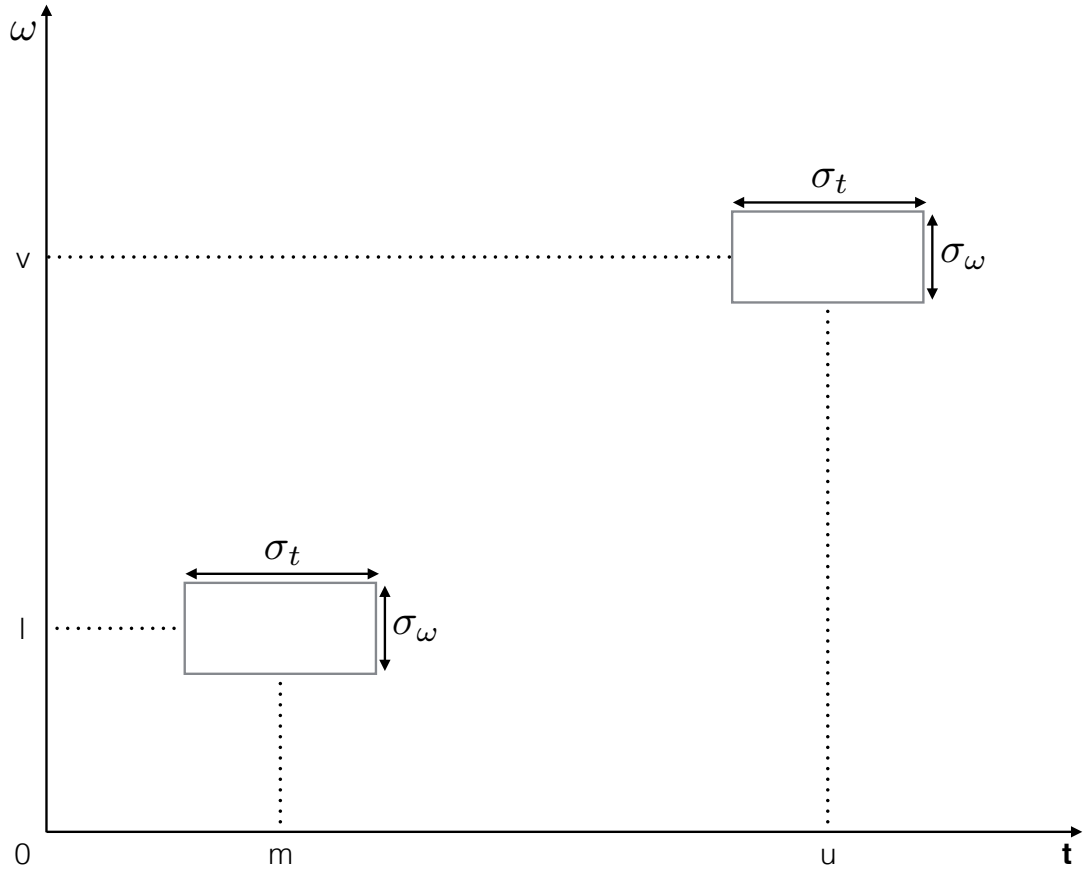

Figure 1.5: A depiction of the time-frequency (Heisenberg) boxes that tile the time-frequency plane. The time width of the box is defined by σ_t while the frequency width is defined by σ_ω. Since the variances are constant, the boxes form a symmetric tiling throughout the plane. This causes issues if we want to investigate low frequencies, which depend on much large time-scales than are resolved by the box.

The key difference and advantage of the wavelet transform over the windowed Fourier transform is in the time-frequency resolution. As for the windowed transform, we must investigate the variances of the integration kernel (in this case the wavelet), σ_t^2, and its conjugate variables variance, σ_ω^2. For the following calculations it will useful to note that the Fourier transform of $\psi_{u,v}(t)$ is

$$\hat{\psi}_{u,s}(\omega) = \sqrt{s}\hat{\psi}\left(s\omega\right) e^{-i\omega u}. \tag{1.12}$$

This states that the $\hat{\psi}_{u,v}$ can be considered a stretched version of $\hat{\psi}$. Since ψ is centred around zero, $\psi_{u,s}$ is centred around u. Using the change of variable $v = \frac{t-u}{s}$ it can be easily shown that in the time domain we have that

$$\sigma_t^2(s) = \int_{-\infty}^{\infty} (t-u)^2 |\psi_{u,s}(t)|^2 dt = s^2 \sigma_t^2, \tag{1.13}$$

with $\sigma_t^2 = \int_{-\infty}^{\infty} t^2 |\psi(t)|^2 dt$. This result corresponds to a stretching of our Heisenberg box. Since $\hat{\psi}(\omega) = \int_{-\infty}^{\infty} \psi(t) dt = 0$, the centre of the spectrum (also called the centre of the passing band), η, is located at some point away from origin. Since the spectrum is even and negative values do not provide additional information, we will limit ourselves to the positive real line. We can find the centre of the passing band by calculating the first moment,

$$\eta = \frac{\int_0^{\infty} \omega |\hat{\psi}(\omega)|^2 d\omega}{\int_0^{\infty} |\hat{\psi}(\omega)|^2 d\omega}. \tag{1.14}$$

Using the results calculated earlier for the Fourier transform of $\psi_{u,v}(t)$, the centre of the frequency is therefore η/s, so we can calculate,

$$\sigma_\omega^2(s) = \frac{1}{2\pi} \int_0^{\infty} (\omega - \frac{\eta}{s})^2 |\hat{\psi}_{u,v}(\omega)|^2 d\omega = \frac{\sigma_\omega^2}{s^2}, \tag{1.15}$$

with $\sigma_\omega^2 = \int_0^{\infty} (\omega - \eta)^2 |\hat{\psi}(t)|^2 d\omega$. These two results show that as we change our scale, we change the shape of the Heisenberg box and the frequency at which we are centred. This can also be thought of as our frequency controlling the scale: as we lower the frequency we increase the range of times considered, and the opposite is true as the frequency is increased. This results in a pyramid shaped tiling over the time-frequency domain which can be seen in Figure 1.6.

By systematically varying the scale, the wavelet transform allows us to tile the time-frequency plane much more efficiently than the windowed Fourier transform, and enables us capture the entire range of frequencies of the original signal. This is the reason we chose to use wavelet analysis over windowed Fourier analysis for our investigation of the sensor data from the cenotes and for the trace element sediment data. Returning to the example first given in Figure 1.4, if we now use the wavelet analysis to compare the two different signals we can readily distinguish the two. In Figure 1.7 we present the same signals from Figure 1.4 where we have taken the wavelet spectrum of the time-series and presented a pseudocolor plot of the time-frequency domain. From this plot we can extract the same frequency information that we could from the Fourier spectrum, (though for wavelet plots it

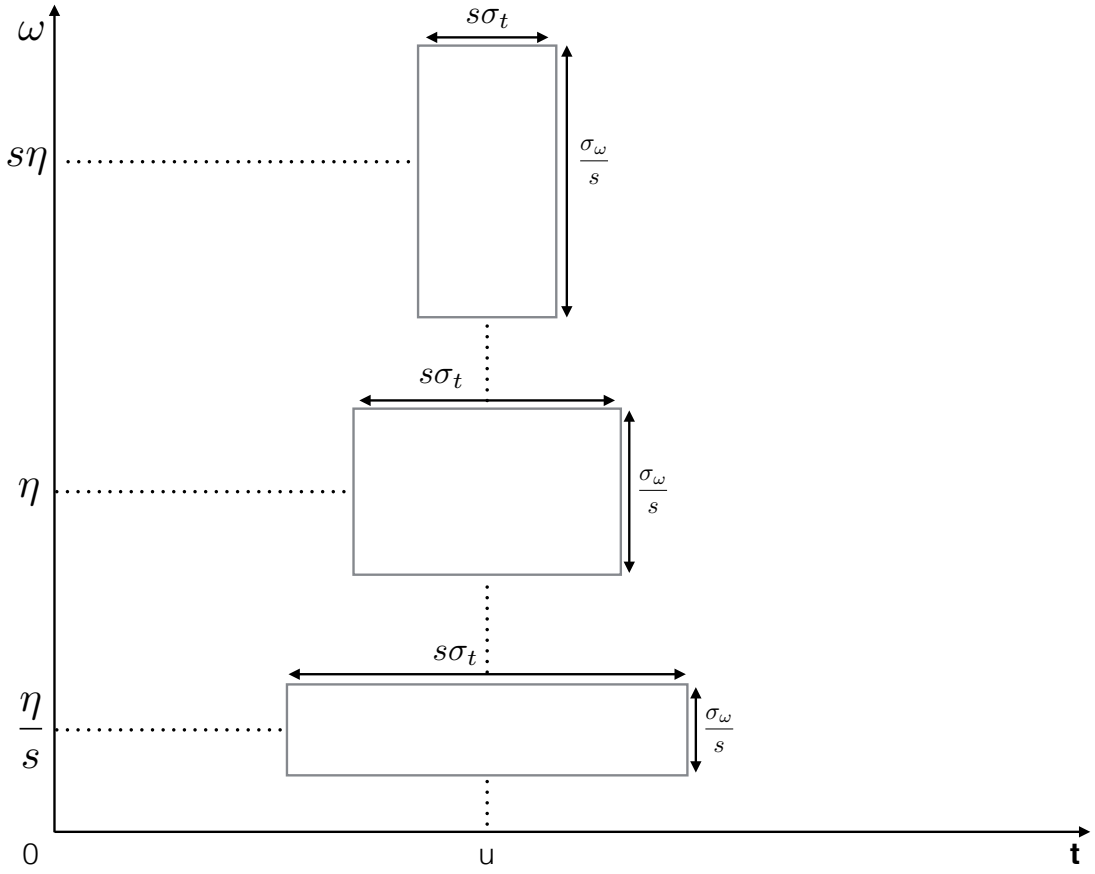

Figure 1.6: A depiction of the time-frequency (Heisenberg) boxes that tile the time-frequency plane using a complex wavelet. The time width of the box is defined by $s\sigma_t$ while the frequency width is defined by σ_ω/s. This means that as we change our scaling we also change our position along the frequency axis and the shape of the time-frequency boxes. This key feature enables sufficient temporal resolution to resolve the corresponding low frequency.

is customary to display the period), but we can also see the duration for which this signal is dominant. In the case of panels (a) and (c) both signals are dominant for the same duration, and are clearly distinct. However in panels (b) and (d) we can now distinguish that in this case the frequencies peak at different times. This is the major gain for wavelet analysis, with the downside being that we have given up some accuracy in calculating the period (the spread of the peaks).

The changes in shape for each peak as we move to lower frequencies are due to the changing size of the time-frequency boxes and the large area of the peaks is due to our uncertainty explained above. The psedocolor plot was interpolated which is why the time-frequency boxes are not visible due to resolution. Also note that the intensity of the smaller period was slightly less than the larger period. This is due to the energy scaling similarly to the support of the frequencies. The other contours shown in Figure 1.7 are important when presenting wavelet spectra and will be used throughout this work. The red diagonal line is called the 'cone-of-influence' and demarcates the region where boundary effects become important. Some daughter wavelets which 'fit' within the centre of time-series, when translated, would extend beyond the limits of the series, causing problems with the transform. Anything above the cone-of-influence has extended beyond boundaries and thus cannot be trusted to be correct. In general we will ignore anything beyond this boundary. The white contour lines indicate a 95% confidence interval for the signal when compared to a background distribution. In these example cases a white noise background is selected, for the field data and in general any physical phenomena, a red noise background is selected. This contour helps to quickly identify regions of interest within the wavelet spectrum and provides evidence that the signal is not a random fluctuation.

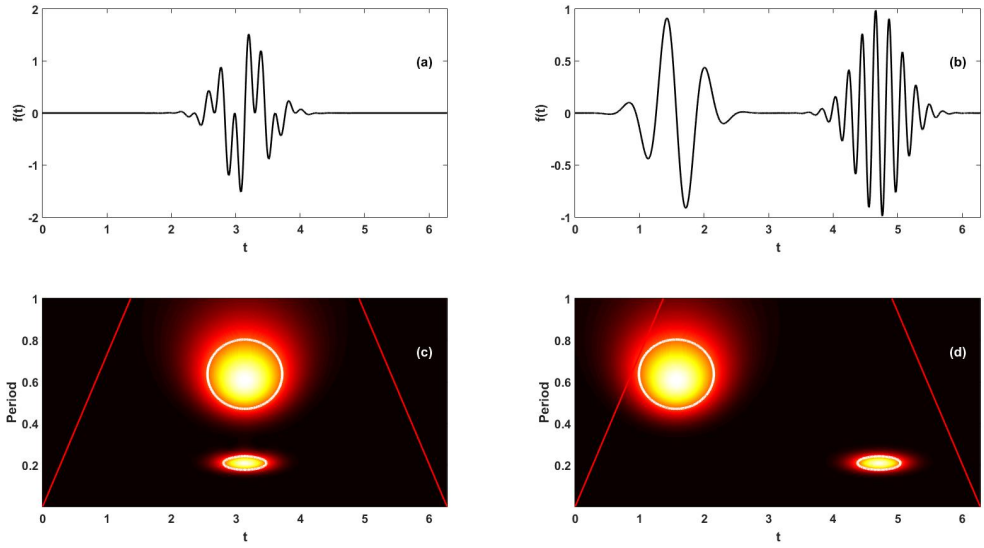

Figure 1.7: The wavelet spectrum for the same signals presented in Figure 1.4. It is important to note is that the y-axis is now given as the period thus the frequencies from Figure 1.4 are of period: $\frac{2\pi}{30}$ (0.2094) and $\frac{2\pi}{10}$ (0.6283). Due to the uncertainty from our time-frequency boxes, the specific frequencies are not exact but rather slightly spread in the time-frequency domain. However, they are sharply spiked at a specific frequency

1.3 Examples of sample signals

Having now discussed the background theory as to how wavelets can be used on a given time series, let us now consider some sample signals and view their response in the time-frequency wavelet plots. We begin by looking at the simplest possible signal, a sine wave. Figure 1.8 shows both the sinusoidal time-series, $y = \sin(t)$, and the corresponding wavelet plot. There is a very clear response in the wavelet plot which is nearing constant for all times within the cone-of-influence. We can also see that the wavelet power is concentrated at the 2π period, which is the period of the time series. There are a few additional important features that we can note from this simple example. The effect of being outside the cone-of-influence are visible near the edges of the plot. Again it should be mentioned that you cannot trust any results from beyond this cone since edges will affect the results. Also notice that the band is not precisely located at 2π but rather there is a frequency width to the band. This is due to the loss of accuracy in the frequency domain that we gave up in order to gain temporal information.

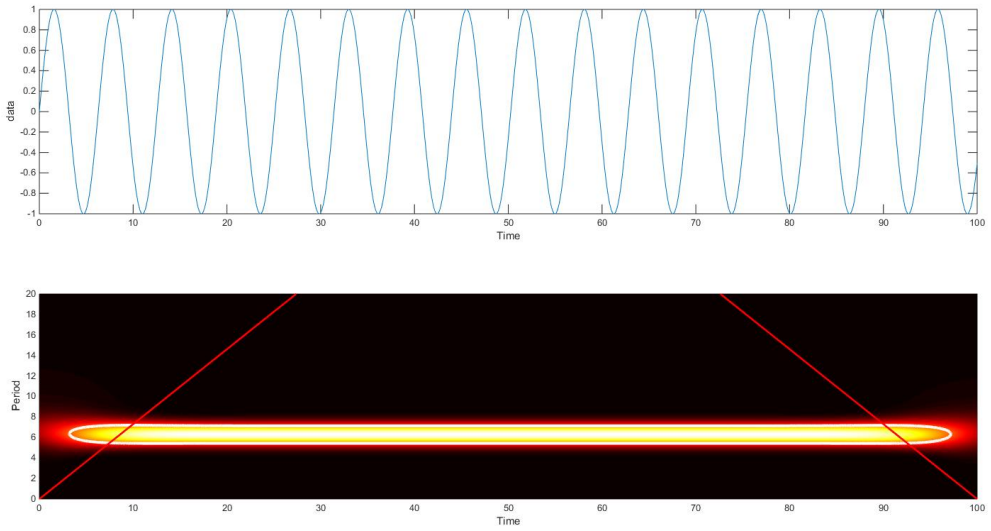

Figure 1.8: A simple sine wave along with its wavelet spectrum. With no changes to the sine wave, there is a clear band at the 2π period mark in the wavelet spectrum plot. This band is consistent for most times within the cone-of-influence.

In the example shown in Figure 1.9 we have now combined (by summation) three different sine waves with different periods (one with the same period as the last figure,

one doubled and one halved). The first feature that is immediately apparent out from this plot is that the magnitudes of the bands of power are much larger for the higher periods. This bias of power at higher periods is because of the width of the wavelet filter in Fourier space. According to the Torrence and Compo (the authors who's wavelet package we use) website FAQ (http://paos.colorado.edu/research/wavelets/faq.html): 'At small wavelet scales (high frequency), the wavelet is very broad in frequency, therefore any peaks in the spectrum get smoothed out. At large wavelet scales, the wavelet is more narrow in frequency, therefore the peaks are sharper and have a larger amplitude ... [I]n mathematical terms, the global wavelet spectrum is an 'efficient' estimator of the 'true' power spectrum, but it is also 'biased'. The bias means that there may be a large difference between the global wavelet spectrum and the 'true' Fourier spectrum. ... [I]f you have sharp peaks in your power spectrum, don't use the global wavelet spectrum to determine the relative magnitude of your peaks'. Therefore, as suggested by Torrence and Compo, we can use the wavelet spectrum to extract time-frequency information from our signal, but in some cases we must combine this with traditional Fourier analysis to compare the magnitudes of different signals. Despite the bias, in Figure 1.9 we can see the distinct bands at the corresponding periods to the sine waves in the time-series.

The next example that we consider is a simple Gaussian function, $y = e^{-\frac{(t-\mu)^2}{2\sigma^2}}$, where μ is the mean, and σ is the standard deviation. In the wavelet spectrum of this Gaussian, seen in Figure 1.10(b), we see a peak in power at 4π which then spreads out in a cone for larger periods. Through modifying the variance (and thus the width of the function) we change both the shape of the cone of periods which are activated and the period of peak power.

The next example, Figure 1.11, we consider is a combination of the two previous; in which we multiply the three sine waves by the Gaussian to create a localized signal. To enable the sinusoids to have enough repetition in time to be differentiated, we have increased the width of the Gaussian by a factor of five. As in Figure 1.9, we can easily distinguish the three sine waves present in the time-series. We also see the differing sizes of the 'bands' for each sine wave created by the variable time-frequency grid. Another striking feature is that the 2π period waves have the highest wavelet power, despite in the non-Gaussian case the highest power was seen at 4π. This is because of the chosen size of the Gaussian envelope and how the daughter wavelets line up. If the width of the Gaussian envelope were to be increased, we would see more power at the larger period oscillation, while the opposite is true for a smaller width. This again enforces the point that the magnitude of the power of the wavelet spectrum cannot be used to compare different periods.

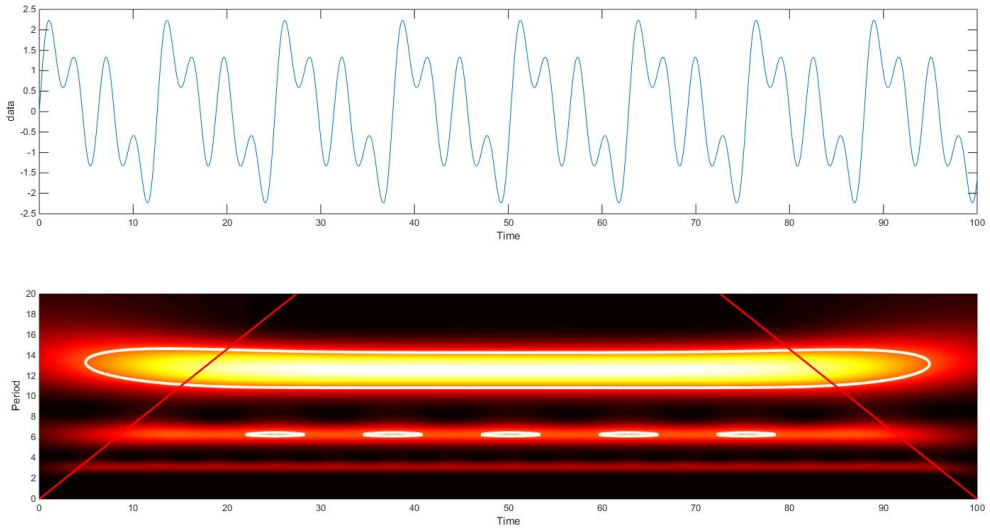

Figure 1.9: Multiple sine waves with different frequencies added together. The time series is shown in the upper panel while the wavelet spectrum is shown in the lower panel. We have the unmodified wave, twice the period and half the period. All three periods can be observed in the wavelet spectrum, however the higher frequencies appear with more power.

The final example we will consider is the same as the example in Figure 1.11, however we have added a background of white noise. The amplitude of this noise has a similar magnitude to the underlying signal. Mathematically this means that the normal distribution used for the noise has a variance of one. The time-series and wavelet plot are presented in Figure 1.12. It is quite evident that despite the addition of significant noise to the signal, the sinusoidal signals are easily extracted from the wavelet plot. Again we must remember that comparing the magnitude of the power across periods is not meaningful, but dominant periods can still be extracted. It is also possible to repeat this with red noise, which is more common for physical applications. The results will be similar, and it is just as easy to extract the sinusoidal signals.

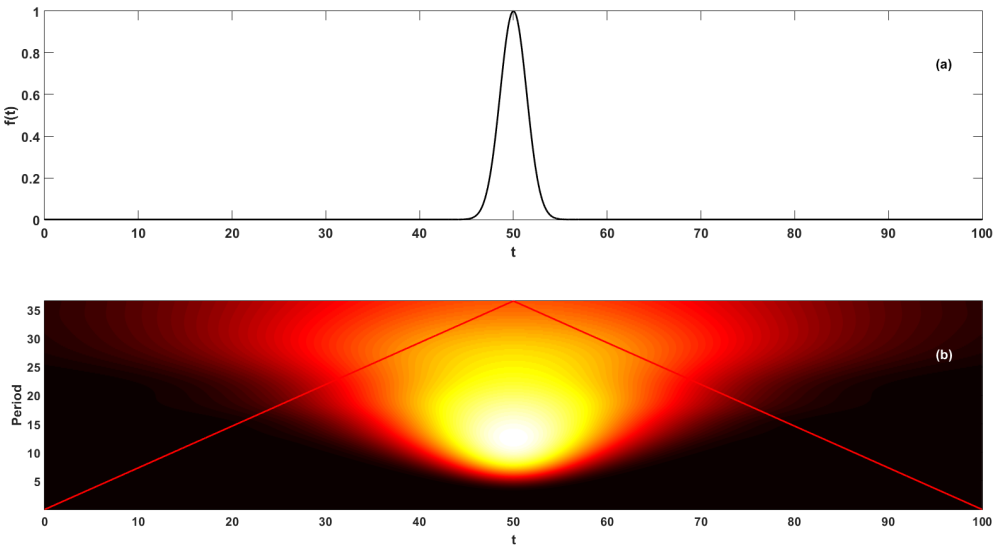

Figure 1.10: The time series of a Gaussian is shown in the upper panel while the wavelet spectrum is shown in the lower panel. The function has a mean of 50 and variance of one. The wavelet power spectra of this function peaks at 4π and then spreads out in a cone from there to larger periods.

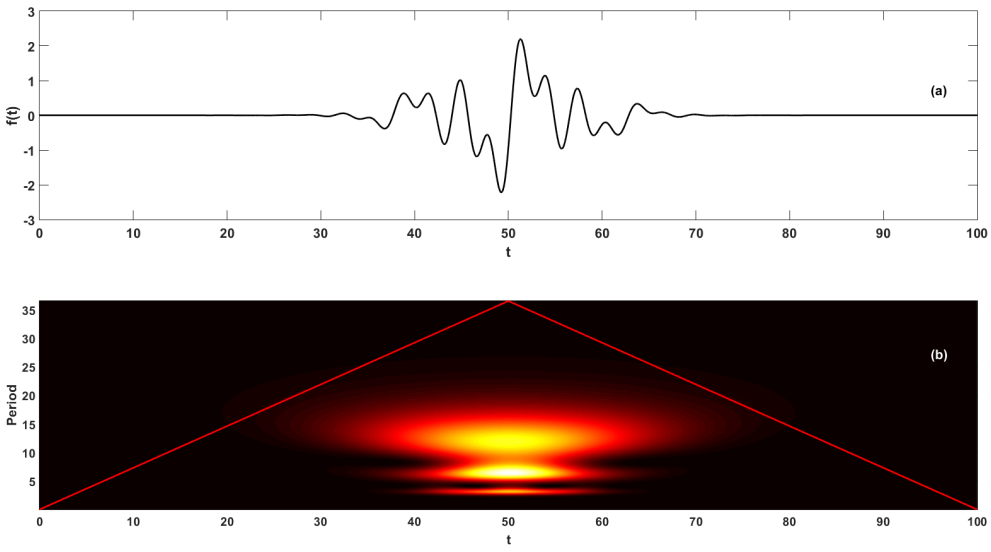

Figure 1.11: The time series of a Gaussian multiplied by the various sine waves shown in Figure 1.9 is shown in the upper panel while the wavelet spectrum is shown in the lower panel. All three periods can be observed in the wavelet spectrum. In this case the 2π period has the highest power.

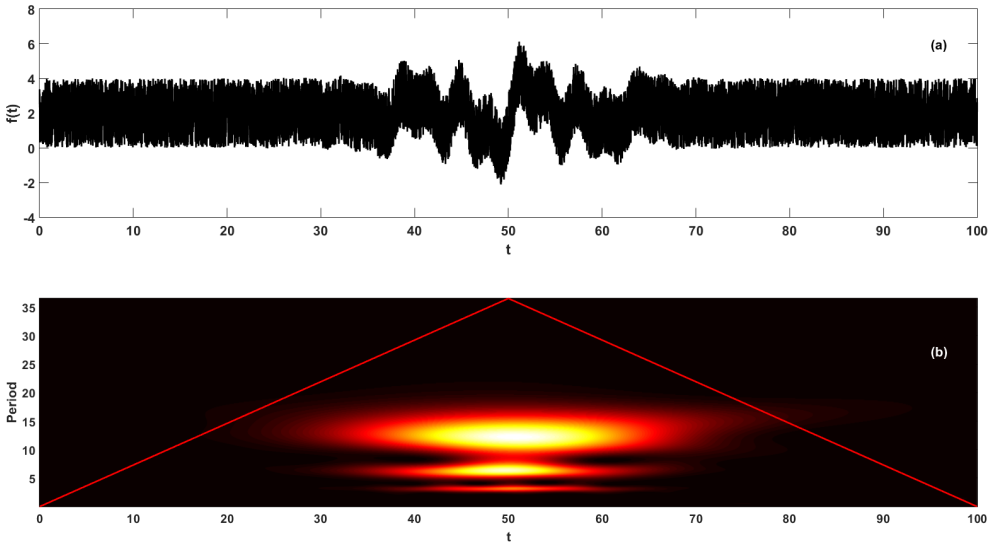

Figure 1.12: The time series of a Gaussian multiplied by the various sine waves shown in Figure 1.9, now with white noise having been added, is shown in the upper panel while the wavelet spectrum is shown in the lower panel. Note that the noise is on the same order as the amplitude of the waves themselves. Again, the three periods can easily be extracted from the wavelet spectrum.

1.4 Hydrodynamic background

The dynamic equations for the systems we will be considering are called the Navier-Stokes equations, under the Boussinesq approximation. These equations are written as

$$\frac{D\boldsymbol{u}}{Dt} = -\frac{1}{\rho_0}\nabla p + \frac{\rho g}{\rho_0}\hat{\boldsymbol{k}} + \nu\nabla^2\boldsymbol{u}, \tag{1.16}$$

$$\nabla \cdot \boldsymbol{u} = 0, \tag{1.17}$$

$$\frac{DT}{Dt} = \kappa\nabla^2 T, \tag{1.18}$$

$$\rho = \rho_0(1 - \alpha(T - T_0)). \tag{1.19}$$

where $\frac{D}{Dt} = \frac{\partial}{\partial t} + \boldsymbol{u}\cdot\nabla$ is called the material derivative, $\boldsymbol{u}$ is the vector of velocities, ρ is the fluid density,ρ_0 is a constant reference density, p is the pressure, g is the acceleration due to gravity, μ is the dynamic viscosity, T is temperature, κ is the thermal diffusivity and α is the thermal expansion coefficient. Equations (1.16) are called the momentum equations and are Newton's law of motion for fluids, with one equation for each component of velocity (see [18], [35], or any standard fluid text for a derivation). Equation (1.17) comes from conservation of mass and the continuity equation. Equation (1.18) is a dynamic equation for temperature. Finally, (1.19) is a simple, linear equation of state. In the following section we will show how this specific set of equations are derived from the full set of more complicated equations to highlight the specific approximations which must be made.

The Boussinesq approximation states that if we assume that density and temperature fluctuations are small in comparison to a reference value, the vertical extent of the fluid is much less than any scale hight (depth $d \ll \rho_0/\frac{d\rho'}{dz}$), and we consider speeds much less than the speed of sound (1482 m/s for water), then density changes can be neglected in all terms except in the gravity term, [29]. To derive this we write

$$\rho = \rho_0 + \rho'(x, y, z), \tag{1.20}$$

where density can be split into a constant part ρ_0 and a small perturbation $\rho'(x, y, z)$ ($|\rho'| \ll |\rho_0|$). Correspondingly, we also assume that we have a constant temperature T_0 added to a small perturbation, T'. We further assume that the background state (ρ_0, p_0) is in hydrostatic balance so that we have,

$$p = p_0 + p'(x, y, z), \tag{1.21}$$

so that

$$\frac{dp_0}{ddz} = -\rho_0 g. \tag{1.22}$$

Substituting this into the full momentum equations,

$$\rho \frac{D\boldsymbol{u}}{Dt} = -\nabla p + \rho g \hat{\boldsymbol{k}} + \mu \nabla^2 \boldsymbol{u}, \tag{1.23}$$

we can cancel out terms and divide through by ρ_0, resulting in

$$\left(1 + \frac{\rho'}{\rho_0}\right) \frac{D\boldsymbol{u}}{Dt} = -\frac{1}{\rho_0} \nabla p' + \frac{\rho'}{\rho_0} g \hat{\boldsymbol{k}} + \nu \nabla^2 \boldsymbol{u}. \tag{1.24}$$

In this equation, ν is the kinematic viscosity and is defined to be $\nu = \mu/\rho_0$, where μ is the dynamic viscosity. If we assume that our density variations are small ($|\rho'| \ll |\rho_0|$), we have for the inertia term that $\left(1 + \frac{\rho'}{\rho_0}\right) \approx 1$. However, applying this on the right hand side of the equation will result in the buoyancy term being eliminated. The buoyancy term is very important and cannot be neglected. Thus we make the assumption that the density perturbation can be neglected except in the buoyancy term. For the continuity equation, we begin with

$$\frac{1}{\rho} \frac{D\rho}{Dt} + \nabla \cdot \boldsymbol{u} = 0. \tag{1.25}$$

It is then argued by scaling the terms, that $\frac{1}{\rho} \frac{D\rho}{Dt}$ will be small compared to the magnitude of the velocity gradients, $\nabla \cdot \boldsymbol{u}$. This will not be true when compressible effects begin to matter, and thus we have a limit on the maximum velocities we can consider. The Boussinesq approximation also effects the general temperature equation,

$$\rho C_v \frac{DT}{Dt} = -k \nabla^2 T + \phi, \tag{1.26}$$

where k is the thermal conductivity of the fluid, C_v is the specific heat at fixed volume and ϕ is the viscous heating. This equation can be derived from the First Law of Thermodynamics, see [18] for details. Using a similar scaling argument to the continuity equation, the magnitude of ϕ in most typical situations will be extremely small ($\approx 10^{-7}$) compared to the other terms and thus can be neglected. The equation of state links thermodynamic variables together (e.g. the ideal gas law). For liquids there is no first-principles equation of state available, rather these relations are empirically derived. In general, the variables in the equation of state could be temperature, salinity, pressure, etc. However in the case of our simulations (Chapter 3) we assume a simple, linear relation between temperature and density,

$$\rho = \rho_0 (1 - \alpha (T - T_0)). \tag{1.27}$$

Equations (1.16)-(1.19) are the equations which were used during our numerical simulation, and indeed are accurate across all scales of motion on which the continuum approximation holds. However, for our specific physical situation, they are not useful across all scales. This is due to the complex boundaries within the karst system. Within the large fractures, caves, and cenotes of the Yucatan peninsula, it is reasonable to assume that the boundaries could be specified in a manner that is representative of the field conditions. However, on the scales of flow through the network of pores within the karst (μm to mm), the Navier-Stokes equations are no longer useful. For such a complicated system it is no longer possible to model the dynamics of the fluid directly but rather we must examine the bulk flow through a sample of karst. In this case we would want to move to a macroscale equation for flow through porous media, that effectively averages the dynamics at the pore level. We would then need to add a preferred flow direction for the many caves and fractures in the landscape which would modify the general flow. The standard, empirical equation for this type of flow is called Darcy's Law and reads

$$\boldsymbol{u} = \frac{1}{\mu}\boldsymbol{K} \cdot \nabla P, \tag{1.28}$$

where the constants and variables have previously been defined except for $\boldsymbol{K}$ which is called the permeability and defines how much fluid is allowed to flow by a given medium. In general it is a second order tensor (this allows for a preferred flow direction within the porous medium), [22].

Chapter 2

Analysis of sensor data

2.1 Overview of sensor placement and methodology

Two types of sensors were used for the measurements within the cave network. The HOBO Fresh Water Conductivity Data Logger from OnSet provided both conductivity and temperature measurements. The Sensus Ultra from ReefNet provided temperature and pressure measurements. OnSet builds the HOBO sensors for research applications, while the Sensus Ultra are a recreational tool for divers. The specifications for the various fields of interest are reported in Tables 2.1, 2.2 and 2.3. Comparing the specifications for the temperature fields reveal that despite the Sensus Ultra not having been designed for research purposes, it can provide roughly the same information as the HOBO sensor. The key difference between the two is that the HOBO sensor is eight times more accurate than the Sensus. Another feature that is important to note is that the Sensus Ultra reports its statistics in depth while the sensor technically takes measurements of pressure. This means that there is some error involved when converting from pressure to depth since the calculation assumes that the density is constant and that there is a constant atmospheric pressure. These approximations do not always hold, especially during storm events where atmospheric pressure may change quite a bit and the density field may get mixed up. However, the effect that these have on the measurements are relatively small, and unless we are looking at centimetre variations they can be considered negligible.

Conductivity	HOBO Conductivity Logger
Range	0 $\frac{\mu S}{cm}$ to 10,000 $\frac{\mu S}{cm}$
Accuracy	3 % of reading, or ± 5 $\frac{\mu S}{cm}$, whichever is greater
Resolution	1 $\frac{\mu S}{cm}$

Table 2.1: Conductivity specifications for the HOBO Conductivity logger that were used in the Yax Chen cave system. (Retrieved from http://www.onsetcomp.com/products/data-loggers/u24-001)

To provide the most relevant information in regards to halocline mixing, the sensors were placed in chains. The two chains that we will be looking at are called the ISOD2 chain and the H2S chain. The location of these chains in the cave system can be seen in Figure 2.1. The chains are close enough to the cenote openings that we expect to see some influence from rain or water flow into the sink hole, but far enough into the cave that the direct dynamic effect of rain falling on the water surface is not seen. Due to the higher price of the HOBO sensors, the Sensus Ultra sensors were used to fill the gaps between the HOBO sensors on the chains. The ISOD2 chain had five Sensus Ultra sensors attached

Depth	Sensus Ultra
Range	1m to 154 m
Accuracy	±0.3048 m
Resolution	0.0127 m

Table 2.2: Depth specifications for the Sensus Ultra sensor that were used in the Yax Chen cave system. (Retrieved from https://reefnet.ca/products/sensus/)

Temperature	Sensus Ultra	HOBO Conductivity logger
Range	−20°C to 40°C	5°C to 35°C
Accuracy	±0.8°C	±0.1°C
Resolution	0.01°C	0.01°C

Table 2.3: Temperature specifications for both the Sensus Ultra sensor and the HOBO conductivity logger that was used in the Yax Chen cave system.

through the length of the cave passage. There were two HOBO sensors attached to the chain, as well. The H2S chain had four Sensus Ultra and two HOBO. The sensor depth and serial numbers are provided in Table 2.4

Type	Serial Number	Location	Depth
HOBO	9960690	ISOD2	10.1 m
HOBO	10039732	ISOD2	9.1 m
Sensus	9208	ISOD2	9.1 m
Sensus	9209	ISOD2	13.6 m
Sensus	9210	ISOD2	12.2 m
Sensus	9218	ISOD2	10.7 m
Sensus	9219	ISOD2	7.6 m
HOBO	10076089	H2S	9.1 m
HOBO	10039736	H2S	10.7 m

Table 2.4: Depth and location information for the sensors used to create the sensor chains. The sensors are split up between two chains, ISOD2 and H2S, the location of which can be seen in Figure 2.1.

The sensors were left in the cave for six-month intervals starting in 2011. To recover

the data, the Sensus Ultra had to be retrieved from within the cave by divers and then manually connected to a computer. For this reason there are gaps within the time series which were interpolated to create a continuous version. A shuttle from OnSet was used for the HOBO sensors which allowed for the data to be collected directly within the cave network and so there was no need to remove the sensors from the chain.

The data that is provided by these sensors can be combined to provide the density through the equation of state. While there are many empirical options for the equation of state, for most saline hydrodynamic applications the UNESCO equation of state is used [34]. This equation computes density using pressure, temperature and salinity. Following Appendix 3 in [9], the density of pure water is given by

$$\begin{aligned}\rho_w = {} & 999.842594 + 6.793952 \cdot 10^{-2} T - 9.095290 \cdot 10^{-3} T^2 \\ & + 1.001685 \cdot 10^{-4} T^3 - 1.120083 \cdot 10^{-6} T^4 + 6.536332 \cdot 10^{-9} T^5. \end{aligned} \tag{2.1}$$

The density when pressure is zero ($p = 0$) is given by

$$\begin{aligned}\rho(S,t,0) = {} & \rho_w + S(0.824493 - 4.0899 \cdot 10^{-3} T + 7.6438 \cdot 10^{-5} T^2 \\ & - 8.2467 \cdot 10^{-7} T^3 + 5.3875 \cdot 10^{-9} T^4) + S^{3/2}(-5.72466 \cdot 10^{-3} \\ & + 1.0227 \cdot 10^{-4} t - 1.6546 \cdot 10^{-6} T^2) + 4.8314 \cdot 10^{-4} S^2. \end{aligned} \tag{2.2}$$

In general we can write that the density is

$$\rho(S,T,p) = \frac{\rho(S,T,0)}{1 - \frac{p}{K(S,T,p)}}, \tag{2.3}$$

where K is the secant bulk modulus and is given by

$$\begin{aligned}K(S,t,p) = {} & 19652.21 + 148.4206 T - 2.327105 T^2 + 1.360477 \cdot 10^{-2} T^3 - 5.155288 \cdot 10^{-5} T^4 \\ & + S(54.6746 - 0.603459 T + 1.09987 \cdot 10^{-2} T^2 - 6.1670 \cdot 10^{-5} T^3) \\ & + S^{3/2}(7.944 \cdot 10^{-2} + 1.6483 \cdot 10^{-2} T - 5.3009 \cdot 10^{-4} T^2) \\ & + p(3.239908 + 1.43713 \cdot 10^{-3} T + 1.16092 \cdot 10^{-4} T^2 - 5.77905 \cdot 10^{-7} T^3) \\ & + pS(2.2838 \cdot 10^{-3} - 1.0981 \cdot 10^{-5} T - 1.6078 \cdot 10^{-6} T^2) \\ & + p^2(8.50935 \cdot 10^{-5} - 6.12293 \cdot 10^{-6} T + 5.2787 \cdot 10^{-8} T^2) \\ & + p^2 S(-9.9348 \cdot 10^{-7} + 2.0816 \cdot 10^{-8} T + 9.1697 \cdot 10^{-10} T^2) + 1.91075 \cdot 10^{-4} p S^{3/2}. \end{aligned} \tag{2.4}$$

While these equations are complicated to write down, pre-made scripts make the equation extremely easy to use numerically. Another important calculation will be converting

from conductivity to salinity. By taking water samples when divers are downloading the data from the sensors, it is possible to calibrate the sensors so as to very precisely calculate the salinity. However, since we are within the range of 2 ppt to 42 ppt we can use an experimentally derived equation for conversion, outlined in [1]. Salinity can be computed from conductivity, C, as

$$\begin{aligned} S(t,C) &= 0.0080 - 0.1692R_T^{1/2} + 25.3851R_T + 14.0941R_T^{2/3} - 7.0261R_T^2 + 2.7081R_T^{5/2} \\ &\quad + \left(\frac{T-15}{1+0.0162(T-15)}\right)(0.0005 - 0.0056R_T^{1/2} - 0.0066R_T - 0.0375R_T^{3/2} \\ &\quad + 0.0636R_T^2 - 0.0144R_T^{5/2}), \end{aligned} \tag{2.5}$$

where R_T is the ratio given as

$$R_T(T,C) = \frac{C}{-0.0267243T^3 + 4.6636947T^2 + 861.3027640T + 29035.1640851} \tag{2.6}$$

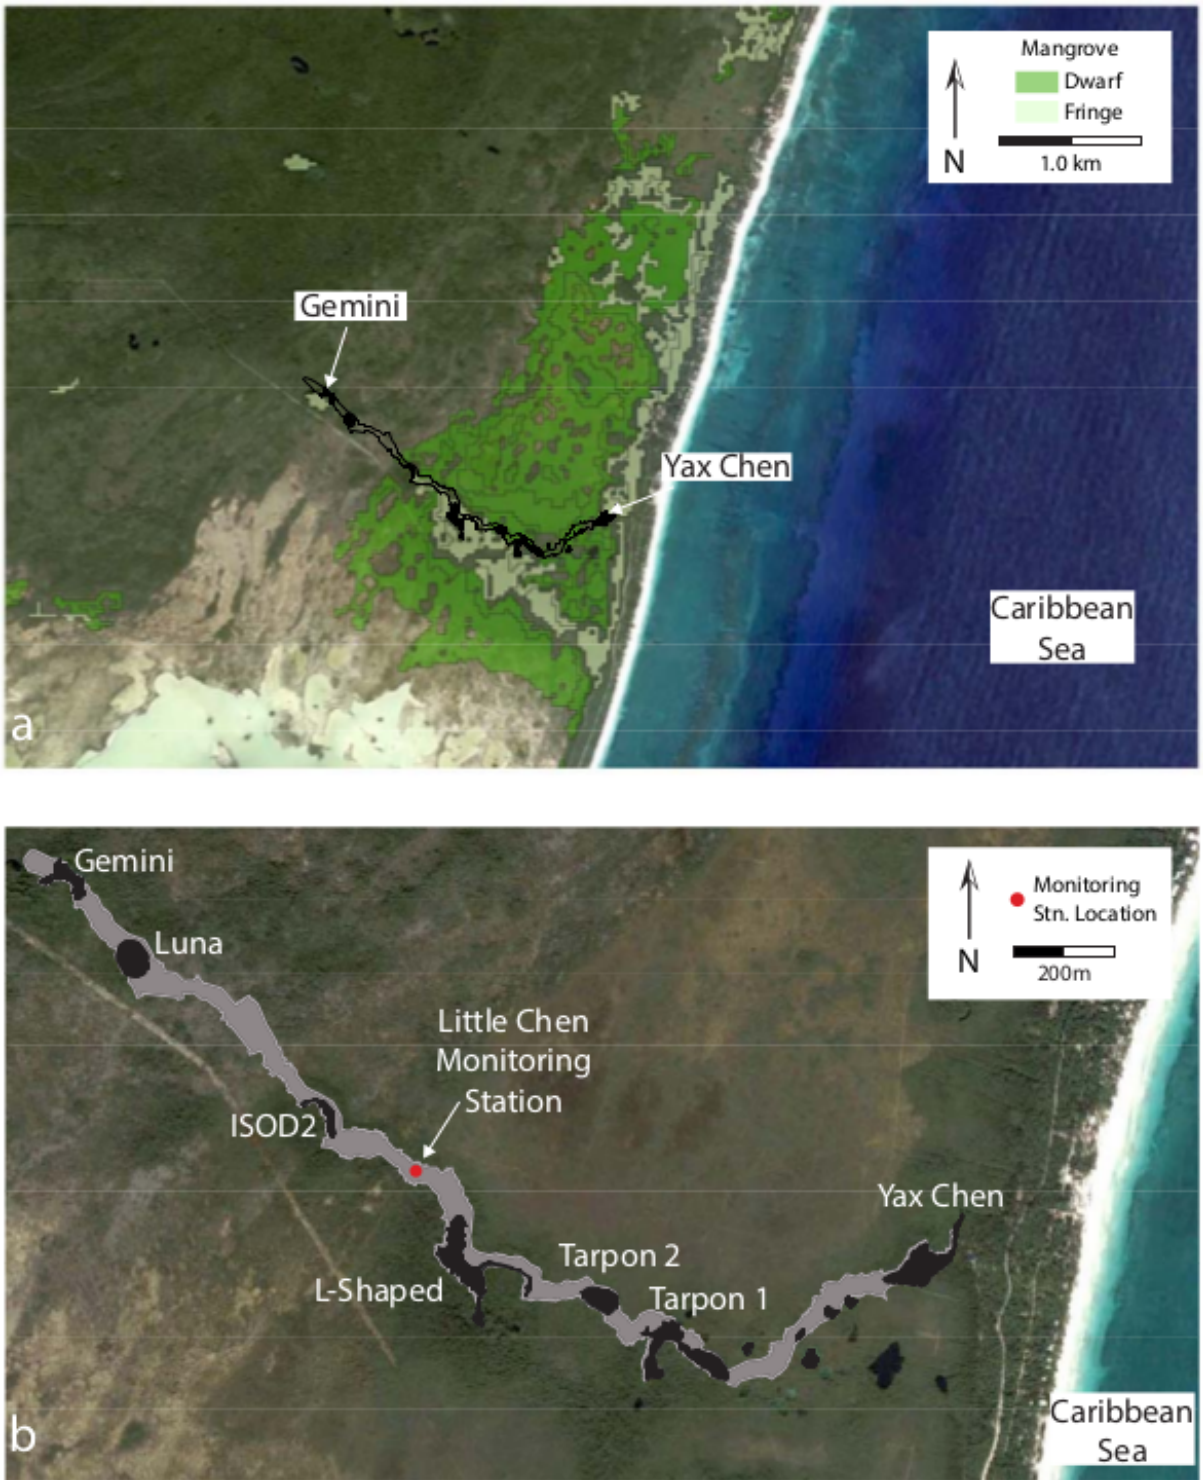

Figure 2.1: A map of the research area courtesy of Shawn Collins. Panel (a) shows a portion of the Yucatan coast, and the location of the Yax Chen cave network. Mangrove cover is shown in green, and the entrance and furthest well mapped point are labelled. Panel (b) shows a close up of the cave system and the names of particular areas within the system.

2.2 Time-series analysis

To allow for easy comparison with figures that do not contain their own precipitation data, we have provided the full rainfall measurements from the Sian Ka'an Biosphere in Figure 2.2. This data indicates several large events, the largest of which (dwarfing all the others) occurs in mid-September 2013, and is a result of Hurricanes Ingrid and Manuel.

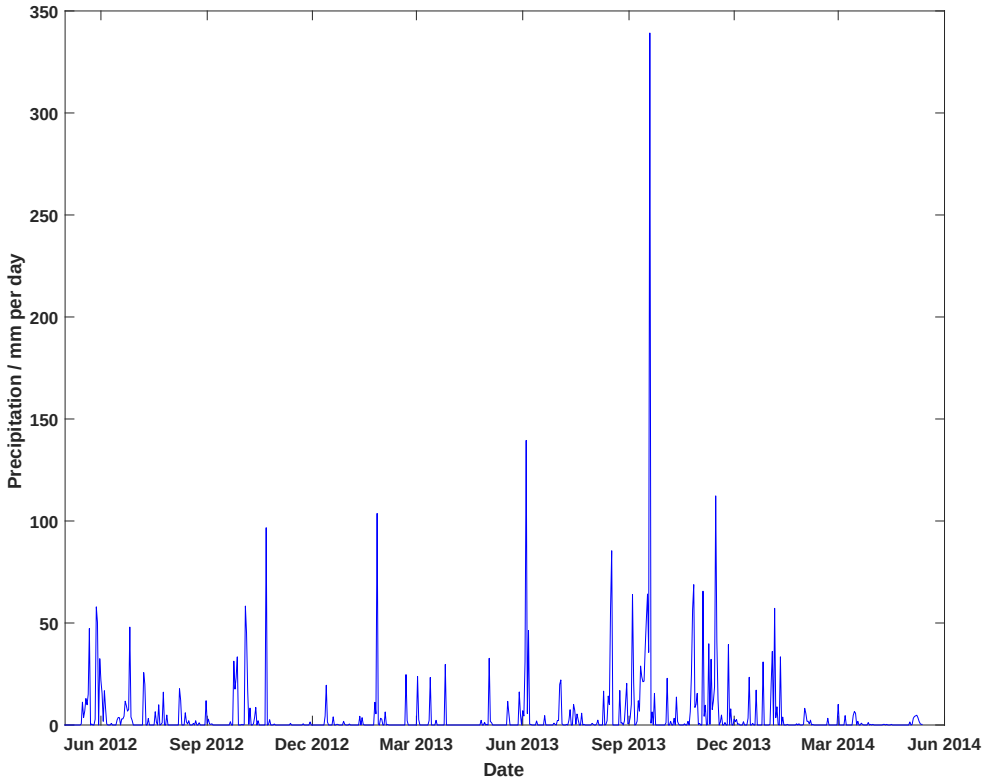

Figure 2.2: Precipitation data gathered from the Sian Ka'an Biosphere relatively close to the sensor site. This is the full extent of the gathered data. The sensor records precipitation in 10 minute intervals, however to make the data more manageable it has been binned into daily amounts. The largest spike in the series correlates with Hurricane Ingrid and Manuel.

Figure 2.3 shows the temperature measurements from the Sensus Ultra sensors on the ISOD2 chain from June 2011 to March 2014. The temperature decreases vertically up the z

axis so that the sensors lines appear as they would on the chain. There are two very striking events visible in this time series, one in September of 2011 and one in September of 2013. There is a clear response, specifically in the deeper sensors at these times. Comparing this time series to weather records for the Yucatan peninsula (Figure 2.2), we can see that these times directly correspond to Hurricane Rina in 2011 and Hurricanes Ingrid and Manuel in 2013. This correlation implies a connection between the hydrodynamics within the caves and large scale weather events. The oscillations of the 10.2 m (red) sensor, appears to alternate between a temperature similar to the sensors above it, and one similar to those below it. This effect can be readily explained by the presence of the halocline at this level, with the oscillations in temperature corresponding to the halocline oscillating about the fixed sensor. This is supported by anecdotal evidence from divers who report that the halocline in this section of the cave is located at roughly a 10 m depth.

Using weather data gathered from the nearby Sian Ka'an Biosphere (this data has been gathered from 2012 onwards), we can compare the temperature time-series with the precipitation data to look for closer correlation between the datasets. Figure 2.4 aligns the temperature time-series with the precipitation data for the time period during and after the 2013 hurricanes. It is very clear that the large drop in temperature of the lower sensors occurs at the same time as the large rainfall event. We can also note that the second largest rainfall event in early November also corresponds to a drop in temperature of the lowest sensor. Notice that the response to both of these events is similar namely a sharp drop in temperature followed by a recovery to a new stable state that is cooler than before.

Figure 2.5 compares the temperature responses to both the 2011 and 2013 hurricanes. The initial drop in temperature looks extremely similar in both cases, especially for the lower two sensors (11.6m and 12.1m). The response afterwards does differ however, though the initial decline in temperature is followed by a quick increase in temperature during the 2013 event, then drops again to a similar level to that seen during the hurricane. The 2011 temperatures return to similar values to those before the lowest sensor slowly declining in temperature. Observing the precipitation data from Figure 2.4, it is apparent that there is a second rainfall event, thus the second decline in temperature is likely a result of this. We would argue that were this second event not to have occurred the temperature response would appear similar to the 2011 hurricane.

While it is possible to observe large events in the raw time-series, we have no information about the periodicity of the signal, and any underlying structure is difficult to readily observe. For this we turn to wavelet analysis. Figure 2.6 shows the time-series, (a), global wavelet power spectrum (GWS), (b), and wavelet power spectrum, (c), for Sensus Ultra ISOD2 sensor located at a depth of 7.2 m. The time period that was analysed spans from September 2012 to September 2013. Neither of the two large hurricane events occur during

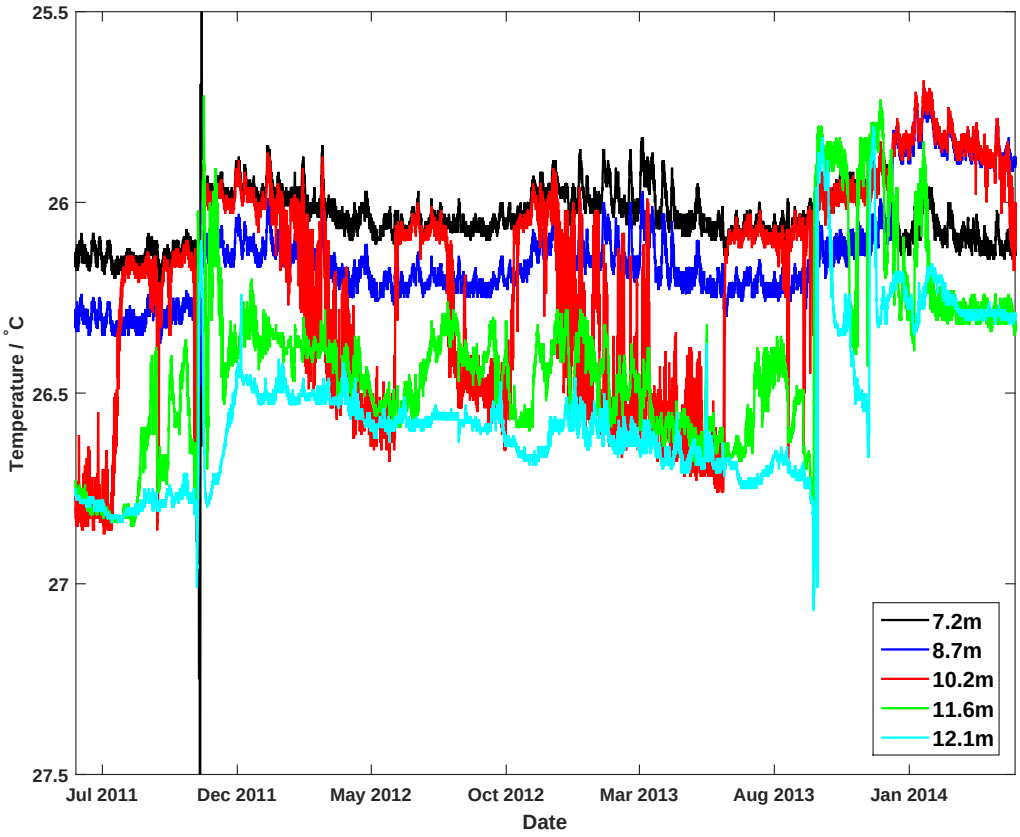

Figure 2.3: The complete time-series of temperature for the Sensus Ultra sensors located at ISOD2, from May 2011 to March 2014. The sensors were organized in a chain, with the corresponding depths given in the legend. The temperature axis has been flipped so that the depths are ascending. Due to the warm surface ocean water entering the system the overlying fresh-water is cooler, however the salinity differences result in a stable stratification.

this time, so the data does not contain any large jumps in temperature. Observing the wavelet power spectrum in Figure 2.6(c), a few features are immediately apparent. Notice that at the current power saturation it is still possible to see a band of power located at the one day mark. This corresponds to the tidal signal and will be considered more closely in Figure 2.7. Focusing on the blue contour representing a 95% confidence in power compared

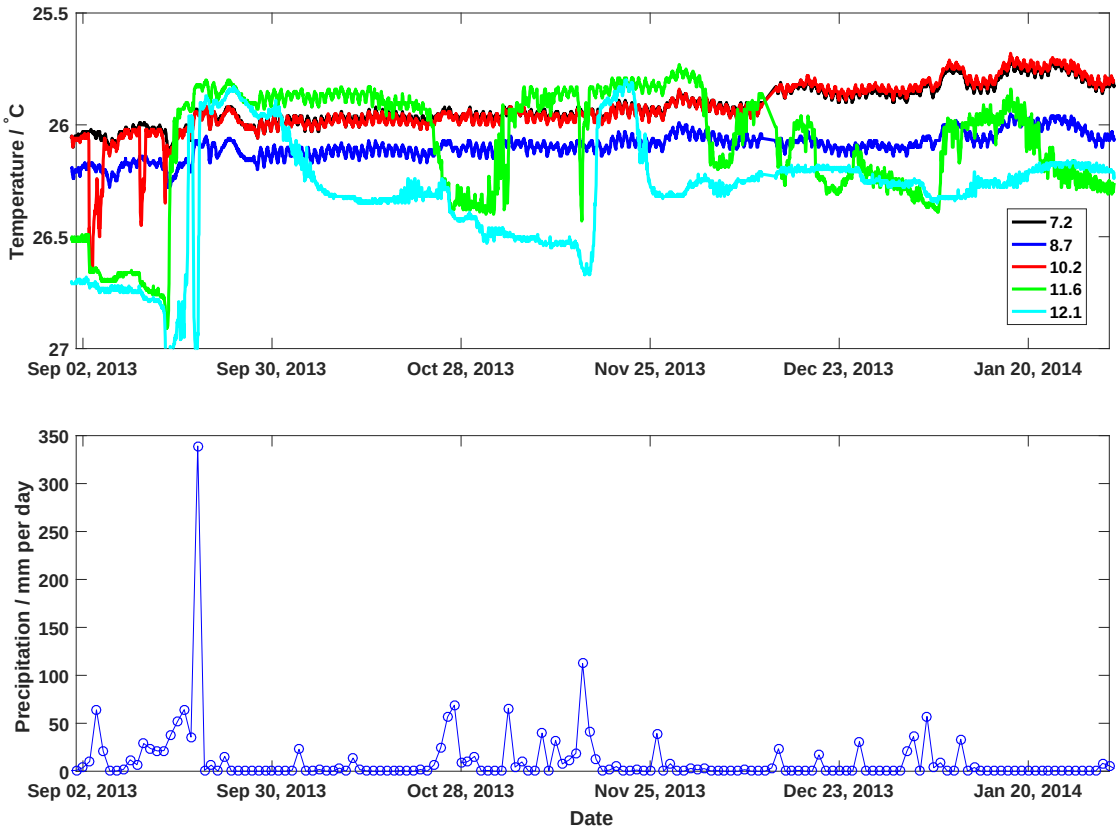

Figure 2.4: A comparison of the Sensus Ultra ISOD2 sensor chain temperature data (top) with the Sian Ka'an weather station precipitation data (bottom). Notice the very strong correlation of temperature drops in the lower sensors during the large rainfall events.

to a background red noise, almost all of the power with a period longer than seven days can be considered significant. The lack of contours surrounding the tidal signal indicate that while it still may be real, the power level at these frequencies is not 95% above the background distribution. The global wavelet spectrum (GWS) is the time-average of the wavelet power spectrum, and can be considered an efficient estimator of the Fourier power spectrum. Using the alignment of the wavelet spectrum and the GWS, Figures 2.6(a) and (c), we can see that the relatively large oscillations between January 2013 to April 2013 appear to correspond to the large peak in power that appears in Figures 2.6(b) and (c)

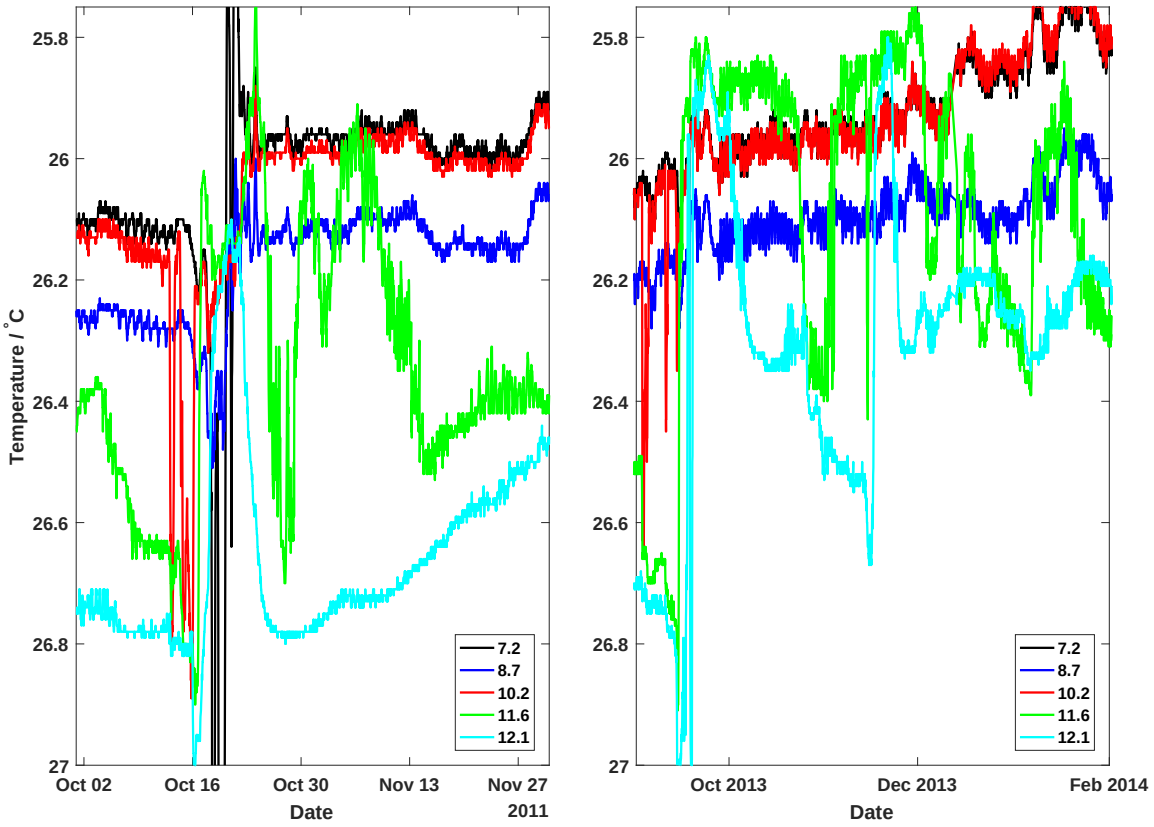

Figure 2.5: A comparison of the 2011 and 2013 hurricane events. There is a very smooth decline in temperature followed by a quick increase to a slightly lower temperature at the 12.1 m sensor in the wake of the 2011 hurricane. The 2013 hurricane appears to follow a similar path at first but then repeatedly and rapidly drops in temperature and then recovers. Observing the precipitation data from Figure 2.4 we can see that there is another noticeable rainfall event after the initial hurricane. We would argue that were it not for this second event, the system would follow a very similar path to the 2011 hurricane.

at roughly the 13-14 day period. The arrows in the GWS indicate peaks in the power of the spectrum using the built-in MATLAB peak finding function the following periods were extracted: 0.5024 days, 1.0048 days, 12.9682 days, 26.1167 days, 56.7639 days and 111.9648 days. The first two peaks are can be readily identified as the semi-diurnal and

diurnal tides. The latter three peaks are more difficult to identify. The 13 day peak may correspond to the luni-solar fortnightly tidal component, to which it is within 5% of the expected period. The next largest most likely corresponds to the lunar month to which again matches within 5%. The last two periods are probably due to seasonal variations, though it is not readily apparent what physical processes correspond to each. There is a general trend within the GWS of the power increasing for longer periods. This is in part caused by the bias of the wavelet analysis, as discussed in the previous chapter, but is also a result of the underlying Brownian noise within the data. Since all physical processes depend on the motion of matter, which is made of atoms and molecules, any measure of these processes will contain this form of noise. The spectrum of Brownian noise is one of exponential decay as frequency increases, thus since period is the reciprocal of frequency we will have exponential growth for larger periods. The wavelet transform can also be considered a broadband filter based on the scale of the particular wavelet that is being used, because of this there is a smoothing effect on the GWS. In comparison to the traditional Fourier spectrum (see Figure 2.9), this makes it much easier to extract the general peaks in power.

Figure 2.7 focuses on the shorter periods of the wavelet power spectrum from Figure 2.6(c). The lower bound of the power spectrum has also been lowered to 2^{-2}, as opposed to $2^0 = 1$, to highlight features which were not apparent in the previous figure. At this saturation level the semi-diurnal and diurnal tidal signals are immediately obvious. It is important to note that the data presented here are temperature measurements, while tidal signals are traditionally measured in water depth. We would argue that these oscillations in temperature are due to the water level fluctuating about the fixed sensor, since the water column should have a stable density profile, as the water level rises the slightly cooler water below the sensor is now detected, the inverse is true during low tide with the warmer water above the sensor. The primary dynamical result from the detection of tidal signals in the sensors is that this portion of the cave network is connected to the oceans. Any modelling of the large scale hydrodynamics of the region must take in to account this connectivity.

To further investigate the periodicity and impact of the tides, we move to the pressure (depth) data. Figure 2.8 shows the same time-domain range as Figure 2.6, but with the pressure data instead of temperature. In this wavelet spectrum the semi-diurnal signal is within the 95% confidence interval, which agrees with intuition given that water depth should change with the tides. In this time-series the spikes in the GWS occur at 0.5130 days, 1.0621 days, 2.4232 days, 8.2643 days, 14.1910 days, 38.2371 days, and 119.1720 days. Comparing the spikes in the GWS and wavelet spectrum to the ones seen in Figure 2.6, we find a strong correlation between the $\approx$0.5 day, $\approx$1 day, $\approx$13 day, and $\approx$112 day events

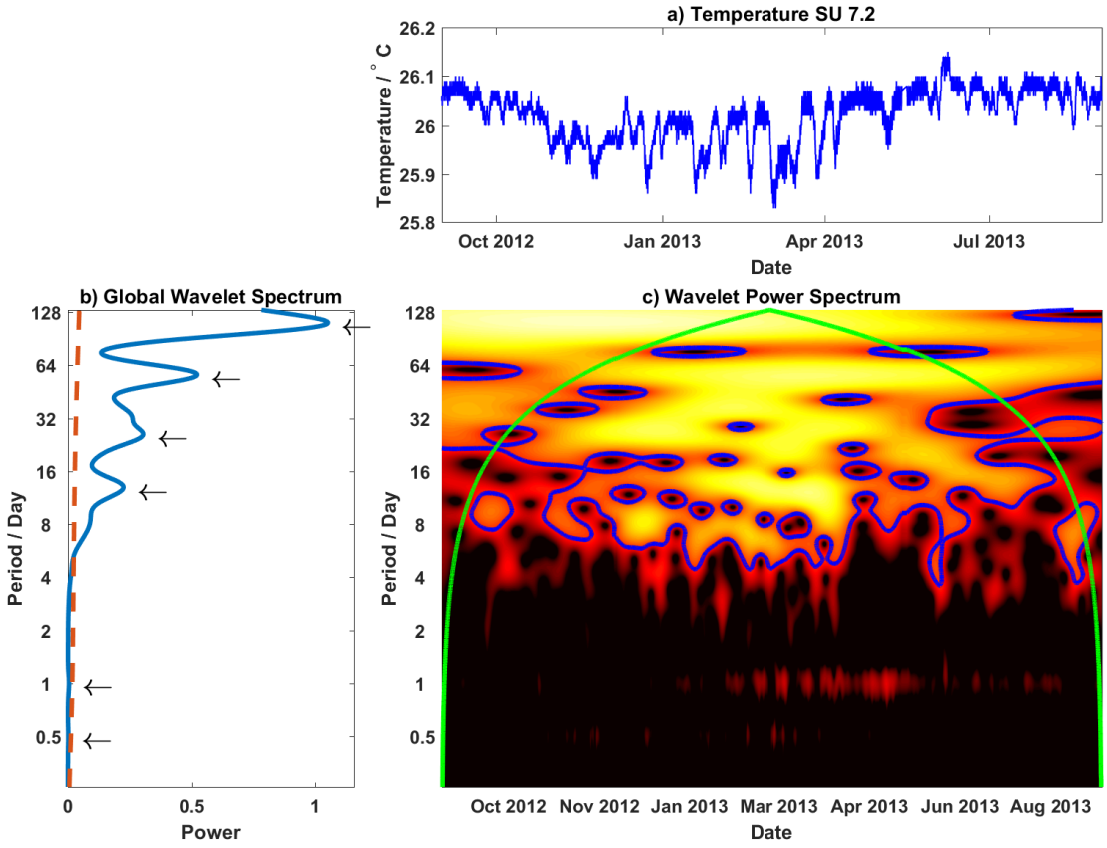

Figure 2.6: (a) the temperature time-series for the shallowest Sensus Ultra sensor at a depth of 7.1 m from September 2012 to September 2013. (b) the global wavelet spectrum (GWS) for the time-series. (c) is the wavelet power spectrum plot, comparing period to time. The GWS is the temporal average of the wavelet spectrum and can be thought of as an estimation of the Fourier spectrum. The dashed line indicates the mean red-noise spectrum and the arrows indicate peaks in power of the GWS. The blue contours indicate a 95% confidence interval and the yellow line indicated the cone-of-influence. Notice a clear band in the wavelet power spectrum at roughly the 60 day mark for most of the time-series.

with the largest error ≈8% occurring at the 13 day event. Comparing the remaining spikes to the wavelet spectrum plot, it is possible to align the 2.5 day spike with the spike that occur in early June 2013. In the pressure series, this corresponds to the relatively large

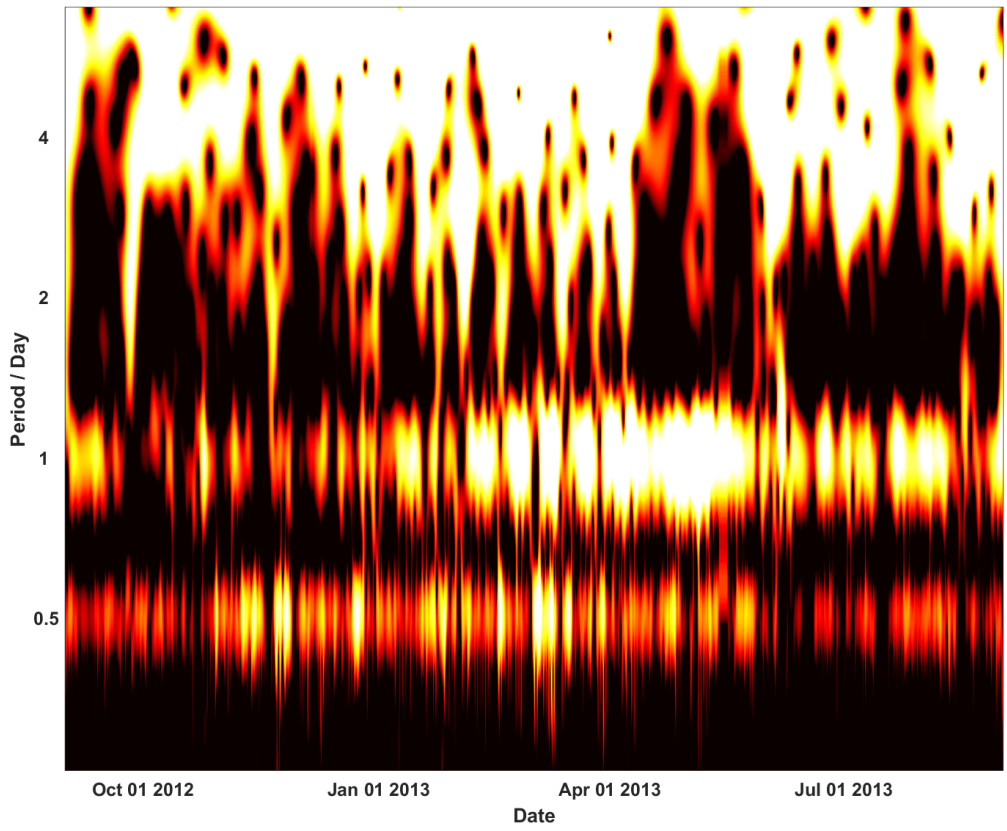

Figure 2.7: A close up of the wavelet power spectrum at low periods for the same time-series as Figure 2.6. By setting the saturation level at (2^{-2}) we can clearly extract a tidal signal from the temperature time-series. Both the diurnal and semi-diurnal tides are visible at this saturation level. This can be used to confirm the expectation that we can extract a tidal signal from any hydrological system that is connected to the ocean.

jump in pressure indicating an increase in water depth. With a quick glance to Figure 2.6, we can notice a spike in temperature as well, though not as strong. The ≈38 day spikes appears to correlate with a large spike in the wavelet spectrum that occurs from June 2013 - September 2013. Again, comparing to the pressure series we can correlate this to the periodicity of spikes at occur in early June, mid-July, and late August. What is interesting is that other than the June event there does not appear to be any corresponding spike in

the temperature data, however, if we compare the dates with Figure 2.2 we can see that these dates do correlate with precipitation events. It is possible that these events were not large enough to impact the temperature at a 7.2 m depth, or that the impact was small enough that it is not noticeable at the current range of the temperature time-series. The remaining spike, at ≈8 days, appears to have contributions throughout the time-series, thus likely has a geophysical cause, but does not correlate with any tidal components or rainfall events.

If we are solely interested in frequency (period) information, wavelet analysis has the disadvantage of losing accuracy in the frequency domain. Thus, in the same cases, it is more effective to use traditional Fourier analysis. Figure 2.9 presents the power spectrum of the pressure data from the 7.2 m Sensus Ultra sensor. The range of dates from September 2012 to September 2013 was used. To properly present the power, the logarithm base 10 of the power is shown. To highlight the diurnal and semi-diurnal signals a magnification from a period of 0.2 to 1.2 is presented as well. As is to be expected due to the red-noise there is a general increasing trend as period increases. An especially notable feature is that the power at the semi-dirnal period (≈0.5 days) is massive when compared to the surrounding data, only for periods larger than 75 days is the power consistently as large as this spike. Comparing the spikes in the magnified plot with the peaks in the GWS from Figure 2.8(b) we can see that rather than being a single spike in power, the general peak is made up of smaller specific peaks. The peaks correspond to the tidal constituents which refer to the specific physical process that results in this excitation. Table 2.5 compares the 'theoretical' values for the eight largest tidal constituents with the largest spikes in spectral power seen in Figure 2.9. The theoretical values, do not imply that there should be a spike in power exactly at that value, but rather that there should be a spike around that value. This is because many factors such as distance from coast, coastal geometry, or geography, can impact the precise period. Comparing the theoretical period to the largest, closest period measured, we see that there is fairly close agreement between all the constituents with the largest error corresponding to a difference of 15 minutes.

We will now move on to analysis of the HOBO sensor data. The HOBO sensors are both more accurate and provide more relevant information since we are interested in how mixing events impact the salinity of the freshwater lens which is used for drinking water. One of the primary motives for this research is to see how large precipitation events affect the salinity within the water column. Measurements of these two fields are presented in Figure 2.10. The salinity measurements are taken at two different depths (9.1 m and 10.1 m) from June 2013 to December 2014. These sensors were located on the ISOD2 chain. In both sensors there is a clear correlation of spikes in salinity with the large precipitation events. There are eight events with precipitation greater than 50 mm/day: June 4, August

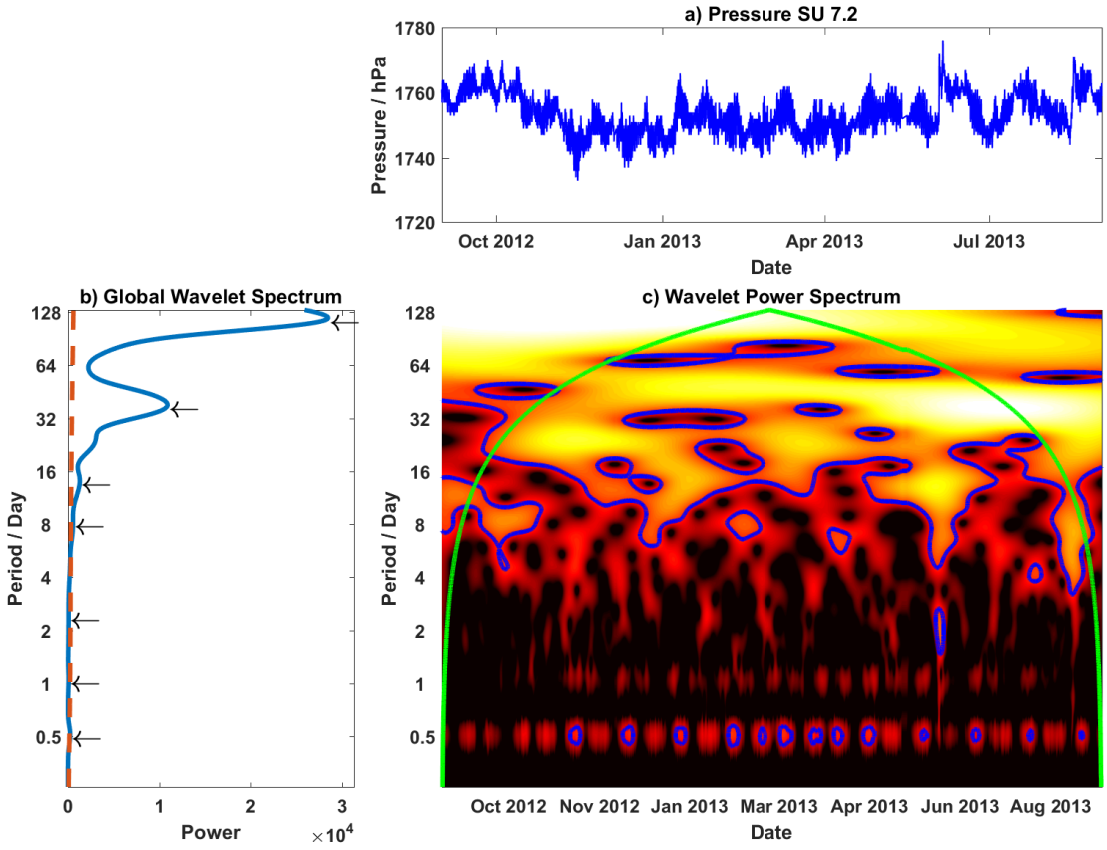

Figure 2.8: (a) the pressure time-series for the shallowest SensusPro sensor at a depth of 7.1m from Sep. 2012 to Sep. 2013. (b) is the global wavelet spectrum (GWS). The dashed line indicates the mean red-noise spectrum and the arrows indicate peaks in the GWS. (c) is the wavelet power spectrum plot, comparing period to time. The blue contours indicate a 95% confidence interval and the yellow line indicated the cone-of-influence.

17, September 4, September 17, September 19, October 27, November 4, and November 15. Each of these events has a corresponding response in the salinity measurements. Though the sensors are located at different levels, it appears that the magnitude of the spikes is the same; the maximum increase occurs during the September 19 hurricane event with an increase ≈2.3 ppt. This data provides strong evidence to support the hypothesis that large precipitation events result in mixing events within the Yax Chen system. The next

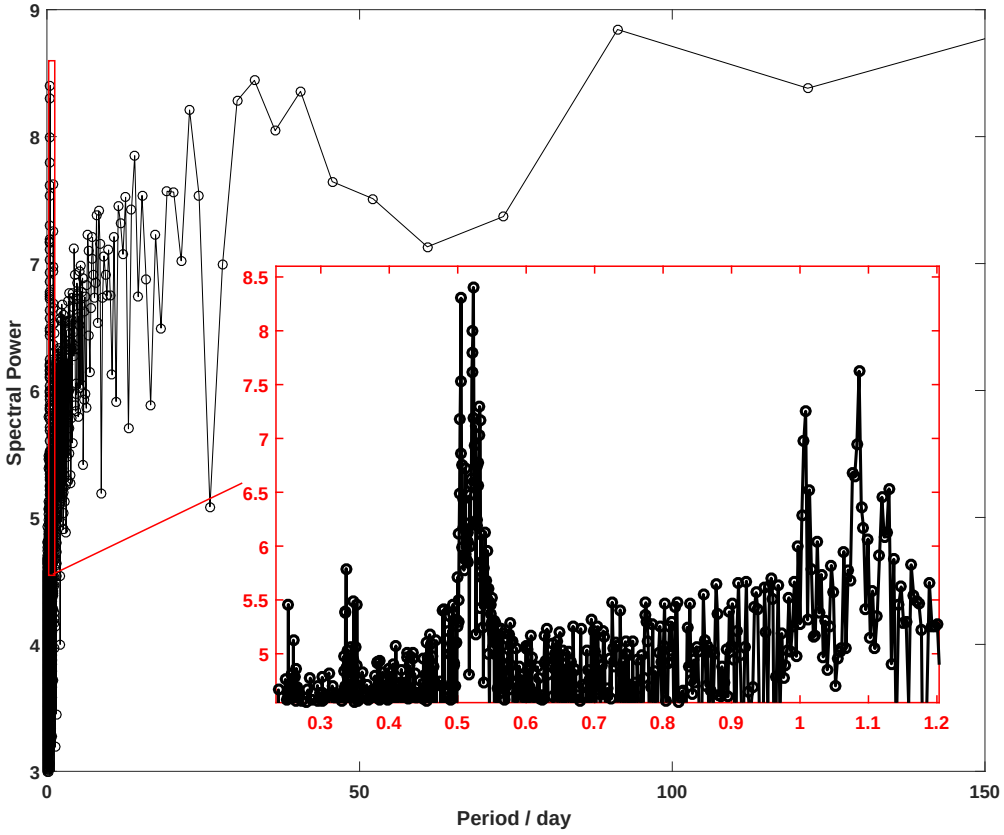

Figure 2.9: The spectral power of the pressure data from the 7.2 m deep Sensus Ultra sensor. A range of dates from September 2012 to September 2013 was used. The horizontal axis is given in period rather than the traditional frequency, which due to the background of red-noise, results in growth in spectral power as period increases. An even extension of the time-series was used so as to make the time-series periodic and avoid errors.

logical question is to what extent this mixing occurs. Currently, with this data, we can say that there is a significant amount of saline water mixing within the first meter or two above the halocline but without sensors at higher locations we cannot make claims about the upper portion of the water column. At this point it is relevant to mention that sensors located near the surface, which Shawn Kovacs is using for study of calcite rafts, have detected spikes in salinity which appear to correlate with precipitation events. Further

Constituent	Symbol	Theoretical Period (h)	Period measured (h)
Shallow water overtides of principal lunar	M_6	4.140	4.076
Shallow water overtides of principal lunar	M_4	6.210	6.062
Shallow water terdiurnal	MK_3	8.177	8.088
Principal solar semi-diurnal	S_2	12	12.120
Principal lunar semi-diurnal	M_2	12.421	12.550
Larger lunar elliptic semi-diurnal	N_2	12.626	12.770
Lunar diurnal	K_1	23.934	24.200
Lunar diurnal	O_1	25.819	26.070

Table 2.5: A comparison of the theoretical values of the first eight most powerful tidal constituents with corresponding spikes in the Fourier spectrum.

study into this data will be needed, but it does support the hypothesis that these mixing events may extend through the majority of the water column. It appears that after the October 27 precipitation events, there is an inversion of the salinity measurements between the two sensors. Before this event the deeper 10.1 m depth sensors consistently measured higher salinities than the 9.1 m sensor. After the event, it appears that we have more saline fluid above less saline fluid. For this configuration to remain stable, there must be a corresponding change in the temperature field as these are two primary variables which control density. Both the salinity and temperature time-series are presented in Figure 2.11. In addition to the two sensors shown in Figure 2.10, we have also presented the 9.1m deep HOBO sensor from the H2S chain. Perhaps the most striking feature of these two plots is that the 9.1m sensors line up almost perfectly in the salinity measurements but are almost a degree apart in temperature. One explanation could be that the H2S chain is located closer to the coast (as seen in Figure 2.1), and so there should be more seawater which is warmer, though this does not explain why the salinity remains the same. Another explanation is error during placement resulted in the sensors not exactly being placed at the depths recorded. In regards to the salinity inversion in late October, we can see that there is no corresponding change in temperature. The temperatures do not event cross, indicating that the salinity inversion should correspond to a density inversion. Using the equations of state from Section 2.1 we can combine the temperature and salinity measurements from Figure 2.11 into one density field which is presented in Figure 2.12. The density profile appears very similar to the salinity data, which is to be expected since the temperature readings remained relatively static during this duration. The second panel shows a positive density difference between the two sensors' indicating that the deeper sensor registers a larger density than the shallower one. The plot has been coloured so that positive values are

in red while negative values are in blue. After the September 2013 hurricane, the difference is predominantly negative for the rest of the time-series. As mentioned previously this is inherently unstable, observing the density difference we can see that change is not a huge amount. The maximum difference appears to be roughly 1 kg/m^3 or roughly 0.01%. We can also see that in the wake of the first declines into a negative difference, the system returns quickly to a stable state, but after several of these events it returns more slowly. We would argue that these inversions are caused by the rainfall induced-mixing the water column above the halocline. Normally these inversions are temporary until the system reverts back to a stable state, but in this case there are a series of rainfall events which continuously causes this mixing and extends the duration of this inversion. Furthermore, due to the large amount of rainwater that gets stored in the mangroves and upper mud layer of the ground, which takes time to drain into the cave system, and the constant out flow of the caves into the ocean, it is possible for this unstable state to exist for such as long period.

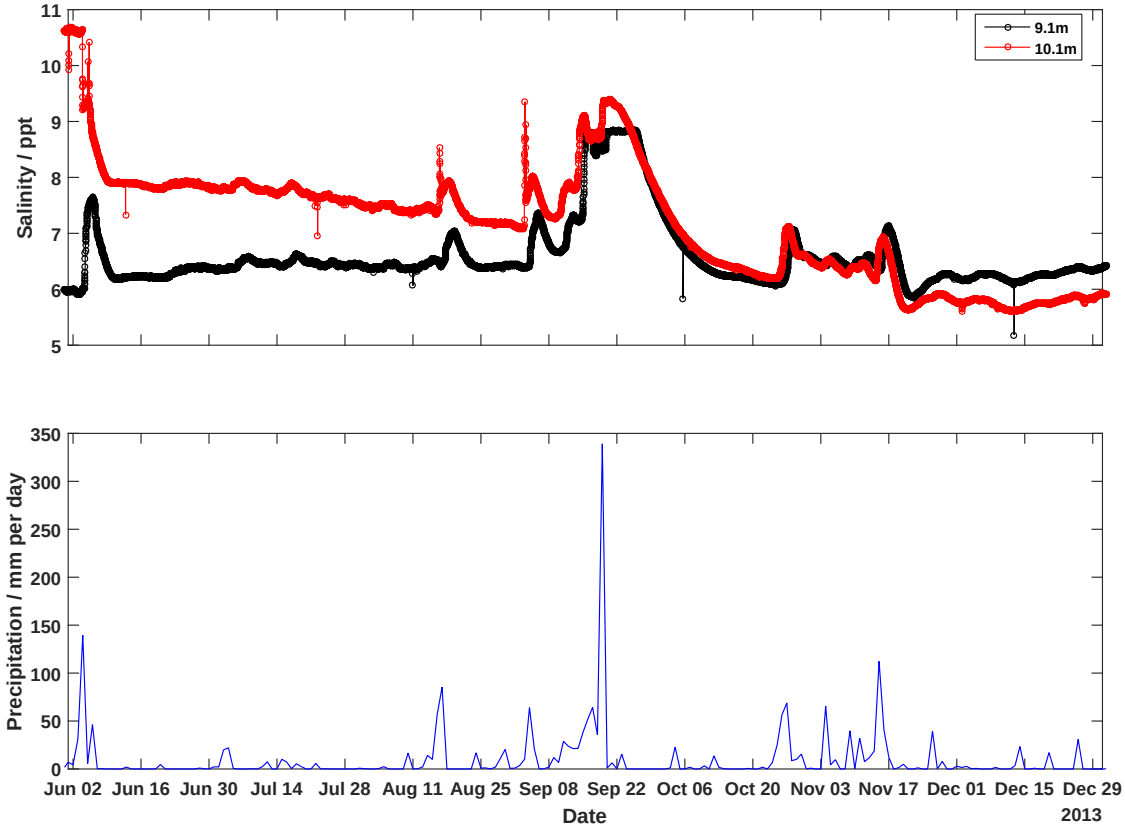

Figure 2.10: The time-series (top) of the salinity data collected by the HOBO sensors located at the ISOD2 chain,from June 2013 to January 2014. The back line corresponds to the 9.1 m deep sensor while the red line corresponds to the 10.1 m one. The bottom panel presents the daily rainfall during the same period. By inspection, it is easy to correlate large rainfall events with large jumps in salinity.

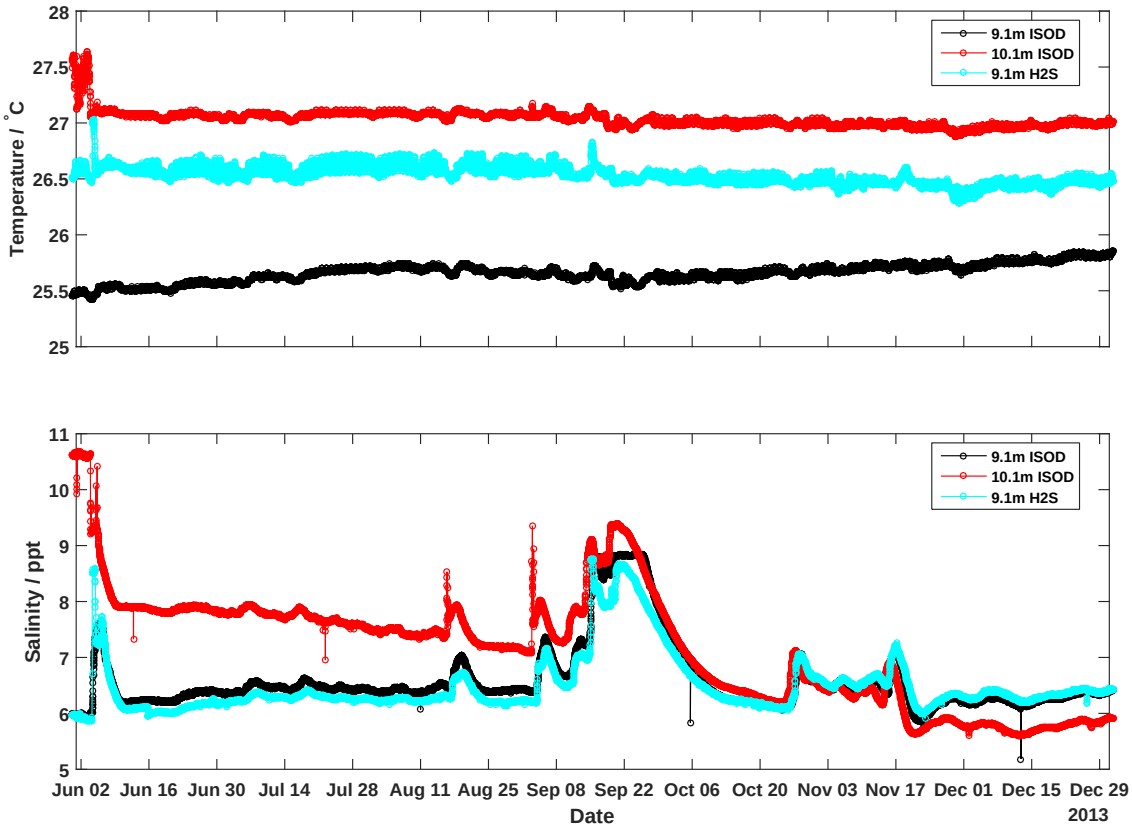

Figure 2.11: The time-series of the temperature data (top) collected by the HOBO sensors located at both the ISOD2 and H2S chains, from June 2013 to Jan 2014. The bottom panel shows the salinity measurements from the same sensors. Though the temperatures remain relatively stable during this half-year, the salinities do not, particularly from November onwards, where it appears that there is salinity inversion with the overlying fluid containing more salt than the underlying fluid. This along, with the temperature remaining constant, implies that the water column over this time period is unstable.

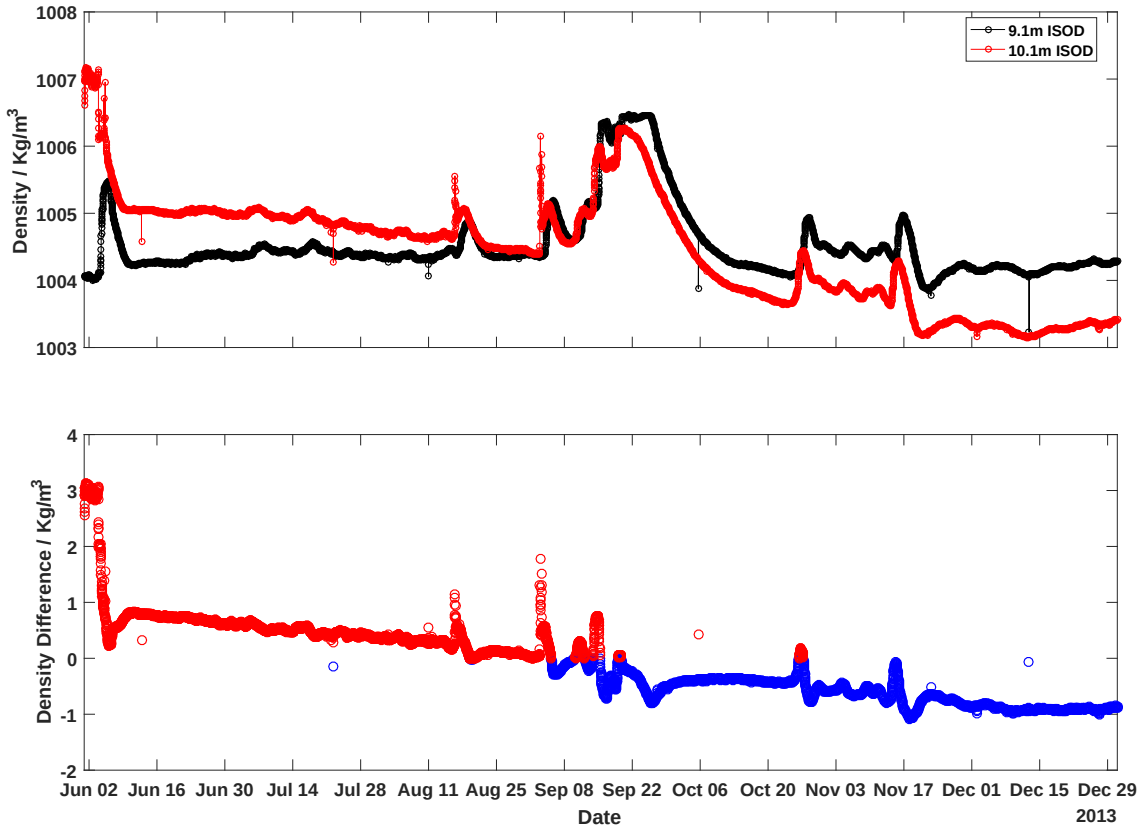

Figure 2.12: A comparison of densities at the two ISOD2 HOBO sensor locations (9.1 m and 10.1 m). The top panel shows the density computed using the UNESCO equation of state and the empirical equation for converting conductivity to salinity. The densities at the two sensors are initially dynamically stable with the deeper sensor registering a larger density. After the September 2013 hurricanes, this inverts and the higher sensor registers the larger density. This is highlighted in the bottom panel, which shows the density difference between the sensors. After the hurricane event, the density difference is consistently negative, albeit the density difference is quite small (a maximum of 1 kg/m^3).

Chapter 3

Numerical simulation

3.1 Reduction of physical experiments

As we saw in the previous chapter, field measurements are an extremely important tool for understanding how hydrological systems such as those seen in the Yucatan peninsula react to various inputs (i.e., hurricanes) and their general dynamics. However, one of the major limitations of field measurements is that there are many degrees of freedom and the systems are extremely complicated. This makes it difficult to concretely conclude that any results observed is uniquely caused by a specific event or mechanism. There is also a finite amount of information that can be gathered through the use of discrete sensors which are placed into the environment. Due to this lack of resolution within the system, many smaller scale features cannot be observed. In our particular case, the sensors only measured pressure, temperature and conductivity, giving no information about velocity or how the water within the freshwater or saltwater layers reacted to the heavy rainfall. Even provided that the sensors could have detected flow rates we still would have extremely sporadic information, which may not be sufficient for understanding the full scope of the phenomenon. To better understand the dynamics, we must thus turn to numerical experimentation. In this manner, we can examine specific cases and fully observe the resulting dynamics. Of course it is not possible to simulate the entire, complicated, real world system and so approximations must be made. One of the primary goals for these simulations is to see how a region that is being forced and mixed directly (e.g., raining on an open cenote) results in mixing away from the forcing site (e.g., within a cave network, for which the mud, mangroves and the karst itself provide shielding from direct rain). The difficulty for a numerical simulation is that 'rain' is an extremely problematic boundary condition to simulate due to its random, time-dependent nature and due to the input of mass and momentum. To see how we can simulate the effect of rain more simply we preformed a series of small scale physical experiments and observed the results. The dimensions of the tanks that were used are shown in Table 3.1. The specific parameters for each experiment are presented in Table 3.2 along with the dimensionless Reynolds and bulk Richardson number. The Reynolds number is a ratio of the inertial force to the viscous force. It is written as $Re = \frac{U}{l\nu}$, where U is the characteristic velocity, l is the characteristic length, and ν is the dynamic viscosity (for water, and thus all cases considered here, $\nu \approx 10^{-6}$). The bulk Richardson number is the ratio of the inertial force to the buoyancy force. It can be written as $Ri = \frac{g'l}{U^2}$, where U is the characteristic velocity, l is the characteristic length, and $g' = g\frac{\rho_1 - \rho_2}{\rho_1}$ is called the reduced gravity with ρ_1 and ρ_2 being the upper and lower layer densities. Our first experiment, shown in Figure 3.1, presents the mixing that occurs for a very small scale event, namely a large beaker. The initial fluid was stratified with a small saltwater layer approximately five times smaller than the freshwater layer in vertical

extent. The rain was created using a funnel with a perforated surface covering the bottom, herein referred to as the 'rainmaker'. When the rain begins, a small amount of mixing occurs, however when more water was added to the rainmaker the momentum of the rain increased slightly (due to the pressure of the increased water depth in the funnel) and this immediately began disturbing the halocline, and resulting in saltwater being transported upwards. The experiment indicates that even with an extremely unrealistic freshwater depth, and at low Reynolds number, that mixing can still occur. Furthermore, from this experiment it appears that the effect of the rain on the halocline is to force the interface it downwards resulting in upward motion due to conservation of mass.

Tank	**Length**	**Height**	**Width**
Small	0.78 m	0.16 m	0.17 m
Large	2.00 m	0.30 m	0.20 m
Beaker	0.10 m	0.30 m	0.10 m

Table 3.1: The dimensions of the various tanks which were used for experimentation.

Experiment Figure	Figure 3.1	Figure 3.2	Figure 3.3	Figure 3.4
Tank	beaker	large	small	small
Freshwater depth	0.28 m	0.10 m	0.11 m	0.08 m
Saltwater depth	0.02 m	0.15 m	0.05 m	0.08 m
Density difference	3%	2%	2%	2%
Characteristic length/velocity	l≈0.01 m, U≈0.01 m/s	l≈0.1 m, U≈0.01 m/s	l≈0.01 m, U≈0.01 m/s	l≈0.01 m, U≈0.01 m/s
Reynolds number	$\approx 10^2$	$\approx 10^3$	$\approx 10^2$	$\approx 10^2$
Bulk Richardson number	≈29.4	≈196	≈19.6	≈19.6

Table 3.2: The experiment figure, tank used, stratification information, and dimensionless parameters for the various experiments shown.

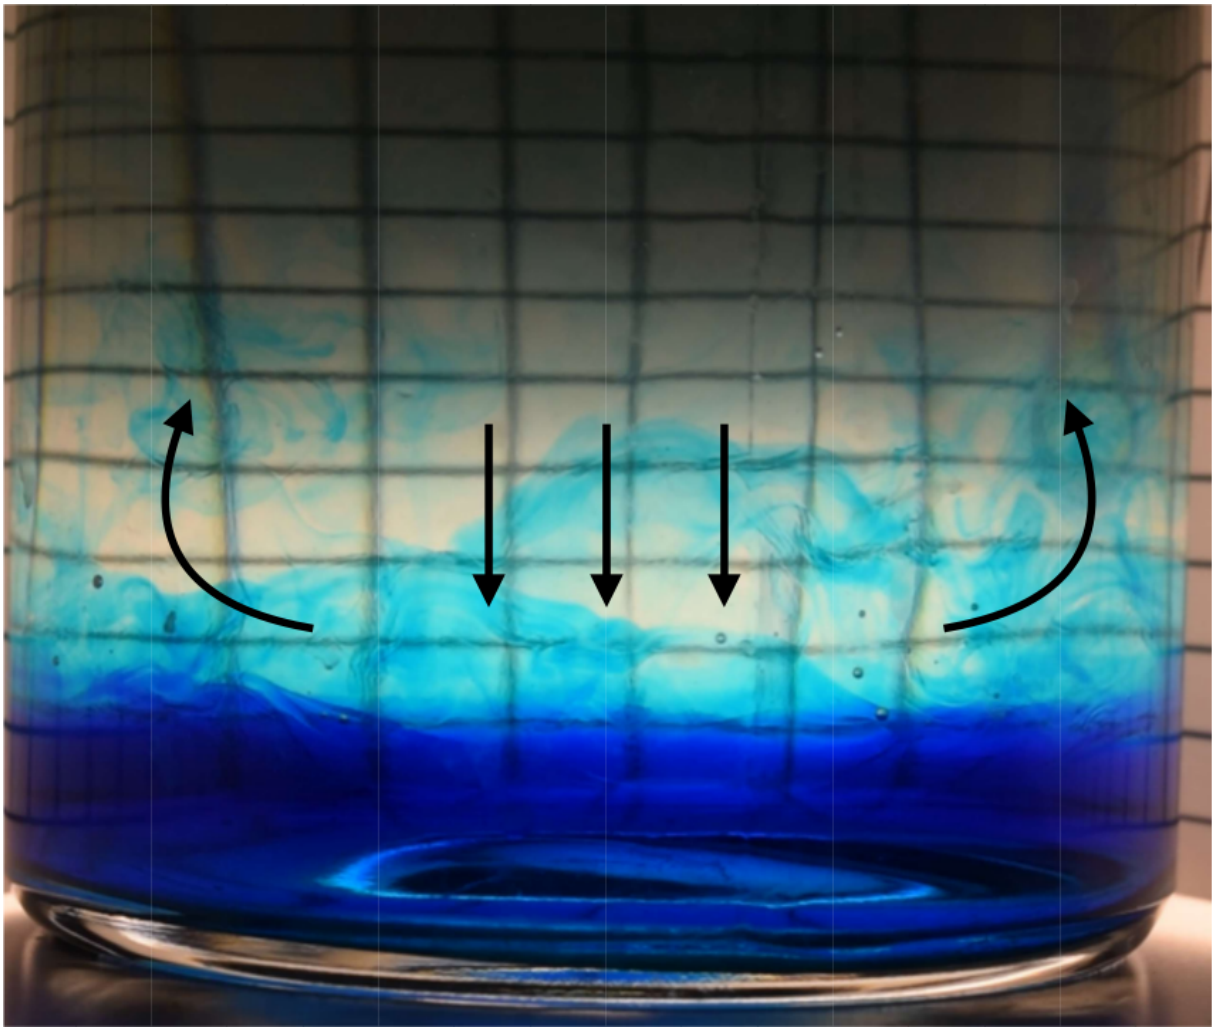

Figure 3.1: A small scale experiment of rain dropping on a stratified fluid. The freshwater lens was made to be approximately five times thicker than the saltwater layer. When the rain first begins falling on the surface the momentum is transferred all the way to the bottom of the fluid and forces the halocline which sends saltwater upwards.

To see the impact that the rain has away from the forcing region we performed a similar experiment in a much larger tank. The halocline was also moved up closer to the mid depth, see the Saltwater depth in Table 3.2. Potassium permanganate dye was placed in columns away from the forcing region to observe the horizontal velocity. Figure 3.2 shows this experiment at two times, one shortly after the forcing begins and one later in the

evolution. At early times the forcing region begins entraining the saltwater at the halocline directly below it, thereby inducing motion elsewhere in the tank. There is a corresponding motion away from the forcing region in the middle of the saltwater and freshwater layers. At later times we can see that the region close to the forcing has been completely mixed. The deformation to the potassium permanganate bands seen at the previous time has become more pronounced, with the region at the halocline moving towards the forcing region, while the mid portion of the density layers moves in the opposite direction. We see the same effect in Figure 3.3, where we have performed a similar experiment, however with a smaller saltwater layer. In this case we can see the extent to which the rain results in the halocline below the forcing region being entrained to the surface. Combining the results across the experiments, the net effect of rain on the halocline appears to be to force the halocline down, which by conservation of mass results in portions being forced up, and eventually getting mixed directly by the rain above the halocline. Over time, as more fluid gets pulled in and mixed, the turbulent region grows and fills the entire domain. Based on these observations, the solution that we came up with to implement in the numerical simulation was a series of jets located at the halocline which would force the halocline apart similarly to the momentum effect of the rain. This will then lead to the propagation of a turbulent front, and systematic transport to and from the forcing region in different sections of the water column. To test the jets idea we performed an experiment where we mimicked the jets by physically mixing the halocline with a potato masher. This experiment is shown in Figure 3.4. This picture was taken 40 seconds after the forcing began. At this point in the experiment, the potassium permanganate has deformed into the familiar wedge shape from the bottom panel in Figure 3.2. This comparison supports our jet idea as a simplification of the net effect of rain. The one major difference apparent from the figures is that there is no rightward flow along the surface. This is visible in the shape of the potassium permanganate column within the freshwater layer. In Figure 3.4 the entire column is being transported towards the forcing region while in Figure 3.2 it is only the portion closest to the halocline. The reason for this difference is that the rain case we are adding freshwater and so there is a net flow across the tank as this water spreads out, this effect is not seen in the jet case. In summary, by performing a series of small-scale experiments and observing the effect that rain has on the halocline we have found an analogue for this effect which is significantly less difficult to implement numerically. By performing an experiment mimicking this jet analogue we have shown that it produce a similar effect as the rain forcing.

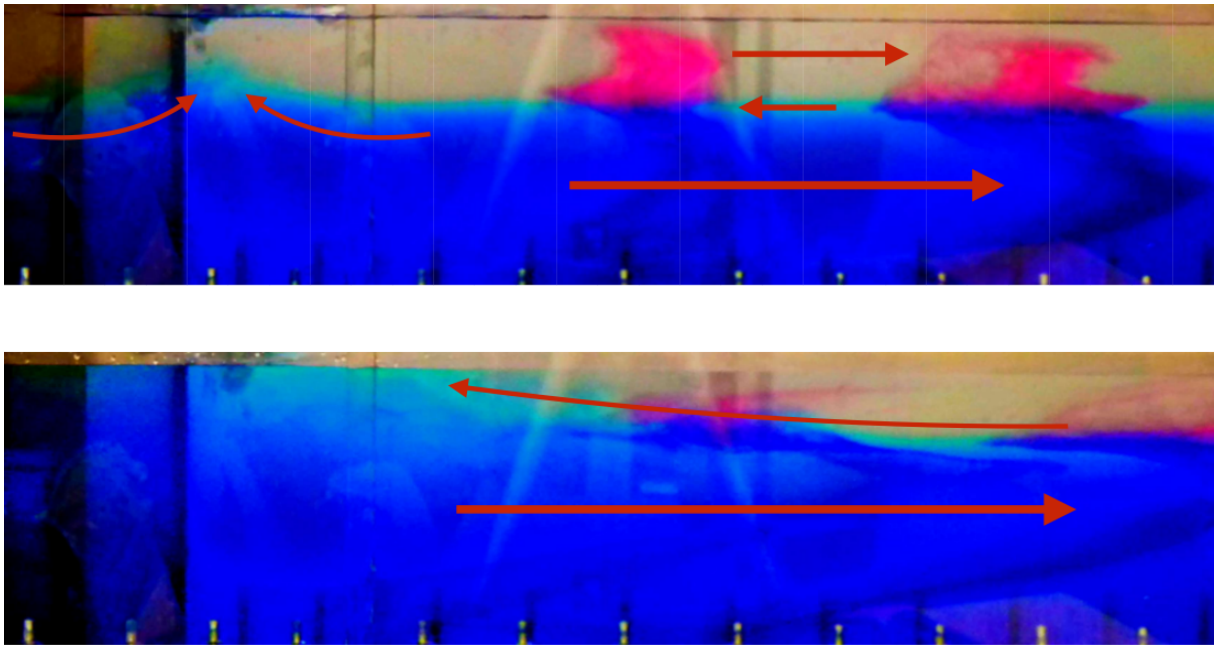

Figure 3.2: A rain-halocline experiment in a large tank at early and late times. Both times show the same general structure, namely motion along the halcline towards the forcing region, and motion away from this region in the middle of each layer. In the forcing region, the halocline is pulled up towards the surface where the rain is falling. This eventually leads to mixing within the entire depth of the tank within this region. Over time, this region grows as more fluid is entrained and is subsequently mixed.

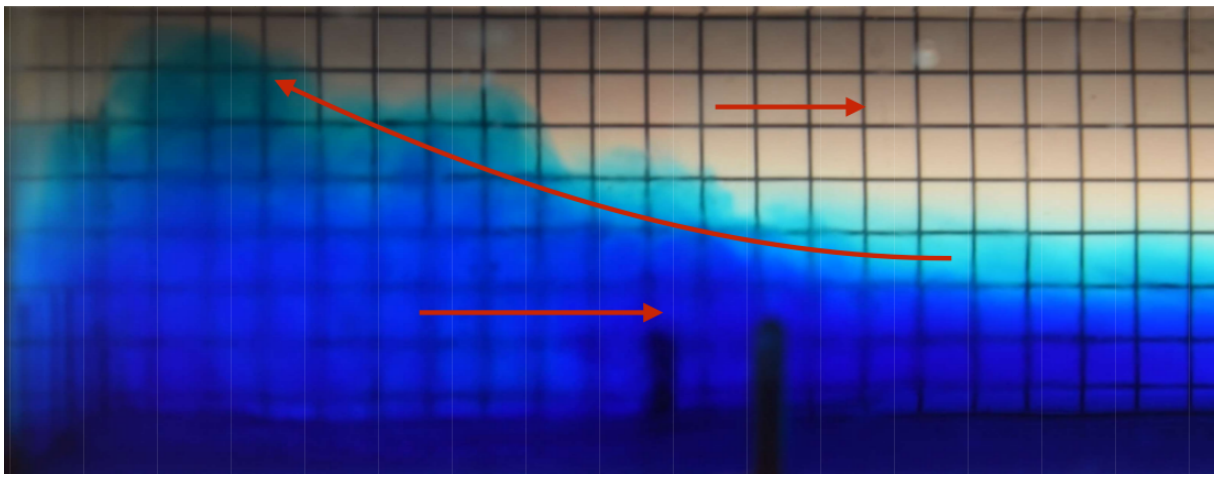

Figure 3.3: A similar experiment to that presented in Figure 3.2, however we have made the saltwater layer smaller, see Table 3.2. We again see the same effect of motion along the halocline towards the forcing and motion away within the rest of the fluid.

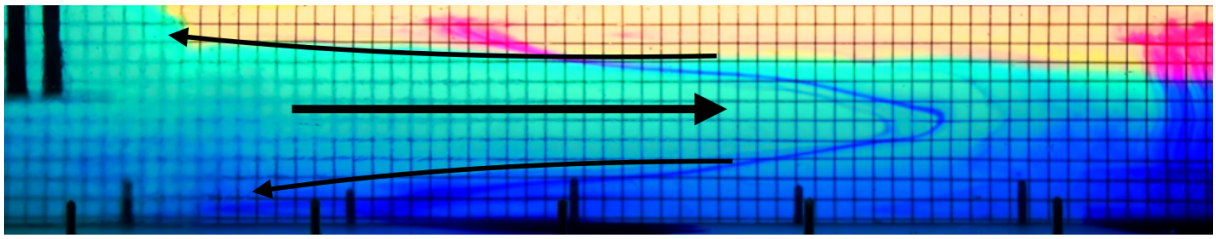

Figure 3.4: To test our jet idea and compare to the actual rain experiment we mimicked the jets by physically forcing the halocline with a potato masher. From the image we see a very similar shape to the previous experiments. We have movement along the halocline towards the forcing region and flow away above and below this level. The similarities are especially clear when compared to Figure 3.2.

3.2 Numerical Methods

As mentioned in Section 1.4, our numerical simulation solve the incompressible Navier-Stokes equations under the Boussinesq approximation. These equations were solved using the in-house, spectral incompressible Navier Stokes equation solver (SPINS) [31]. For temporal discretization, the code uses a third-order variable length timestepping method. A pseudo-spectral collocation method for spatial discretization is implemented, with the choice of spectral expansion depending on the desired boundary conditions. A Fourier expansion is used for periodic boundary conditions, a sine/cosine expansion is used for free-slip boundaries, and Chebyshev polynomials are used for no-slip conditions. See [4] for more details. The problem with using a spectral collocation method for the nonlinear advection term is that it introduces aliasing errors, wherein high-wavenumber waves, beyond the Nyquist frequency, alias onto lower wavenumber, well resolved, waves. To remedy this, an exponential filter with a cutoff is used to suppress these high wavenumber waves. The most computationally difficult part of the code is solving the Poisson problem for pressure, since there is no evolution equation for pressure. To solve this, a finite difference operator is used to precondition the Poisson operator before using the generalized minimum residual method (GMRES) to solve the problem. Together, the result of these methods is a fully-parallelized code which can solve the incompressible Navier-Stokes equations with spectral accuracy (i.e. the accuracy of the simulation scales with the number of grid points [32]).

For our specific numerical simulation we imposed free-slip boundary conditions along all boundaries because of our interest in the dynamics of the pycnocline, as opposed to details of the boundary layer dynamics. This was also done to speed up computation time, as full 3D simulations are extremely computationally intensive. In accordance with the small-scale laboratory experiments, the physical dimensions of the tank were chosen to be 1m $\times$ 20cm $\times$ 20cm ($L_x \times L_y \times L_z$). Since we are expecting mixing to occur we chose to use a high resolution, namely 4096 $\times$ 256 $\times$ 256 (same ordering as the physical dimensions). This number of points provides a resolution of $2.44 \cdot 10^{-4}$m in the x-direction and $7.81 \cdot 10^{-4}$m in both y and z. This is sufficient for a direct numerical simulation (DNS) where we have resolved all scales and do not require a separate model for turbulence below the grid scale. A calculation of the smallest Kolmogorov length scale, $\eta = (\nu^2/2e_{ij}e_{ij})^{1/4}$, where e_{ij} is the strain rate tensor, at 80s results in $\eta \approx 3.7 \cdot 10^{-4}$m which is resolved in the x direction and nearly resolved in y and z. This is sufficient to show that we are resolving all the physical scales and can treat the simulation as a 'numerical experiment'.

The initial conditions for the density are defined as

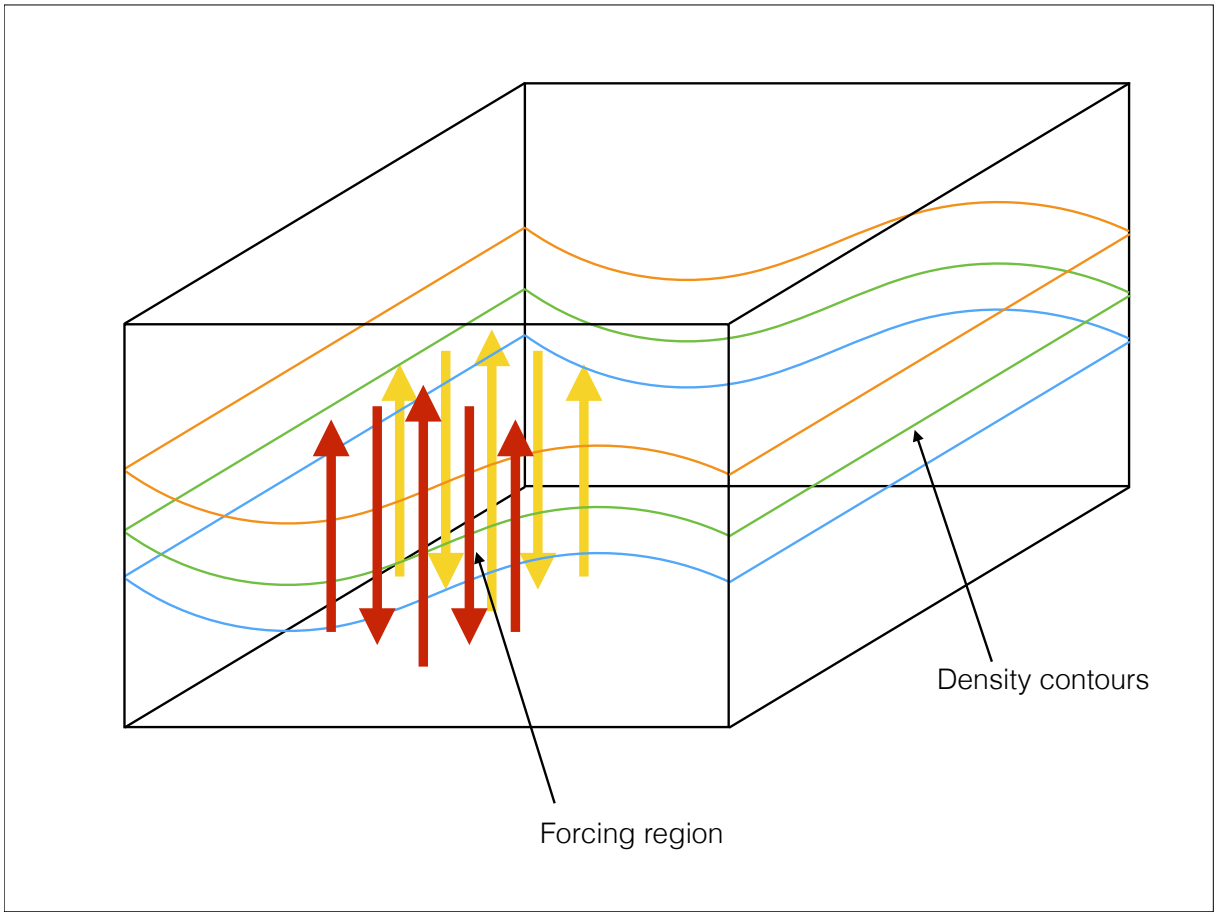

Figure 3.5: A schematic of the numerical simulation set up. The forcing region is located on the left hand side of the tank, centred at $x = 0.25$. The pycnocline is located exactly at the mid-depth, and corresponds to a 1% density change. All the boundaries were given free-slip boundary conditions since we were interested in the effect that the jets had on the pycnocline.

$$\rho = \rho_0 \left(1 - 0.005 \tanh\left(\frac{z - 0.5L_z}{0.05L_z}\right)\right),$$

where $\rho_0 = 1000\,\mathrm{kg/m^3}$. The pycnocline was fixed at the mid-depth to minimize the effect of the boundaries. Though not fully comparable to known measurements ([3] measured a $\Delta\rho \approx$ 2-3%), we decided to implement a 1% density change to allow for weaker

jets, due to lower inertia. As explained in the previous section, we have decided to mimic the effect of rain forcing by using alternating jets located within the pycnocline. This jet forcing has the functional form,

$$\boldsymbol{F}(x,z,t) = 0.2\sin(2\pi t)\exp\left(\frac{x-0.25}{0.05}\right)^2 \sin\left(\pi\frac{z}{L_z}\right)\sin\left(2\pi\frac{x}{0.04}\right)\hat{\boldsymbol{k}}.$$

This function forces the w-momentum equation directly, and forms jets which alternate direction and weaken through x, and pulsate in time with a period of one second. Since the outputs occur every two seconds, on the second, the forcing will appear to be constant through outputs. The horizontal extent of the forcing region is 0.1m centred at $x = 0.25$ m. The forcing is constant through the y direction. We have placed a passive 'dye' column at $x = 0.65$ m (0.1 m wide) to observe how this column is entrained by the forcing, similar to the potassium permanganate in the physical experiments.

This simulation was performed on the high performance computing (HPC) network SHARCNET over a period of 32 days using 128 processors. Due to the size of the simulation and the amount of data that it produced, many new issues arose during the analysis and post processing phase of the project that would not have occurred for smaller-sized simulations. In computing, these issues are known as the field of big data. While our simulation was not at the scale for which all previous analytical techniques become impossible, it was still necessary to change almost all of the previous techniques for efficiency (e.g., the computation of derivatives in post-processing) and in some cases to change the technique completely (e.g. for 3D graphics the use of Matlab became impossible and the open source software Visit [6], was used instead).

3.3 Velocity and density

Line plots can often provide more quantitative information than 2D (e.g. contour) or 3D (e.g. isosurface) plots. For examining the general evolution of this simulation, the vertically and spanwise averaged kinetic energy provides a clear distinction between the different regions that form within the tank and how they evolve. Figure 3.6 shows this averaged kinetic energy for a sequence of times from 40 s to 400 s. The x-axis begins at $x = 0.25$ m (the centre of the forcing region), since we are primarily interested in how the forcing impacts the development of a turbulent front that moves into the quiescent region. For all times, we can see that from $x = 0.25$ m to $x = 0.35$ m, the kinetic energy is completely dominated by the jets, with their shape visible in the kinetic energy. As we

evolve in time we can see a local maximum (not due to the jets) within the kinetic energy that propagates rightward through the tank. This is especially visible from 40s (at $x \approx 0.45$ m) to 200s (at $x \approx 0.85$ m). To the right of this local maximum the kinetic energy drops off, though at early times there are oscillations that occur near the far end of the tank. These oscillations appear to correspond to internal waves which were created when the jets began disturbing the pycnocline. This also explains their lack of appearance at later times since by then any waves have reflected back into the turbulent region. Figure 3.6 allows us to demarcate clear distinct regions based upon their kinetic energy profile. The three regions can be separated as jet-dominated (from $x = 0.25$ m to $x = 0.35$ m for all times), turbulence dominated (from $x = 0.35$ m to the turbulent front defined as the local maxima in kinetic energy) and quiescent (from the front to the end of the tank).

To see the general impact of the jets over time, Figure 3.7 shows a 3D volume plot of the kinetic energy. The panels correspond to 100 s, 200 s, 300 s and 400 s respectively. A key result of this figure is that ,throughout the simulation, the areas of large kinetic energy (seen in white in these figures) do not appear to propagate away from the jet forcing region. The colorbar has been kept the same between the different panels. Another key aspect from this figure is the location of the high kinetic energy areas. Other than the jets, which are obviously high kinetic energy areas, for all times, the regions near the upper and lower boundaries to the sides of the jets contain high kinetic energy. This can be attributed to the flow generated as the jet-forced fluid impacts the upper and lower boundaries of the tank. Comparing these 3D figures to the 1D Figure 3.6, we can see that while the 1D figure provided general information about the different regions in the simulation. it completely lacked information about the structure within these regions. For example, in Figure 3.7(d), there is an extrusion of high kinetic energy along the close wall which is not revealed in the corresponding 1D plot.

Figure 3.8 shows the spanwise averaged horizontal velocity u in the x-z plane. Panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. The panels have been saturated at 1 mm/s to highlight the distinct regions. Focusing on the right side of the jets, we can see that there is a change in dynamics between 100 s and 200 s. For the earlier times, we can see that there is a positive horizontal velocity, u, within the pycnocline and primarily negative u above and below this. At the later outputs this reverses and the pycnocline has a negative u with regions of positive u above and below it. The structure of this velocity field indicates that there are numerous areas of high shear between the different flow directions which will produce large amounts of vorticity via tilting and stretching. This is particularly evident in the region near the boundaries to the sides of the jets, which corresponds to the regions of high kinetic energy visible in Figure 3.7. As mentioned earlier in regards to the 1D plots, the problem of averaging

through a dimension is that we lose any information about the 3D structure, which given the turbulent nature of this simulation is assuredly there. One way to gain this information through 2D figures is to not average, but rather take a slice of the field at a given time. In Figure 3.9, we present slice in the x-y plane at (a) $z = 0.12$ m, (b) $z = 0.14$ m, (c) $z = 0.16$ m, and (d) $= 0.18$ m. As we move through the region above the pycnocline we can see a multitude of 3D structures throughout the spanwise dimension. This is particularly evident in Figure 3.9(c), with a large structure of positive u completely surrounded by negative u.

The most useful field to compare to physical simulations, and to field data is density (since in the lab we can dye one density one colour and we can directly measure this in the field). Figure 3.10 shows the time evolution of density using a 3D volume plot. Panels (a), (b), (c) and (d) correspond to 100 s,200 s, 300 s, and 400 s, respectively. At early times, (a), we can see how the jets immediately begin deforming the pycnocline, however, it is important to note that the fluid being forced is primarily of a density from the outer layers of the pycnocline, or the weakly stratified region. It is not until later times, (c) and (d), where we begin to see fluid from a more intermediate density begin to move. With how we have chosen the opacity of these plots we can see all the 3D structure that appears within the turbulent region. We can also see how this region entrains fluid from the quiescent region towards it. It appears that the fluid being entrained is entirely from the outermost region of the pycnocline. Furthermore based on the shapes visible within the turbulent region, the entrainment occurs via vortices within the region ripping away at the outer layers of the pycnocline. By the last panel (d) the shape of the pycnocline appears to be extremely similar to the shape seen in experiments, as typified by the first panel of Figure 3.2.

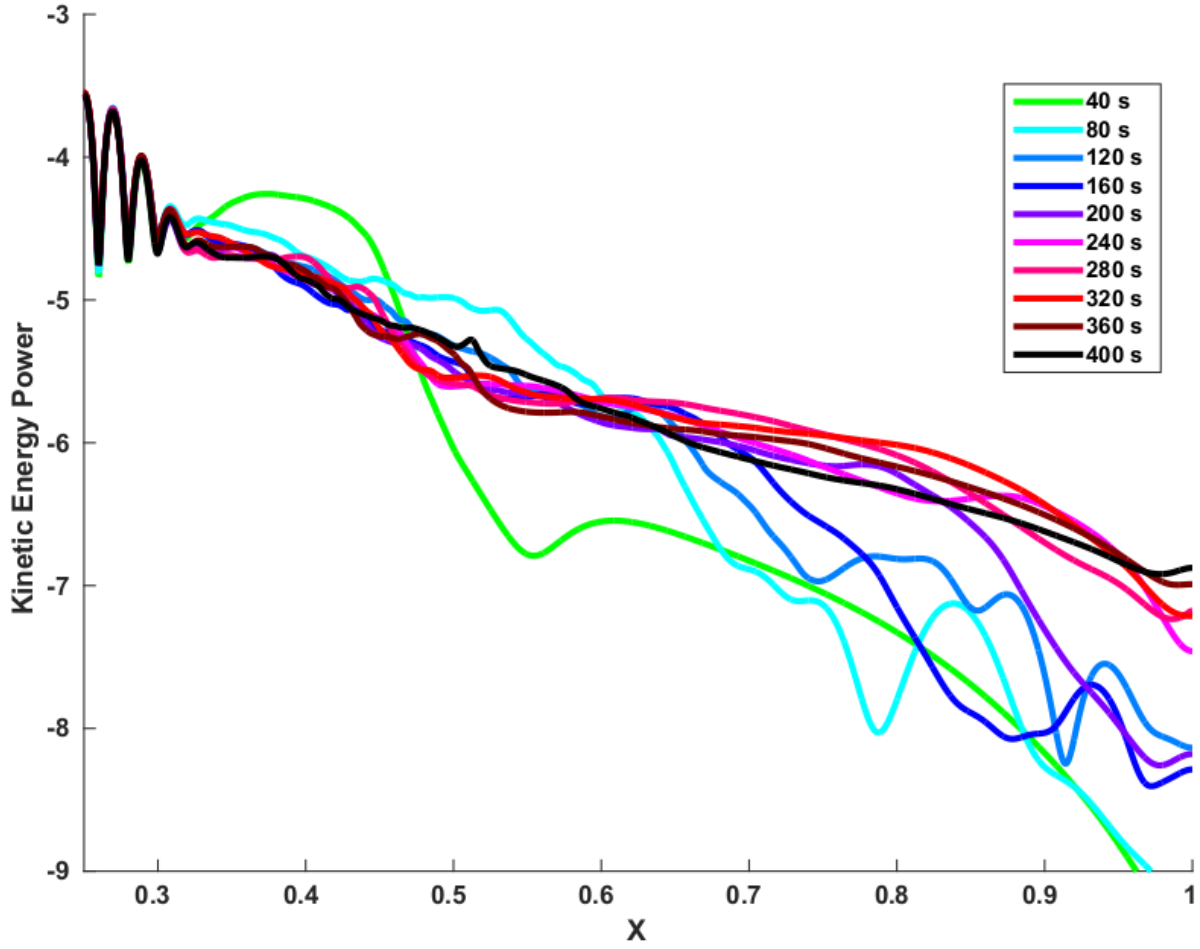

Figure 3.6: Vertically and spanwise averaging the 3D kinetic energy reduces the full 3D scalar field into a more manageable 1D vector. This method also highlights the extent to which the turbulent entrainment has moved through the tank. There are three distinct regions to the right of the mixing region, which are highlighted by the kinetic energy. The first region is dominated by the motions induced by the forcing and extends to $x = 0.45$ m. The second region is characterized by a relatively slow decline in kinetic energy and is where the entrainment occurs. Over time, the rightward edge of this region has moved through the tank. The third region is the relatively quiescent fluid that has not been disturbed yet on the far right. This region shrinks as the turbulent entrainment moves through the tank. This figure shows that by 80 s the kinetic energy has equilibrated and follows a steady pattern afterwards.

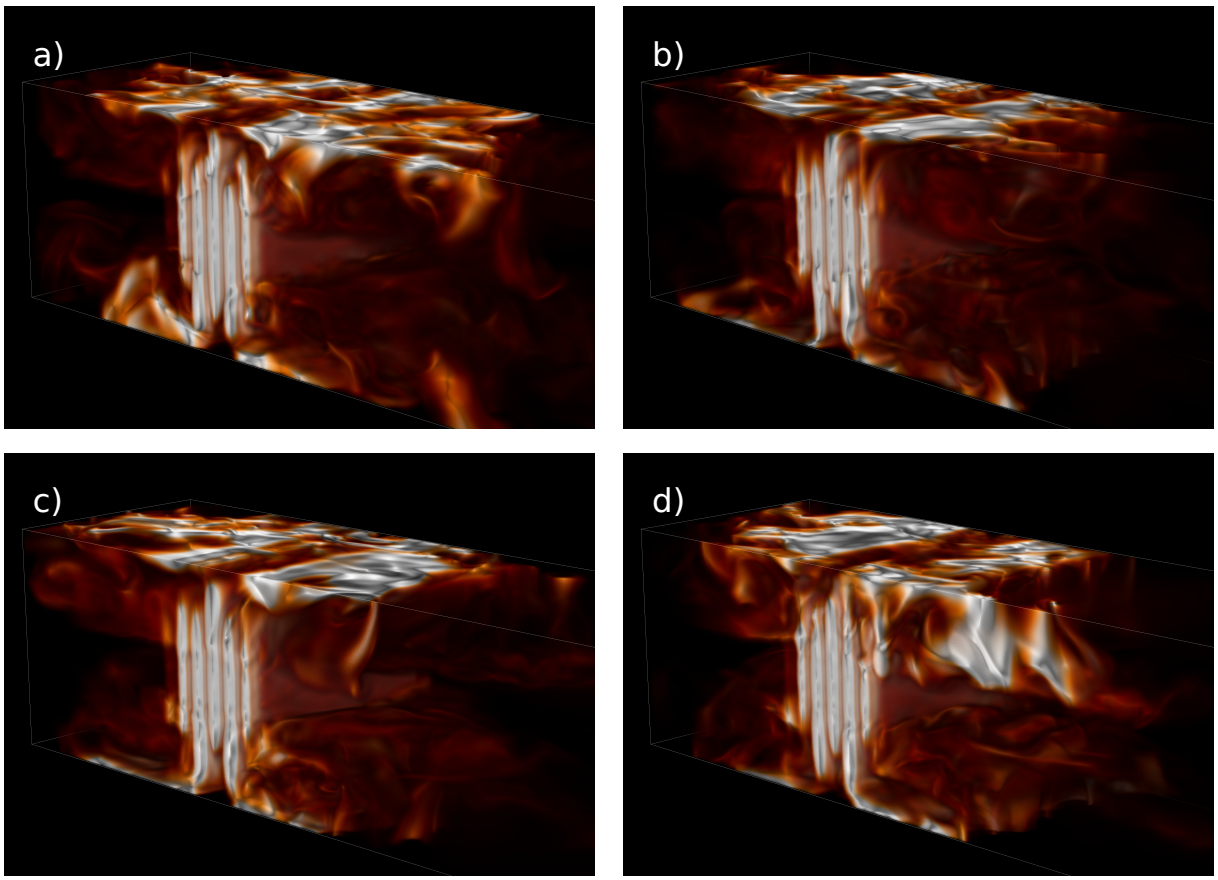

Figure 3.7: A 3D volume plot of kinetic energy made using a ray casting method. Panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. The colour saturation remains the same between all the panels. The regions of high kinetic energy (white) are concentrated in the jets and along the boundaries to the sides of the jets.

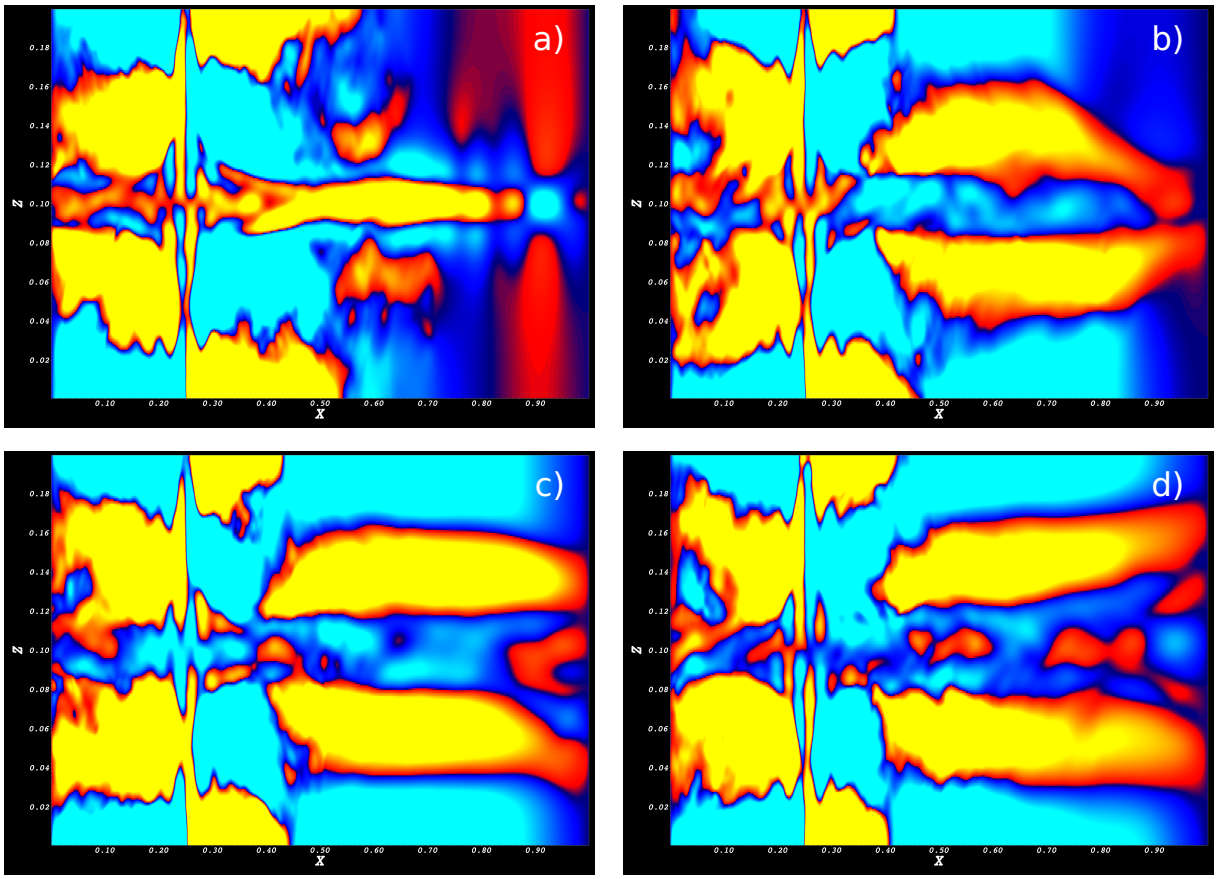

Figure 3.8: 2D spanwise-averaged horizontal velocity pseudocolour plots. The panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. The panels have been saturated at ±1mm/s to highlight the distinct regions. There is a switch in the dynamics between panel (a) and the later ones, with the regions within and around the pycnocline switching flow direction.

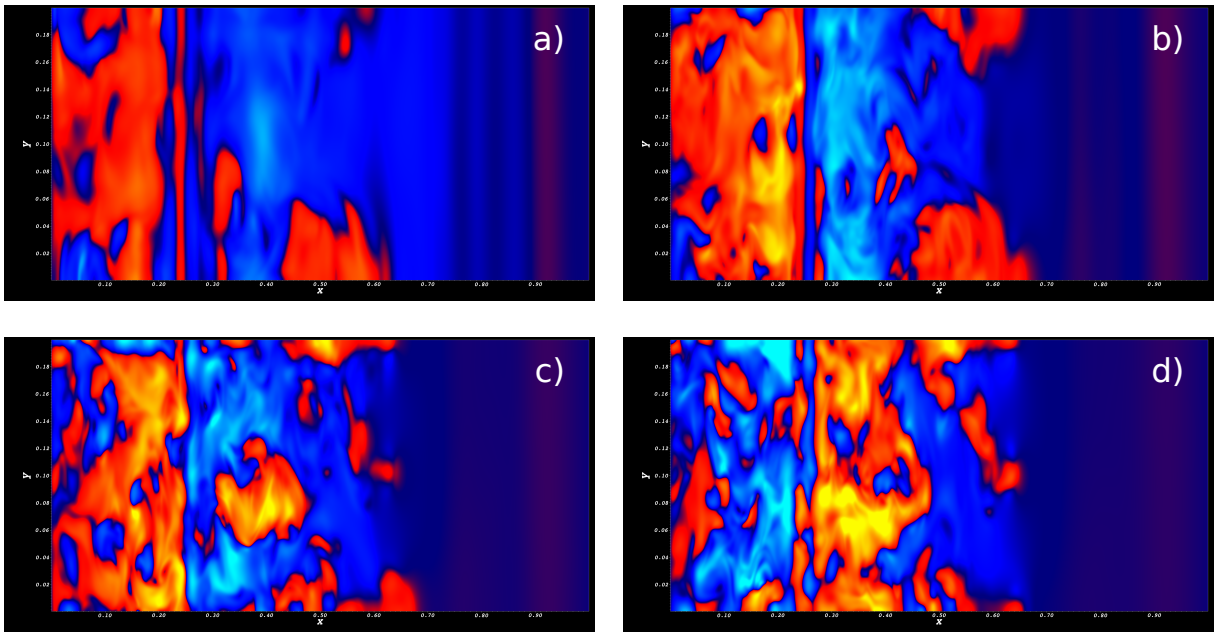

Figure 3.9: 2D slices of horizontal velocity in the $x-y$ plane. The slices were taken at (a) 60%, (b) 70%, (c) 80%, and (d) 90% of the vertical extent. To allow for more structure to be seen within positive and negative regions, we have increased the maximum and minimum velocity values to $\pm$5mm/s. As we move through the upper layers of the tank we can see a plethora of 3D structures which are not visible when we average in the spanwise.

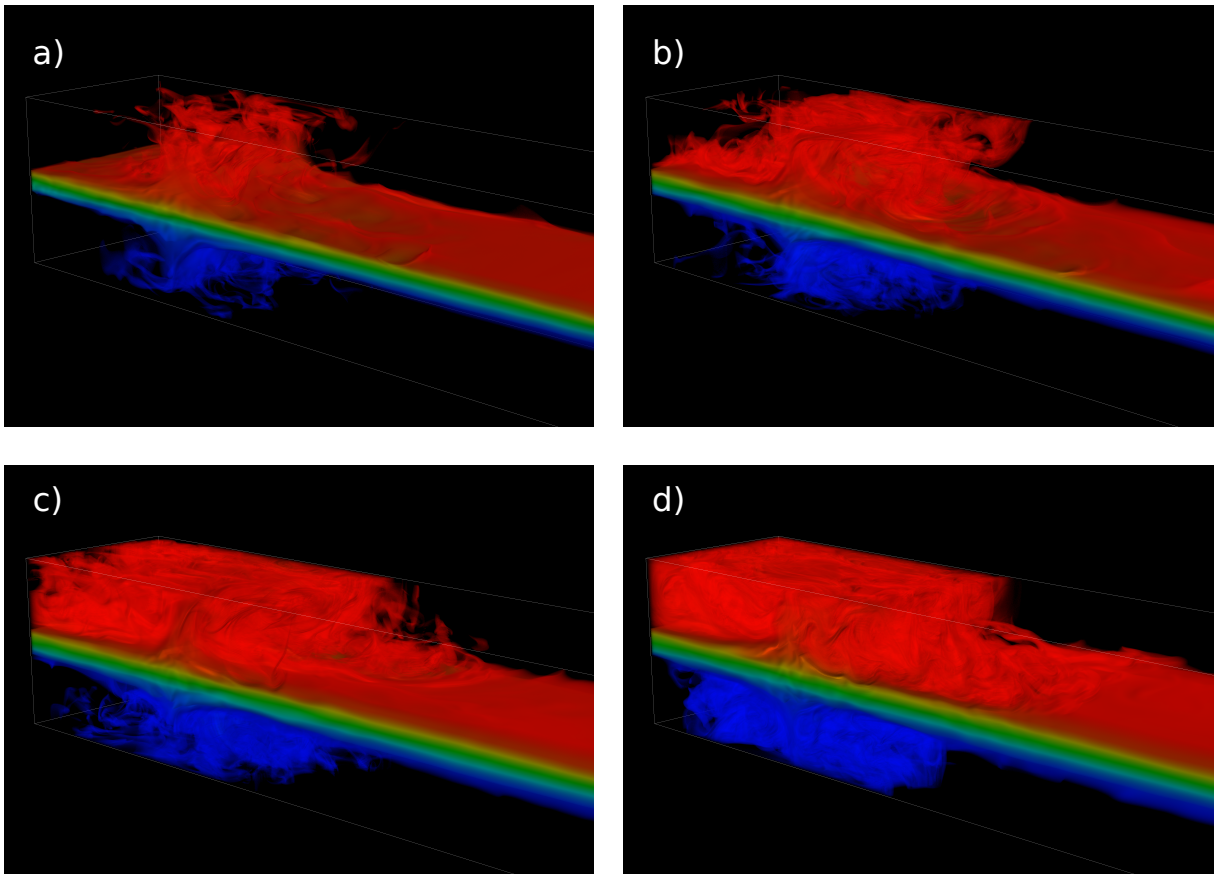

Figure 3.10: A 3D volume plot of density made using a ray casting method. Panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. We can see that, for early times, the jets only move the outermost fluid of the pycnocline, but at later times, (d) especially, we begin to see more intermediate density fluid being forced away. We can also see how the generation of the turbulent front entrains fluid from the quiescent region, especially from the outermost layers of the pycnocline.

3.4 Dye

As mentioned in the numerical methods section, we added a column of dye in the simulation. Figure 3.11 shows the time evolution of dye using a 3D volume plot. Panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. An interesting feature visible when comparing panels (a) and (b) is a change in the shape of the dye column near the pycnocline. Where in panel (a) the dye at the centre was forced to the right, above and below this to the left, we have an inversion of this in panel (b) with the centre being pulled to the left, and above and below to the right. This is consistent with the change in horizontal velocities visible in Figure 3.8. In the later panels, (c) and (d), we can see that the dye above and below the pycnocline is forced to the right end of the tank, while the portion that was pulled in from the middle of the pycnocline has reached the left end of the tank. We can also see the complex 3D structures which are created in the turbulent region, this is particularly visible in panel (d). In general, it appears that the portions of the column in the middle of the pycnocline, and along the top and bottom, get entrained the most, eventually reaching the jet forcing region. While dye in between these regions is advected in the opposite direction.

One analysis which can be performed on the passive tracer in this simulation is locating, and separating, mixing from stirring. This is important to quantify the extent to which actual mixing is occurring and its localization. This can be applicable to trace contaminants being released into this system. Mixing is the irreversible homogenization of a specific tracer (eg., ρ or dye), while stirring is the mechanical motion of the tracer. For two fluids this increases the interface surface area. Following the derivation in [25], consider a general passive tracer θ. The evolution of this tracer follows a simple advection-diffusion equation (where κ is the diffusion coefficient for the tracer),

$$\frac{D\theta}{Dt} \equiv \frac{\partial \theta}{\partial t} + \boldsymbol{u} \cdot \nabla\theta = \kappa \nabla^2 \theta.$$

One way to measure the spatial variability of the tracer is the gradient $\nabla\theta$. Over the entire domain we can define a useful measure,

$$C = \frac{1}{2} \iiint_V \nabla\theta \cdot \nabla\theta \, dV.$$

To see how this measure evolves in time, we next consider the time derivative of this term. By using the evolution equation for the tracer and the no-outflow boundary conditions we can write,

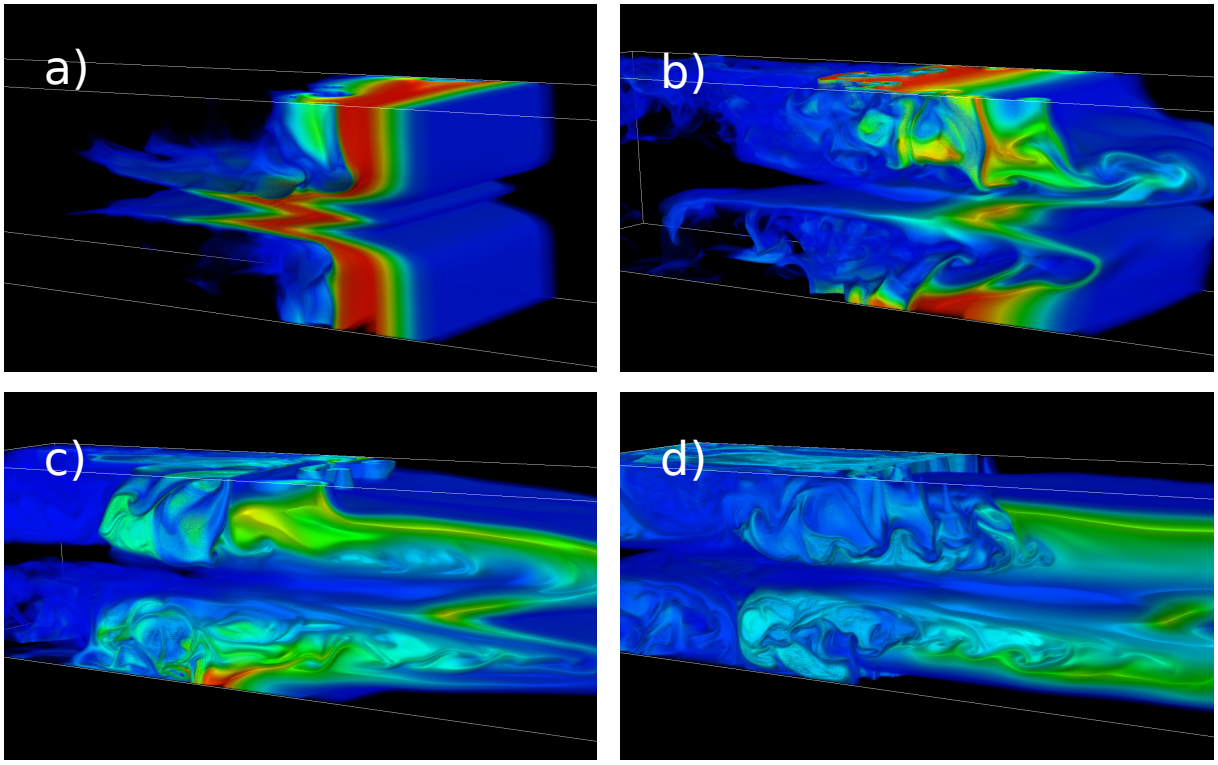

Figure 3.11: A 3D volume plot of dye concentration using a ray casting method. Panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. At early times, only the left side of the dye column is deformed, but by panel (d) the entire column has been ripped apart. We can also see the shift in horizontal velocities that was visible in Figure 3.8 by comparing panel (a) to (b). The zigzag shape from (a) becomes inverted and deformed in the next panel, with part of it being pulled all the way into the jet region. Also interesting with this change in shape, is that part of the dye above, and below the pycnocline actually gets pushed away from the jets into the quiescent region.

$$\frac{dC}{dt} = \iiint_V \left((\boldsymbol{u} \cdot \nabla\theta)\nabla^2\theta - \kappa \left(\nabla^2\theta\right)^2 \right).$$

In this equation, the first term inside of the equation represents stirring while the second represents mixing. Mixing is always negative and decreases the variability in C, this corresponds with physical intuition, as mixing homogenizes the fluid. Stirring can be either sign (though in most cases it is positive since it's easier to stir then unstir). An indirect

result of positive stirring is to make the mixing more efficient since stirring will increase the surface area over which mixing occurs. Together, these effects result in C increasing to a maximum at which point it decreases rapidly. Without mixing, C would solely increase over time, while without stirring, C would decrease though at a slower rate than both together.

Figure 3.12 shows the spanwise-average of this mixing parameter for the dye. The log base 10 of the colour axis has been taken due to the disparate scales of mixing. The panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. In panel (a) the mixing appears most dominant within the dye column, this is consistent with the panel (a) from Figure 3.11 since the column has not been deformed by that time. At intermediate times, in panels (b) and (c), the mixing reaches its maximum within the turbulent region. This is consistent with the turbulent region deforming the dye which results in increased mixing. At late times, in panel (d), the dye has become diffused throughout the tank, and so the intensity of the mixing has declined, though the areal extent has increased. An interesting feature primarily visible in panels (b) and (c) is an asymmetry between mixing above and below the pycnocline. There appears to be significantly more mixing occurring below the pycnocline than above. It is likely that this is due to large scale motions which can break the symmetry. Particularly visible in panels (c) and (d), are two extrusions of mixing to the right side of the tank above and below the pycnocline. Comparing these to panel (d) of Figure 3.11, these likely correspond to the dye that has been forced to the right side of the tank.

The partner of the mixing term, stirring, is shown in Figure 3.13. The panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. A quick glance at the panels reveals a very strong correlation between the two fields. When stirring is high it greatly deforms the tracer interface, which results in higher mixing. This correlation between the fields is particularly evident in panels (b) and (c). There appears to be little stirring occurring within the pycnocline and the jets. Not appearing in the jets can be explained physically since it is very hard for the dye to move against the forcing. However, within the pycnocline, we might expect some stirring given the presence of mixing and the 3D plots. It is possible that due to the spanwise averaging we are missing some of these features, however it is also possible that the motion within the pycnocline is not deforming the dye shape, but solely advecting it, there is evidence for this in Figure 3.11. By not deforming the dye it is not stirring and thus does not appear.

As when we considered the horizontal velocities, by computing the spanwise average we have lost any information about 3D structure within each field. This structure can be seen by considering slices as opposed to averages. Figure 3.14 shows two slices, at $y = 0.05$ m and $y = 0.15$ m, of dye density (shown in purple), along with both stirring and

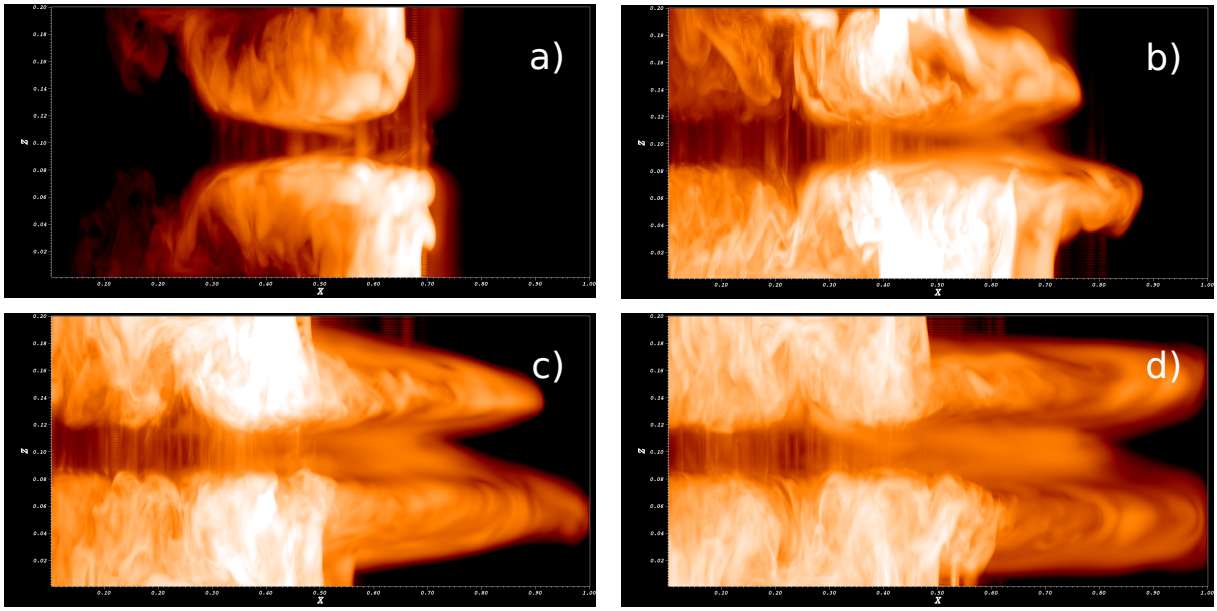

Figure 3.12: 2D spanwise averaged mixing parameter pseudocolour plots for dye. The panels (a), (b), (c) and (d) correspondto 100 s, 200 s, 300 s, and 400 s, respectively. At early times in panel (a), the dye mixing is concentrated around the dye column since there has been little deformation of the column. For intermediate times (b) and (c), the maximum mixing occurs within the turbulent region, and especially near the front. At late times (d), the dye has been diffused out and so less intense mixing occurs, albeit over a larger area.

mixing. The mixing parameter is presented in blue coloured areas, and a single contour of stirring is overlain in green. Panels (a) and (b) correspond to $y = 0.15$ m, while (c) and (d) correspond to $y = 0.05$ m. Panels (a) and (c) have been saturated to show only the highest regions of mixing, stirring and dye concentration (maximums correspond to 20000 for mixing and stirring and 0.25 for dye concentration), while panels (b) and (d) have a lower threshold (maximums correspond to 2000 for mixing and stirring and 0.1 for dye concentration). These slices were taken at 200 s. As in the previous figures, it is very clear that the mixing and stirring are very closely correlated, almost every location of high mixing is encircled by a contour of stirring. It is also quite clear that the majority of the mixing is occurring along the dye that is within the turbulent region. In all the panels the mixing primarily takes place between $x = 0.35$ m and $x = 0.55$ m, which corresponds to the interface of the dye with the turbulent region. Comparing the saturated ((b) and (d)) panels, with the unsaturated ((a) and (c)) panels, the highest stirring and mixing occurs

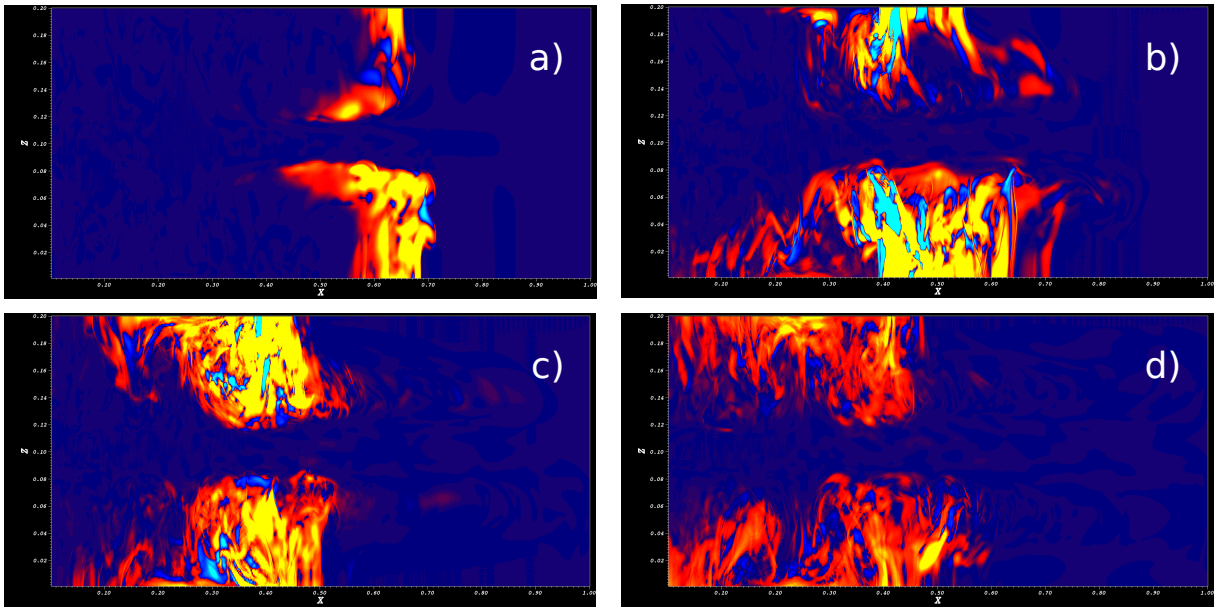

Figure 3.13: 2D spanwise averaged stirring parameter pseudocolour plots for dye. The panels (a), (b), (c) and (d) correspond to 100 s, 200 s, 300 s, and 400 s, respectively. Mimicking what we saw in Figure 3.12 we can correlate locations of high stirring with those of high mixing. As in the previous figure the stirring reaches a maximum in panel (b) and (c) with the stirring concentrated within the turbulent region. It is interesting that there are some regions of negative stirring, though the majority is positive.

along the edges of dye filaments with steep concentration gradients, mathematically this is seen by the $\nabla^2\theta$ term in both terms.

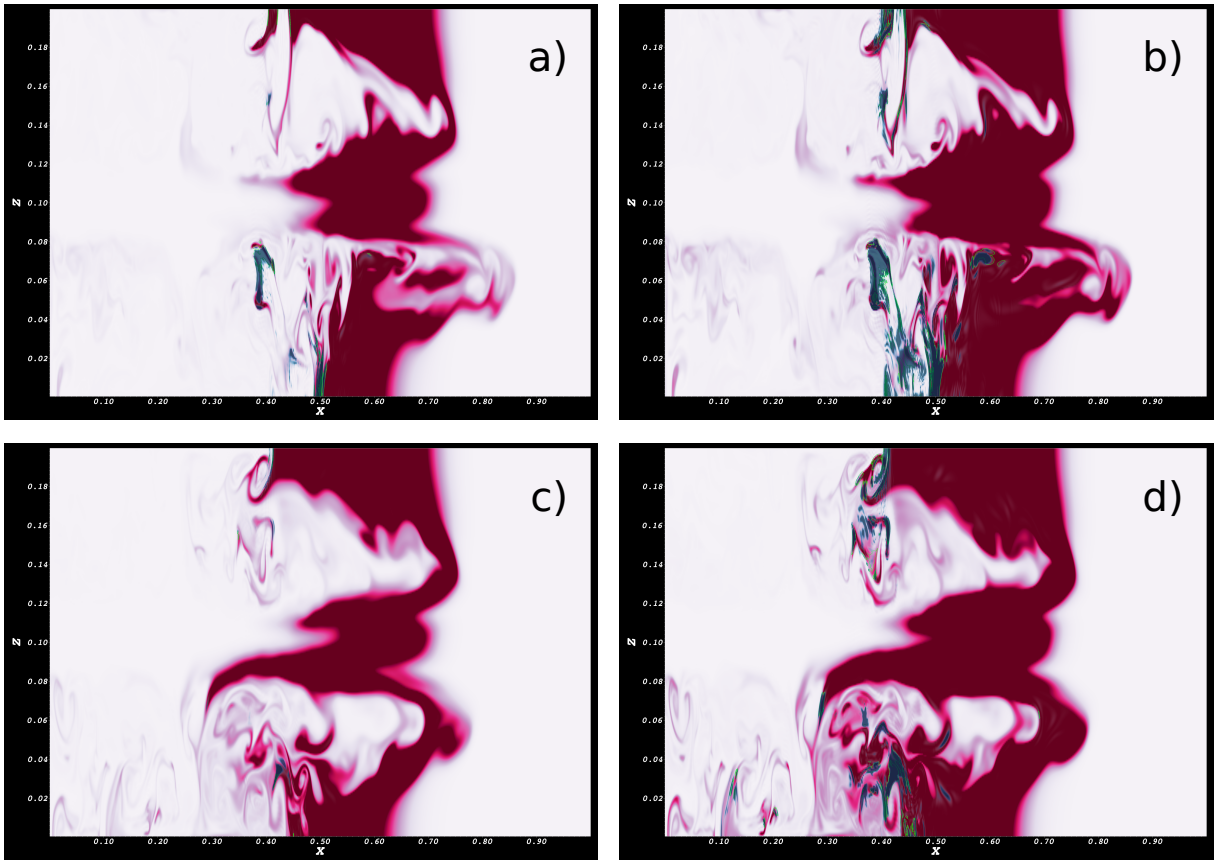

Figure 3.14: Slices of dye concentration shown in purple, with mixing overlain in blue, and contours of stirring in green. Panels (a) and (b) correspond to $y = 0.15$ m and (c) and (d) to $y = 0.05$ m. The panels are then horizontally organized by saturation with (a) and (c) corresponding to maximum values of 20000 for mixing and stirring, and 0.25 for dye concentration, while panels (b) and (d) correspond to maximum values of 2000 for mixing and stirring, and 0.1 for dye concentration. This change in saturation results in only the highest areas of mixing and stirring being shown in the first column. In all of the panels it is clear that the stirring and mixing occurs along the thin filaments of dye which are within the turbulent region.

Chapter 4

Conclusions

4.1 Linking numerical simulation to physical observations

While obviously greatly simplified, and not to scale, we can use some of the results from the numerical simulation and relate them to the full physical system. One key result is the development of a turbulent front from the forcing which propagates away from the initial disturbance entraining quiescent fluid. If we consider the direct forcing region to be an open-air cenote we would have these fronts forming within the water column and then propagating through the karst system. Assuming that the top mud layer and mangrove cover prevents the rainwater from directly permeating into the karst, on a regional scale we would expect the most mixing to occur within the centoes and there to be measurable fronts moving within the cave networks. Furthermore as the front moves away from the initial site there should be a measurable flow above and below the halocline entraining new fluid and mixing. A feature that is uniquely suited to analysis through the numerical simulation is the dye column and the mixing parameter. While possible (and has been done [21]) releasing dye or other tracers into this system can at best provide point wise measurements of the dye concentration and no dynamic information. Furthermore to observe the results during a heavy rainfall event the researcher would have to actively be in the field to release the dye and perform measurements as the event is occurring. This can be extremely dangerous in remote locations and during extreme weather. Quantifying mixing would be practically impossible in the field but is an important measure to understanding the dynamics. This type of information is extremely relevant in the event of a contaminant leak into the system during a rainfall event. Quantifying the location of maximum mixing and how the contaminated water interacts and spreads into the quiescent fluid can provide key information for managing such a situation.

4.2 General field data

The major result from the field data is the confirmation of the hypothesis that heavy rainfall results in halocline mixing. There is extremely high correlation between the known rainfall events and changes in temperature and salinity above the halocline. Within the data there is evidence of a sustained density inversion which is inherently unstable. Other than errors in the sensors, a possible explanation could be a series of rainfall events which prolong this period of instability. Using the wavelet and Fourier analysis it was possible to very accurately match up the first eight theoretical tidal components to corresponding spikes in the water level data. The wavelet analysis also found longer scale trends on the

order of ≈ 120 days in both the temperature and the water level data. Further investigation is required to verify this signal and find a physical cause. More recent wavelet analysis focusing on the large rainfall events has also found that these events locally suppress the tidal signal for the days following a large event.

An important use for this field data is relating these observations to the current geological research being done within the same cave networks. Preliminary trace element measurements from sediment cores gathered within the cave network (shown in Figure 4.1) shows an increase in the salinity of the groundwater during the classical Mayan period and a corresponding decrease in salinity at the end of the Classical period. Historical records of rainfall show that during the preclassical period of the Maya there was an abundance of rainfall which helped support their civilization [7], while their decline was characterized by a rapid series of droughts. We can interpret these historical results through the current results found from the sensors to argue that during these prolonged periods of rain the halocline was continuously mixed resulting in an increase in salinity of the groundwater. The opposite would have happened during prolonged drought with any density differences in the water having enough time to settle out.

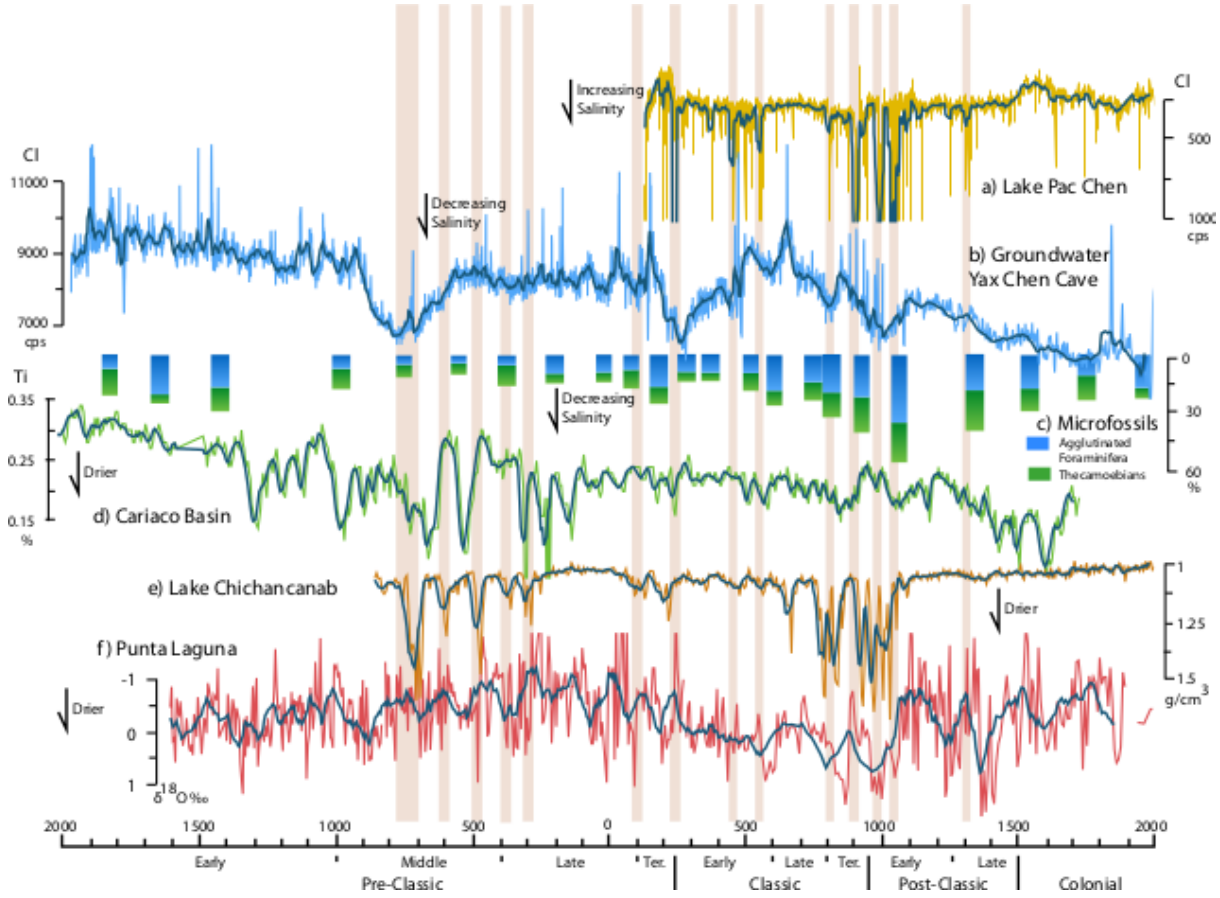

Figure 4.1: A figure courtesy of Dr. Reindhardt. A comparison of Chlorine trace element data from a sediment core from Yax Chen to other standard historical records. The chlorine is used as a proxy for the salinity of the groundwater and the shaded areas show a correspondence between the different records.

4.3 Future work

Currently we are extending the work presented in this thesis, both the numerics and the field measurements. During the last trip to the field site in May 2016, the HOBO sensors were repositioned to a higher location in the water column. All HOBO sensors were set to this constant depth (4m) forming a cross section inland through the cave network. This will allows us to observe spatial differences in mixing and observe if said mixing events extend to this level in the water column. We are also planning on placing a small

weather station on site which will provide much more accurate rainfall readings and allow for accurate atmospheric pressure readings to calibrate water level data. Current analysis of the numerical simulation includes a number of different parameters such as λ_2 (for identifying coherent vortices), Q-R (for identifying dominant turbulent structures) and viscous dissipation (for quantifying energy loss). Through these parameters we can extract more useful information from the current simulation and link our results to the classical literature.

Immediate future work on the numerical simulation will include more realistic parameters, such as jet strength, and the addition of temperature and salinity, combined through the UNESCO equation of state for density. Another avenue is to provide a more realistic forcing by moving the jets above the pycnocline. We will also apply classical theory of turbulent fronts (e.g., [33]) to our simulations, and investigate to what extent such theory is applicable. To achieve a simulation which more physically matches the field it will be necessary to move to a different model. Simplifications will have to be made to enable computation and would not provide the dynamic accuracy that the current model provides. However moving to a hydrogeological model would allow for an analogue of the entire physical system to be modelled and allow for direct comparison with sensor measurements. Future work in the field will include the addition of flow sensors to accurately measure and compare with simulations. In general an overall increase in the number of sensors will provide better resolution within the system and will enable a deeper understanding of the dynamics. A major area of future work will be the application of the time series analyses shown here to trace element data from sediment cores. In particular wavelet analysis will allow us to identify long term high-power climactic signals within the data. This can then be matched and correlated to known climate signals.

As mentioned in the introduction, this research is important to many of the companies developing the coastline. As their demand for water increases this will put a strain on the existing supply and particularly the underdeveloped infrastructure throughout the region. This research is also relevant for waste management as current common procedure is to pump waste below the halocline, however we have shown that during heavy rainfall events this under-layer can mix with the potable water above. A key component of future work should look to collaborate with some of these companies to allow for better management decisions to be made.

References

[1] A. P. H. Association. *Standard methods for the examination of water and wastewater 20th Edition*. American Public Health Association, Washington, DC, 1999.

[2] P. Bauer-Gottwein, B. R. N. Gondwe, G. Charvet, L. E. Marin, M. Rebolledo-Vieyra, and G. Merediz-Alonso. Review; The Yucatan Peninsula karst aquifer, Mexico. *Hydrogeology Journal*, 19(3):507–524, 2011.

[3] P. A. Beddows, P. L. Smart, F. F. Whitaker, and S. L. Smith. Decoupled fresh-saline groundwater circulation of a coastal carbonate aquifer: Spatial patterns of temperature and specific electrical conductivity. *Journal of Hydrology*, 346(1-2):18–32, 2007.

[4] J. P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover Publications, second edition, 2000.

[5] J. C. Chatters, D. J. Kennett, Y. Asmerom, B. M. Kemp, V. Polyak, A. N. Blank, P. a. Beddows, E. Reinhardt, J. Arroyo-Cabrales, D. a. Bolnick, R. S. Malhi, B. J. Culleton, P. L. Erreguerena, D. Rissolo, S. Morell-Hart, and T. W. Stafford. Late Pleistocene human skeleton and mtDNA link Paleoamericans and modern Native Americans. *Science*, 344(6185):750–4, 2014.

[6] H. Childs. VisIt: An End-User Tool for Visualizing and Analyzing Very Large Data, 2013.

[7] J. H. Curtis, D. A. Hodelle, and M. Brenner. Climate Variability on the Yucatan Peninsula (Mexico) during the Past 3500 Years , and Implications for Maya Cultural Evolution. *Quaternary Research*, 47(1):37–47, 1996.

[8] D. Gabor. Theory of communication. Part 1: The analysis of information. *Electrical Engineers-Part III: Radio and ...*, 93(26):429–441, 1946.

[9] A. E. Gill. Adjustment under gravity in a rotating channel. *Journal of Fluid Mechanics*, 77(03):603–621, apr 1976.

[10] A. H. González González, C. Rojas, A. Terrazas, M. Benavente Sanvicente, W. Stinnesbeck, J. Aviles, M. De los Ríos, and E. Acevez. The Arrival of Humans on the Yucatan Peninsula: Evidence from Submerged Caves in the State of Quintana Roo, Mexico. *Current Research in the Pleistocene*, 25:1–24, 2008.

[11] A. Grossmann and J. Morlet. Decomposition of Hardy functions into square integrable wavelets of constant shape. *SIAM journal on mathematical analysis*, 15(4):723–736, 1984.

[12] R. Guenther and J. Lee. *Partial differential equations of mathematical physics and integral equations*. Courier Corporation, 1988.

[13] R. Haberman. *Applied partial differential equations: with Fourier series and boundary value problems*. Pearson, fourth edition, 2004.

[14] G. H. Haug, D. Gunther, L. C. Peterson, D. M. Sigman, K. a. Hughen, and B. Aeschlimann. Climate and the collaspe of Maya Civilization. *Science*, 299(5613):1731, 2003.

[15] P. N. Kambesis and J. G. Coke. Overview of the Controls on Eogenetic Cave and Karst Development in Quintana Roo, Mexico. *Coastal Karst Landforms*, 5(2002):347–373, 2013.

[16] D. J. Kennett and Et.al. Development and Disintegration of Maya Political Systems in Response to Climate Change. *Science*, 338(6108):788–791, 2012.

[17] P. Kumar and E. Foufoula-Georgiou. Wavelet analysis in geophysics: An introduction. In *Wavelets in geophysics*, chapter Wavelet An, pages 1–43. Academic Press, 1994.

[18] P. K. Kundu, I. M. Cohen, and D. R. Dowling. *Fluid Mechanics*. Academic Press, fifth edition, 2011.

[19] L. J. Lucero, J. D. Gunn, and V. L. Scarborough. Climate Change and Classic Maya Water Management. *Water*, 3(2):479–494, 2011.

[20] S. Mallat. *A wavelet tour of signal processing*. Academic Press, London, second edition, 1999.

[21] Y. H. Moore, R. Stoessel, and D. Easley. Fresh water/sea-water relationship within a groundwater system, northeastern coast of the Yucatán Peninsula. *Ground Wat.*, 30(3):343–350, 1992.

[22] D. A. Nield and A. Bejan. *Convection in porous media*, volume 24. Springer, third edition, 2013.

[23] W. H. Organization. Guidelines for drinking-water quality. Technical report, World Health Organization, 2011.

[24] E. Perry, L. Marin, J. McClain, and G. Velazquez. Ring of Cenotes (sinkholes), northwest Yucatan, Mexico: its hydrogeologic characteristics and possible association with the Chicxulub impact crater, 1995.

[25] R. Salmon. Lectures on geophysical fluid dynamics. *Oxford University Press*, 1998.

[26] V. L. Scarborough. Ecology and Ritual: Water Management and the Maya, 1998.

[27] P. Schulte, L. Alegret, I. Arenillas, J. A. Arz, P. J. Barton, P. R. Bown, T. J. Bralower, G. L. Christeson, P. Claeys, C. S. Cockell, G. S. Collins, A. Deutsch, T. J. Goldin, K. Goto, J. M. Grajales-Nishimura, R. A. F. Grieve, S. P. S. Gulick, K. R. Johnson, W. Kiessling, C. Koeberl, D. A. Kring, K. G. MacLeod, T. Matsui, J. Melosh, A. Montanari, J. V. Morgan, C. R. Neal, D. J. Nichols, R. D. Norris, E. Pierazzo, G. Ravizza, M. Rebolledo-Vieyra, W. U. Reimold, E. Robin, T. Salge, R. P. Speijer, A. R. Sweet, J. Urrutia-Fucugauchi, V. Vajda, M. T. Whalen, and P. S. Willumsen. The Chicxulub Asteroid Impact and Mass Extinction at the Cretaceous-Paleogene Boundary. *Science*, 327(5970):1214–1218, 2010.

[28] P. L. Smart, P. A. Beddows, J. Coke, S. Doerr, and F. F. Whitaker. Cave Development on the Caribbean coast of the Yucatan Peninsula, Quintana Roo, Mexico. *Geological Society of America*, 2404(10):105–128, 2006.

[29] E. A. Spiegel and G. Veronis. On the Boussinesq Approximation for a Compressible Fluid. *The Astrophysical Journal*, 131:442, 1960.

[30] R. K. Stoessell. Dampening of Transverse Dispersion in the Halocline in Karst Limestone in the Northeastern Yucatan Peninsula. *Ground Water*, 33(3):366–371, 1995.

[31] C. J. Subich, K. G. Lamb, and M. Stastna. Simulation of the Navier-Stokes equations in three dimensions with a spectral collocation method. *International Journal for Numerical Methods in Fluids*, 73(2):103–129, sep 2013.

[32] L. N. Trefethen. *Spectral Methods in Matlab*, volume 10. Siam, 2000.

[33] J. S. Turner. *Buoyancy effects in fluids*. Cambridge University Press, 1973.

[34] I. UNESCO and I. SCOR. Tenth report of the joint panel on oceanographic tables and standards. *UNESCO Technical Papers in Marine Science*, 36:24, 1981.

[35] G. K. Vallis. *Atmospheric and Oceanic Fluid Dynamics*. Cambridge University Press, 2006.

[36] G. Veni. Maya utilization of karst groundwater resources. *Environmental Geology and Water Sciences*, 16(1):63–66, 1990.

[37] A. E. Weidie. Part I: Geology of Yucatan Platform. In *Geology and hydrogeology of the Yucatan and Quaternary geology of northeastern Yucatan Peninsula.*, pages 1–19. New Orleans Geological Society, 1985.