Title
Dispersion and Tidal Dynamics of Channel-Shoal Estuaries

Author
Holleman, Christopher Dean

Publication Date
2013

Peer reviewed|Thesis/dissertation

Dispersion and Tidal Dynamics of Channel-Shoal Estuaries

by

Christopher Dean Holleman

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering – Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Mark T. Stacey, Chair
Professor Fotini K. Chow
Professor Philip S. Marcus

Fall 2013

Dispersion and Tidal Dynamics of Channel-Shoal Estuaries

Abstract

Dispersion and Tidal Dynamics of Channel-Shoal Estuaries

by

Christopher Dean Holleman

Doctor of Philosophy in Engineering – Civil and Environmental Engineering

University of California, Berkeley

Professor Mark T. Stacey, Chair

Estuaries, and the varied ecosystems they support, are affected by human action in many ways. One of the fundamental environmental questions pertaining to estuaries is how material is mixed into ambient waters and transported within and beyond the estuary. Characterizing transport and mixing is essential for tackling local environmental questions, such as gauging the water-quality impacts of a wastewater outfall, or how dredging may alter estuarine circulation. In many estuaries the dominant physical dispersion mechanism is shear dispersion. A fundamental and idealized approach to the scalar transport problem is undertaken, in which canonical shear dispersion regimes are augmented with a new, intermediate regime. This mode of shear dispersion, in which the longitudinal variance of the plume increase with the square of time, is likely to occur in channel-shoal basins where bathymetric transitions can create sheared velocity profiles. A key tool in the engineer's arsenal for pursuing studies of scalar transport is the numerical model, but care must be taken to separate physical mixing effects from numerical artifacts. Unstructured models are particularly well-suited to estuarine problems as the grid can be adapted to complex local topography, but the numerical errors of these models can be difficult to characterize with standard methods. An in-depth error analysis reveals a strong dependence of numerical diffusion on the orientation of the grid relative to the flow. Flexible grid generation methods allow for optimizing the grid in light of this dependence, and can decrease across-flow numerical diffusion by a factor of two.

The most far-reaching impact of our actions, though, must certainly be global climate change. A period of rising sea level is being ushered in by climate change, and the final research-oriented chapter seeks to further our understanding of how rising sea levels will affect basin-scale tidal dynamics. Numerical and analytic approaches show that as basins get deeper tidal amplification becomes more effective and tidal range increases. The crux of the analysis, though, is the inclusion of inundation within these scenarios. Inundated areas dissipate incident tidal energy, countering the added amplification due to basins becoming deeper. The net effect, in the case of San Francisco Bay, is that tidal amplification under

sea level rise actually decreases, such that a particular rise in mean higher high water in the coastal ocean is predicted to raise mean higher high water within the bay by a lesser amount.

To my parents, my brothers, my friends.

Contents

List of Figures

List of Tables

Acknowledgments

I am indebted to Mark Stacey for his ever-positive outlook on the course of my research. The denizens of 202 and 205 O'Brien have been invaluable. Many thanks also to Oliver Fringer and members of the Stanford Environmental Fluid Mechanics Laboratory. This work has been financially supported by the California Coastal Conservancy and the National Science Foundation. Computational resources were provided by NSF XSEDE, Berkeley EFM, and Stanford EFML.

Chapter 1

Introduction

Estuaries are the bodies of water where tidally drive, saline ocean waters meet freshwater rivers. The interaction of tides and river flow and the mixing of freshwater and saltwater make estuaries ecologically important and dynamically complex. The tidal, brackish habitats offered by estuaries are unique and important for a wide range of species; the tidal marshlands such as those found in the margins of San Francisco Bay support shorebirds, waterfowl, young salmon and steelhead trout, harbor seals and a host of other birds, mammals and fish. Some species live their entire life cycle in estuarine waters, while for other species the unique habitat supports a specific life phase. For some animal species estuaries offer a shallow, protected environment for rearing their young. For migratory birds, flyways often coincide with large estuaries and these areas serve as seasonal stopping points during migration.

For the human population estuaries provide protection in the form of calm waters where ocean-going trade vessels can find refuge from wind and waves. In addition to shelter from ocean conditions, estuaries are also, by definition, the junction of rivers and the ocean and thus the gateway to river-based commerce. It is no surprise that historically estuaries around the world have been a popular backdrop for settlement, development and increasingly large population centers on their shores. The hidden caveat, though, is that this development comes with many new problems, both in terms of how we have an impact on the estuary and how we come to depend on a tame and predictable waterway.

With concerns for both the continued ecological function within estuaries and the support and protection of population centers adjacent to estuaries, environmental engineers are tasked with quantifying the key physical processes in these water bodies. Not only must we characterize the natural state of these systems, but also how human actions, past, present and future, affect these water bodies. The effects are sometimes intentional, dredging navigational channels or leveeing shorelines for flood protection, but often as not there are unintentional side effects. Construction of causeways decreases exchange, allowing pollutant concentrations to peak. In portions of San Francisco Bay upstream mining operations introduced mercury which has accumulated in bed sediments. The devastation of oil spills speaks for itself.

A broad category of environmental issues surrounding estuaries concerns the introduc-

tion and dispersal of pollutants. With both large population centers and agricultural regions often located near estuaries, there are numerous situations in which a contaminant is introduced into the waterways and we must determine the spatial and temporal extents of its effects. One example is agricultural lands where water runoff tends to have unnaturally high nutrient content. Such nutrient loading can lead to algal blooms, which in turn deplete the water of oxygen, threatening fish. Similarly, in urban settings treated wastewater is often discharged into adjacent waters, and though the wastewater treatment process nearly purifies the wastewater, there is some non-negligible, residual nutrient and contaminant loading. In all of these cases, the size of the affected water mass and its ultimate fate are important for understanding ecological implications and making related management decisions. Quantifying the mixing and dispersal of these sources allows informed decisions ranging from where to site a wastewater treatment outfall to selecting the most appropriate treatment methods.

The available methods for characterizing transport and mixing processes in the estuarine context can be roughly categorized into analytic descriptions, lab- or field-based methods, and numerical modeling. Analytic descriptions offer the purity of a mathematical approach which distills the problem down to the barest necessities and establishes functional relationships between the parameters of the problem and the resulting transport and mixing. While real world situations introduce many complexities which are generally beyond the scope of an analytic approach, purely analytic descriptions are highly useful for quick or preliminary analyses, and as a tool to inform decisions on how to proceed with a more in-depth analysis. For example, in the case of a pollutant plume introduced into a tidally oscillating flow, an analytic solution could offer an estimate of how quickly the plume will grow based on a a minimal description of the basin. Such a description may then help rule out or highlight particular concerns, such as whether a fragile ecosystem some distance away is at all likely to be affected the plume. Analytic models also inform how more in-depth observational or model-based techniques should be applied. Questions such as how big the model domain should be, where measurements should be taken and what time scales are relevant can be answered by leveraging the simplified mathematical description of the problem. These solutions also provide hints for how a particular process might be parameterized when included in a larger-scale analysis.

One of the most efficient large-scale mixing processes in rivers and estuaries is shear dispersion, also known as Taylor dispersion due to the original description in [74]. A vertical or cross-flow velocity shear in the along-stream direction, coupled with a lateral mixing process such as small-scale turbulent mixing, are shown to lead to a highly amplified longitudinal mixing rate. The typical shear dispersion analysis focuses on the long term evolution of a scalar field, i.e. over time scales longer than the lateral mixing time. This is particularly appropriate for flows such as pipe flow (the application of [74]), and river flows where bank-to-bank width is small compared to the length of a reach. Estuarine basins, however, are often much larger than the width of the in-flowing river, and basins with a planform aspect ratio in the range of 1–10 are common. In light of this geometry, significant evolution of a scalar plume occurs before the assumptions of lateral mixing are met. This has led to the analysis of transient dispersion regimes, the focus of chapter 2. This chapter explores the

evolution of a plume as it transitions from a spatial scale smaller than the length scale of velocity variations (e.g. the width of a shear layer separating a channel and a shoal), through an intermediate scale for which lateral boundaries have not begun to limit the lateral expansion of the plume, and finally the limit of Taylor dispersion with complete lateral mixing. The first and last periods have been well-studied, and a plume in these cases can be shown to exhibit growth in its along-channel variance of $\sigma_x^2 \sim t^3$ and $\sigma_x^2 \sim t$, respectively. The intermediate regime is introduced and is shown to fill in the gap between the two previously described regimes, having a variance evolution proportional to t^2. In terms of the time rate of change of the variance, the linear case has a constant rate, $d\sigma_x^2/dt \sim 2K$ (with dispersion coefficient K). For the quadratic and cubic regimes, though, there are still terms of t in the expression for the time rate of change. These factors indicate the presence of additional state in the plume, and this chapter explores the origin and quantification of this additional state.

When the details of the problem are known and there is a need for quantitative information, numerical models are often the best tool. Analytic approaches and scaling arguments help enumerate which physical processes are relevant, but as the number of processes involved grows, interactions between competing processes can be difficult to characterize through the simplifying lens of scaling laws and analytic relationships. It is at this point that one must turn to numerical models. Such models, though, come with their own limitations, and estuarine hydrodynamic models are not yet sophisticated enough (and may never be) to be blindly applied to an arbitrary problem. Rather, the user must take care in selecting a model which includes the relevant physics and resolves the range of scales expected to be important. This decision is often complicated by the nuanced line between numerical artifacts and resolved physical physical processes. In the context of scalar transport and mixing, the distinction can be particularly difficult, and effective application of numerical model includes quantifying and taking steps to reduce numerical errors.

High-order numerical schemes with minimal numerical mixing are widely available in models based on rectilinear and curvilinear grids, where larger stencils are straightforward and well-behaved. The simple geometry of most river reaches and the broad open space of coastal ocean domains lend themselves to structured grids. When attempting to apply these models in estuarine domains, though, one must confront the complicated geometry of many estuaries, which is difficult to shoehorn into a structured grid. Irregular shorelines and the desire to resolve different parts of the domain with different spatial resolutions motivate the use of unstructured grid models in these situations. Unfortunately, the high-order numerical schemes available on structured grids are difficult to translate to unstructured grids. However, unstructured grids do offer some advantage in allowing the grid orientation to be locally adapted to the flow. Chapter 3 develops an estimate for the numerical diffusion on triangular grids, similar to grids used in several unstructured hydrodynamic models. The error estimate, in the form of a numerical diffusion tensor, depends on the alignment of the flow relative to faces of the grid and quantifies the benefits of designing a flow-aligned grid. Grid alignment approaches often rely on manual specification of the alignment direction, but in complicated geometries this is time-consuming and prone to errors. An automated algorithm

for generating a flow-aligned grid is presented, based on a new, optimizing, orthogonal grid generator. Dominant flow directions are extracted from an unaligned, reference simulation of the same domain, then used to govern the placement of triangular cells during generation of the aligned grid. Along-flow diffusion is only mildly decreased in the aligned grid, but cross-flow diffusion is very sensitive to grid alignment. In the case of a perfectly regular triangular grid the cross-flow diffusion can be eliminated entirely, and reduced by 50% in a realistic case.

Chapters 2 and 3 focus on scalar mixing, tracking the composition of the water. The applications of scalar transport most commonly cover the fate of materials introduced into the estuary, originating from the shorelines or upstream rivers. But the other defining feature of estuaries is the direct connection to ocean waters, and in thinking broadly of anthropogenic effects on an estuary one eventually must consider changes in the ocean forcing. This, of course, includes one of the most significant anthropogenic changes in our environment: climate change. From the perspective of estuarine dynamics, the ocean conditions are taken as a given, and from the climate modeling community we know that higher sea levels and a shifting distribution of extreme weather events are possible. Understanding how these coastal conditions correlate to conditions inside the estuary is the subject of chapter 4. At first glance it may seem sufficient to examine present day high water levels and extrapolate that a specific rise in coastal high water elevation would lead to the same rise in high water within an estuary. This over-simplified view is problematic, though, as even under present day conditions there are significant variations in mean sea level and tidal range throughout a typical estuary. Frictional damping of the tidal wave decreases amplitudes as one moves farther from the ocean. At the same time, many basins have a converging geometry, getting narrower and shallower as one moves upstream. This convergent geometry has an effect opposite to the frictional damping, leading to amplification of the tidal range in the upstream reaches of an estuary. The nonlinear aspects of these mechanisms mean that a simple increase in the tidal range or mean sea level at the ocean boundary will not necessarily translate to the same increase of tidal range of mean sea level within the basin. Clearly one must account for the basin geometry and the role of friction in accurately predicting a basin's response to a change in the tidal forcing.

Even with a physically-based model of a basin with reasonable skill at predicting present-day conditions, sea level rise scenarios still present a significant challenge. Areas which are currently above the reach of tidal action will become tidal in the future, and their contribution to the wetted area of the basin should be considered. Furthermore, tidal dynamics and inundation are inherently coupled. Any increase or decrease of the tidal range caused by newly inundated areas will in turn increase or decrease the inundated area. Newly inundated areas are expected primarily to attenuate the tidal range, as these areas will be shallow and friction-dominated. At the same time, sea level rise will cause the interior of the basin to become deeper, lessening the effects of bed friction. In most basin geometries deepening leads to an increase in the tidal range. These competing effects, amplification in response to deepening and attenuation in response to inundation, further motivate the need for numerical modeling to characterize the effects of sea level rise.

Capturing the extent of inundation comes with its own challenges, too. In some cases the land area susceptible to flooding is presently broad marsh lands which will presumably transition to mud flats as sea level rises. In many cases, though, these low-lying areas have already been extensively alterated by human action, with levee systems and dredged channels partitioning the landscape into tidal and non-tidal areas. These complex landscapes have hydrodynamically relevant features at spatial scales down to tens of meters, where a single narrow levee may segregate expansive tracts of low-lying marsh from the action of the tides. Special care must be taken to include as much information as possible about connectivity between basins, ponds, sloughs and rivers. Chapter 4 introduces a method for preserving a significant degree of hydrodynamic connectivity even when the physical features are smaller than the nominal grid resolution. A static inundation analysis of the grid with and without this additional process demonstrates the efficacy of the method, where the original grid shows extensive flooding under present-day conditions while the connectivity-preserving method more accurately reflects the known present-day conditions.

A series of tidal simulations are used to explore the coupling between sea level rise, tidal amplification and inundation. The study area comprises two basins, with predominantly M2 tides which are near a standing wave in one basin, and a partially progressive wave in the other basin. Simulating a wide range of scenarios combining various shoreline configurations and sea level rise conditions, the analyses explore the tidal response and role of inundation. For the hypsometry of San Francisco Bay, the results show that the additional amplification from a deeper basin is more than offset by the damping in the inundated areas. The response of a specific basin depends not only on the details of the forcing but also the tidal wave characteristics in that basin. The spatially varying characteristics of the tidal wave in each basin are analyzed in terms of energy flux and change in tidal phasing. Energy flux analysis shows a general shift towards standing wave tides under conditions of sea level rise without inundation. When inundation is allowed, energy fluxes show increased dissipation of the incident M2 wave on perimeter sloughs and rivers. Modulation of the M2, though, accounts for only a fraction of the change in high water inside the bay. Analysis of M2 overtides show complex and significant differences in overtide generation between the scenarios, and these differences make up the bulk of the remaining differences in high water across scenarios. Though the specific quantitative details of the sea level rise simulations depend on the details of San Francisco Bay, the broader conclusion is that the dynamics of inundation are an important factor when studying sea level rise within enclosed basins.

Estuaries will always be important ecologically due to their unique habitats. Though we have made great progress in reducing pollution from population centers adjacent to most estuaries, there are still significant present day sources of concern as well as the lingering presence of past contaminants. Sea level rise, the continued expansion of developed areas around estuaries, and an increasing willingness to engage in restoration projects all point to the continued importance of understanding how our actions have an impact on these water bodies.

Chapter 2

Transient Shear Dispersion Regimes

Abstract

Characterizing scalar dispersion is a key concern in a wide variety of applications, including both steady-state and time-dependent studies of wastewater outfalls, salinity distribution in estuaries, and the spreading of pollutants from industrial spills. As the size of a scalar plume grows with respect to the size of the containing water body, the effective dispersion varies, from the well-known $\sigma_x^2 \sim t^3$ behavior for a plume smaller than the lateral extent of the shear, to the $\sigma_x^2 \sim t$ behavior at the limit of a laterally well-mixed plume. An additional regime is introduced in which the plume extends across the full range of the available shear, but is not significantly affected by the lateral bounds of the water body. An analytical treatment shows that this regime exhibits a $\sigma_x^2 \sim t^2$ behavior, independent of lateral mixing coefficient. Particle tracking results in an idealized, tidal channel-shoal basin demonstrate this regime as particle clouds straddle the channel-shoal interface. Quantitative analysis of spatial moments as plumes transition between regimes show good correlation between the observed parameters and parameters predicted by the analytical framework.

2.1 Introduction

Scalar dispersion, in the most general sense, describes any of a numerous set of physical processes by which features of the flow lead to the spreading out of a scalar species. When considered in Fickian terms the effective dispersion rate is often much greater than the molecular diffusion rate. This rate of spreading is often of fundamental interest in a wide range of applications. In environmental flows, questions related to spatial distributions of nutrient abundance, the spreading of algal blooms, larval transport, and the zone of influence

Also in press as: Holleman, R. C. and M. T. Stacey. Transient Dispersion Regimes. *J. Fluid Mech.*, 2013.

of engineered sources such as wastewater treatment outfalls all trace back to questions of scalar dispersion. In ocean flows dispersion rates are essential for understanding the mixing effects of eddies, formation and break down of patches of heightened biological activity, and the mixing action of passing and breaking internal waves [85].

Of particular relevance to riverine and estuarine environments is the case of linear shear dispersion, also known a Taylor dispersion due to the first mathematical analysis of [74]. Given the primary application to estuarine and riverine systems, the term longitudinal is used to refer to the along-flow axis of the domain (x), and the term lateral to refer to the perpendicular horizontal axis (y). Taylor demonstrated that lateral gradients in the longitudinal velocity interacting with lateral diffusion leads to Fickian longitudinal dispersion, albeit with a highly amplified dispersion coefficient. Though the original analysis was targeted at laminar pipe flow, subsequent work in [75] and [2] expanded the analysis to cover a broad range of domains under both laminar and turbulent conditions, summarized in [23]. Central to Taylor's analysis and many subsequent works is the description of dispersion in terms of the growth of the second central moment of the tracer distribution, σ_x^2, equivalently labeled the spatial variance. In the common Fickian model in which scalar flux is proportional to the concentration gradient, σ_x^2 grows with a linear dependence in time $\sigma_x^2 \approx 2K_x t$, with longitudinal dispersion coefficient K_x.

Scalar dispersion in environmental flows may result from the combined action of multiple dispersive processes, such as turbulent mixing, flow asymmetries giving rise to tidal pumping [69], tidal trapping [53], and various interactions of velocity shear in one direction with dispersion in an orthogonal direction. The mixing due to these often complex dispersive processes is at least superficially similar to the homogenizing effects of Fickian diffusion, and Fickian models are often used to describe the bulk mixing effects. However, the underlying processes do not necessarily obey a Fickian model [22], and at short time scales plume spreading may significantly depart from the Fickian description. These departures from the Fickian model are often termed *anomalous diffusion*, and when considered in terms of a power law relating tracer variance and time, $\sigma_x^2 \sim t^\beta$; the term *subdiffusion* indicates $\beta < 1$ and *superdiffusion* indicates $\beta > 1$ [84].

[2] presented a general framework for evaluating shear dispersion, showing that lateral diffusion coupled with lateral shear in the longitudinal velocity leads to longitudinal dispersion with a Fickian behavior at asymptotically long time scales. [61] extended the analysis of Aris to the case of a plume released at ground level into a semi-infinite shear layer above the plane of the ground, showing that in such a case the variance grows with a cubic dependence on time compared to the linear dependence of Taylor's and Aris' scenarios. [3] refined the analysis of Aris, clarifying the mathematical assumptions of the analysis, and for a broad subset of flows presented a method for distilling the PDEs into a simpler set of ODEs. Solving the more tractable set of ODEs leads to analytical solutions, valid at all time scales, for the moments of a scalar cloud. In introducing considerable mathematical rigor, the approaches of Barton become difficult to apply to environmental flows. For example, while scaling analyses may easily reach the results of [61], a semi-infinite flow is incompatible with the approach of [3], such that one could not arrive at a relationship directly describing the cubic time

dependence in unbounded, constant shear cases. Scaling laws and corresponding qualitative regimes are quite helpful and broadly applied in understanding environmental flows where the precise global flow field may not be known, which in turn motivates the more schematic methods described below. Other recent dispersion mechanism studies include [68], considering shear dispersion in a laterally bounded flow but with a non-zero Lagrangian time scale, demonstrating that a non-zero Lagrangian time scale can increase the dispersion coefficient.

The present work is concerned with three regimes of shear dispersion in environmental flows and the transitions there between. The initial shear dispersion regime is essentially that described by [61], among others, in which a constant shear velocity profile completely envelopes the scalar plume. As the plume expands laterally the range of velocities sampled by the parcels within the plume also increases, leading to an ever-increasing velocity scale associated with the longitudinal dispersion. The straining of the plume by this ever-increasing velocity scale has been shown to result in cubic in time variance growth, $\sigma_x^2 \sim t^3$, and will generally be referred to as the "cubic" regime.

A second regime, introduced more fully in section 2.2, describes the variance growth of a plume which is not bounded in the lateral direction, but for which the range of velocities sampled by its parcels is essentially constant in time. Schematically this could correspond to a narrow (relative to the lateral extent of the plume) shear layer separating two regions of smaller velocity gradients. Physically, a plume may experience such a flow when the discrete spatial features leading to velocity gradients are large compared to plume extent. One particular example is the geometry of a channel-shoal basin, in which shear is maximal over the sloping region joining the channel and shoal, with comparatively little shear in the channel or shoals. A simple scaling approach to this scenario starts with a plume centered on a shear layer with velocity $+U$ to one side of the shear layer and $-U$ to the other side. The centroid of the plume is stationary and the longitudinal variance a function of the typical displacement Δx of a parcel. While on a particular side of the shear layer, $\Delta x \sim \pm Ut$ and $\sigma_x^2 \sim U^2t^2$. The variance growth is then $d\sigma_x^2/dt \sim U^2t^2/T$, where T is the characteristic time for a parcel to return to the shear layer. T scales with the lateral size of the plume as $T \sim \sigma_y^2/K_y$, and in the absence of lateral boundaries $\sigma_y^2 \sim K_yt$, leading to an overall scaling $d\sigma_x^2/dt \sim U^2t$ or $\sigma_x^2 \sim U^2t^2$. The application of the Aris method of moments to this flow in section 2.2 supports this scaling and provides the exact coefficients of proportionality.

In the case of a plume subject to shear but also constrained by lateral boundaries, the variance growth asymptotes to the linear shear dispersion limit $\sigma_x^2 \sim t$ after a time sufficient for the plume to become laterally well-mixed. Only in this final regime does the plume assume a Fickian behavior.

Section 2.3 moves to the inverse problem of taking a time series of spatial moments describing the evolution of a plume and estimating the relative importance and parameters of each dispersion regime. Following a description of the inverse problem, idealized flows and plume releases, described in section 2.4, serve as a testbed for the inverse problem.

2.2 Analytical Derivation

Following the approach of [2], the advection-diffusion equation can be transformed into an evolution equation for the spatial moments of the scalar distribution. The derivations of this section assume a rectangular channel aligned with the x axis, with lateral coordinate $y \in [-B/2, B/2]$ and vertical coordinate $z \in [-H, 0]$. Define pth moment in x for a filament extending along a given line of constant y and z as $c_p(y, z, t)$:

$$c_p(y, z, t) \equiv \int_{-\infty}^{\infty} x^p C(x, y, z, t)\, dx. \tag{2.1}$$

The pth moment of the overall scalar cloud, m_p, is then defined as the integral of c_p over the cross-section of the channel

$$m_p(t) \equiv \int_{-B/2}^{B/2} \left[\int_{-H}^{0} c_p(y, z, t)\, dz \right] dy. \tag{2.2}$$

Assuming a constant but anisotropic diffusivity described completely by the axis-aligned diffusivities K_x, K_y and K_z, lateral and vertical velocities equal to zero, and along-channel velocity $u = u(y, z)$, the advection-diffusion equation can be written

$$\frac{\partial C}{\partial t} + u\frac{\partial C}{\partial x} = K_x \frac{\partial^2 C}{\partial x^2} + K_y \frac{\partial^2 C}{\partial y^2} + K_z \frac{\partial^2 C}{\partial z^2}. \tag{2.3}$$

(2.3) is then multiplied by x^p and integrated over $x \in (-\infty, \infty)$ to get

$$\int_{-\infty}^{\infty} \left(x^p \frac{\partial C}{\partial t} + u x^p \frac{\partial C}{\partial x} \right) dx = \int_{-\infty}^{\infty} \left(K_x x^p \frac{\partial^2 C}{\partial x^2} + K_y x^p \frac{\partial^2 C}{\partial y^2} + K_z x^p \frac{\partial^2 C}{\partial z^2} \right) dx. \tag{2.4}$$

Pulling the time derivative outside the spatial integral and rearranging the advection term leads to

$$\frac{\partial}{\partial t}\left[\int_{-\infty}^{\infty} x^p C dx \right] + u \int_{-\infty}^{\infty} \left[\frac{\partial}{\partial x}\left(x^p C\right) - p x^{p-1} C \right] dx = K_x \int_{-\infty}^{\infty} x^p \frac{\partial^2 C}{\partial x^2} dx + K_y \int_{-\infty}^{\infty} x^p \frac{\partial^2 C}{\partial y^2} dx + K_z \int_{-\infty}^{\infty} x^p \frac{\partial^2 C}{\partial z^2} dx. \tag{2.5}$$

Applying the boundary conditions that C and its derivatives approach zero as $x \to \pm\infty$, (2.5) can be further simplified and expressed in terms of c_p, c_{p-1} and c_{p-2} as

$$\frac{\partial c_p}{\partial t} - u p c_{p-1} = K_x p(p-1) c_{p-2} + K_y \frac{\partial^2}{\partial y^2} c_p + K_z \frac{\partial^2}{\partial z^2} c_p. \tag{2.6}$$

Define an initial condition for the scalar distribution $C(\mathbf{x}, t = 0) \equiv C_0(\mathbf{x})$ from which $c_{p,0}(y, z)$ can be directly evaluated, and Neumann boundary conditions of zero flux at the y

and z boundaries: $K_{ij}\frac{\partial C}{\partial x_i}\cdot\hat{n}\Big|_{\delta\Omega}=0$, where $\delta\Omega$ denotes the boundaries of the domain, and $\hat{n}$ the unit vector normal to the boundary. Over the cross-section, the lateral and vertical diffusion integrate to zero, such that substituting the definition of the moment m_p, one arrives at an evolution equation for arbitrary moments,

$$\frac{dm_p}{dt}=p\overline{uc_{p-1}}+K_x p\left(p-1\right)m_{p-2}., \tag{2.7}$$

where the overbar denotes integration over the cross-section. As a final step, if the coordinate reference frame is chosen such that $m_1=0$ for all time, i.e. the reference frame of the centroid of the distribution, then the longitudinal spatial variance of the plume may then be described by

$$\frac{d\sigma_x^2}{dt}=\frac{1}{m_0}\frac{dm_2}{dt}. \tag{2.8}$$

Linear and cubic time dependence

[2] developed the above method and applied it to the case of a channel with an arbitrary lateral velocity profile $u=u\left(y\right)$. Taking the simplest shear velocity profile, a linear profile with constant shear $S=\partial u/\partial y$, $u=Sy$, with lateral boundaries at $y=\pm W/2$ (illustrated in panel c of figure 2.1), the spatial variance is described by:

$$\frac{d\sigma_x^2}{dt}=\frac{W^4S^2}{60K_y}+2K_x \tag{2.9}$$

where σ_x^2 has a linear time dependence - the familiar shear dispersion. If the lateral boundaries are absent but the constant shear velocity profile is retained (panel a of figure 2.1), the evolving spatial variance is instead described by

$$\frac{d\sigma_x^2}{dt}=2K_yS^2t^2+2K_x \tag{2.10}$$

such that σ_x^2 has a cubic time dependence [61].

Quadratic time dependence

In the linear case described above both the velocity range and the lateral domain are bounded, leading to a linear growth rate. The cubic case has an unbounded velocity range as well as a lack of lateral boundaries, and the method of moments predicts a cubic time dependence. At a qualitative level, a scenario in which exactly one of the velocity range or lateral boundaries are finite would be expected to yield a quadratic time dependence. Construction of a finite width channel with an infinite range of velocities is distinctly non-physical, but an unbounded channel with finite shear is well-posed and potentially physical, such as the limiting case of a plume which is large relative to local variation in the velocity field but

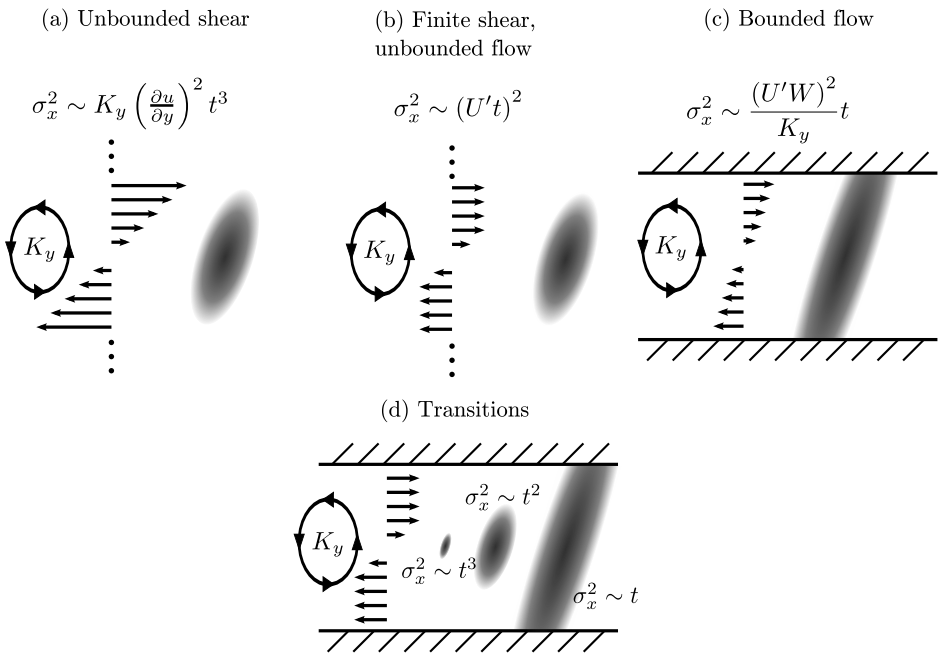

Figure 2.1: Schematic of the three dispersion regimes

small relative to the width of the channel. This section applies the method of moments to the scenario of an infinitely wide domain with a finite velocity range. In order to avoid introducing an additional length scale into the model corresponding to the thickness of the shear layer, an infinitely thin shear layer is imposed, with velocity field $u(y,z) = U\mathrm{sgn}(y)$, where U is a constant velocity scale.

The zeroth per-filament moment c_0 is purely a function of lateral diffusion and takes the form of a simple Gaussian distribution:

$$c_0(y,t) = \frac{m_0}{H\sqrt{4\pi K_y t}} \exp\left(-\frac{y^2}{4K_y t}\right). \tag{2.11}$$

Assuming that the plume is vertically well-mixed, solving for the per-filament first moment requires a solution for c_1 satisfying the equation

$$\frac{\partial c_1}{\partial t} = U\mathrm{sgn}(y)\, c_0 + K_y \frac{\partial^2 c_1}{\partial y^2}. \tag{2.12}$$

The form $c_1 = f c_0$ is assumed, where $f = f(y,t)$ gives the longitudinal per-filament centroid location as a function of time and lateral coordinate. Substituting $f c_0$ into the evolution equation for c_1 produces

$$\frac{\partial f}{\partial t} = u + K_y \frac{\partial^2 f}{\partial y^2} - \frac{y}{t}\frac{\partial f}{\partial y}. \tag{2.13}$$

This can be further simplified by non-dimensionalizing f by Ut, and assuming a similarity variable $\eta = y/\sqrt{4K_y t}$, such that $f/Ut = g(\eta)$ and the PDE is reduced to a second order nonlinear ODE in η:

$$g + \frac{\eta}{2}g' - \frac{g''}{4} - \operatorname{sgn}(\eta) = 0. \tag{2.14}$$

By symmetry $g(y = 0, t) = 0$. As y approaches $\pm\infty$ the motion of the centroid is independent of the shear at $y = 0$ and carried only by the uniform flow, giving boundary conditions for f:

$$\lim_{y\to\pm\infty} f = \pm Ut \tag{2.15}$$

or equivalently

$$\lim_{\eta\to\pm\infty} g(\eta) = \pm 1 \tag{2.16}$$

The solution to (2.14), subject to the constraints at $y = 0$ and (2.16) is then

$$g(\eta) = \sqrt{\pi}\eta \exp\left(\eta^2\right) \operatorname{erfc}|\eta|. \tag{2.17}$$

g is both continuous and differentiable at $\eta = 0$, however the second derivative is undefined at $\eta = 0$. While further manipulations are required to demonstrate that g is formally a weak solution to (2.14), g has exactly the expected behavior, and precisely matches a numerical integration of (2.14).

Returning to the method of moments,

$$c_1 = c_0 Utg(\eta) \tag{2.18}$$

is substituted back into the cross-sectional integral of (2.7) to obtain an equation for the evolution of the second moment m_2. Applying (2.8), the time dependence of σ_x^2 for the infinite domain, finite shear case is reached:

$$\frac{d\sigma_x^2}{dt} = U^2 t + 2K_x. \tag{2.19}$$

Integrating in time from an initial point release recovers the original quadratic scaling:

$$\sigma_x^2 = \frac{1}{2}U^2 t^2 + 2K_x t. \tag{2.20}$$

As expected, the time dependence falls between the two previously discussed cases, with variance growing quadratically in time. Additionally, the quadratic case includes no dependence on the lateral dispersion coefficient K_y.

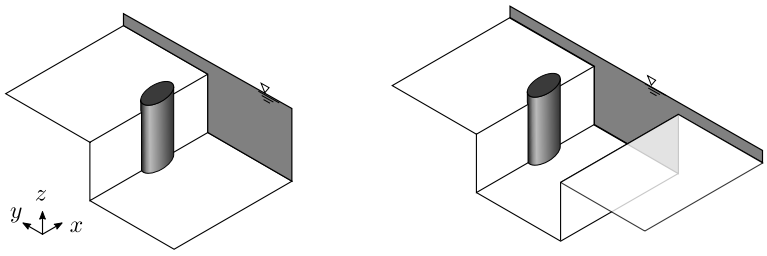

(a) shelf bathymetry (b) channel-shoal bathymetry

Figure 2.2: Schematic of plume interaction with bathymetry variability

Effects of Bathymetry

In physical shoal-channel flows, lateral variation of the velocity field is typically correlated to variation in depth. Sloping regions of the lateral bathymetry profile act as partial lateral boundaries, and a scalar plume which extends over sloping areas would be partially bounded at depth. In this partially bounded state one might assume that the realized growth of σ_x^2 is then best described by a combination of regimes. For the specific case of a vertically mixed plume centered on a break in the lateral bathymetry, such as in panel a of figure 2.2, the growth of variance is shown to remain purely quadratic in time, though with an attenuation factor which is a function of the ratio between the two depths. The derivation of this behavior relies on the observation that when the plume is centered on the step then solutions for the per-filament moments c_p are the same as for the flat-bottom case. (2.7) is applied to get the first and second moments of the total tracer distribution. $\mu_x = m_1/m_0$ describes the centroid of the plume, and in the step bathymetry case this is no longer stationary but instead follows

$$\mu_x = \frac{1-\alpha}{2\left(\alpha+1\right)} Ut \tag{2.21}$$

where $\alpha \equiv H_s/H_c$, the ratio of the two depths. In order for m_2/m_0 to represent the second central moment, (2.7) is evaluated in a coordinate system centered on μ_x, leading to

$$\frac{d\sigma_x^2}{dt} = \left[\frac{4\alpha}{\left(\alpha+1\right)^2}\right] U^2 t + 2K_x. \tag{2.22}$$

In essence the step bathymetry leads to an attenuated quadratic dispersion, where the attenuation factor depends only on the ratio of the two depths. This result only holds in the very specific case of a plume centered on the bathymetric break — in general one would expect that bathymetry, acting as a submerged boundary, would have a significant effect on the lateral tracer concentration distribution and the resulting dispersion would be a mix of regimes.

Transitions

The scaling arguments and results of the method of moments clearly delineate behavior of σ_x^2 in the simplified, "pure" scenarios, but for a tracer cloud which passes through each regime there is the additional question of how one regime transitions to the next. Further examination of (2.9), (2.10) and (2.19) reveals that these transitions are characterized by a continuous $d\sigma_x^2/dt$.

Examining the evolution of σ_x^2 in the cubic case, (2.10), $\sigma_y^2 = 2K_y t$ describes the lateral variance and $\sigma_y = \sqrt{2K_y t}$ gives a linear scale for the lateral extent of the plume. When coupled with the shear this gives a scale for the range of velocities $U(t) \approx \partial u/\partial y \sqrt{K_y t}$. (2.10) could then be rewritten $d\sigma_x^2/dt \approx U^2(t)\, t$, similar to the first term in the quadratic evolution described by (2.19). The transition from cubic to quadratic is contained in the time-dependency of the velocity scale, where the cubic regime implies a linear growth in U, and as the range of velocities asymptotes to a constant value the plume moves to the quadratic regime.

The transition from quadratic to linear dispersion can be examined in terms of the time scale for lateral mixing $T \sim W^2/K_y$. For $t > T$, quadratic regime must be shut down and the dispersion reverts to the linear regime. Evaluating (2.19) with $t = T$ returns exactly the linear scaling, implying that the variance growth rate during the linear regime is also the maximum growth rate achieved by the quadratic regime.

More broadly, the generic scaling

$$\frac{d\sigma_x^2}{dt} \sim U^2(t)\tau(t), \tag{2.23}$$

may be considered, where $U(t)$ is a velocity scale describing the range of velocities sampled by the plume and $\tau(t)$ is a time scale describing how long it takes for a parcel to sample that range of velocities. The three regimes can then be categorized by the time dependency in $U(t)$ and $\tau(t)$. In the cases where the lateral extent of the plume is not constrained by domain boundaries, the lateral mixing time grows with the plume size such that $\tau(t) \sim t$, but as the plume becomes large enough to "feel" the constraints of the boundaries the time scale τ asymptotes to a constant value W^2/K_y, where W is the width of the domain. During periods when the plume has not yet sampled the full range of velocities in the flow, the lateral spreading and an assumed linear shear profile lead to $U^2(t) \sim t$. However, any finite flow has a finite range of velocities such that eventually U must also asymptote to a constant value. Combining the velocity and time scales, the resulting evolution of σ_x^2 is then clearly continuous, with time dependence ranging from t^2 when neither U nor τ have reached their asymptote, to a constant value in the linear shear dispersion limit of a constant U and τ.

2.3 Estimation of Regimes and Parameters

The next analysis considers a number of particle plumes released in an idealized flow and analyzed through the lens of the regimes discussed above. The aim is to discern, , by careful

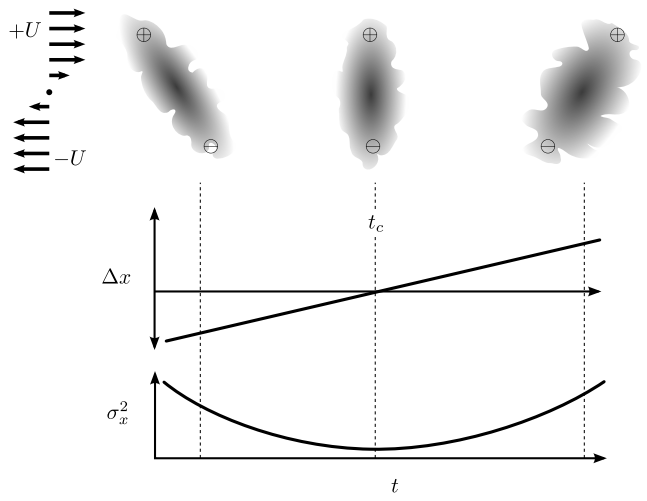

Figure 2.3: Accumulation of strain by a plume in the finite shear, infinite domain case

evaluation of the time series of spatial moments, (i) the relative importance of each regime, and (ii) the degree to which the parameters of the simulated flow predict the evolution of the plume. The goal is then to describe the evolution of $\sigma_x^2(t)$ in terms of the three regimes and the relevant parameters including the shear S, finite velocity range U, and lateral dispersion coefficient K_y. While the parameters of the regimes are generally changing in time, both due to the changing plume size and the local hydrodynamics driving the dispersion, the problem is simplified by considering the parameters to be constant over short windows of time. For each time window the time rate of change of the variance can be modeled by a second degree polynomial

$$\frac{d\sigma_x^2}{dt} \approx \alpha_2 t^2 + \alpha_1 t + \alpha_0 \tag{2.24}$$

with the three coefficients corresponding conceptually to the three dispersion regimes. The presence of t and t^2 in (2.24) imply a time origin for the cubic and quadratic regimes. Since the plume behavior should not depend on the choice of an absolute time origin, one must assume that t is instead relative to one or more time origins related to the plume state. One possible choice is to take all times relative some t_y at which $\sigma_y^2 = 0$. However, as a plume shifts between regimes the x-y covariance of the plume of the plume may not evolve in lockstep with the lateral variance σ_y^2 and, as will be shown below, the covariance plays a key role in the time dependence of variance evolution.

Consider a pair of particles, such as in figure 2.3, in the finite shear flow of the quadratic regime, where one particle $\oplus$ is in the $y > 0$ half-plane and advected with a velocity $+U$, and the other, $\ominus$, is in the $y < 0$ half-plane advected with velocity $-U$. Ignoring for the moment effects of K_x and K_y, the longitudinal separation between the particles is then expected to evolve as $\Delta x = 2U(t - t_c)$, where t_c indicates the time when the particles are aligned across the flow (and more generally the time at which the covariance is zero). In the quadratic and

cubic cases the particles are free to diffuse laterally without limit and on average never mix back across the shear region, which allows the approximation $\sigma_x \approx \Delta x/2$ (as opposed to the laterally bounded, linear regime where particles are periodically mixed back to the center of the domain and this scaling does not hold). The variance then evolves as the square of time, and the time rate of change goes as $d\sigma_x^2/dt \sim U^2\,(t - t_c)$. Combining the t_c dependence from the present scaling argument with (2.19), the adjusted form for the evolution of the variance is then

$$\frac{d\sigma_x^2}{dt} \approx U^2\,(t - t_c)\,. \tag{2.25}$$

Understanding the independent effects of t_y and t_c in the cubic regime requires a return to the method of moments. The previous analysis is modified in the choice of initial conditions for c_0, which now describes a laterally spreading Gaussian plume for which $\sigma_y^2\,(t = t_y) = 0$,

$$c_0\,(y,t) = \frac{m_0}{H\sqrt{4\pi K_y\,(t - t_y)}} \exp\left(-\frac{y^2}{2K_y\,(t - t_y)}\right). \tag{2.26}$$

For $t_y = t_c = 0$, (2.26) reduces to (2.11), the lateral mass distribution in both the cubic and quadratic analyses. The longitudinal concentration distribution is also modified such that $\mathrm{cov}_{xy}\,(t = t_c) = 0$. $t_y < t_c$ indicates an offset in time between the start of lateral spreading and the start of straining. Subject to the initial and boundary conditions, and satisfying (2.6), the per-filament centroid is described by

$$c_1\,(y,t) = \frac{c_0 Sty}{2}\left(1 + \frac{t_c - t_y}{t - t_y}\right). \tag{2.27}$$

Applying (2.7), the evolution of the longitudinal variance is then given by

$$\frac{d\sigma_x^2}{dt} = 2K_y S^2\,(t - t_c)\,(t - 2t_y + t_c) + 2K_x. \tag{2.28}$$

The schematic of figure 2.4 depicts the relationship between t_c, t_y, cov_{xy}, σ_y^2 and σ_x^2. Conceptually, $\sigma_y^2\,(t = t_c)$ dictates the initial range of velocities sampled by the plume and the covariance reflects an accumulation of strain in the plume, the same role as in the quadratic analysis above.

In moving from the pure, analytical cases of section 2.2 to the present forms which aim to be applicable to transitional states, it also becomes necessary to distinguish between the internal mixing of parcels within the plume, given by the dispersion coefficient K_y, and the lateral spreading of the plume, $d\sigma_y^2/dt$. In the unbounded case these are related by a factor of two, but as a plume becomes partially bounded or in the linear limit fully constrained by lateral boundaries $d\sigma_y^2/dt$ and K_y diverge. K_y continues to describe mixing within the plume but $d\sigma_y^2/dt$ tends towards zero. In the infinite shear case it is the lateral spreading of the plume which is relevant, and which should be attenuated when lateral boundaries play a

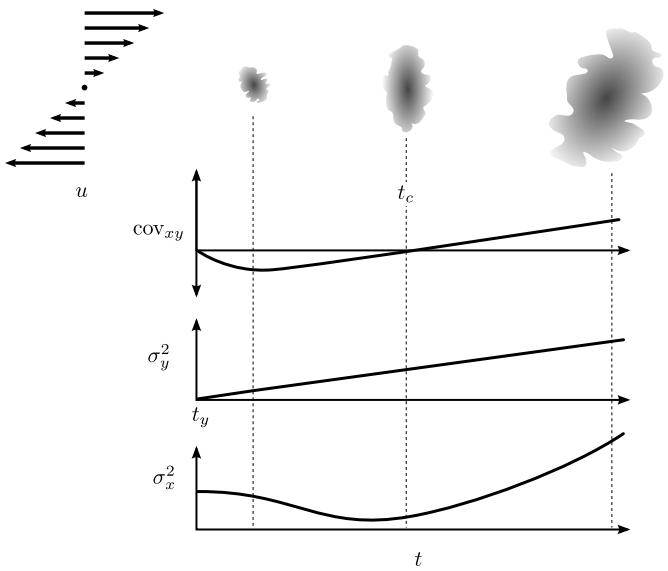

Figure 2.4: Relationship between σ_x^2 and the time origins for covariance and lateral variance in the case of infinite shear.

role. Combining the time origins and replacing K_y with the more directly descriptive $d\sigma_y^2/dt$, the infinite shear regime is then characterized by the form

$$\frac{d\sigma_x^2}{dt} \approx \left(\frac{\partial u}{\partial y}\right)^2 \frac{d\sigma_y^2}{dt} (t - t_c)(t - 2t_y + t_c). \tag{2.29}$$

In the case of linear shear dispersion, for which the plume is laterally well-mixed, the lateral variance is by definition constant. The covariance in the linear dispersion case also asymptotes to a constant value with a time scale proportional to the lateral mixing time, after which the straining action of the shear is balanced by the homogenizing effect of lateral diffusion, shown schematically in figure 2.5. Corresponding to the lack of any evolving state in the linear case other than σ_x^2, there is no additional time dependence and no need to include t_c or t_y at times beyond the lateral mixing time, and the evolution of σ_x^2 remains

$$\frac{d\sigma_x^2}{dt} \approx \frac{U^2 W^2}{\gamma K_y} \tag{2.30}$$

where γ is a constant related to the specific velocity profile.

The simplest representation of the combined effects of the three regimes is then a simple summation of the individual contributions, taking into account the time origins for the higher order regimes:

$$\frac{d\sigma_x^2}{dt} \approx \left(\frac{\partial u}{\partial y}\right)^2 \frac{d\sigma_y^2}{dt} (t - t_c)(t - 2t_y + t_c) + U^2 (t - t_c) + \frac{U^2 W^2}{\gamma K_y} \tag{2.31}$$

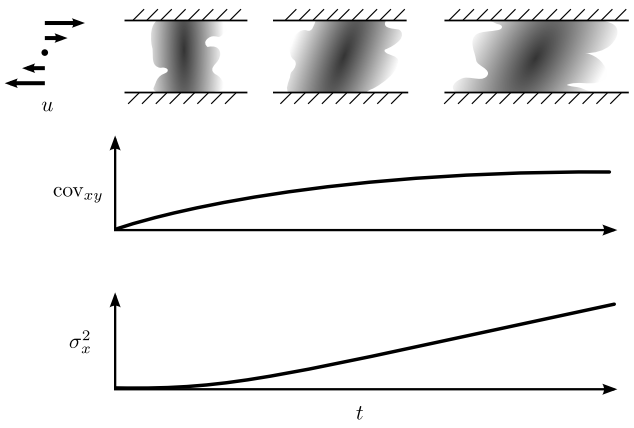

Figure 2.5: Asymptotic behavior and lack of accumulated state for linear shear dispersion. The upper panel shows a sample tracer distribution, with corresponding time series of covariance and longitudinal variance in the lower panels.

In order to make meaningful comparisons between observed coefficients of a polynomial fit as in (2.24) and (2.31) it becomes necessary to estimate t_c and t_y. Considering the effect of the time origins on each term, the terms of (2.31) are connected with the polynomial coefficients of (2.24). The highest order coefficient is unaffected by the time origins:

$$\left(\frac{\partial u}{\partial y}\right)^2 \frac{d\sigma_y^2}{dt} \approx \alpha_2. \tag{2.32}$$

Taking into account the time origins of the cubic dispersion term, the coefficient for t can then be approximated by

$$\left[U^2 - 2\left(\frac{\partial u}{\partial y}\right)^2 \frac{d\sigma_y^2}{dt} t_y\right] \approx \alpha_1. \tag{2.33}$$

Similarly, time origins for both the cubic and quadratic dispersion terms are relevant for the constant term of (2.24):

$$\frac{U^2 W^2}{\gamma K_y} - U^2 t_c + \left(\frac{\partial u}{\partial y}\right)^2 \frac{d\sigma_y^2}{dt}\left(2t_c t_y - t_c^2\right) \approx \alpha_0. \tag{2.34}$$

Note that (2.24), while including terms quadratic, linear and constant in time, does *not* map directly to the cubic, quadratic and linear shear dispersion regimes due to the time origins t_c and t_y. To more cleanly separate the influence of each regime, (2.32)-(2.34) are used to define

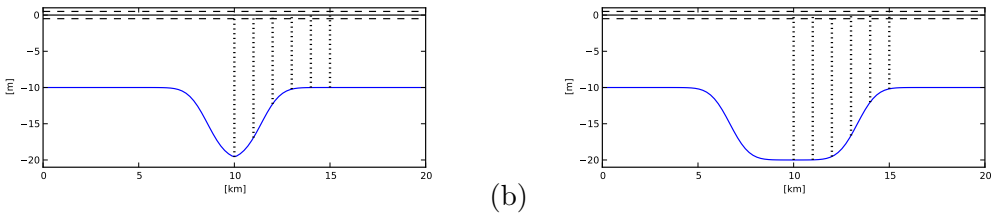

(a) (b)

Figure 2.6: Bathymetry profiles idealized domains, (a) narrow channel and (b) wide channel. Horizontal, dashed lines bracketing $z = 0$ denote the range of the tidal boundary condition. Vertical, dotted lines denote particle release locations.

a set of coefficients α_i^*, in terms of α_i, which describe the role of each individual regime:

$$\alpha_2^* = \alpha_2 \approx \left(\frac{\partial u}{\partial y}\right)^2 \frac{d\sigma_y^2}{dt} \tag{2.35}$$

$$\alpha_1^* = \alpha_1 + 2\alpha_2 t_y \approx U^2 \tag{2.36}$$

$$\alpha_0^* = \alpha_0 + \alpha_1^* t_c - \alpha_2^* \left(2t_c t_y - t_c^2\right) \approx \frac{U^2 W^2}{\gamma K_y} \tag{2.37}$$

With the effects of the time origins removed from each coefficient, what remains is a series of coefficients with one-to-one relationships with the three dispersion regimes and a prediction of each based on parameters of the flow.

2.4 Tidally-forced Channel-Shoal Domain

In a step towards applying the above analysis to real world conditions this section considers a more realistic but still idealized flow. Two domains are utilized, both 100km long, and 20km wide, varying in the lateral bathymetry profile as shown in figure 2.6. The $x = 0$ end of the channel is forced by 12 hour periodic tides, with a peak-to-peak amplitude of 1.0m, while at $x = 100$km the landward end of the channel is closed. The hydrodynamic simulations utilize the SUNTANS RANS model [28], run in 2-D mode. The primary goal of these simulations is to drive a moderately complex but tractable two dimensional flow field in which the three dispersion regimes may be simultaneously observed. Dimensions of the basin, tidal amplitudes and the lateral bathymetry profile fall within the range of typical physical values. The focus is on the subtidal dispersion caused by interaction of the subtidal flow field with shorter time scale turbulent mixing and tidal stirring. The residual flow field is essentially the landward portion of the flow described by [43], in which the channel carries a residual landward flow and the shoals have a residual seaward flow.

Particle tracking model

FISH-PTM [32] simulates the transport and dispersion of passive particles by the predicted hydrodynamic flow field output by SUNTANS. This particle tracking model includes both deterministic transport by the Eulerian velocity field as well as a non-deterministic parameterization of turbulent dispersion. The non-deterministic component, a carefully formulated random walk, is calculated at each time, for each particle, according to the formulation of [37] with a constant eddy diffusivity of $K_0 = 0.1\ \mathrm{m^2\,s^{-1}}$. The diffusivity has been chosen to be relatively small so that transitions between regimes are sufficiently separated in time to allow a detailed analysis of each regime. While the full formulation of [37] is quite involved, the bulk of the mathematical machinery is due to treatment of gradients and discontinuities in K_0. In the present case the diffusivity is constrained to be constant, uniform and isotropic, such that the particle tracking integration is greatly simplified. Under these conditions, each particle location is updated according to the split time-stepping scheme

$$X_i^{n+1/2} = X_i^n + u_i\left(\mathbf{X}^n\right)\Delta t \tag{2.38}$$

$$X_i^{n+1} = X_i^{n+1/2} + R\sqrt{2r^{-1}K_0\Delta t} \tag{2.39}$$

where i denotes the coordinate dimension, X_i^n the particle's location in the ith dimension at step n, $u_i\left(\mathbf{X}\right)$ the velocity from the hydrodynamic model, interpolated in time and space as needed, R a random number evenly distributed over $[-1, 1]$, and $r = 1/3$. Both the advective step and the diffusive step are subcycled as needed. The advective step is subcycled as particles cross an edge into a new cell of the hydrodynamic grid, and the diffusive step is subcycled according to the time step constraints detailed in [60]. The diffusion step may lead to a particle location which falls outside the domain. In this case the particle location is reflected back into the domain.

Particle clouds are released, evenly distributed in the vertical, at several locations in both the lateral and longitudinal directions, shown in figures 2.6 and 2.7. The initial lateral position of the plume ranges from the center of the channel to halfway between the channel center and the lateral boundary of the domain. Each plume is tracked over 100 tidal periods. At the start of each tidal cycle the first and second moments are calculated as:

$$\mu_j\left(t\right) = \frac{1}{N}\sum_{p=1}^{N} r_{p,j}\left(t\right), \text{ and} \tag{2.40}$$

$$\sigma_{jk}\left(t\right) = \frac{1}{N}\sum_{p=1}^{N}\left[r_{p,j}\left(t\right) - \mu_j\left(t\right)\right]\left[r_{p,k}\left(t\right) - \mu_k\left(t\right)\right], \tag{2.41}$$

where $r_{p,j}\left(t\right)$ is the position of particle p along axis j at time t, N is the total number of particles and the indices j and k range over the coordinate axes x and y. With these definitions, $\sigma_x^2 \equiv \sigma_{xx}$, $\sigma_y^2 \equiv \sigma_{yy}$ and $\mathrm{cov}_{xy} \equiv \sigma_{xy}$.

In addition to tracking the evolution of each plume, the particle tracking model is also used to estimate a residual Lagrangian velocity field. This field is estimated by releasing

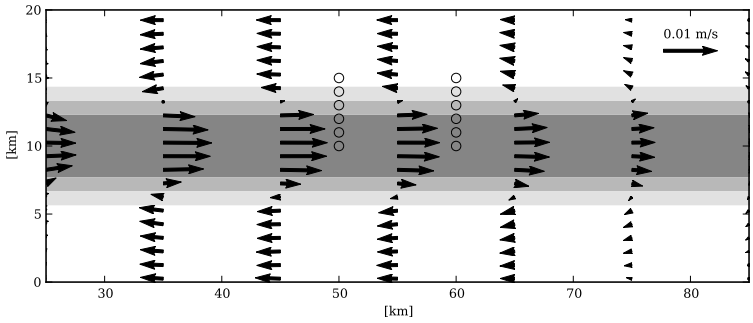

Figure 2.7: Plan view of the central portion of the wide channel idealized domain. Locations of particle plume releases are shown by empty circles, and the Lagrangian residual field is shown by arrows. Greyscale depth breaks shown at depths of 11, 15 and 19 m.

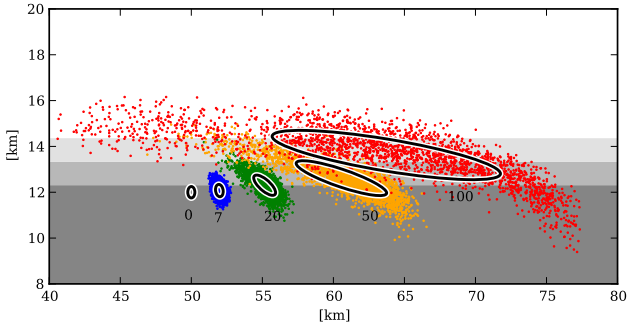

Figure 2.8: Instantaneous plume distributions after 0 (gray), 7 (blue), 20 (green), 50 (orange) and 100 (red) tidal periods. Greyscale depth breaks at 11, 15, and 19 m, with the lower portion of the plot falling within the channel region and the upper half in the shoal region. Ellipses are drawn according to the 2-D covariance matrix.

particles on a uniform, 200m grid throughout the domain, and extracting the displacement of each particle after being advected for exactly one tidal period with the dispersive random walk disabled. The displacements define the residual Lagrangian velocity, anchored in space at each respective initial position. The residual velocity field for the wide channel domain is shown by the arrows in figure 2.7 domain.

A sample plume trajectory is shown in figure 2.8, showing the interaction of the plume with slope-generated residual shear. The initial release is within the channel, close to the edge of the slope region. The initial residual transport carries the plume towards the closed end of the domain, while a mix of dispersion and lateral advection lead to a growing portion of the plume which samples the slope and shoal regions. Straining of the plume across the slope region leads to a significant quadratic growth in the variance. The time series of σ_x^2 is shown in figure 2.9, partitioned into contributions from the three dispersion regimes as

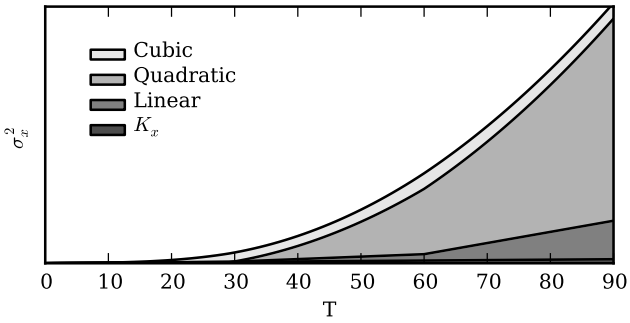

Figure 2.9: Cumulative growth of σ_x^2 for the same particle cloud as depicted in figure 2.8. The time series of σ_x^2 has been partitioned into the cumulative contributions of the three shear dispersion regimes, based on the observed values for α_i^*, plus the constant contribution of the imposed K_0.

well as the constant contribution from the imposed background dispersion rate K_0. The partitioning is according to the values of α_i^* and estimates for t_c and t_y.

Comparison of fit parameters to direct estimates

In order to assess the relevance of the analytical regime descriptions to cases where the net dispersion is an evolving combination of multiple regimes, the adjusted observed polynomial coefficients α_i^* is compared to predicted quantities according to (2.35)-(2.37). The predicted quantities depend on approximations of parameters describing the subset of the flow sampled by the plume during each time window. Within a single tidal cycle the internal dispersion is dominated by the imposed K_0 and for the estimates in this section $K_x = K_y = K_0$ is assumed. While the particle tracking results include tidal variations, the present analysis consider only the subtidal signals, obtained by extracting parameters only at integer tidal periods $t_n = nT$ where $T = 12$h is the tidal period. The width scale $W(t_n)$ is approximated by $\sigma_y(t_n)$. The velocity scale $U(t_n)$ is extracted from the residual Lagrangian velocity field by taking the standard deviation σ_u of the residual along-channel velocity at each particle location at time t_n. Per-window quantities W and U are then taken as averages of $W(t_n)$ and $U(t_n)$ over all t_n in the time window. Approximation of the quadratic coefficient in (2.35) is simplified to a direct estimate of the rate of change of the range of velocities sampled by the plume. Assuming a uniform, constant shear $\partial u/\partial y$ in the vicinity of the plume (the base assumption for the cubic regime),

$$\left(\frac{\partial u}{\partial y}\right)^2 \frac{d\sigma_y^2}{dt} = \frac{d}{dt}\left(\frac{\partial u}{\partial y}\sigma_y\right)^2 = \frac{d\sigma_u^2}{dt} \tag{2.42}$$

where the right hand side is evaluated by fitting a line to the time series of σ_u^2 within each time window. The time origins t_c and t_y are estimated by a least-squares linear fit to time series of

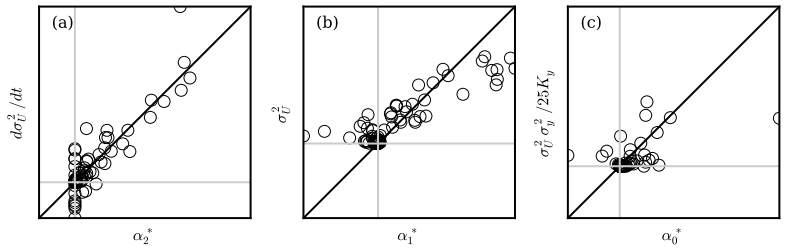

Figure 2.10: Correlation between predicted and observed growth coefficients. Diagonals show the 1:1 lines and grey perpendiculars intersect at $(0, 0)$. Panels show correlation for (a) the infinite shear regime, (b) the finite shear / unbounded regime, and (c) the linear, Fickian limit of bounded shear dispersion. Correlation coefficients are $R = 0.890$, 0.867, and 0.565 for α_2^*, α_1^* and α_0^*, respectively.

both cov_{xy} and σ_y^2 over the time window and solving for the time when the line crosses zero. The polynomial fit (2.24) is also constrained to a non-negative leading quadratic coefficient. While it is possible for a plume to exhibit growth described by a negative coefficient on t^2, the scaling relationships presently discussed are not applicable to such flows. Flow features including longitudinal convergence or mean vorticity which rotates longitudinal variance into the lateral axis could lead to decreasing σ_x^2 but are beyond the present analysis.

A comparison of the measured α_i^* and corresponding predicted coefficients is shown in figure 2.10. Since the velocity shape factor γ is not one of the estimated parameters, it is given a nominal value of 25 for the purposes of the plotting, noting that the slope of the correlation in the right-hand panel is arbitrary. The comparison for α_2^* (figure 2.10a) shows good agreement for larger observed values, but significant clustering around 0, largely due to the limitation that α_2 is required to be non-negative. Windows for which α_2 would have been negative if not constrained are most often times when the plume encounters some longitudinal variation in the flow. For example, plumes in the center of the channel move slowly to the closed end of the basin and eventually encounter a significant longitudinal convergence, $\partial u/\partial x < 0$, compressing the plume and decreasing σ_x^2. A similar process occurs for plumes in the shoals which advect far enough towards to the open end of the domain to encounter the "null" in the residual field. Both cases fall outside the range of regimes covered by the method of moments as applied in section 2.2, and a broader, fully two-dimensional approach would be required to adequately treat these cases. The lower correlation for α_0^* (figure 2.10c), relative to α_2^* and α_1^* (figure 2.10a,b), is due in part to variability of the velocity profile which is not captured by the constant value of γ. The fact that α_0^* relies on all three values of α_i, as well as linear and nonlinear terms in t_c and t_y undoubtedly contributes to the higher error relative to the higher order α_2^* and α_1^*.

2.5 Discussion

The flows presented here have been purposely distilled down to cases where higher order dispersive regimes may be observed and quantified. In observations or physically realistic simulations a broad range of processes lead to a more complex picture, in particular in the temporal and spatial variability of K_y and K_x. In the context of subtidal dispersion, these dispersion coefficients parameterize the combined effects of numerous processes including turbulent mixing, longitudinal-vertical shear dispersion, tidal stirring and lateral-vertical shear dispersion. Additional non-Fickian processes may also contribute to the overall plume dynamics, including residual density-driven transport in both longitudinal and lateral axes. Depending on the time scales involved, these processes may add to or mask high-order dispersion regimes. However, there remains the fundamental notion that the velocity scale U present in shear dispersion relations cannot be assumed to asymptote to a constant value in the same time period as the lateral extent of the plume asymptotes to a constant value.

Note that the quantitative estimates for the importance of quadratic dispersion do not necessarily exclude periods of cubic dispersion, which when observed for a short period may exhibit a quadratic time dependence. Further examination of periods of significant quadratic dispersion, however, show little correlation between the measured quadratic and cubic terms, suggesting that dispersion in these periods is best classified as quadratic. Specifically, comparing α_2^* and α_1^* the correlation coefficient is $R = -0.07$. A similar comparison, but between the predicted coefficients shows a correlation coefficient of $R = 0.20$. Both comparisons reinforce the conceptual model of a plume quickly growing beyond the lateral scale of the shear and then reverting to a quadratic dispersion regime.

Of particular interest in the high order regimes is that the overall dispersion, in terms of cross-sectionally averaged gradients and fluxes, is not Fickian. Operational 1-D models (e.g. [62]) must make many simplifying assumptions, and the use of Fickian dispersion coefficients is one of them. While Fickian dispersion is easily expressed in terms of either plume growth ($\sigma^2 \approx 2Kt$) or a differentially defined flux ($F \approx -K\frac{\partial C}{\partial x}$), high order dispersion regimes cannot be expressed in terms of local gradients. For the simplest, infinitely long, steady state case one can argue that all time scales for lateral mixing are achieved and all dispersion reverts to linear Taylor dispersion. In more realistic applications, though, embayments are finite and often exhibit significant variation along the axis. One can imagine a situation such as figure 2.11 where a residual circulation transports salt through the embayment. In terms of a one dimensional salt balance, though, gradients of the cross-sectionally averaged salinity $\langle s \rangle$ are negligible, even though the salt flux is quite significant. This situation is analogous to the quadratic dispersion regime in the sense that straining of the concentration field leads to significant dispersion or flux before lateral mixing is able to homogenize the cross-section.

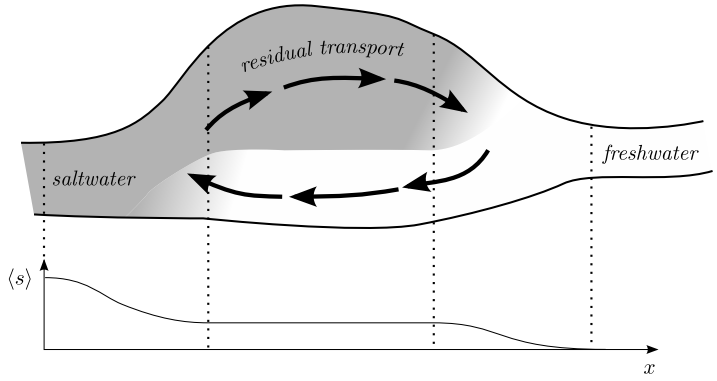

Figure 2.11: Flux view of quadratic dispersion in an estuarine reach which is not laterally well-mixed. Significant salt flux occurs even though gradients in the cross-sectionally averaged salinity $\langle s \rangle$ are negligible.

2.6 Conclusions

Motivated by the structure of environmental flows, this chapter has introduced a dispersion regime with quadratic dependence on time, to augment the previously studied cubic and linear dispersion regimes. This regime may be prevalent in flows where the lateral extent of the velocity shear is smaller than the lateral extent of the domain, such as channel-shoal systems in which the slope features are narrow compared to the width of the channel or the total width of the basin. A plume with a lateral length scale falling between the width of the shear and the width of the domain is expected to evolve according to this quadratic regime. While simple scaling arguments quickly arrive at this behavior, a more rigorous and precise description is obtained by way of the method of moments. In a suitably constructed idealized flow all three regimes occur and the net evolution of the plume variance can reasonably be predicted by parameters extracted from the flow.

Chapter 3

Numerical Diffusion on Unstructured, Flow-aligned Meshes

Abstract

The benefits of unstructured grids in hydrodynamic models are well understood but in many cases lead to greater numerical diffusion compared to methods available on structured grids. The flexible nature of unstructured grids, however, allows for the orientation of the grid to align locally with the dominant flow direction and thus decrease numerical diffusion. This chapter investigates the relationship between grid alignment and diffusive errors in the context of scalar transport in a triangular, unstructured, 3-D hydrodynamic code. Analytical results are presented for the 2-D anisotropic numerical diffusion tensor and verified against idealized simulations. Results from two physically realistic estuarine simulations, differing only in grid alignment, show significant changes in gradients of salinity. Changes in scalar gradients reflect reduced numerical diffusion interacting with the complex three-dimensional structure of the transporting flow. Also described is a method for utilizing flow fields from an unaligned grid to generate a flow-aligned grid with minimal supervision.

3.1 Introduction

Unstructured grid models are used in a wide variety of computational fluid dynamics applications, including extensive use in hydrodynamics [11, 28, 12]. Unstructured models allow seamless transitions from large scales to small and efficient adaptation to complex domain geometry. Along with these advantages come challenges related to grid generation, higher-

Previously published by Wiley and Sons: Holleman, R. C., O. B. Fringer and M. T. Stacey. Numerical Diffusion on Unstructured, Flow-aligned Meshes. *Int. J. Numer. Meth. Fluids*, 2013.

order methods and error analysis. In particular, many higher-order advection schemes are not directly applicable to unstructured grids, leaving these models to use diffusive low-order advection schemes. Characterizing this numerical diffusion and finding ways to reduce its impact on global model error are important facets of utilizing unstructured models.

A standard mathematical tool for evaluating discretization errors is modified equation analysis [33, 82], which can be applied to structured grid methods with relative ease. This analysis is still possible in the case of unstructured grids, but irregular stencils and coupling between coordinate dimensions complicate matters. Special care must be taken to account for heterogeneities in the grid and the fact that neighboring cells may not be characterized by the same modified equation and may not be independently consistent with the original PDE being modeled. For example, Bouche [5] investigated the upwind finite volume method applied to the scalar advection equation and observed an apparent loss of consistency when analyzing the method on unstructured grids. Even discretizations on structured but nonuniform grids can lead to inconsistency, such as the nonuniform curvilinear meshes studied by Turkel [77]. They observed that irregular meshes led to inconsistent discretizations if the spacing was not quasi-uniform and the staggering placed the dependent variable at cell centers. In the specific case of finite volume, first-order upwind on triangular unstructured grids, the focus of this chapter, these heterogeneities arise on a uniform triangular grid due to asymmetries between two classes of cells, distinguished by the number of in-flow faces.

In practice first-order accurate advection schemes are often too diffusive to accurately predict scalar gradients [15]. While there do exist higher-order methods that are commonly implemented on unstructured grids, such as TVD and FCT, these methods are generally formulated as corrections or extensions to an underlying low-order method [20]. Monotonic higher-order schemes are necessarily nonlinear, as Godunov's theorem states that a linear monotone numerical scheme is at most first-order accurate [30]. This nonlinearity makes a generalized error analysis difficult and motivates an investigation of the diffusive behavior of the underlying low-order method, for which a complete analysis is possible in the case of a uniform mesh.

The bulk of analysis for numerical methods is focused on (i) consistency (with particular attention paid to the order of the error terms) and (ii) stability. Together, consistency and stability imply convergence by way of the Lax equivalence theorem [42], and for some applications it is sufficient to identify only the order of the error terms and the presence of stability, without regard to the coefficients on the error terms. Several aspects of hydrodynamic applications, however, call for a more quantitative analysis of numerical diffusion. Firstly, hydrodynamic applications often have well defined principal velocity directions, whether due to river flows, prevailing winds, or the major axis of the tidal ellipse. This attribute of marine flows motivates the desire to understand not just the worst case behavior of an advection operator, but also the expected behavior given a specific velocity distribution. Additionally, hydrodynamic applications often need to distinguish between lateral, or cross-flow, diffusion and longitudinal, or streamwise, diffusion. Secondly, oceanographers are often interested in quantifying physical mixing due to processes ranging from small-scale turbulent diffusion to diffusive processes arising from the evolution of the structure of the large-scale flow. Non-

physical diffusive errors are not just a frustration, but significantly raise the noise floor when the ultimate goal is to quantify diffusion arising from physical processes.

Both physical and numerical diffusion can interact with spatial variations in the velocity field to modify large-scale mixing and transport. For instance, shear in the velocity field will cause the spreading of a scalar distribution as fluid parcels sample different portions of the velocity profile. Diffusion perpendicular to the flow will cause a single parcel to sample a broader swath of the velocity profile. In the limit of infinite lateral diffusion, all parcels will sample all parts of the velocity profile equally, such that the mean transport of each parcel is identical and the distribution of parcels in the direction of the straining will remain constant. This interaction of shear and lateral diffusion where longitudinal diffusion is inversely proportional to lateral diffusion (*i.e.* shear dispersion, [74]) is an example of the potentially hidden effects of numerical diffusion, and motivates the need to differentiate lateral and longitudinal numerical diffusion.

There have been many methods proposed for reducing numerical diffusion in unstructured finite volume schemes while retaining monotonicity. TVD methods for one-dimensional problems (see [71]) are naturally extensible to multi-dimensional structured methods by way of operator splitting, and have more recently been extended to unstructured multi-dimensional problems, *e.g.* [48, 18]. Another approach, the Reservoir method [1], decreases numerical diffusion by time-shifting fluxes to achieve an effective local Courant number near unity. Grid alignment has long been known to affect the accuracy of results which has resulted in a number of approaches for adapting a grid to flow features. Adapting a grid to the flow may occur before the simulation, such as in applications focused on resolving boundary layers that use a static grid which envelopes boundaries with layers of highly anisotropic, boundary-oriented cells (such as the advancing layers method discussed in [54]), and the methods described in the current chapter. Many adaptive methods rely on runtime error estimates to target areas for refinement, such as the methods described in [34] applied to quadtree and octree meshes. Adaptive methods for fully unstructured grids such as [73] are able to apply anisotropic stretching and refining of grid elements in order to better resolve oriented features of a flow. Another perspective on adaptive grid methods is the Vofire method of [39] in which triangular cells are subdivided parallel to the local velocity direction at each step, decreasing the cross-flow component of numerical diffusion. These runtime adaptive methods require significant modifications to or complete reimplementation of the numerical core of a simulation code. Additionally, due to the nonlinear character of most of the schemes numerical diffusion is both difficult to predict *a priori*, and may be either positive or negative (such as the over-sharpening of gradients seen with the superbee limiter[18]).

For an important class of finite volume C-grid hydrodynamic solvers (*e.g.* [11, 28]) the computational grid must be orthogonal, meaning that the line joining adjacent cell centers must be perpendicular to their common edge, which is equivalent to requiring that cell centers are the circumcenters of the cells. Additionally they require that the cell center falls within the cell itself, a condition which is violated for triangular cells which are not acute. This constraint both limits the utility of anisotropic mesh approaches, and makes re-gridding

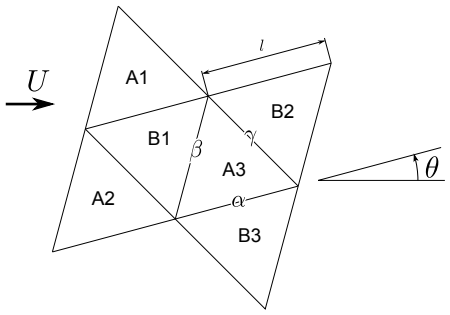

Figure 3.1: Schematic for derivation of recurrence relation

generally too expensive to be done dynamically.

Section 3.2 presents the details of the modified equation analysis for the linear advection equation on an idealized triangular grid. The analysis focuses on a first-order upwind, finite volume discretization of the scalar advection equation. Numerical diffusion is first analyzed in each of two classes of grid cells independently to demonstrate that such an independent analysis is insufficient to capture the coupling of errors between the two classes of cells. Consequently the analysis is extended to consider the aggregate errors due to the interaction of the two classes of cells, leading to analytical expressions for the components of the 2-D anisotropic numerical diffusion tensor. Section 3.3 introduces a series of idealized scalar advection test cases which are used to calculate empirical numerical diffusion coefficients. These empirical estimates are shown to verify the analytical expressions derived in Section 3.2. Section 3.4 presents an algorithm for generating flow-aligned grids and a simple test case applying this algorithm and evaluating the effects on numerical diffusion in the test case. Section 3.5 describes a physically realistic pair of simulations of San Francisco Bay, where the first case uses an unaligned grid and the second employs a grid with cells aligned to the principal velocity directions. Comparisons of salinity model data to observations show significant improvements in salinity gradients when using the aligned grid, consistent with decreased non-physical lateral diffusion. While the first-order upwind advection scheme utilized in these simulations is not state-of-the-art, it is effective at highlighting the role of grid alignment in numerical diffusion and the effect of anisotropic numerical diffusion on scalar distributions, both of which are present in higher-order advection schemes.

3.2 Analytical Grid Diffusion

As a point of reference for modified equation analysis [82], this section briefly considers the modified equivalent partial differential equation for one-dimensional, first-order, forward in time, backward in space scalar advection. This scheme can be written

$$A_j^{n+1} = CA_{j-1}^n + (1 - C)\, A_j^n \tag{3.1}$$

where A_j^n is the scalar concentration in cell j and time step n, and $C = U\Delta t/\Delta x$ denotes a volumetric Courant number, equivalent to the fraction of each cell's mass displaced by advective fluxes during each time step. A Taylor series expansion of all scalar values about time n and location j leads to the modified equivalent equation

$$A_t + UA_x = \frac{\Delta x^2}{2\Delta t} C\left(1 - C\right) A_{xx} + O\left(\Delta x^2, \Delta t^2\right) \tag{3.2}$$

where subscripts denote the corresponding partial derivative, and the coefficient on A_{xx} is interpreted as numerical diffusion. The stability of this scheme can be related to the sign of the diffusive term, leading to the conclusion that this method is stable for $0 \leq C \leq 1$.

In the context of two-dimensional structured grids this analysis may be applied directly, but a difficulty arises in applying a modified equation analysis to an unstructured triangular grid, even when the flow is uniform and the grid regular. The key hurdle to direct application of modified equation analysis is the fact that for a uniform and constant flow, half of the cells have one inflow face and two outflow faces (type A cells in Figure 3.1), while the other half have two inflow faces and one outflow face (type B cells). The direct, decoupled analysis is illustrated in the following section which derives the truncation error independently for type A and type B cells. The result is unrealistic, but motivates an extension of the analysis that considers the coupling of type A and B cells.

In the derivation below, following the schematic in Figure 3.1, the scalar concentration in type A cells at time step n is denoted A_i^n, and for type B cells B_i^n, to reinforce the notion that the two scalar fields are not interchangeable. Scalar values are cell averages, and for the purpose of gradients are located at the circumcenter of each cell. The orientation of the grid relative to horizontal is given by $0 \leq \theta \leq \frac{\pi}{3}$, and the three types of edges are denoted by α, β, and γ. The free-stream velocity U is taken parallel to the x axis, and the components of the velocity normal to the faces are given by $U_\alpha = U\sin\theta$, $U_\beta = U\sin\left(\theta + \frac{\pi}{3}\right)$, and $U_\gamma = U\sin\left(\theta + \frac{2\pi}{3}\right)$. Since the grid is congruent over rotations of $n\pi/3$, the given range of θ is sufficient to cover all possible orientations of the grid relative to the flow. The edge length is l, and the cell area is a, which for the case of equilateral triangles is $a = \frac{\sqrt{3}}{4}l^2$.

Independent analysis of A and B cells

The discretization of the scalar advection equation for axis-aligned flow, $A_t + UA_x = 0$, for a type A cell using a first-order upwind, finite volume discretization is

$$a\left(A_3^{n+1} - A_3^n\right) = l\Delta t\left(U_\beta B_1^n - U_\alpha A_3^n - U_\gamma A_3^n\right). \tag{3.3}$$

Defining a finite volume analog to the Courant number for each face $C_\zeta = U_\zeta l\Delta t/a$, and noting that, from continuity, $C_\beta = C_\alpha + C_\gamma$, (3.3) is rewritten:

$$A_3^{n+1} = \left(1 - C_\beta\right) A_3^n + C_\beta B_1^n \tag{3.4}$$

Applying Taylor series expansions in time and space, centered around A_3^n, the modified equation for a type A cell is

$$\begin{aligned} A_t &+ \frac{4}{3}\cos^2\left(\theta-\frac{\pi}{6}\right)UA_x + \frac{4}{3}\cos\left(\theta-\frac{\pi}{6}\right)\sin\left(\theta-\frac{\pi}{6}\right)UA_y \\ &= \left[\frac{2}{3\sqrt{3}}\cos^3\left(\theta-\frac{\pi}{6}\right)Ul - \frac{1}{2}\Delta t U^2\right]A_{xx} \\ &+ \frac{2Ul}{3\sqrt{3}}\cos\left(\theta-\frac{\pi}{6}\right)\sin^2\left(\theta-\frac{\pi}{6}\right)A_{yy} \\ &+ \frac{4Ul}{3\sqrt{3}}\cos^2\left(\theta-\frac{\pi}{6}\right)\sin\left(\theta-\frac{\pi}{6}\right)A_{xy} + O\left(\Delta t^2, l^2\right) \end{aligned} \tag{3.5}$$

The analysis of type B cells follows in the same manner, with the resulting modified equation:

$$\begin{aligned} B_t &+ \frac{4}{3}\left[\sin^2\left(\theta+\frac{2\pi}{3}\right)+\sin^2\theta\right]UB_x + \frac{2}{3}\left[\sin\left(2\theta+\frac{\pi}{3}\right)-\sin\left(2\theta\right)\right]UB_y \\ &= \left[\frac{2l}{3\sqrt{3}}\left(\sin^3\left(\theta+\frac{2\pi}{3}\right)+\sin^3\theta\right) - \frac{1}{2}U\Delta t\right]UB_{xx} \\ &+ \frac{2l}{3\sqrt{3}}\left(\cos\left(\theta+\frac{\pi}{6}\right)\sin^2\left(\theta+\frac{\pi}{6}\right)+\cos^2\theta\sin\theta\right)UB_{yy} \\ &+ \frac{4Ul}{3\sqrt{3}}\left[\cos^2\left(\theta+\frac{\pi}{6}\right)\sin\left(\theta+\frac{\pi}{6}\right)-\cos\theta\sin^2\theta\right]B_{xy} + O\left(\Delta t^2, l^2\right) \end{aligned} \tag{3.6}$$

Denoting the coefficients on A_x and B_x as the effective advective speeds U_A and U_B, the coefficients of A_{xx}, A_{yy}, and A_{xy} as numerical diffusivities K_A^x, K_A^y, and K_A^{cross} (and similarly for B cells), the modified equations for each type of cell have an apparent loss of consistency, due to $U_A \neq U_B \neq U$. Similar observations of loss of consistency have been made for discretizations on irregular structured grids, for example [77]. However, in this regular, unstructured case the expected advective speed U can be recovered by assuming a simple average of the coefficients in (3.5) and (3.6), which also allows the cross-stream advection coefficients on A_y and B_y to cancel. Similarly, averaging the diffusion coefficients and simplifying gives

$$K_{mean}^x = \left[\frac{1}{3\sqrt{3}}\sin^3\theta + \frac{1}{4}\cos\theta\right]Ul - \frac{1}{2}U^2\Delta t \tag{3.7}$$

$$K_{mean}^y = \left(\frac{1}{3\sqrt{3}}\cos^2\theta\sin\theta + \frac{1}{12}\cos\theta\right)Ul \tag{3.8}$$

$$K_{mean}^{cross} = \left(\frac{1}{6} - \frac{2}{3\sqrt{3}}\cos\theta\sin\theta\right)\sin\theta Ul \tag{3.9}$$

While the cross-stream advection coefficients cancel, coefficients on $\frac{\partial}{\partial x^2}$ and $\frac{\partial}{\partial y^2}$ in the Taylor series expansion are always positive and do not cancel, such that the resulting diffusion coefficients will overpredict both lateral and longitudinal diffusion. Note that the cross term

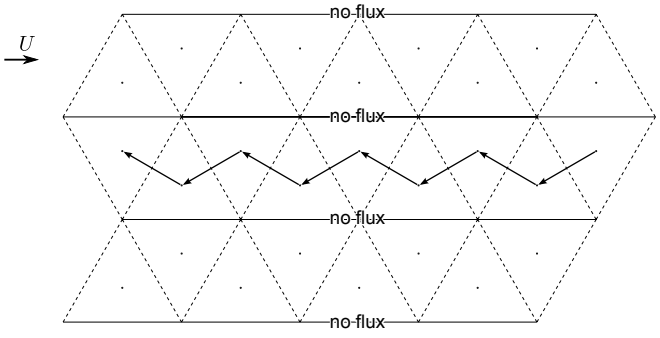

Figure 3.2: Schematic showing vanishing of lateral diffusion for $\theta = 0$. Solid lines denote zero-flux faces of the cells, and dashed edges have nonzero flux. The arrows along the center strip show the vector $(\Delta x, \Delta y)$ for which the Taylor series expansion is evaluated for each cell's upwind neighbor. The nonzero y component of the vectors leads to an overpredicted K^y unless the two types of cells are treated simultaneously.

only provides information on the orientation of diffusion - the magnitude of the total diffusion depends only on K^x and K^y. This result is intuitively evident from examining a strip of triangular cells, parallel to the flow, such as in Figure 3.2. Cell centers of the two types of cells are offset from each other in the cross-flow direction, such that the Taylor series expansion in the lateral direction includes positive coefficients on A_{yy} and B_{yy}. The positive coefficients lead directly to a nonzero K^y in the averaged modified equation. Geometrically, though, any diffusive spreading is limited to this single-width triangle strip, as the flow-aligned faces have zero flux. The source of this discrepancy is that a Taylor series expansion about one type of cell ignores the biases in the other type of cell.

Combined analysis

One method for amending the analysis is to expand the values B_i^n for type B cells in terms of values in type A cells, using a recurrence relation rather than a Taylor series expansion. For the schematic in Figure 3.1, the update equation for cell B_1 at time step n, analogous to (3.4) but for a type B cell, is

$$B_1^n = (1 - C_\alpha - C_\gamma)\, B_1^{n-1} + C_\alpha A_1^{n-1} + C_\gamma A_2^{n-1} \tag{3.10}$$

For notational convenience $\lambda \equiv 1 - C_\alpha - C_\gamma$. Successive application of the recurrence relation replaces B_1^{n-j} with an expression containing B^{n-j-1} and A_i^{n-j-1}, such that the result of k expansions is

$$B_1^n = \lambda^k B_1^{n-k} + C_\alpha \sum_{j=1}^{k} A_1^{n-j} \lambda^{j-1} + C_\gamma \sum_{j=1}^{k} A_2^{n-j} \lambda^{j-1} \tag{3.11}$$

The upwinded values from A_i at past time steps can be expanded as a Taylor series in time around step n, yielding A_i^{n-j} in terms of A_i^n and its time derivatives. Since the purpose of the analysis is to quantify numerical diffusion, the expansion are limited to the second derivative. Assuming an n large enough that B_1^n is fully decoupled from the initial condition B_1^0, *i.e.* $n\Delta t \gg l/U$ and $\lambda^n \approx 0$, allowing k to go to infinity results in

$$\begin{aligned} B_1^n &= \sum_{j=0}^{\infty} \lambda^j \left[C_\alpha A_1^n + C_\gamma A_2^n - (j+1)\,\Delta t \left(C_\alpha A_{1,t}^n + C_\gamma A_{2,t}^n \right) \right. \\ &\quad \left. + \frac{(j+1)^2\,\Delta t^2}{2} \left(C_\alpha A_{1,tt}^n + C_\gamma A_{2,tt}^n \right) + O\left(\Delta t^3\right) \right] \end{aligned} \tag{3.12}$$

for which the infinite sums can be replaced with closed expressions:

$$\begin{aligned} B_1^n &= \frac{1}{1-\lambda} \left[C_\alpha A_1^n + C_\gamma A_2^n \right] - \frac{\Delta t}{(1-\lambda)^2} \left[C_\alpha A_{1,t}^n + C_\gamma A_{2,t}^n \right] \\ &\quad + \frac{(1+\lambda)\,\Delta t^2}{2\,(1-\lambda)^3} \left[C_\alpha A_{1,tt}^n + C_\gamma A_{2,tt}^n \right] + O\left(l^3\right) \end{aligned} \tag{3.13}$$

This expression for B_1^n is now in terms of values at time step n evaluated only at the centers of type A cells. For the purposes of the modified equation analysis A_i^n and corresponding derivatives are then expressed as Taylor series expansions around A_3^n and its derivatives. After substitution, the modified equation (now dropping the subscript on A_3, as all values have been expanded about this point) reads

$$A_t + U_{eff} A_x = \left[2K^x\right] A_{xx} + \left[2K^y\right] A_{yy} + \left[2K^{cross}\right] A_{xy} + O\left(\Delta t^2, l^2\right) \tag{3.14}$$

The simplified expressions (see appendix A for details of the simplification) for the anisotropic numerical diffusion coefficients are:

$$\begin{aligned} 2K^x &= \left[\frac{2\sqrt{3}}{3} \left(\cos^2\theta \cos\left(\theta + \frac{\pi}{6}\right) + \sin\theta \cos^2\left(\theta - \frac{\pi}{3}\right) \right) \right. \\ &\quad \left. - \frac{\sqrt{3}}{4\sin\left(\theta + \frac{\pi}{3}\right)} \right] Ul - U^2 \Delta t \end{aligned} \tag{3.15}$$

$$2K^y = \frac{\sin(3\theta)}{2\sqrt{3}} Ul \tag{3.16}$$

$$2K^{cross} = \left[\frac{4}{\sqrt{3}} \sin\theta \sin\left(\theta - \frac{\pi}{6}\right) \sin\left(\theta - \frac{\pi}{3}\right) \right] Ul \tag{3.17}$$

$$U_{eff} = U$$

The factor of 2 arises from the fact that the analysis includes truncation errors for both the update of a B cell from its upwind A neighbors and for an A cell updated from its upwind B

neighbor. The effective mean numerical diffusion coefficients are consequently one-half the coefficients that appear in the modified equation.

For the special case of a flow-aligned grid $\theta = 0$ and the longitudinal numerical diffusion simplifies to

$$K^x = \frac{1}{4}Ul - \frac{1}{2}U^2\Delta t \tag{3.18}$$

Employing the fact that the longitudinal spacing of cell centers is $\Delta x = l/2$ when $\theta = 0$, and the Courant number $C = U\Delta t/\Delta x$, the longitudinal diffusion can be further simplified to $K_x = \frac{1}{2}\left(U\Delta x - U^2\Delta t\right) = \frac{\Delta x^2}{2\Delta t}C\left(1 - C\right)$. Thus in the case of exactly aligned flow, the unstructured method reduces to the same behavior as a one-dimensional structured grid. Likewise, the lateral diffusion coefficient goes to zero for an aligned grid, consistent with the schematic of Figure 3.2.

Since the update equation in the aligned-grid case is identical to that for a regular 1-D discretization, the same conditions apply for stability, namely that $0 \leq C \leq 1$. For arbitrary θ the condition for monotonicity, which is a sufficient condition for stability [20], is

$$0 \leq C_\beta = \frac{4U\Delta t \sin\left(\theta + \frac{\pi}{3}\right)}{\sqrt{3}\, l} \leq 1 \tag{3.19}$$

The worst-case scenario, which occurs at $\theta = \frac{\pi}{6}$, requires $0 \leq \frac{U\Delta t}{l} \leq \frac{\sqrt{3}}{4}$. Stability may also be investigated in terms of the diffusion coefficients. Since stability is related to the fastest growing Fourier mode, which may be anisotropic and not necessarily aligned with the axes, the two-dimensional case requires consideration of the signs of the eigenvalues of the diffusion tensor

$$\underline{\underline{K}} = \begin{bmatrix} K^x & \frac{1}{2}K^{cross} \\ \frac{1}{2}K^{cross} & K^y \end{bmatrix} \tag{3.20}$$

rather than the total diffusion $K^x + K^y$ or the separate axis-aligned diffusion coefficients K^x and K^y. The condition that the smallest eigenvalue of $\underline{\underline{K}}$ be non-negative can be shown to be equivalent to the monotonicity condition (3.19), though the lengthy and generally unenlightening algebra is omitted from this chapter.

For comparison, the diffusion coefficients for the same first-order upwind advection scheme on a regular Cartesian grid can be similarly (and more simply) derived. The discretization is

$$A^{n+1} = \left(1 - C_W - C_N\right)A^n + C_W A_W^n + C_N A_N^n \tag{3.21}$$

Where subscripts W and N denote the respective quantities for the upwind cells to the "west" (negative x direction when $\theta = 0$) and "north" (positive y when $\theta = 0$) direction. The Courant numbers are defined as:

$$C_N = \sin\theta\frac{U\Delta t}{l} \tag{3.22}$$

$$C_W = \cos\theta\frac{U\Delta t}{l} \tag{3.23}$$

Where θ takes on the same sense as in the triangular case, namely a counterclockwise rotation of the grid relative to a $+x$ directed mean flow. The modified equation analysis follows the same conventions as for the triangular grid, although with only a single type of cell the derivation is straightforward and as such has been omitted for brevity. The anisotropic numerical diffusion coefficients for a Cartesian grid then read:

$$K^x = \frac{1}{2}\left[\sin^3\theta + \cos^3\theta\right] Ul - \frac{1}{2}U^2\Delta t \tag{3.24}$$

$$K^y = \frac{1}{2}\left[\cos\theta\sin^2\theta + \cos^2\theta\sin\theta\right] Ul \tag{3.25}$$

$$K^{cross} = \left[\cos^2\theta\sin\theta - \cos\theta\sin^2\theta\right] Ul \tag{3.26}$$

For the Cartesian case the edge length l is equal to the cell-center spacing Δx, and θ is constrained to the interval $\left[0, \frac{\pi}{2}\right]$. As expected, the flow-aligned ($\theta = 0$) behavior is identical to the one-dimensional case and the two-dimensional flow-aligned triangular case. The condition for monotonicity is $0 \leq 1 - C_W - C_N \leq 1$, which for arbitrary θ leads to the condition

$$0 \leq \frac{U\Delta t}{l} \leq \frac{1}{\sqrt{2}\sin\left(\theta + \frac{\pi}{4}\right)} \tag{3.27}$$

The worst case scenario for monotonicity in the Cartesian case occurs at $\theta = \frac{\pi}{4}$, for which $\frac{U\Delta t}{l} \leq \frac{1}{\sqrt{2}}$. For equivalent edge lengths l, the Cartesian grid permits a longer time step by a factor of $\frac{2\sqrt{2}}{\sqrt{3}} \approx 1.63$. For equivalent cell areas (*i.e.* equivalent spatial resolution), the Cartesian grid still permits a longer time step, though by only a factor of $\frac{\sqrt{2}}{\sqrt[4]{3}} \approx 1.075$.

3.3 Idealized Simulations

Empirical estimates of numerical diffusion are used to corroborate the modified equation analysis of Section 3.2. The estimates come from a series of idealized scalar transport test cases simulating uniform flow in a flat-bottomed channel. The uniform nature of the test domain and the aim of quantifying the full 2-D diffusion tensor motivate the use of the method of moments [2] for estimating diffusion coefficients from the time rate of spreading of a Gaussian scalar plume. The domain, shown in Figure 3.3, extends $L = 50$km in the x dimension and $W = 20$km in the y dimension, with a constant depth of 10m. In all cases the northern and southern boundaries are closed, and a uniform velocity $u = 0.5$ m s^{-1} is imposed throughout the domain. Simulations are carried out using the discretization given by equations (3.4) and (3.10).

The simulations sample a portion of the parameter space relevant to physical simulations of estuarine domains, where the parameters are l, the side length of the cells ranging from 75m to 500m, θ, the smallest positive counter-clockwise angle between grid edges and the x axis, ranging from 0° to 60°, and the time step. The grid is created by taking a tessellation of equilateral triangles, rotating by θ, and removing all triangles that fall outside the rectangular

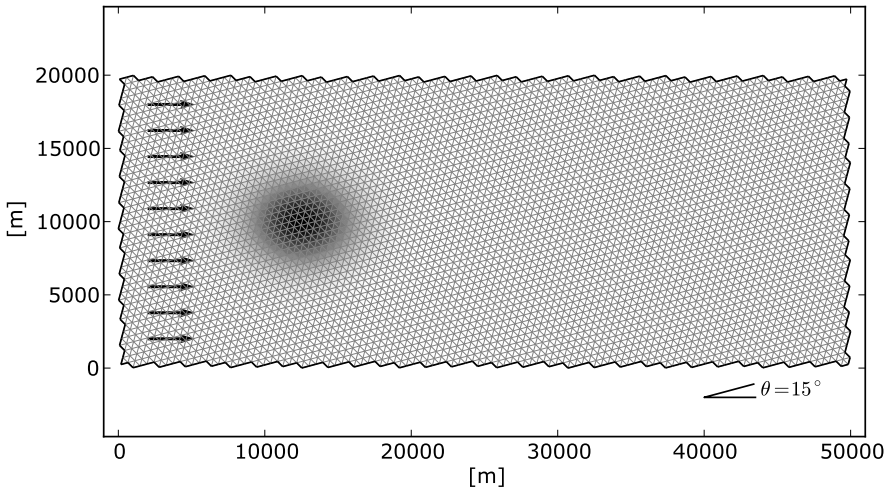

Figure 3.3: Idealized domain with initial condition for passive scalar. $l = 500$m, $\theta = 15°$.

boundary (leaving the grid with two or four "ragged" edges). The time step ranges from 40s for the highest resolution runs to 120s for the low resolution runs. The corresponding range of the volumetric Courant number is from 0.80 to 0.24. The passive scalar field is initialized with a Gaussian plume, centered at $(L/4, W/2)$ with $\sigma_x = 0.05L$ and $\sigma_y = 0.1W$ as shown in Figure 3.3. These dimensions ensure that the plume is sufficiently free from wall effects as it advects through the domain, while keeping computational costs reasonable. An additional constraint on the size of the plume is that it should cover a sufficient number of cells to reasonably approximate a Gaussian distribution. The simulations run until the domain-integrated scalar mass falls to 99% of its original value, at which point the trailing 25% of the simulation period is additionally discarded to avoid outflow boundary effects, leaving the period $t \in [t_0, t_f]$. The tracer field is output every 20 minutes, and each snapshot of the tracer distribution is processed to find the centroid, second central moments in x and y, and the spatial covariance. The centroid of the scalar plume is computed at each output time step according to

$$\mu = \frac{\sum_i \mathbf{r}_i V_i A_i}{\sum_i V_i A_i} \tag{3.28}$$

where $\mathbf{r}_i = (x_i, y_i)$ denotes the location of the ith cell center, V_i its volume, and A_i its tracer

concentration. The second central moments and spatial covariance are then computed as

$$\sigma_x^2 = \frac{\sum_i (x_i - \mu_x)^2 V_i A_i}{\sum_i V_i A_i} \tag{3.29}$$

$$\sigma_y^2 = \frac{\sum_i (y_i - \mu_y)^2 V_i A_i}{\sum_i V_i A_i} \tag{3.30}$$

$$\mathrm{cov}_{xy} = \frac{\sum_i (x_i - \mu_x)(y_i - \mu_y) V_i A_i}{\sum_i V_i A_i} \tag{3.31}$$

Values of the time-averaged diffusion coefficients K^x, K^y and K^{cross} are estimated as one half the average rate of growth for the second central moments and covariance [2], specifically

$$K^x = \frac{1}{2}\overline{\left(\frac{d\sigma_x^2}{dt}\right)} = \frac{1}{2}\left(\frac{\sigma_x^2|_{t_f} - \sigma_x^2|_{t_0}}{t_f - t_0}\right) \tag{3.32}$$

$$K^y = \frac{1}{2}\overline{\left(\frac{d\sigma_y^2}{dt}\right)} = \frac{1}{2}\left(\frac{\sigma_y^2|_{t_f} - \sigma_y^2|_{t_0}}{t_f - t_0}\right) \tag{3.33}$$

$$K^{cross} = \overline{\left(\frac{d\mathrm{cov}_{xy}}{dt}\right)} = \left(\frac{\mathrm{cov}_{xy}|_{t_f} - \mathrm{cov}_{xy}|_{t_0}}{t_f - t_0}\right) \tag{3.34}$$

where t_0 corresponds to the time when the initial Gaussian plume is introduced, and t_f the cutoff time as described above. Correlation between predicted and measured numerical diffusion, using the results of the independent modified equation analysis (Section 3.2), are shown in Figure 3.4. Both K^x and K^y are consistently overpredicted, and for the case of K^y with $\theta = 0$ it is apparent that measured diffusion is nearly zero despite the prediction being only mildly attenuated relative to the unaligned case. As shown in Figure 3.5, the analytical predictions for numerical diffusion using the combined approach of Section 3.2 correlate very well with measured diffusion from simulations, with slightly more noise in the correlation for K^{cross} than K^x or K^y.

3.4 Automated Grid Alignment in an Idealized Flow

The analytically derived diffusion coefficients (3.15)-(3.17) quantify how much the orientation of grid cells affects grid diffusion, confirming that cells aligned with the velocity field lead to zero lateral diffusion. This motivates the desire to create computational grids locally aligned with the dominant flow, a common but labor-intensive practice. This section describes an algorithm for automating this process, demonstrated on an idealized, rigid-body rotation flow field, and evaluates the efficacy of aligning the grid in terms of the measured and predicted numerical diffusion coefficients. The steps to creating a flow-aligned grid are (i) run a reference simulation on an unaligned grid for a representative period of time (ii) extract principal velocity directions from the reference simulation, (iii) construct principal

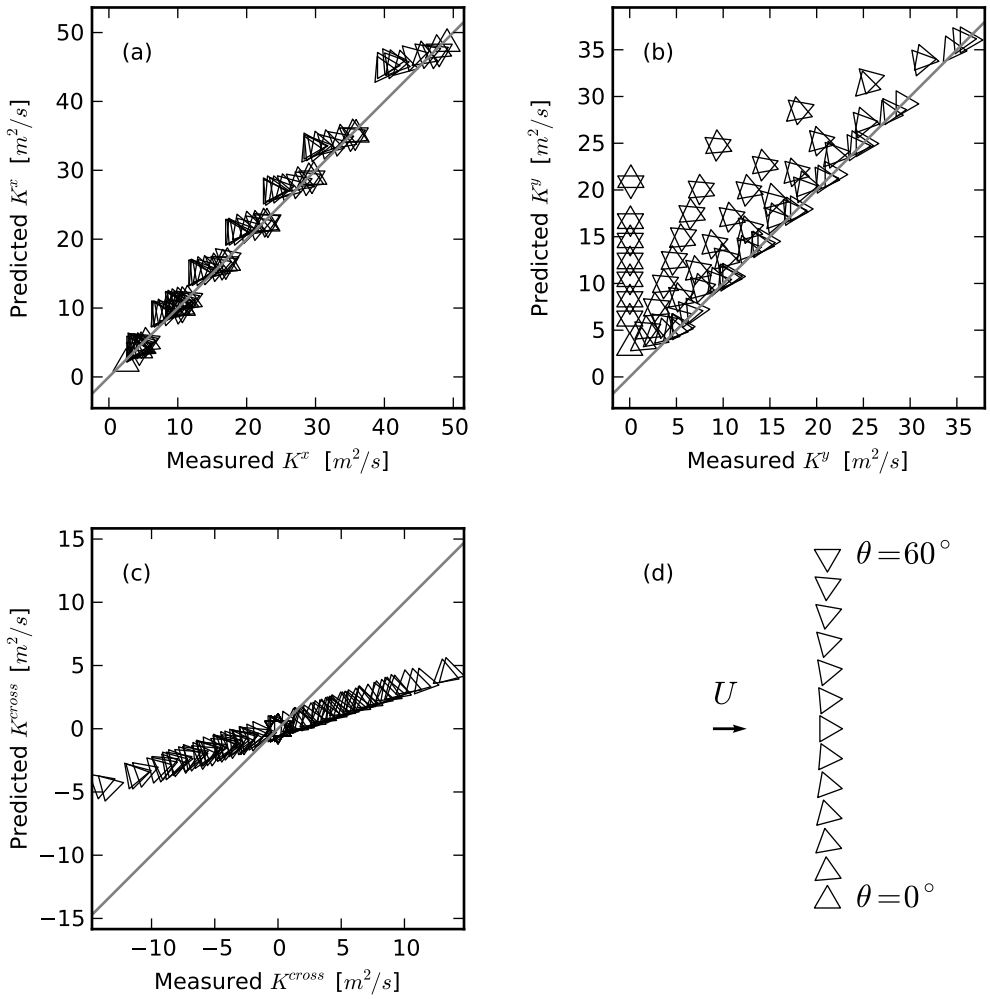

Figure 3.4: Measured versus predicted numerical diffusion, using the average of the independent modified equations. **(a)** Longitudinal diffusion **(b)** Lateral diffusion. **(c)** Cross diffusion. Grid alignment in (a)-(c) is shown by the orientation of the triangular markers, relative to a left-to-right flow as illustrated in **(d)**.

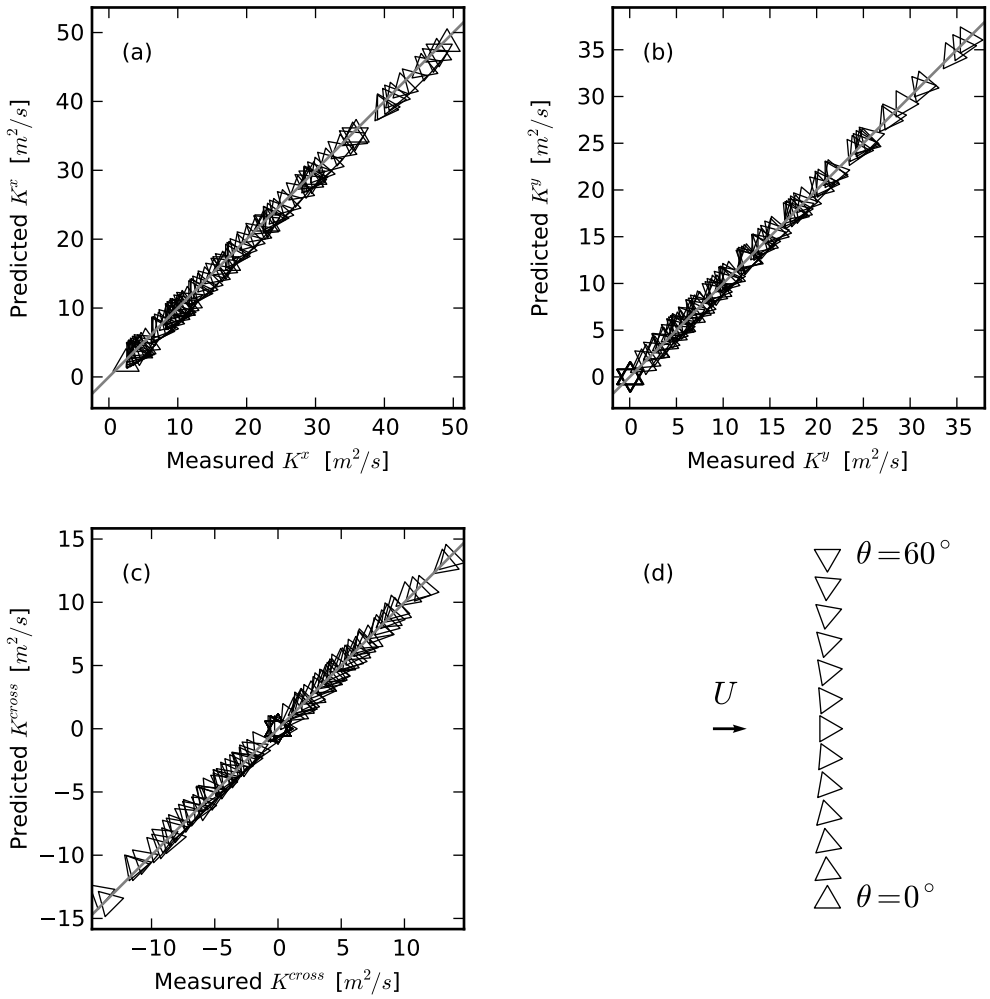

Figure 3.5: Measured versus predicted numerical diffusion using a recurrence relation to unify the two types of cells. **(a)** Longitudinal diffusion **(b)** Lateral diffusion **(c)** Cross diffusion. Grid alignment in (a)-(c) is shown by the orientation of the triangular markers, relative to a left-to-right flow as illustrated in **(d)**.

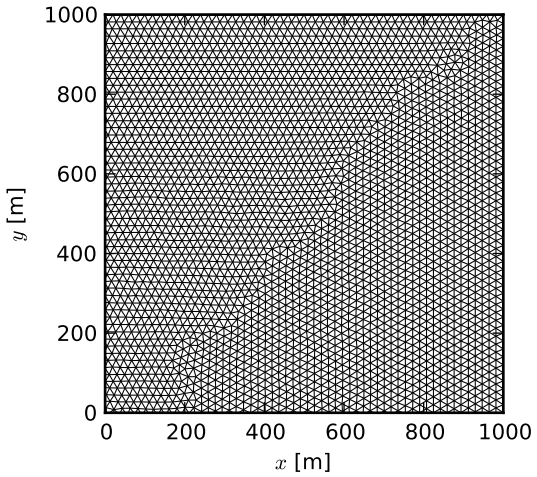

Figure 3.6: Unaligned grid for rigid-body rotation test case, upper-right quadrant

streamlines by integrating trajectories in the principal velocity field and (iv) generate a new grid in which the principal streamlines are embedded in the grid as cell edges.

The flow for the idealized flow test case is steady rigid body rotation,

$$u = \frac{2\pi}{T} y \tag{3.35}$$

$$v = -\frac{2\pi}{T} x \tag{3.36}$$

such that a full revolution is achieved over a time period T. The domain is the region $[-L, L]^2$, chosen to be square such that the boundaries do not coincide with streamlines. The unaligned grid is denoted G, with a velocity field $u_i^m(t_n)$ where $i \in \{1, 2\}$ indicates components of the 2-D depth-averaged velocity field, m an index unique to each cell, and t_n the time at step n of the simulation. For the steady rigid body rotation case, u_i^m is clearly independent of t_n, but in anticipation of the tidal case described in Section 3.5, the present description of the algorithm retains the time variability. The upper right quadrant of the unaligned grid G for the test case is shown in Figure 3.6, with $L = 1000$ m and a nominal edge length l of 20 m. In the more general case, the nominal edge length, for both the unaligned and aligned grids, is potentially variable in space, prescribed by the target edge length $l = l(x, y)$.

The reference simulation on the unaligned grid is run for a time period long enough to obtain accurate principal velocities from the flow, such as a spring-neap cycle for a tidal flow. In the steady test case the reference simulation is run for 12 h after which principal velocities are extracted from the simulation results. For the purposes of the principal velocities the

velocity field is treated as constant within each cell. The velocity correlation matrix for each cell m, over N time steps of the unaligned simulation, is computed as

$$C_{ij}^m = \overline{u_i^m u_j^m} = \frac{1}{N} \sum_{n=0}^{N} u_i^m \left(t_n\right) u_j^m \left(t_n\right). \tag{3.37}$$

λ_1^m and λ_2^m denote the eigenvalues of C^m, with $\lambda_1^m \geq \lambda_2^m$. The principal velocity direction $\hat{u}_i^m$ is then defined as the eigenvector corresponding to λ_1^m, and the ratio $e^m \equiv \lambda_2^m / \lambda_1^m$ is a measure of how well-defined the principal velocity direction is. A direct analog to a tidal ellipse can be defined with the orientation of the major axis parallel to the principal velocity direction $\hat{u}_i^m$, and the lengths of the major and minor axes defined to be λ_1^m and λ_2^m respectively, though the principal velocity captures variability at all frequencies, while tidal ellipses are specific to each tidal constituent. For the steady flow of the test case, the principal velocity directions coincide with the steady velocity field itself, within a constant factor (since the principal directions are eigenvectors).

A collection of flow-aligned curves, termed *principal streamlines*, are created by integrating trajectories in the principal velocity field. These contours are similar to streamlines, but differ in that they come from an integration of the principal velocity directions rather than the instantaneous velocity field. Two aspects of the creation of principal streamlines for grid alignment present a challenge relative to calculating streamlines in an instantaneous velocity field. The first issue is that since $\hat{u}_i$ is defined as an eigenvector, it is only determined within a constant factor. This issue is handled by assuming that when a principal streamline crosses from cell m to cell p, $\hat{u}_i^p$ is scaled by a factor $\alpha = \pm 1$ such that $\hat{u}_i^m \alpha \hat{u}_i^p > 0$ (utilizing Einstein summation notation). For a well-resolved velocity field this heuristic has been robust and all cases for which $\hat{u}_i^m \hat{u}_i^p \approx 0$ also exhibited $e^m \approx 1$, indicating that the principal velocity direction in cell m was poorly defined and the local orientation of the grid irrelevant. The second issue is that only a limited number of curves can be embedded in the aligned grid. The collection of principal streamlines should be dense enough to force most cells to be aligned but not so dense as to over-constrain the grid generation process which would lead to poor quality cells. A lower limit on the proximity between shorelines and principal streamlines, defined by a constant factor ϵ_l times the local target edge length l. enforces reasonable spacing. Tests show that $\epsilon_l = 3$ acheives a reasonable trade-off between grid quality and alignment, meaning that there is space for three to four cells between any two principal streamlines or shorelines.

An algorithm, described below, has been implemented to evenly distribute principal streamlines throughout the domain while maintaining this lower bound on proximity between curves, T is the constrained Delaunay triangulation of the shoreline, incrementally updated with principal streamlines as they are calculated. All edges, both from the shoreline and the principal streamlines, are sampled such that adjacent vertices are separated by at most the local grid scale l. $m\left(\mathbf{x}\right)$ is the grid cell in G containing point $\mathbf{x}$, $\mathbf{c}\left(t\right)$ is the circumcenter of Delaunay triangle t, and $r\left(t\right)$ the circumradius of Delaunay triangle t. T' is defined as the set of triangles $t \in T$ for which $m\left(\mathbf{c}\left(t\right)\right) \neq \emptyset$, and $e^{m(\mathbf{c}(t))} < \epsilon_e$. These are the triangles of T

which fall inside the computational domain and for which the circumcenter lies in a region where the principal velocity direction is well-defined. Observed distributions of e^m suggest a nominal choice of $\epsilon_e = 0.2$. Each principal streamline is traced according to the algorithm:

1. Find the triangle $t \in T'$ with the greatest circumradius relative to the local scale, $r(t)/l(\mathbf{c}(t))$.
2. If $r(t) < \epsilon_l l(\mathbf{c}(t))$ then terminate
3. Add the point $\mathbf{c}(t)$ to the triangulation T
4. For each α_0 in $\{+1, -1\}$:
   a) Initialize $i \leftarrow 0$, $\mathbf{x}_0 \leftarrow \mathbf{c}(t)$, $\alpha = \alpha_0$
   b) Loop:
      i. Calculate the minimum δ such that $m\left(\mathbf{x}_i + \delta\alpha\hat{\mathbf{u}}^{m(\mathbf{x}_i)}\right) \neq m(\mathbf{x}_i)$, *i.e.* take a step in the principal velocity direction to reach another grid cell.
      ii. $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_i + \delta\alpha\hat{\mathbf{u}}^{m(\mathbf{x}_i)}$
      iii. If the distance between $\mathbf{x}_{i+1}$ and any constrained edge of T is less than $\epsilon_l l$ then end loop
      iv. Insert the point $\mathbf{x}_{i+1}$ into T, with a constrained edge $\overline{\mathbf{x}_i\mathbf{x}_{i+1}}$
      v. If $\hat{\mathbf{u}}^{m(x_i)} \cdot \hat{\mathbf{u}}^{m(x_{i+1})} < 0$ then $\alpha \leftarrow -\alpha$
      vi. $i \leftarrow i + 1$

This process is repeated until the termination condition in step 2 is achieved. The set of all edges added in step 4(b)iv form the principal streamlines which will be used as constraints in the grid generation step. In the rigid-body rotation test case, the principal streamline tracing generates a set of concentric rings as shown in Figure 3.7, with a gap of approximately $3l$ between each ring, and a similarly sized gap in each ring. Finally, the set of scale-compatible principal streamlines are included as additional constraints for the generation of an aligned grid, following the method described in section D.2. The upper-right quadrant of the aligned grid for the rigid-body rotation test case is shown in Figure 3.8.

To quantify the effects on diffusion of orienting the grid with the principal streamlines, a passive tracer field is advected in both the unaligned and aligned domains. The tracer field is initialized with a Gaussian concentration distribution centered at $(L/2, 0)$ with standard deviation $\sigma = L/10$, as shown in Figure 3.9. The rigid-body rotation flow advects this tracer distribution for one full revolution period T. Comparison of the second central moments between the initial and final concentration distributions allow an estimate of the grid diffusion during one revolution. While the plume spans some finite range of angles relative to the center of the domain, for the sake of simplicity the analysis uses the Cartesian moments σ_x^2 and σ_y^2 to approximate the radial diffusion coefficient

$$K^r \approx \frac{\sigma_x^2(T) - \sigma_x^2(0)}{T} \tag{3.38}$$

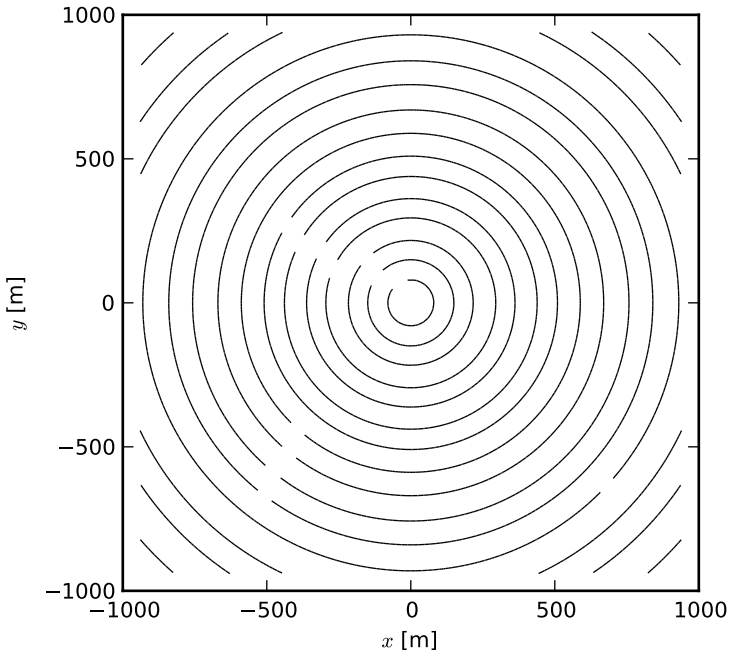

Figure 3.7: Principal streamlines extracted from the unaligned simulation

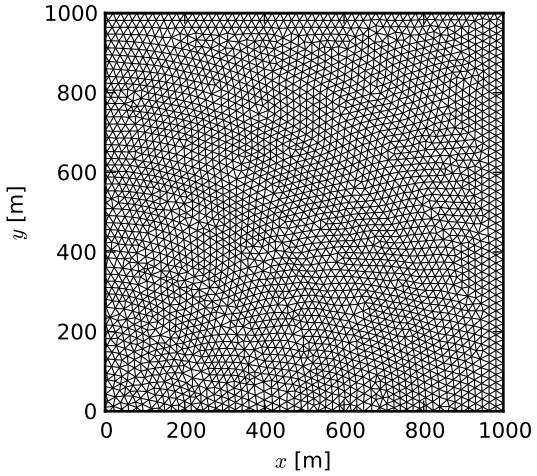

Figure 3.8: Upper-right quadrant of the aligned grid for the rigid-body rotation test case

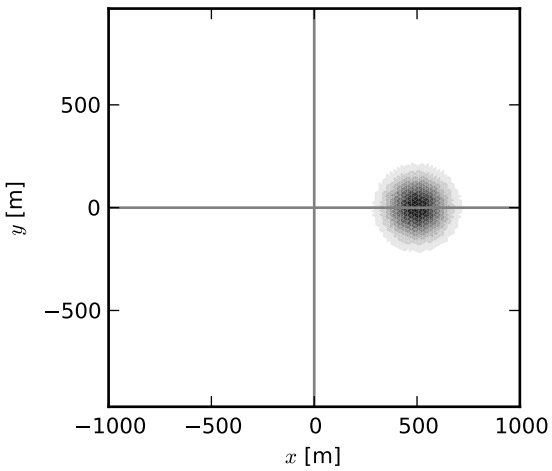

Figure 3.9: Initial scalar field on the unaligned grid. Grayscale graduations indicate 0.1 contours, with a maximum concentration of 1.0

and tangential diffusion coefficient

$$K^{\theta} \approx \frac{\sigma_y^2(T) - \sigma_y^2(0)}{T}. \tag{3.39}$$

In the case of the unaligned grid, this leads to an estimate of $K^r = 0.081 m^2 s^{-1}$, and $K^{\theta} = 0.145 m^2 s^{-1}$. For the aligned grid $K^r = 0.020 m^2 s^{-1}$ and $K^{\theta} = 0.134 m^2 s^{-1}$. Evaluating the analytical diffusion coefficients at the center of the release, a perfectly aligned grid should have zero radial diffusion and $K^{\theta} = 0.162 m^2 s^{-1}$. A "perfectly" unaligned grid, taking $\theta = \pi/6$, leads to $K^r = 0.105 m^2 s^{-1}$ and $K^{\theta} = 0.137 m^2 s^{-1}$. The along-flow diffusion coefficient K^{θ} is comparable in both cases, and falls within the expected range based on the analytical coefficients. While theory suggests that along-flow tangential diffusion should be slightly greater on the aligned grid, the results show a slightly attenuated tangential diffusion. The major difference is in the radial, or cross-flow diffusion. Here the observed diffusion on the unaligned grid is 77% of the predicted maximum, while the aligned grid sees only 19% of the predicted maximum.

3.5 Application to a Physical System

To further explore the effects of grid alignment and grid diffusion, this section describes a physically realistic tidal simulation of San Francisco Bay, with and without grid alignment, and a comparison of the resulting salinity fields. The SUNTANS [28] code in hydrostatic

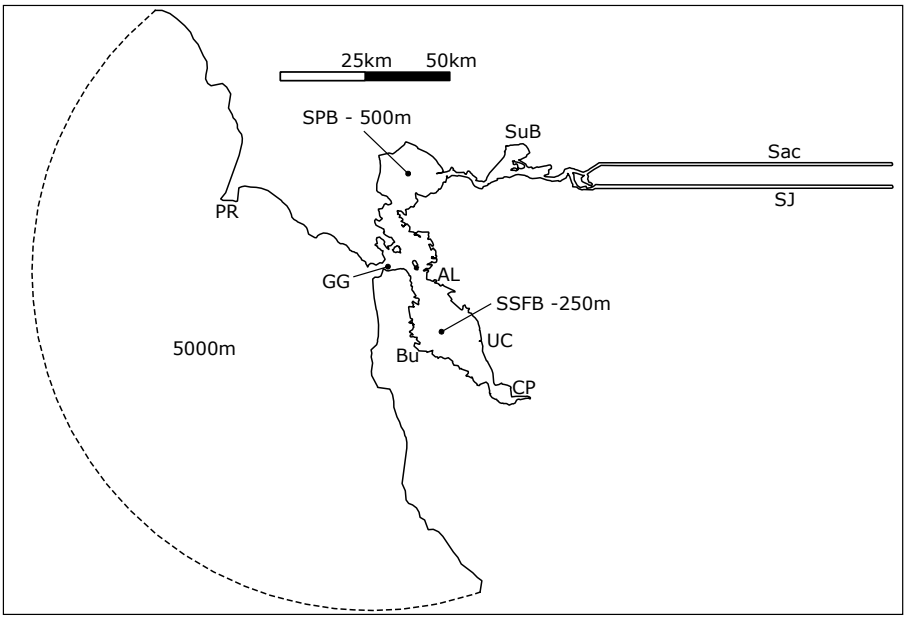

Figure 3.10: Computational domain for the San Francisco Bay model runs. Nominal grid edge-length is 5000m in the coastal ocean, 500m in North San Francisco Bay, and 250m in South San Francisco Bay. For computational efficiency and boundary condition simplicity the Sacramento-San Joaquin Delta has been replaced by a pair of long dissipative channels. The dashed line represents the tidal ocean boundary condition.

mode is utilized to solve the 3-D Reynolds-averaged primitive equations with the Boussinesq approximation in a rotating frame of reference. In accordance with the hydrostatic approximation, the vertical velocity field w is solved by continuity. Horizontal numerical diffusion is assumed to be on the order of or larger than the physical horizontal eddy viscosity and the latter is omitted from the equations of motion. The equations of motion are then

$$\frac{\partial u}{\partial t} + \nabla \cdot (\mathbf{u}u) - fv = -\frac{1}{\rho_0}\frac{\partial p}{\partial x} + \frac{\partial}{\partial z}\left(\nu_V \frac{\partial u}{\partial z}\right) \tag{3.40}$$

$$\frac{\partial v}{\partial t} + \nabla \cdot (\mathbf{u}v) + fu = -\frac{1}{\rho_0}\frac{\partial p}{\partial y} + \frac{\partial}{\partial z}\left(\nu_V \frac{\partial v}{\partial z}\right) \tag{3.41}$$

$$\nabla \cdot \mathbf{u} = 0 \tag{3.42}$$

Where u is eastward velocity, v northward velocity, $f = 2\omega \sin\phi$ the Coriolis parameter with angular velocity of the Earth ω and latitude ϕ, ν_V the vertical eddy viscosity derived from a two-equation turbulence model (Mellor-Yamada 2.5 [47]), and ρ_0 a constant reference density ($\rho_0 = 1000$ kg/m^3). The hydrostatic pressure p is defined by $\frac{\partial p}{\partial z} = -(\rho_0 + \rho)\, g$, where $\rho = \beta s$ is the density anomaly due to salinity s and coefficient of contractivity β. These equations are discretized on a prismatic finite-volume grid comprising unstructured triangles in the horizontal and structured z-levels in the vertical.

Two nearly identical simulations of San Francisco Bay are carried out, differing only in the alignment of grid cells. Model outputs from both simulations are compared to observations of depth-averaged salinity in the study area, South San Francisco Bay. The model domain, shown in Figure 3.10, extends to the south through South San Francisco Bay (SSFB) as far as Calaveras Point (CP), beyond which the basin devolves into a network of tidal sloughs. In the northern end of the bay, the domain includes San Pablo Bay (SPB), Suisun Bay (SuB), and a pair of false deltas each with a width of 1km. The false deltas are surrogates for the Sacramento (Sac) and San Joaquin (SJ) Rivers, and have hypsometric curves matched to accurate Delta bathymetry, evenly split between the two channels. The ocean boundary is an arc of radius 100km, centered on the Golden Gate (GG). The spatial resolution is nominally 5km in the coastal ocean, gradually decreasing to 250m in South San Francisco Bay, leading to approximately 26,000 2-D cells and 430,000 3-D cells. The vertical dimension is discretized with 40 z-levels, with cell thicknesses ranging from 0.5m at the free-surface to 10m at a depth of 100m. Depths below 100m are lumped into a single layer. The time step is 60s, and the scalar transport routines utilize substepping such that the overall hydrodynamic time step is limited only by the propagation speed of wetting and drying fronts and internal waves.

The ocean boundary is forced with zero velocity and a free-surface height set by observed tides from Point Reyes (PR). In order to reach phase and amplitude agreement at the Golden Gate, the ocean tides are advanced in time by 17.7 minutes and attenuated by a factor of 0.97. Winds in the model were spatially uniform, as measured at Alameda (AL). Wind stress is calculated using the drag law of Large and Pond [41]. Freshwater inputs are assembled from a wide range of data sources, including Delta flows from the Delta Dayflow [21] measurements, USGS gauged streams, and ungauged streams correlated to nearby gauges based on drainage area. Discharges from wastewater treatment plants are also included, based on daily data for the larger sources and monthly data otherwise. Evaporation and precipitation are spatially uniform, with evaporation based on a monthly climatology at Burlingame (Bu) and hourly precipitation from a gauge in Union City (UC). In order to avoid introducing a bias towards the validation of one grid over the other, calibration is limited to the already mentioned phase and amplitude tuning at the Golden Gate. The simulations are initialized with a quiescent velocity field and a salinity distribution extrapolated from salinity observations [78] along the thalweg of the bay. Following a spin-up period starting on August 1, 2008 and running for 90 days, the analysis period is October 30, 2008 through January 22, 2010.

While the model domain includes the majority of the San Francisco Bay system, the study area for the simulations is South San Francisco Bay (Figure 3.11), a 45km long basin forming one of two branches in the larger San Francisco Bay system. Most of this embayment exhibits a channel-shoal morphology, with channel depths of 15-20m and shoal depths of 0-4m. The pronounced channel-shoal geometry leads to significant lateral shear and lateral salinity gradients.

In the present case of San Francisco Bay, a tidally dominated estuarine and coastal ocean domain, the depth-averaged principal velocities are essentially stationary over time-scales greater than the spring-neap cycle, such that the representative period for the reference

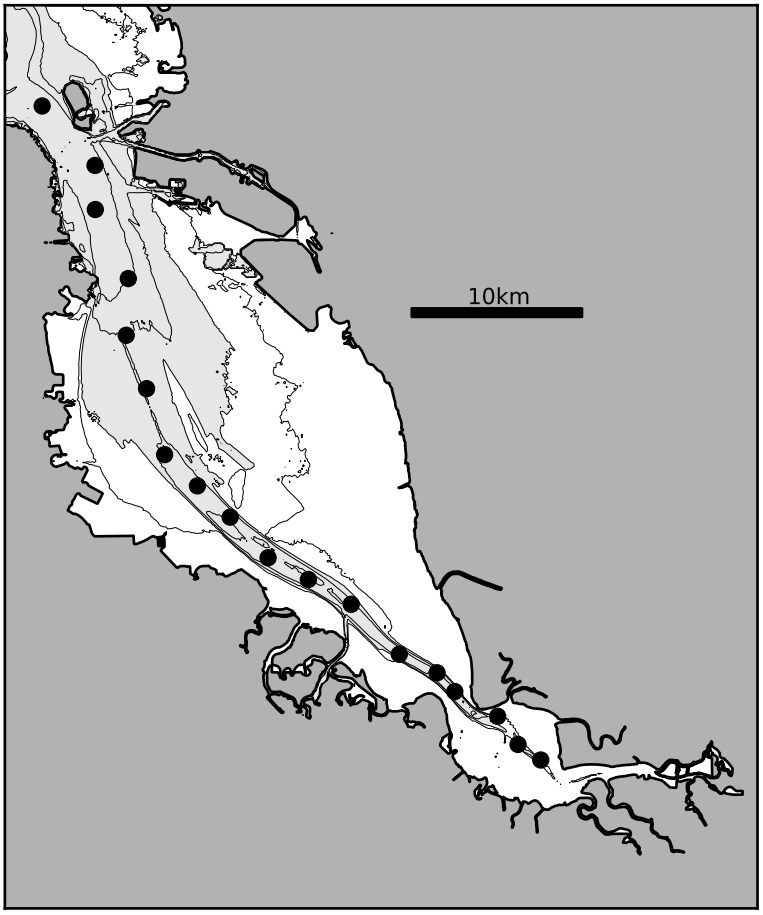

Figure 3.11: South San Francisco Bay bathymetry used in the physically realistic model runs. Contours at -15, -10, -5 (shaded) and -2.5m (NAVD88). Transect sampling locations are shown as black dots, with the start of the transect (0km) at the southern end.

simulation should span at least one spring-neap tidal cycle. Model output from the unaligned simulations over the first 90 days of the analysis period is used to establish a representative principal velocity field. Figure 3.12 shows the principal velocity ellipses for a portion of the domain, resampled to a coarse regular grid for clarity as well as normalized such that the length of the major axis is constant. Principal streamlines, consistent with the grid resolution function $l(x, y)$, are depicted in Figure 3.13. The flow-aligned grid, with identical resolution and shoreline data, is then constructed such that the flow-aligned contours are embedded in the grid as edges of cells (shown in Figure 3.15). Section D.2 gives further details on the mesh generation algorithm. Figure 3.14 compares distributions of cell-center spacing and cell skewness between the original unaligned grid and the aligned grid. Cell skewness is defined as $1 - A_{eq}/A$ where A is the area of the cell and A_{eq} is the area of an equilateral triangle with the same circumcircle, such that 0.0 indicates a perfectly equilateral cell and 1.0 a zero-area

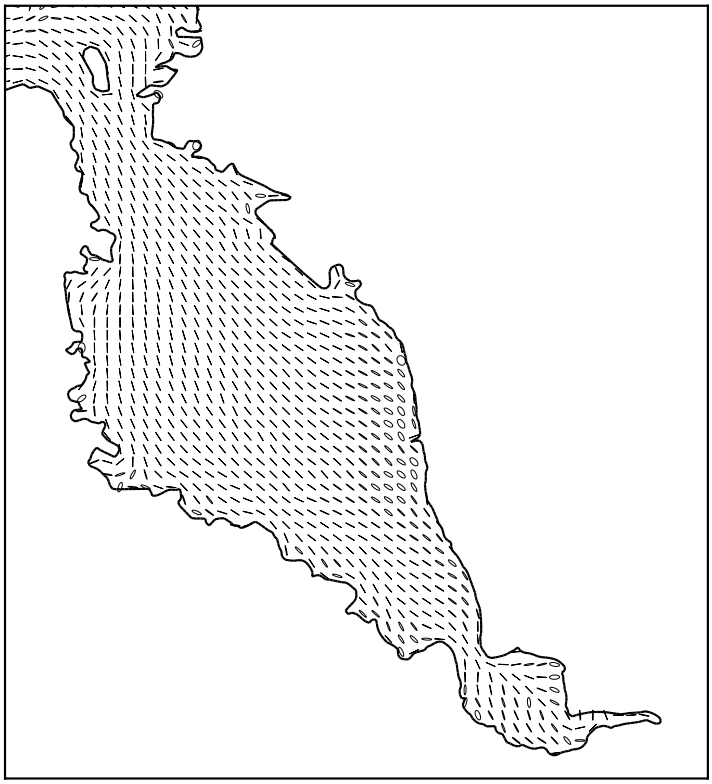

Figure 3.12: Principal velocity ellipses in the vicinity of the study area, derived from model output from the unaligned grid. Ellipses are calculated for all cells, but are shown interpolated onto a coarse regular grid and normalized by major axis magnitude for clarity.

sliver. These distributions reflect the entire domain, but since the majority of the cells are in the high-resolution South San Francisco Bay area, it is this area with a nominal 250m resolution that dominates the histograms. Cell-center spacings are similar between the two grids, but skewness shows an increase of moderately low-quality cells in the aligned grid, most likely due to the greater number of constraints on cell placement.

To quantify the effects of grid alignment in the San Francisco Bay case, (3.15)-(3.17) are evaluated to obtain K^x, K^y, and K^{cross} for each grid cell, depth-averaged and averaged over the first 90 days of the analysis period. While the method of Burchard and Rennau [7] would likely give a more accurate measure of the total numerical mixing for a specific scalar field, our approach directly evaluates the numerical diffusion coefficients rather than calculating decay of scalar variance on a per-scalar basis (though a diffusion coefficient may be obtained

Figure 3.13: Flow-aligned contours, traced from the principal velocities, ready for input to the re-gridding process. The separation between adjacent contours scales with the grid resolution, leading to sparser contours in the north and west.

by normalizing by the gradient of the scalar concentration). Our method also retains the distinction between lateral and longitudinal diffusion, which is helpful in interpreting the effects of numerical mixing on the synoptic evolution of the scalar field. The alignment angle θ was taken from the cell edge most closely aligned with the instantaneous velocity vector, modulo $\pi/3$. Since the expressions of Section 3.2 give the diffusion tensor in a coordinate system parallel to the instantaneous velocity, it is necessary to rotate each instantaneous diffusion tensor into a common coordinate system before averaging over time and depth. The principal velocity direction of each cell is used as the common coordinate system, such that the resulting diffusion tensor gives K^x as longitudinal diffusion, parallel to principal flow and K^y as lateral diffusion, perpendicular to principal flow. Figure 3.16 shows the resulting spatial distributions of longitudinal K^x and lateral K^y. Volume- and time-averaged diffusion coefficients for the study area, roughly defined as the southern two thirds of the region shown in Figure 3.16, over days 91-180, are $K^x = 12.9^{m^2}/_s$, $K^y = 6.5^{m^2}/_s$, and $K^{cross} = 0.25^{m^2}/_s$ for the unaligned grid. For the aligned grid, the same calculation gives

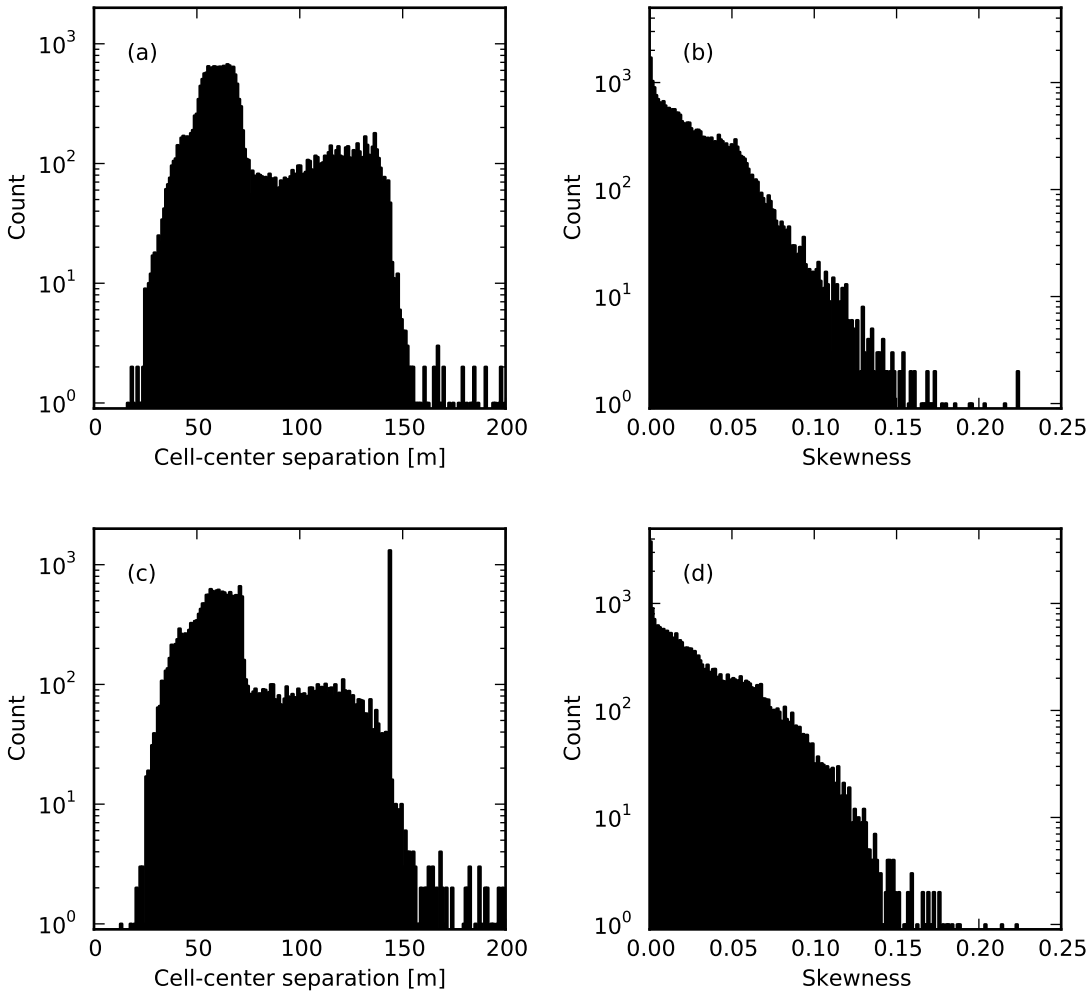

Figure 3.14: **(a)** Distribution of cell-center spacing for the unaligned grid. **(b)** Distribution of skewness for the unaligned grid. **(c)** Distribution of cell-center spacing for the aligned grid. **(d)** Distribution of skewness for the aligned grid.

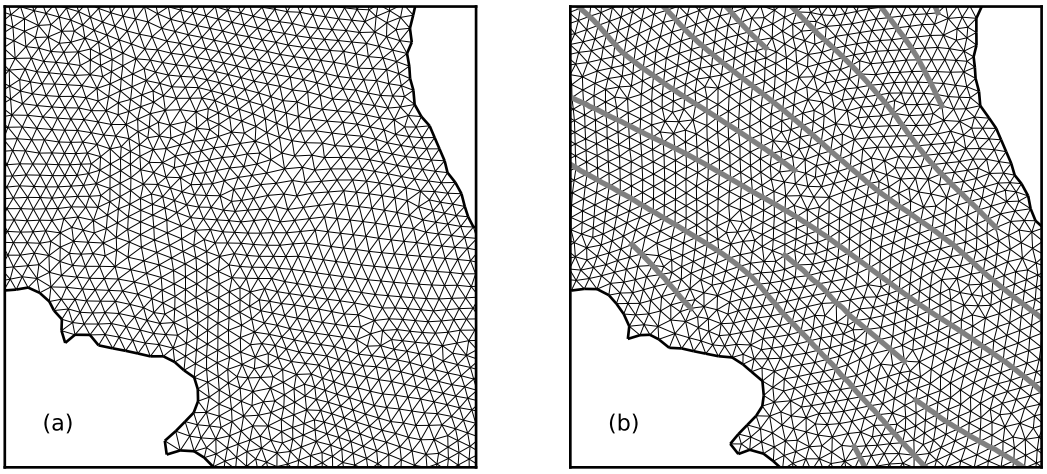

Figure 3.15: Details of the computational grids, towards the southern end of the study area. **(a)** The unaligned grid. **(b)** The aligned grid with flow-aligned contours superimposed.

$K^x = 12.8 m^2/s$, $K^y = 3.2 m^2/s$, and $K^{cross} = 0.1 m^2/s$, showing a significant decrease in lateral numerical diffusion while the longitudinal component is essentially unchanged. While the mean lateral diffusion is improved in the aligned grid, the spatial distribution of the diffusion coefficient shows much more variability than the unaligned grid. Low-diffusion regions occur along embedded contours and high-diffusion regions between the contours. This small-scale heterogeneity may lead to anomalous results when the length scale of a tracer distributions is small relative to the inter-contour length scale.

Validation of unaligned and aligned model output

While the focus of the comparisons is the difference between the two models rather than specific validation against observations, included below is a comparison of both models to observed salinity in South San Francisco Bay. A full treatment of model predictive skill would necessarily include evaluations of forcing data, the turbulence closure and bathymetry, and is beyond the scope of this chapter. Model-data comparisons of the longitudinal salinity field are shown in Figures 3.17 and 3.18. Each set of salinity observations was collected by the USGS Polaris [78] over the course of a particular day, starting at the southern end of San Francisco Bay. Model output has been interpolated in time to match the respective time of each observation along the transect. Both the model and observations resolve the water column at approximately 1m resolution, but have been depth-averaged for simplicity.

Considering only the portion of the transects that fall within the high resolution portion of the grid (the region depicted in Figure 3.11), bias (mean error), the Pearson product-moment

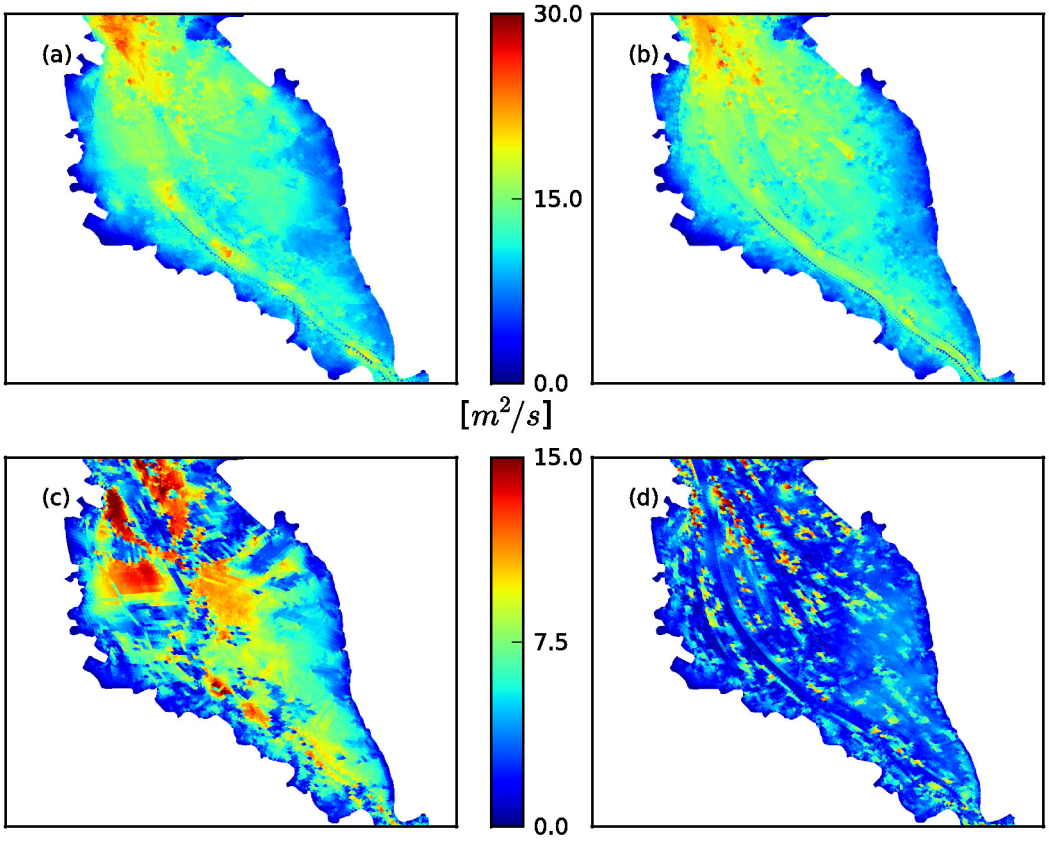

Figure 3.16: Distribution of depth- and time-averaged longitudinal and lateral diffusion in $m^2 s^{-1}$. **(a)** K^x for unaligned grid. **(b)** K^x for aligned grid, showing minimal decrease in numerical diffusion relative to the unaligned grid. **(c)** K^y for unaligned grid. **(d)** K^y for aligned grid, showing a significant decrease in numerical diffusion.

correlation coefficient, the index of agreement [83], and the RMS error in the longitudinal salinity gradient are calculated over all transect observations available for the 450 day analysis period. These results and definitions of the respective metrics are summarized in Table 3.1. While the unaligned grid has a marginally smaller bias, both the unaligned and aligned simulations under-predict salinity and the difference in bias is small relative the standard deviation of the errors. The aligned grid, however, has better correlation with observations and a slightly better prediction of longitudinal gradients.

The observed transects, such as those shown in Figures 3.17 and 3.18, are available only once a month and are timed to the phase of the tide. As such, there is the possibility of a sampling bias relative to the true mean salinity distribution. Figure 3.19 shows predicted long-term average salinity distributions sampled every 20 minutes. The difference in longi-

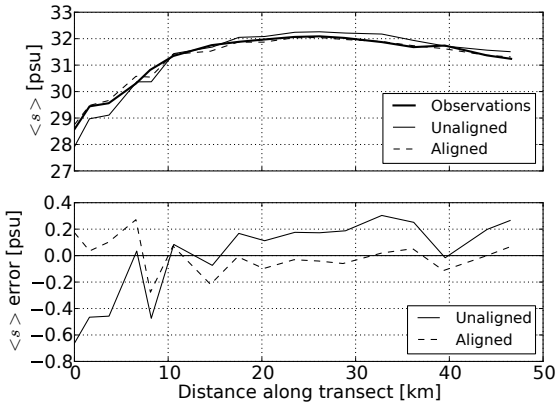

Figure 3.17: Comparison of unaligned and aligned model output and observations of salinity along a longitudinal transect in South San Francisco Bay (shown in Figure 3.11), from November 5, 2008. **(upper)** Depth-averaged salinity **(lower)** Absolute error of model output compared to observations.

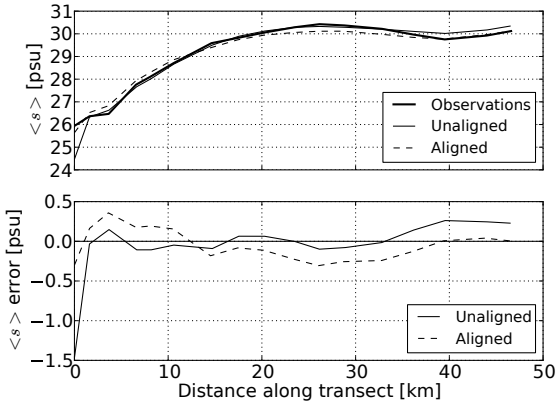

Figure 3.18: Comparison of unaligned and aligned model output and observations of salinity along a longitudinal transect in South San Francisco Bay (shown in Figure 3.11), from January 1, 2009. **(upper)** Depth-averaged salinity **(lower)** Absolute error of model output compared to observations.

	Definition	Unaligned	Aligned
Mean error [psu]	$\langle s_i^n - \hat{s}_i^n \rangle$	-0.36 ± 0.69	-0.42 ± 0.49
Correlation coefficient (r)	$\frac{\mathrm{COV}(s,\hat{s})}{\sigma_s \sigma_{\hat{s}}}$	0.945	0.958
Index of agreement	$1 - \frac{\sum\left(\hat{s}_i^n - s_i^n\right)^2}{\sum\left(\left\|\hat{s}_i^n - \langle s \rangle\right| + \left\|s_i^n - \langle s \rangle\right|\right)^2}$	0.966	0.971
RMS gradient error [psu]	$\left\langle \left[\left(s_i^n - \hat{s}_i^n\right) - \left(s_{i-1}^n - \hat{s}_{i-1}^n\right) \right]^2 \right\rangle^{\frac{1}{2}}$	0.29	0.22

Table 3.1: Error statistics for comparison to longitudinal salinity transects.

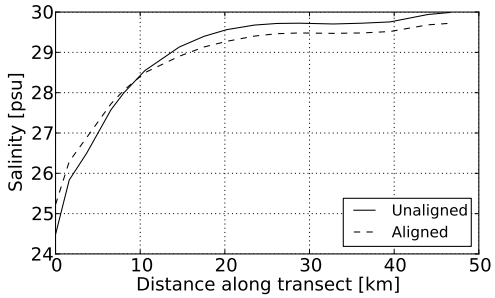

Figure 3.19: Time and depth averaged salinity for the unaligned and aligned grids, along the thalweg of South San Francisco Bay. Decreased longitudinal gradients in the aligned grid model are consistent with an increase in longitudinal diffusion.

tudinal salinity gradients between the two simulations is persistent and significant, though of similar order to the error relative to observations.

3.6 Discussion

A priori estimates of the diffusion tensor have been both verified for idealized simulations and estimated for physically realistic simulations. A distinct advantage to estimation of the diffusion tensor compared to other error estimates is that it communicates the error back to the modeler in the same language as the original governing PDE. In this way the effects of numerical diffusion and how it compares to and interacts with physical dispersion processes is clearer.

Perhaps the greatest advantage of using grid alignment to combat numerical diffusion is that it is code agnostic. Any simulation code which uses a prismatic triangular grid can potentially benefit from an aligned grid, though of course the specifics of the numerical diffusion will vary by choice of numerical scheme. The fact that grid generation in a non-adaptive simulation is decoupled from the computation itself allows for greater reuse of the grid generation methods across simulation frameworks. Of course a key step in the application of these methods to realistic simulations is the creation of a high quality grid

which enforces cell alignment. While the method detailed in section D.2 is able to produce reasonably aligned grids, there are always trade-offs. In particular, the price of enforcing cell alignment is a slight decrease in cell quality in terms of cell-center spacing, which in turn affects the maximum scalar advection time step. 90 day simulations on the unaligned grid average 20.5h wall-time, compared to 23.1h for the aligned grid (with both simulations using 16 cores of an HPC cluster). Still, given the linear dependence of K_{lat} on the grid size l, a simple reduction of grid scale to achieve the same 50% reduction in lateral numerical diffusion would require four times more cells and potentially twice as many time steps. While switching to higher-order advection schemes is almost certainly a more efficient way to achieve the same decrease in numerical diffusion, the two approaches are complementary, and the results of using a flow-aligned grid in conjunction with a high-resolution advection scheme can be expected to perform better than either method alone. Another trade-off is in heterogeneity of the lateral diffusion coefficient. Grid alignment, as implemented here, increases quality of many cells at the expense of a small number of cells. The misaligned cells tend to fall into lines between the embedded contours, creating streaks of high lateral diffusion. Small scalar plumes may experience erratic diffusion until they are large enough to sample regions of both high and low lateral diffusion.

The primary observation from comparing the salinity gradients predicted in the unaligned and aligned simulations is an effective *increase* in longitudinal diffusion in the aligned grid simulations. While estimates of the longitudinal numerical diffusion are essentially equal between the two grids, the interaction of lateral shear and lateral diffusion leads to an effective increase in the longitudinal diffusion, *i.e.* the shear flow dispersion scaling of [23], $K_{lon} \sim U^2h^2/K_{lat}$. In physical systems this mechanism is often responsible for the bulk of longitudinal dispersion, and comparisons between the two simulations show this scaling at work, translating the decreased lateral dispersion of the aligned grid into a more efficient longitudinal shear dispersion process. Enhanced longitudinal diffusion is evidenced by the decreased salinity gradients in Figure 3.19, most notably around 10km along the transect, and in fact the salinity gradients in the aligned case more closely track the observations. Note, however, that comparisons between either simulation and the observations should be viewed through a wide-angle lens, as the quality of the validation is affected by a broad set of factors, ranging from the quality of available data for freshwater flows to the treatment of unstable stratification by the turbulence closure. It is the comparison between the two simulations, in particular the difference in mean gradients shown in Figure 3.19, which best demonstrates the effect of grid alignment.

Only velocity fields where the principal velocities are statistically stationary have been considered. This is a valid assumption for strongly tidal embayments, but there is a clear extension to non-stationary flows with an adaptive re-gridding process. Compared to related forms of adaptive mesh refinement, such as the algorithm proposed by Marcuzzi *et al* [46] in which directionality in the error field adaptively controls anisotropy in a finite element mesh, the numerical schemes with which this chapter is concerned place much tighter constraints on the geometry of the grid. These constraints translate to more computationally intensive grid generation procedures, and while our current implementation would be inefficient for

frequent, global re-gridding steps, localized re-gridding in areas with non-stationary principal velocities would be feasible even with the current implementation. In applications such as persistent, meandering ocean currents or morphological hydrodynamic models, the re-gridding costs could be amortized over relatively long computational periods.

3.7 Conclusions

Numerical diffusion for first-order upwind advection on a regular triangular grid has been analytically derived and shown to agree with estimates from idealized simulations. Two distinct classes of cells, when analyzed independently, lead to two distinct modified equations. The sum of these equations recovers consistency but overestimates numerical diffusion. An analysis that takes into account the coupling of the two cell types results in a numerical diffusion tensor that correctly quantifies the known behavior that flow-aligned grids eliminate lateral numerical diffusion.

In the context of a physically realistic simulation, we have presented a largely unsupervised process for creating a flow-aligned computational grid, and compared model results from an unaligned grid and an aligned grid. The mean predicted lateral numerical diffusion within the study area, South San Francisco Bay, was reduced by a factor of two. Longitudinal salinity gradients were smaller on the aligned grid, consistent with an increase in longitudinal diffusion. Considering the interaction of lateral shear and diffusion, this increase in longitudinal diffusion is in turn consistent with a decrease in lateral diffusion. Model-observation comparisons confirm that the salinity gradient is better predicted by the aligned grid model.

Chapter 4

Coupling of Sea Level Rise, Tidal Amplification and Inundation

Abstract

With global sea level rising it is imperative to quantify how the dynamics of tidal estuaries and embayments will respond to increased depth and newly inundated perimeter regions. With increased depth comes a decrease in frictional effects in the basin interior and altered tidal resonance. When inundation is considered, sea level rise also causes an increase in planform area, tidal prism, and potentially significant frictional effects in newly inundated areas. A high resolution 2-D hydrodynamic model of San Francisco Bay, California, US is employed to investigate the coupling between ocean forcing, tidal dynamics and inundation. This domain comprises two basins with distinct tidal characteristics. Multiple shoreline scenarios are simulated, ranging from a fully leveed scenario, in which tidal flows are limited to present day mean higher high water shorelines, to a freely inundating simulation, in which all topography is allowed to flood. These simulations show that increasing the mean sea surface elevation while preserving original shorelines leads to additional tidal amplification. However, the inclusion of inundated areas, flooded by the increased sea surface elevation, introduces frictional, intertidal regions around the perimeter of the basins which serve as energy sinks for the incident tidal wave. This dissipation more than compensates for the reduction in friction due to a deeper basin interior, and net tidal amplification in most areas is ultimately lower in the sea level rise scenarios. Details of the changes in dynamics are analyzed in terms of energy flux and tidal phasing, which show a shift to a more progressive-wave, dissipative environment with rivers and tidal sloughs becoming major energy sinks. The near standing-wave southern arm of San Francisco Bay couples more strongly back to

Under consideration for publication in *J. Phys. Oceanogr.*: Holleman, R. C. and M. T. Stacey. Coupling of Sea Level Rise, Tidal Amplification and Inundation.

the central portion of the bay, in contrast to the progressive-wave northern reach of the bay. Overtide generation, specifically the M4 wave, is also found to vary significantly between scenarios and is a non-negligible factor in determining net changes to high water elevations.

4.1 Introduction

Among the many concerns related to recent and predicted climate change is the trend of rising sea levels. At global scales studies such as [19] show global sea level rising approximately 0.20 m in the past 100 years, and predictions for sea level rise in the the next 100 years range from 0.5 m to 1.4 m [58]. Adding to trends in global mean sea level, observations in some regions also show a pattern of increasing tidal amplitudes and increasing non-tidal variations in sea surface height. The combined effects of sea level rise and potentially increasing tidal ranges will have far-reaching impacts on coastal inundation as many low-lying areas either become uninhabitable or require massive mitigation measures to fend off higher sea levels. As inundation is a consequence of peak sea level, not mean sea level, it is essential to consider both coastal ocean trends in mean sea level and how those trends will couple with local tidal dynamics to affect the peak sea surface height adjacent to areas in danger of flooding. At the same time, inundated areas add to the available tidal prism and the overall tidal energy dissipation, such that one must consider the whole system in order to accurately capture the coupling of tidal dynamics and inundation.

Nearshore regions are also influenced management decisions, which in turn rely on predictions of flooding and sea level rise. Relevant management actions fall into two main categories. The first category, shoreline "hardening", describes construction of hydrodynamic barriers such as concrete sea walls or levees. These projects may be motivated by flood risks, "reclamation" of shallows into dry land, or creation of ponds for salt harvesting. In many areas shoreline hardening is widespread and significantly alters the dynamics of the basin, such as in San Francisco Bay, California where upwards of 85% of historic marshlands have been filled or fundamentally altered [17]. Shoreline hardening decreases the tidal prism and often leads to greater tidal amplification. The second, generally opposite, category of shoreline modifications could be labeled shoreline "softening", but since it is often attempting to reverse the effects of earlier hardening projects, these actions may also be termed restorations. Typical restoration projects include breaching old levees or dredging new channels. Returning tidal action to these areas serves a number of purposes including reestablishing highly productive marsh ecosystems, improving water quality, and even mitigating flood risks. These projects often increase the area available to tidal action, and introduce softer, natural shorelines and slough networks which are effective at dissipating tidal energy. Between the growing number of restoration projects and the potential for widespread sea level rise mitigation efforts, it will be important in the coming century to quantify the range of shoreline modifications and the effects those changes will have both on localized inundation and basin-wide tidal dynamics.

Tidal amplification

Variations in tidal range within a basin come primarily from four physical processes: standing wave resonance from reflection of the incident tidal wave, frictional effects, converging geometry (i.e. cross-sectional area decreases upstream), and inertial effects [79]. Resonance and converging shoreline geometry lead to an increase in tidal amplitude away from the open ocean boundary of a basin, while diverging shorelines and friction lead to attenuation. Inertial effects are typically negligible and are ignored in most analyses which do not target shallow macrotidal systems.

Standing wave resonance is easily understood in terms of a simple prismatic channel, in which the tides can be described by the superposition of an incident wave and a reflected wave. Standing tidal waves occur when the incoming tidal wave is fully reflected, such as in a non-frictional basin with a non-dissipative landward boundary, and in this case the superposition of the two waves simplifies to

$$\eta = \frac{H}{2} \cos kx \cos \omega t, \tag{4.1}$$

where η is the space- and time-varying free surface perturbation, x the distance from the close end of the basin, ω the angular frequency of the tidal forcing, and H the tidal range at the close end of the estuary. Given the length L from the closed, reflective landward boundary of a basin to the open, ocean-forced mouth, the amplification is simply $\alpha = \sec kL$. In systems where L is near a quarter-wave node the amplification approaches a resonant peak, such as the famous tides of the Bay of Fundy. The resonant period of a basin depends on the phase speed of the tidal wave; if the resonant period is altered towards the dominant period of the tidal forcing one would expect that the net result would be an increase in tidal range. However, in most basins, locations which are close enough to the closed end of the basin for the reflected wave to retain significant amplitude (and thus behave as a standing wave system) are also shorter than the quarter-wave resonance length. In these cases an increase in the tidal wave phase speed will push a system farther from resonance. A fundamental parameter characterizing the degree to which a standing wave is present is the velocity phase lead ϕ, which is defined here as the phase offset between peak flood velocity and peak high water for a specific tidal constituent. A progressive wave tide in which reflected energy is vanishingly small will have a phase lead approaching zero, while a standing wave system will see $\phi \approx 90°$. The velocity phase lead provides a useful local estimate of basin-wide tidal energy dynamics. Aside from diagnosing standing wave or progressive wave dynamics, ϕ is also relevant for residual scalar transport and sediment dynamics, since Stokes transport is greatest for $\phi = 0$ and negligible when $\phi = 90°$.

The geometry of a basin can also lead to amplification or attenuation through converging or diverging shorelines. [79] investigated the competing roles of convergence (both in depth and width), friction and reflection. He found that in sufficiently long, deep and converging estuaries the amplifying effects dominate and tidal amplitude increases towards the head of the estuary. Shallow converging channels are dominated by friction, resulting in an attenuated tidal range landward of the mouth. In broad terms, he found that the reflected

wave, if one exists, affects roughly the landward third of the basin, as the reflected wave is both dissipated by friction and attenuated by diverging shorelines as it travels seaward. In strongly converging channels the phase lead of the peak flood velocity ahead of high water approaches 90°, independent of the presence of a closed landward boundary. [66] and [79] term this condition an *apparent* standing wave.

The friction term is often the most challenging term in analytic treatments of tidal basins due to nonlinearity of the quadratic drag relation. Using a linear combination of two approaches to the friction term, [8] derived an analytic model of basin amplification applicable over a wide range of basin parameters. The resulting expressions allow a classification of basins based on how the actual depth compares to the ideal depth (i.e. producing zero amplification) and the critical depth (i.e. producing maximum amplification). The model also includes prognostic equations for the phase lead and non-dimensional amplification factor, as a function of basin geometry, tidal forcing, friction, and mean depth. The flexibility of the input parameters and wide range of behaviors which can be predicted make this model particularly relevant for sea level rise forecasts. Section 4.54.5 applies it to a portion of the study area and compares analytic and numerical predictions for the M2 tide in order to understand the degree to which the analytic approach captures the necessary physics.

Inundation and tides

In regions with considerable inundatable area, the effects of inundation on tidal dynamics must also be considered. Higher sea surface heights allow the tides to access a greater tidal prism. The inundated areas are almost universally very shallow, and while the increased tidal prism may increase tidal velocities seaward of the inundated region, the shallow expanses have an overall dissipative effect on the tides. This additional dissipation tends to decrease reflection and mitigate some fraction of the sea level rise. While there is a well-established body of work on inundation resulting from storm surge, the literature on energetics of inundation coupled with tides is relatively sparse. Depending on the characteristics of the newly wetted area, the amount of dissipation of tidal energy at the perimeter may decrease or increase. At one end of the continuum one could imagine a basin with shear vertical walls at the original mean higher high water (MHHW) contour. As the sea surface rises shallows which were originally intertidal become subtidal and less frictional. Overall, the perimeter becomes more reflective, leading to a greater tidal range. At the other end of the continuum, one can imagine that the region that was originally supertidal is instead flat and littered with drag-inducing features. In this case, the newly inundated areas are dissipative and tend to absorb the energy of the incoming tidal wave. Flows within the perimeter would shift towards a frictional regime, and flows in the interior of the basin would shift towards a progressive wave as incident tidal energy is absorbed in the intertidal areas.

Despite continual progress in analytic solutions to tidal propagation such as [40, 66, 79], the complications of real world tidal basins limit the application of such models. Spatially varying friction and reflection, and geometries which do not fall cleanly into straight, exponentially converging, or steadily sloping beds still frustrate analytic treatment and dictate

the need for numerical approaches. Adding two-way coupling between inundation and tidal energetics, the problem is most thoroughly treated with numerical approaches. Recent work in tide-inundation coupling includes [52], who implemented a wetting and drying scheme in the Princeton Ocean Model (POM) which was then applied to modeling the dynamics of the wetting and drying on the extensive mudflats of Cook Inlet, Alaska. Their results showed up to a 20% increase in tidal range when wetting and drying were included, as well as a slowing of the tidal wave, reducing phase angles by up to 10%. The increase in tidal range with inundation appears to contradict expectations based on frictional dissipation in intertidal areas. The exact comparison in [52], though, is between a case with wetting and drying allowed in the intertidal, and a case where the would-be intertidal area is numerically "dredged" to become subtidal. [64] used the same wetting and drying adapted version of POM in a model sensitivity study of an idealized semi-submerged seamount. They found that tidal ranges in the simulations with wetting and drying were higher on the seaward side of the island and lower on the landward side, compared to non-wetting simulations. Contributions to the momentum balance from bottom stress and barotropic pressure gradients were twice as great in the wetting and drying simulations, further emphasizing the role of friction in inundation studies.

Modeling effects of sea level rise

Over the next century anticipated sea level rise will require a wide range of managed responses, with potentially far-reaching and complexly interacting side-effects. The practical approaches for studying tide-inundation coupled dynamics range widely in complexity. The most basic sea level rise analyses completely ignore all dynamics, treating the basin as a bathtub with equal peak sea surface height everywhere. Inundation is then a matter of simply drawing a constant contour at the anticipated ocean MHHW elevation. Even with this simple approach, care must be taken in evaluating the sea level at which regions become tidally active. In many situations levees isolate low-lying ponds from tidal action and rising sea level, and evaluating hydrodynamic connectivity requires resolving small levee and channel features. Appendix B compares two methods for preparing elevation data for model bathymetry, and applies this basic approach to inundation in San Francisco Bay to demonstrate the importance of accounting for connectivity and levees below the nominal grid scale.

Studies such as [35] add considerable but necessary complexity to the analysis of inundation. A hydrodynamic model predicts future conditions in the interior of the basin, based on rising sea level, altered tidal ranges, and predictions of storm effects. The predicted high water marks at the basin perimeter are then extended over high resolution topography adjacent to the basin to ascertain the extent of inundation. Resolving hydrodynamics within a basin at this level can be accomplished with a very efficient model, allowing a wide range of ocean forcing conditions to be sampled. However, this approach omits the coupling of inundation back to tides.

Present goals

The aim of this chapter is to investigate how sea level rise in the coastal ocean modifies the coupled tidal-inundation dynamics, in hopes of informing future mitigation and restoration efforts and anticipating their consequences. In an attempt to capture the complexities of a physical system, while maintaining broad applicability to other systems, San Francisco Bay, California has been chosen as the domain for the numerical experiments. San Francisco Bay has moderate, mixed tides, representative of a wide area [6], and without particular anomalies which would make an analysis irrelevant for other basins. One advantageous feature of San Francisco Bay is its pair of dynamically distinct channels: a short, reflective, convergent channel to the south and a longer, progressive-wave channel to the north leading to a dissipative inland river delta. A single numerical experiment is thus able to reproduce a wide variety of responses and interactions.

San Francisco Bay is also a prime example of the range of management actions which affect and are affected by inundation dynamics. Multiple large restoration projects will be returning previously nontidal salt ponds to tidal action. In addition to anticipated restoration efforts, sea level rise mitigation projects are also likely to alter significant reaches of shoreline in the next 50–100 years, with two airports and numerous transportation corridors within reach of rising bay waters. An important question for planners is how far-reaching the effects of a particular mitigation effort are. This research is motivated by the need to answer questions such as whether the hardening of a stretch of shoreline by additional levees will increase inundation risk for neighboring soft shorelines. At larger spatial scales one may ask whether hardening shorelines around one embayment alters the tidal signal in another embayment. To this end, multiple shoreline scenarios are modeled, with leveed reaches of shoreline inserted into the model bathymetry to simulate shoreline hardening. Understanding the interplay between tidal dynamics, sea level rise, tidal marsh restoration and the resulting inundation is essential for achieving the goals of these coastal engineering projects at the same time as predicting and mitigating inundation hazards.

4.2 Physical Domain

San Francisco Bay has one of the longest continuous tidal records on the Pacific Ocean [72], showing sea level trends of 0.22 m rise per century [24]. The trends are not limited to mean sea level, but also include tidal and non-tidal variation, such as the 64 mm increase in diurnal tide range over 1900-1998 [24]. Interestingly, San Francisco saw a decrease in the diurnal tide range at a nearby coastal ocean sea level gauge (Point Reyes), but at the mouth of the bay diurnal tide range has been increasing, a trend which gets more pronounced at gauges landward of the mouth. Considering the power density across the full spectrum of sea level variation, records also show that decadal, annual and semiannual variability are increasing. [6] focused on measures of storminess, and found oscillatory and increasing signals of storminess on the US west coast. Other recent analysis [59] has found a reversal

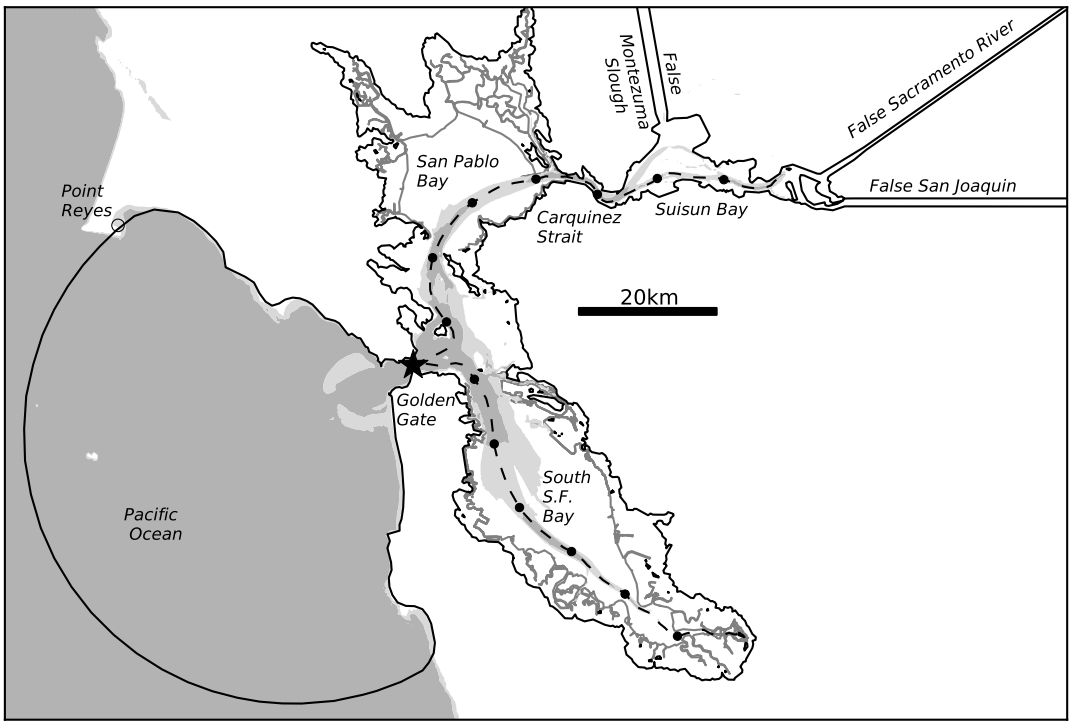

Figure 4.1: Overview of model domain. Shaded depth contours are shown at -5 m and -12 m NAVD88. The solid gray line outlines the present day MHHW shoreline in South San Francisco Bay and San Pablo Bay. The solid black outlines the model domain including three false deltas. Dashed black line is the end-to-end thalweg, with the origin at the Golden Gate marked by $\star$ and 10 km intervals marked by $\bullet$.

in the trend of sea level off the coast of California since 1997, though they attribute this to the Pacific decadal oscillation, and note that in due time the trend of increasing sea level is likely to return. These trends all underline the importance of understanding how dynamics will evolve within a basin such as San Francisco Bay as coastal sea level evolves.

In addition to long term trends in ocean forcing, landward forcing has also seen significant changes, including sediment loading, total river flows and shifts in the annual cycle of peak flows, all changes which are likely to continue into the future. Given the strong anthropogenic controls on river and watershed forcings, and the difficulty in projecting how those controls will evolve, the simulations described below take river inputs as constant, with no further consideration of changes in the landward forcing.

San Francisco Bay is a bifurcated, mesotidal estuary. Tides at the mouth are mixed diurnal and semidiurnal, with a great diurnal range at the mouth of 1.8 m [51]. The mouth of San Francisco Bay is the Golden Gate ($\star$ in figure 4.1), a 100 m deep, constricted channel, connecting Central Bay to the Pacific Ocean. Along channel distances throughout this

chapter are referenced to the Golden Gate with negative distances denoting the southern transect and positive denoting the northern transect. The southern branch of the bay, typically referred to as South San Francisco Bay, has little freshwater inflow, is roughly funnel-shaped, and is characterized by a single channel 12-20 m deep and broad shoals tapering from 5 m deep to intertidal. Tides in South Bay are close to a standing wave, with a velocity phase lead for the M2 constituent in the channel of approximately $\phi_{M2} \approx 75^\circ$ (i.e. peak flood velocity leads the peak sea surface elevation by 75°). The perimeter of South San Francisco Bay, particularly to the south and southeast, is dominated by tidal sloughs and salt ponds. Ongoing restoration projects are returning many of the salt ponds to tidal action by levee breaches.

The northern branch of San Francisco Bay connects through two largely self-contained bays, consecutively San Pablo Bay and Suisun Bay, before reaching the Sacramento-San Joaquin River Delta. A large fraction of the northern borders of both San Pablo Bay and Suisun Bay are tidal marshland joined to the respective bays via networks of tidal sloughs. A large number of wetland restoration projects, in various phases of planning or completion, are also targeted at the northern perimeter of San Pablo Bay.

4.3 Numerical Model

To quantify the tidal processes and contributions to variation in high water across the domain, the SUNTANS hydrodynamic code [28, 80] is employed to model a range of sea level and shoreline configurations in San Francisco Bay. The model is run in depth-averaged mode and solves the shallow water equations,

$$\frac{\partial u}{\partial t} + \mathbf{u} \cdot \nabla u - fv = -g\frac{\partial \eta}{\partial x} - \frac{1}{2h} C_D u \|\mathbf{u}\| \tag{4.2}$$

$$\frac{\partial v}{\partial t} + \mathbf{u} \cdot \nabla v + fu = -g\frac{\partial \eta}{\partial y} - \frac{1}{2h} C_D v \|\mathbf{u}\| \tag{4.3}$$

$$\frac{\partial \eta}{\partial t} = -\nabla \cdot (h\mathbf{u}) \tag{4.4}$$

where u is eastward velocity, v northward velocity, $f = 2\Omega \sin \Phi$ the Coriolis parameter with angular velocity of the Earth Ω and latitude Φ. Elevation of the free surface is given by η, measured as a departure from the NAVD88 (North American Vertical Datum of 1988) geopotential surface, and h gives the total height of the water column. While estuarine circulation and other 3-D processes will undoubtedly be altered by sea level rise, the model is run in the less computationally expensive depth-averaged mode since baroclinic effects are likely to have only minimal effects on inundation and the barotropic tidal response. These equations are discretized on a prismatic, finite-volume grid comprising unstructured triangles in the horizontal. The drag coefficient C_D is derived from a bottom roughness z_0 which in turn is calculated from the water column height based on relationships from previous modeling efforts in San Francisco Bay [31]. At each time step the bottom roughness for each

Table 4.1: Bottom roughness as a function of water column height (values taken from [31]).

h [m]	z_0 [m]
0.1	0.00123
0.6	0.00229
2.0	0.00116
6.5	0.00025
8.5	0.00010

cell is found by linear interpolation over the values in table 4.1. Following [45], roughness throughout the false delta regions is set to a minimal 10^{-5}m. At each time step n an effective drag coefficient $C^n_{D,j}$ is calculated from the edge-local bed roughness $z^n_{0,j}$ as

$$C^n_{D,j} = \left(\frac{\kappa}{\log\left(\frac{h^n_j}{2z^n_{0,j}}\right)} \right)^2 \tag{4.5}$$

where h^n_j is the height of the water column on edge j at time step n, and $\kappa = 0.42$.

Wetting and drying in the model is handled by deactivating cells for which the water column height falls below a threshold height. A threshold of 5 mm is used, chosen to be small enough to avoid significant artificial storage, but large enough that the model is stable at a reasonable time step of 10 s.

The ocean boundary of the model domain is approximately 50 km beyond the Golden Gate, coinciding with a long-term tidal gauge at Point Reyes (see figure 4.1). The focus of the present study is South San Francisco Bay and San Pablo Bay. Within these basins the model domain extends up to the 3.5 m NAVD88 contour to accommodate tidal amplification and sea level rise. Upstream of San Pablo Bay and seaward of the Golden Gate the model extends up to the present day MHHW shoreline. Beyond Suisun Bay, the Sacramento-San Joaquin Delta is represented by a pair of hypsometry-matched false deltas, as is the slough and marsh network north of Suisun Bay. For each of the three false deltas, hypsometry (the relationship between planform area and free surface elevation) is extracted from a 10 m digital elevation model (DEM), excluding areas already accounted for in the original grid. A length for each false delta is determined by the along-channel length of the primary channel in each region. The width is then a consequence of this length and the maximum area found in the hypsometry calculations. A regular triangular grid is constructed to match these dimensions, with a resolution of approximately 400 m. The hypsometry is then binned by depths such that each successive depth bin corresponds to an increase in the planform area equivalent to the area of one cell. These ordered depths are then assigned to cells starting at the seaward end of the false delta, proceeding first along a strip of cells to the landward end, and then proceeding laterally to the next strip of cells until all cells have been assigned a

depth. This simple approach ensures that a subtidal channel extends the length of the false delta (consistent with known river features), with a bed which slopes up in the landward direction and has a degree of lateral bathymetry variation. By matching hypsometry and length the false deltas approximate the tidal response of the more complex physical channel network with a substantial reduction in the number of grid cells compared to a fully resolved delta.

For the purpose of this study, the portion of the domain at elevations between 0 m and 3.5 m NAVD88 is the most relevant as this is the intertidal range across the imposed sea level rise scenarios of 0 m – 1.0 m. The grid resolution in the area between these contours is set to a uniform nominal length scale of 50 m (resolving the broad intertidal areas at finer resolution than this runs into practical computational limits). Additionally, five dynamically important channels outside this region are also given increased resolution: the Golden Gate (resolved at a scale of 125 m), Carquinez Strait and Suisun main channel (100 m), Suisun Cutoff (100 m), and New York Slough (200 m). In all other areas, the grid resolution is allowed to increase at a rate of 10%, up to a maximum grid scale of 3 km at the open ocean boundary. The resulting grid has 937,759 cells.

Bathymetry data were derived from a range of sources covering subtidal, intertidal, and supertidal areas up to the 3.5 m NAVD88 contour. The base elevation data source was a 10 m seamless topography-bathymetry product designed for inundation studies [9]. Bathymetry at 10 m resolution for the Sacramento-San Joaquin Delta was taken from [25]. Given the impact of small levee and slough features on inundation and hydrodynamic connection, special care was taken to assemble up-to-date and high resolution intertidal topography. This included the preprocessed 2 m bathymetry from [26], as well as gridded bare earth LiDaR datasets from [27], [49], and [50]. Missing data in the LiDaR datasets in small regions were filled via interpolation from nearby LiDaR data, or, in cases where the LiDaR was missing data over a span greater than 10 m, data were filled in from [9].

Though the intertidal regions are resolved at 50 m, the length scale of many channels and levee features, essential for the inundation characteristics of the marshes, is 5–25 m. As in previous studies such as [4], simple averaging of the DEM along each edge was insufficient to resolve either channels or levees robustly. [4] proposed a solution for a similar application, but they explicitly state that small coherent structures such as levees are not adequately captured by their statistical analysis. Additionally, the grid orthogonality constraints of the finite-volume code utilized in the present study preclude the flexible manipulation of grid geometries which they are able to leverage. The importance of narrow channels and levees in quantifying inundation motivated development of a bathymetry approach which calculates the overtopping elevation for each edge. The method is described in appendix B along with a comparison to a simple averaging approach.

Calibration and validation of the model (see appendix C) has been performed with observed tides and winds (scenario sObs). Periodic tides are used for all subsequent analysis in order to avoid the need for spring-neap duration runs of each scenario and to allow the analysis to focus on the individual effects of a single tidal constituent. The ocean free surface is forced with an M2 period (12.42h) sinusoidal signal with a peak-to-peak amplitude of 1.64

Table 4.2: Naming convention for numerical experiments. Obs. denotes observed tides, M2 denotes 12.42h periodic tides with amplitude matched to spring range of observed tides. ("N" denotes hardening in the Northern Reach; "S" in the Southern.)

	Sea level	Hardened shorelines			
Tides	rise [m]	none	SPB	SSFB	both
Obs.	0.0	sObs			
M2	0.0	s0	hN0	hS0	hNS0
M2	0.6	s60	hN60	hS60	hNS60
M2	1.0	s100	hN100	hS100	hNS100

m. The amplitude was chosen to match the great diurnal range observed at the Golden Gate. The imposed M2 amplitude, large compared to the measured M2 constituent, allows for a range of inundation similar to the combined tides. This avoids the complications of nonlinear interactions between constituents, but retains and resolves the interplay between tidal dynamics and inundation regions, which is not affected by these interactions, and allows the development of higher frequency harmonics. The tidal boundary condition is applied with uniform phase and amplitude across the entire ocean boundary; tidal dynamics outside the Golden Gate are not expected to be realistic.

The numerical experiments cover three variations in sea level rise and four shoreline configurations, with a naming convention outlined in table 4.2. The range of ocean boundary conditions comprises (i) present day mean sea level, (ii) an increase of 0.6 m and (iii) an increase of 1.0 m. These values were chosen to roughly cover the range of predictions over the next 50-100 years [58]. Multiple shoreline configurations are used to simulate the effects of mitigation efforts such as the construction of levees at present day MHHW shorelines, and how shoreline hardening in one portion of the domain affects tidal range in other portions of the domain. The first shoreline scenario, "soft" (s), does not include any explicit shoreline protection, only present day topography and bathymetry. The completely hardened scenarios (hNSx) limit flows in both San Pablo Bay and South San Francisco Bay to present day MHHW shorelines. Two additional scenarios represent shoreline hardening limited to either San Pablo Bay (hN) or South San Francisco Bay (hS). In all cases except sObs, the model is allowed to spin up for four days before data is extracted for a single M2 period.

4.4 Energy Flux and Tidal Phase Analysis

Energy flux and tidal phasing provide fundamental information about the spatially variable dynamics in each basin. While the one-dimensional tidal propagation problem is well-described by the velocity phase lead ϕ, in two dimensions the direction of wave propagation and the sense of landward or seaward can become ambiguous. Switching to a comparison of

depth-averaged tidal energy flux and an effective tidal phase lag derived from the energy flux allows for a more robust description. The eastward and northward energy fluxes are defined as the pressure work done by the flow on an imaginary surface normal to the respective coordinate direction, averaged over the mean depth of the water column [36]:

$$F_i = \left\langle \int_{-h}^{0} p u_i dz \right\rangle \tag{4.6}$$

where u_i is the velocity component in the ith coordinate direction (assumed constant in the vertical), $p = -\rho g z$ is the hydrostatic pressure, h is the time varying height of the water column, and $\langle \circ \rangle$ denotes a tidal average. For a tidal constituent with frequency ω_c, the harmonic velocity and depth are given by

$$u_c = U_c \cos\left(\omega_c t + \phi_{x,c}\right) \tag{4.7}$$

$$v_c = V_c \cos\left(\omega_c t + \phi_{y,c}\right) \tag{4.8}$$

$$h_c = H + \eta_c \cos \omega_c t. \tag{4.9}$$

where H is the mean depth. Substituting (4.7) – (4.9) into (4.6), the energy flux for constituent c is then

$$F_{x,c} = \frac{1}{2} \rho g U_c H \eta_c \cos \phi_{x,c} \tag{4.10}$$

$$F_{y,c} = \frac{1}{2} \rho g V_c H \eta_c \cos \phi_{y,c}. \tag{4.11}$$

Consistent with the energy flux estimates, a measure of the tidal phase can then be determined as the inverse cosine of the ratio of the energy flux to the maximum possible flux for a fully progressive wave with the same H, U_c, V_c and η_c:

$$\phi_c = \cos^{-1}\left[\frac{\sqrt{(U_c \cos \phi_{x,c})^2 + (V_c \cos \phi_{y,c})^2}}{\sqrt{U_c^2 + V_c^2}}\right]. \tag{4.12}$$

In the case of unidirectional tidal flow aligned with the x coordinate, (4.12) simplifies to $\phi_c = \phi_{x,c}$. However, compared to considering the tides to be unidirectional along the principal axis, (4.12) is robust to amphidromes and rotary tides where the principal axis is poorly defined.

The tidal energy in higher harmonics is relatively small in the majority of the domain and our initial analysis is focused on the M2 constituent. The numerical experiments do predict a significant M4 overtide, which is later considered in section 4.6. From each of the periodic scenarios in table 4.2, M2 phase and amplitude of the sea surface height, eastward velocity and northward velocity were extracted by a least squares approach over exactly one tidal period. Changes in M2 energy flux between pairs of scenarios elucidate how the tides change in response to shoreline hardening and sea level rise, and how this response differs between the two bays.

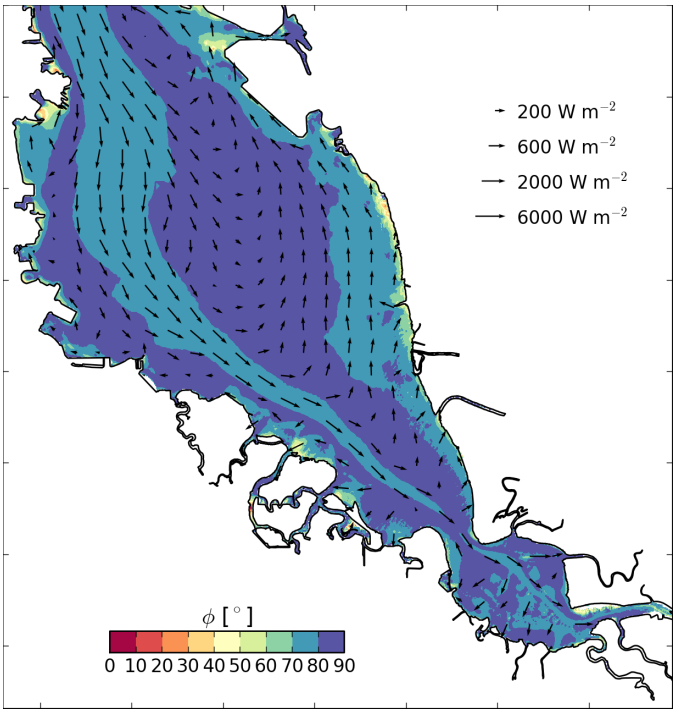

Figure 4.2: South San Francisco Bay: ϕ_{M2} and tidal energy flux for hNS0.

South San Francisco Bay

Figure 4.2 shows the M2 energy flux and tidal phase for South San Francisco Bay. The dominantly standing-wave tidal dynamics are clear, with the majority of the embayment showing $\phi_{M2} > 80°$. The channelized portion of the bay to the southwest is slightly more progressive, but still much closer to a standing wave than the channel in San Pablo Bay (discussed below). Though the tidal phase is generally close to standing throughout the bay, the energy flux magnitudes are still substantial due to the large tidal range.

Interestingly, in much of the bay the easternmost portion of the shoals shows slightly "overstanding" tides with a seaward energy flux. At the most eastward margins of the bay the seaward-directed progressive component is sufficient to reduce ϕ_{M2} below 80°. This seaward energy flux has been observed in South San Francisco Bay [38] and is reminiscent of the seaward residual transport in [43]. In that study, analytic solutions to the residual transport in a short, V-shaped basin yield landward transport in the channel and seaward transport in the shoals. To better understand the overstanding wave, the analytic, intratidal solution of [44] (which served as the intratidal solution on which the subtidal solution of [43] was based) is applied to the circulation in this stretch of South San Francisco Bay. The analytic solution describes the time-varying velocity and free surface in terms of complex amplitudes U, V and A for the longitudinal velocity, lateral velocity and free surface anomaly. Each

complex amplitude describes a magnitude and phase in the form (with variables renamed for consistency within the chapter)

$$h(x,y,t) = \Re\left\{\eta(x,y)e^{i\omega t}\right\} + H(y) \quad (4.13)$$
$$u(x,y,t) = \Re\left\{U(x,y)e^{i\omega t}\right\} \quad (4.14)$$
$$v(x,y,t) = \Re\left\{V(x,y)e^{i\omega t}\right\} \quad (4.15)$$

where $H(y)$ is the mean depth and a function only of the lateral position for the purposes of the analytic solution. The analytic solution is then described by

$$\beta = \frac{8C_D U_0}{3\pi} \quad (4.16)$$
$$\mathcal{F} = -\int_0^D \frac{gH}{i\omega + \beta/H} dy \quad (4.17)$$
$$\sigma^2 = \frac{i\omega D}{\mathcal{F}} \quad (4.18)$$
$$\eta = \eta_0 \frac{\cos\left[\sigma\left(x-L\right)\right]}{\cos\sigma L} \quad (4.19)$$
$$U = \frac{g}{i\omega + \beta/H}\frac{\eta_0\sigma}{\cos\sigma L}\sin\left[\sigma(x-L)\right] \quad (4.20)$$
$$V = -\frac{\eta}{H}\left(i\omega y + \int_0^y \frac{gH\sigma^2}{i\omega+\beta/H} dy\right) \quad (4.21)$$

where $C_D = 0.0025$ is the drag coefficient, $D = 11.8$ km is the width of the cross-section, $U_0 = 0.83$ m s^{-1} a representative velocity scale, $i = \sqrt{-1}$, $\eta_0 = 0.4$ m the tidal amplitude at the open boundary, and $L = 30$ km the length of the basin. The form of $H(y)$ (figure 4.3a) was extracted from gridded bathymetry data at -38 km from Golden Gate. The results of the analytic model are shown in figure 4.3b, where the velocity phase lead has been plotted as a function of the cross-channel distance. The deep channel is the only portion of the cross-section with a landward energy flux, with the greatest seaward energy flux in the far shoals, in general agreement with figure 4.2. Note that the analytic model can only describe rectangular basins, and since the standing wave of South San Francisco Bay is largely due to channel convergence, it is not surprising that this analysis predicts more progressive features compared to figure 4.2. Nonetheless, the qualitative agreement in the velocity phase lead is a good indicator that a channel-shoal bathymetry profile and a reflective basin are the drivers of the overstanding wave in the shoals.

Analysis of how the present day M2 dynamics are altered by sea level rise proceeds in two steps. First, to approximately separate the effects of deepening from inundation, the changes due to sea level rise with hardened shorelines are considered throughout the region. The change in M2 phase and energy flux between the hNS0 and hNS100 scenarios is shown in figure 4.4. The choice of the 1.0 m scenarios was motivated by the characteristics of inundation in South San Francisco Bay, where a majority of the inundation occurs above

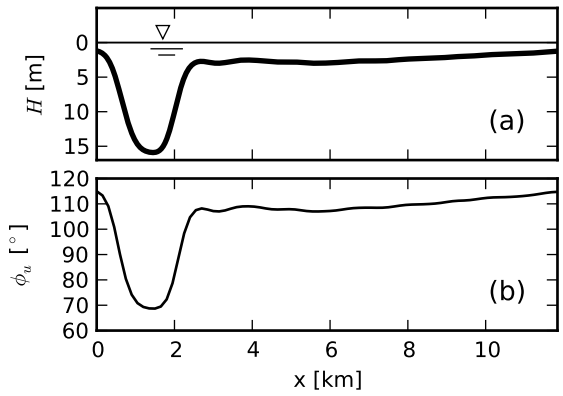

Figure 4.3: Application of [44] analytic intratidal solution to South San Francisco Bay. (a) Bathymetry profile extracted from DEM at -38 km from Golden Gate. (b) Predicted velocity phase lead for the along-channel velocity. The subscript u denotes that velocity phase lead was calculated just for the longitudinal velocity, such that $\phi_u < 90°$ describes a progressive wave with landward energy flux, and $\phi_u > 90°$ a progressive with seaward energy flux, i.e. an overstanding wave.

a sea level rise of 0.6 m. The phase changes in the bulk of the bay are minimal. The landward end of the bay has a convergent geometry and is nearly closed, such that, absent any dissipation from inundation, tidal energy has nowhere to go and the landward energy flux is constrained to be near zero. The southern half of the bay shows only small and scattered changes in the energy flux, while the northern half shows a distinct seaward shift in the energy flux. This shift is consistent with a deeper South Bay that is less frictional and more reflective. Portions of the eastern shoals become more progressive, departing slightly from the bulk of the bay, but notably the change in energy flux is actually seaward, showing that the overstanding tidal phasing of the shoals is accentuated by the deeper sea level and greater tidal range in the far south end.

The effects of inundation on M2 dynamics are shown in figure 4.5, comparing scenarios hNS100 and s100. The greatest changes are at the mouths of sloughs, which function as gateways to the increased tidal prism when inundation is permitted. South San Francisco Bay is ringed by numerous tidal sloughs, connecting pond and slough networks to the main body of the bay. The sloughs are typically small, but when considered in the aggregate they are a significant sink of tidal energy in the M2 band. The change in energy flux is everywhere landward, consistent with a basin transitioning towards a progressive wave. The eastern shoals actually show an increase in ϕ_{M2}, towards a standing wave. In the scenarios with hardened shorelines, these shoals had a seaward energy flux, such that as the basin as a whole becomes more progressive, those areas which were originally overstanding (i.e. a

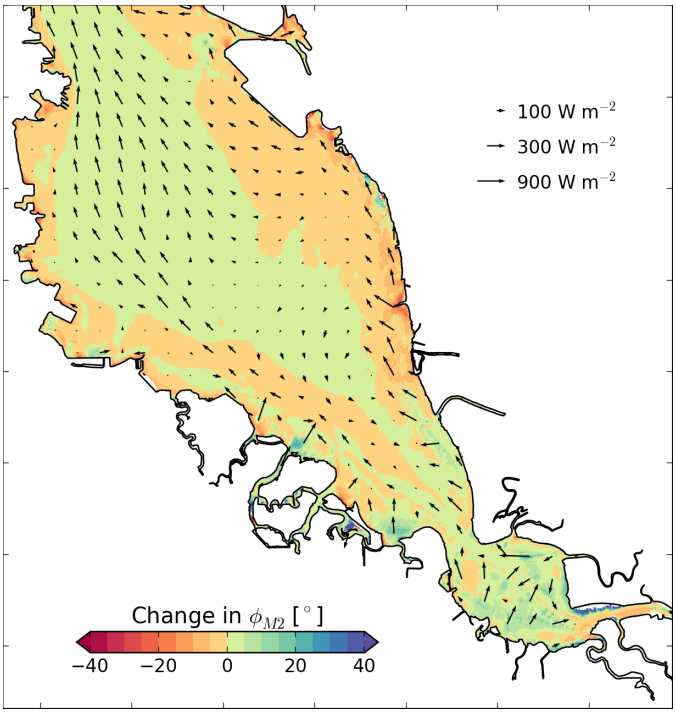

Figure 4.4: South San Francisco Bay: change in depth-averaged M2 energy flux between scenarios hNS0 and hNS100. Change in ϕ_{M2} is contoured in the background.

seaward tidal energy flux) shift towards a standing wave. The flux changes of figure 4.5 are quite uniform across the width of the bay, especially compared to the baseline energy flux of figure 4.2 which has significant lateral variation. This stems from the fact that in South San Francisco Bay the majority of the inundation occurs in the southern portion of the bay. The primary effect of inundation for most of the bay is directly related to the amplitude of the reflected wave, rather than local or lateral dynamics.

In South San Francisco Bay the competing effects of deepening and inundation are clearly demonstrated. Deepening reduces frictional effects in the subtidal basin interior, and allows energy to reflect back out of the basin, shifting the dynamics towards a standing wave. Inundation counters this shift by attenuating the incident wave in the southern end of the basin. Less energy is reflected and the phase shifts towards a progressive, non-amplifying wave.

San Pablo Bay

Figure 4.6 shows ϕ_{M2} throughout the interior of San Pablo Bay, under present day conditions (scenario hNS0). San Pablo Bay has somewhat progressive tides along the main channel in the south, and a partially standing wave across the shoals. Separated from the main channel

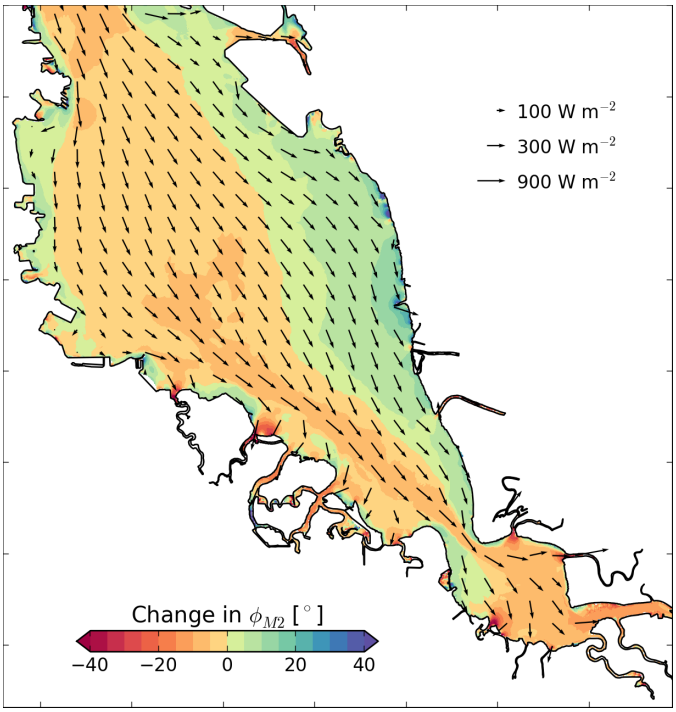

Figure 4.5: South San Francisco Bay: change in depth-averaged M2 energy flux between hNS100 and s100. Change in ϕ_{M2} is contoured in the background.

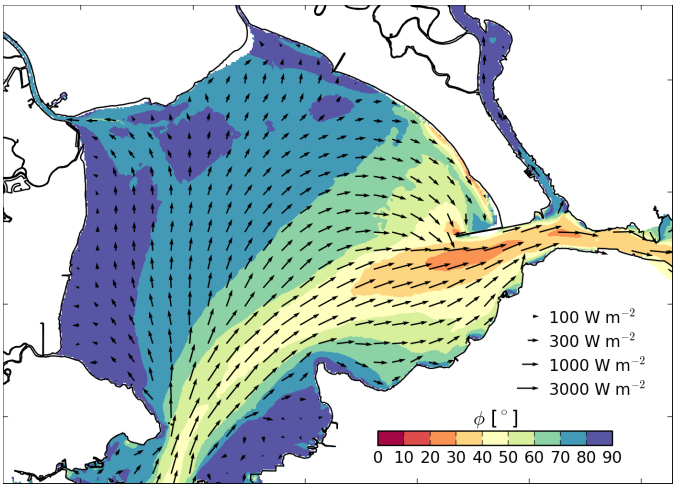

Figure 4.6: San Pablo Bay: ϕ_{M2} and tidal energy flux for the hNS0 scenario.

by the broad shoals, two significant river/slough mouths connect to the northern shore of San Pablo Bay: the Petaluma River to the west and Sonoma Creek to the west. Additionally, the Napa River joins Carquinez Strait shortly upstream of San Pablo Bay. Under present day conditions all three of these riverine features are significant sinks of tidal energy. However, the ranges of their respective influence are localized, as indicated by the spatial extent of the depression in ϕ_{M2} near each mouth (figure 4.6).

As with South San Francisco Bay, scenarios hNS0 and hNS100 are compared first in order to isolate the effects of a deeper basin interior, without significant change in inundation or tidal prism. Figure 4.7 shows the change in energy flux and phase between these two scenarios. Unlike South San Francisco Bay where most inundation occurs above a 0.6 m increase in sea level, in San Pablo Bay the incremental change in inundated area between s0 and s60 is similar to the change between s60 and s100. The landward energy flux shifts from the channel to a proportionally greater flux in the shoals. The present day mean depth of the off-channel areas of San Pablo Bay is quite shallow, making it highly frictional and a high impedance path for tidal propagation. A 1 m increase in mean sea level has a proportionally greater effect on the role of friction in the shoals, allowing a greater fraction of the tidal energy to propagate via the shoals. Another feature of the change in energy flux is the significant increase of energy leaving San Pablo Bay by way of Carquinez Strait to the east. This is likely due to a combination of less energy being lost in San Pablo Bay, and that the upstream prism is greater, with more area and more dissipation which attenuates the return wave and allows a greater net energy flux landward. In terms of the tidal phasing the trend is clear that almost all of the bay shifts towards a progressive wave. The most striking feature is the northeastern shoreline, which, under the higher sea level, permits a southeasterly energy flux with a much more progressive wave, feeding tidal energy to Carquinez Strait. In both hNS0 and hNS100 this progressive, along-shore wave is evident, but in the hNS100 case the greatly increased flux through the shoals means that there is that much more tidal energy available to then be transported south to Carquinez Strait.

Figure 4.8 shows the incremental change in energy flux and phase between scenarios hNS100 and s100, depicting the effect of soft shorelines under 1.0 m rise in sea level. The bulk effect in San Pablo Bay is an increase in tidal energy entering the bay from the south, and a decrease in tidal energy leaving the bay in the east (note that figure 4.8 shows depth-averaged fluxes; in the south, the incoming flux is in twice the depth as the outgoing flux). Based on these changes in energy fluxes it is clear that the inundation of the soft shorelines leads to greater dissipation and the bay has become a greater sink of tidal energy. The hot spots of energy flux and progressive phase at the mouths of all three rivers (west to east, Petaluma River, Sonoma Creek, and Napa River) show that the bulk of the newly inundating areas are not directly connected to the main body of the bay but are instead connected via river and slough features. While in figure 4.7 the eastern shore had a pronounced shift towards a progressive wave, this same region partially reverts towards a standing wave once the additional tidal prism is available via the two nearby river mouths. In the hard shorelines case the tidal wave propagates through the shoals but with no outlet to the north it is focused on the eastern shore and then routed to the south. The soft shorelines case introduces highly

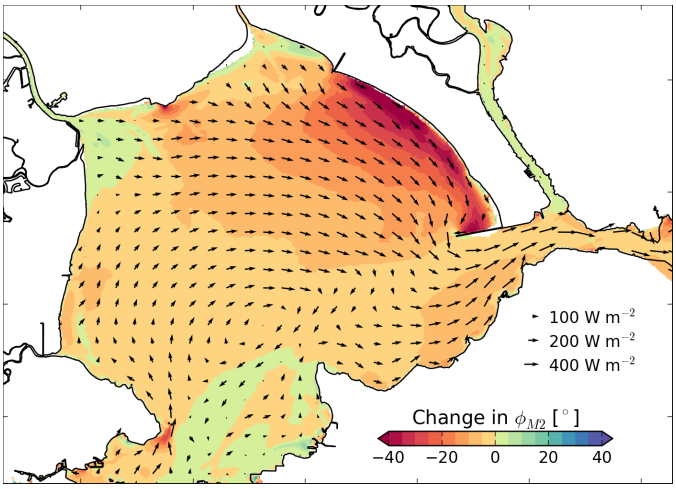

Figure 4.7: San Pablo Bay: change in depth-averaged M2 energy flux between hNS0 and hNS100. Change in ϕ_{M2} is contoured in the background.

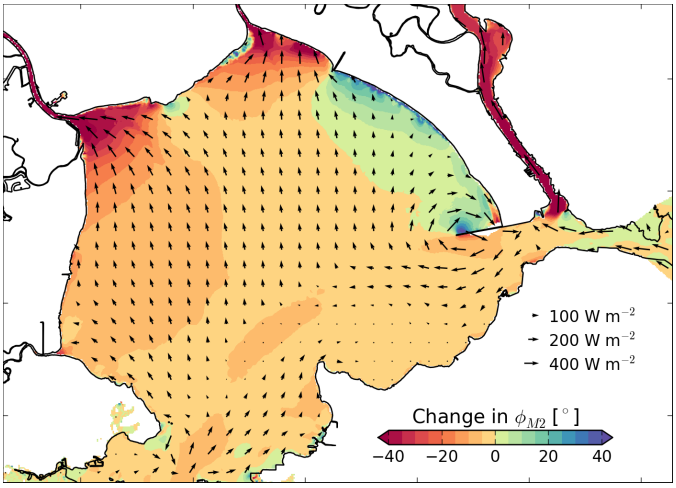

Figure 4.8: San Pablo Bay: change in depth-averaged M2 energy flux between hNS100 and s100. Change in ϕ_{M2} is contoured in the background.

dissipative river features, such that tidal energy is partially redirected and dissipated in the sloughs rather than being focused on the eastern shore and routed south to Carquinez Strait.

4.5 Tidal Amplification and Damping

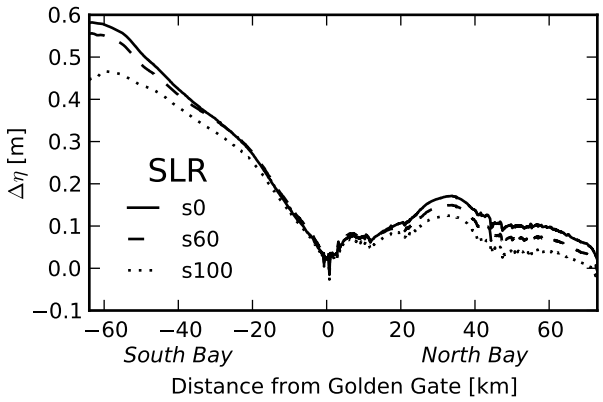

Figure 4.9: Transect of relative high water for soft shoreline scenarios s0 – s100. Negative longitudinal distance indicates South San Francisco Bay, positive is San Pablo Bay and beyond.

The analysis has so far examined only the M2 wave, but to more concisely address potential for inundation, the following section turns to a direct measure of the peak sea surface height. The high water elevation includes the combined effects of the M2 wave and its harmonics, as well as constant offsets from processes such as the Stokes transport increasing mean free surface setup. Changes in high water within the basin are a combination of changes in the ocean forcing and changes in dynamics within the basin. Capturing only the effects of dynamics within the basin, the relative high water is defined as

$$\Delta\eta\left(\mathbf{x}\right) = \max_t \eta\left(\mathbf{x}\right) - \max_t \eta_{BC}. \tag{4.22}$$

This is the height by which the high water level at a point in the domain exceeds or falls shy of a reference high water level at the coastal ocean boundary. Lateral variations in $\Delta\eta$ are typically small, allowing comparisons between scenarios simply along the central thalweg of the bay, as shown in figure 4.9.

The highly reflective and converging South San Francisco Bay is evidenced by the tidal amplification on the left side of the plot. At its most extreme, high water at the far south end of the bay exceeds coastal ocean high water by nearly 0.6 m in the s0 scenario. Interestingly, the highest *rate* of amplification in the southern reach occurs in the first 20 km, which is still essentially part of Central San Francisco Bay. The amplification rate slumps for the reach

between 20 km and 40 km, consistent with the divergent geometry of the bay here, with a slight uptick south of 40 km as the shorelines converge.

The more transmissive and dissipative northern reach of the bay sees mild amplification for the first 35 km, up to the transition from San Pablo Bay to Carquinez Strait. The seaward half of that stretch has particularly complex geometry and bathymetry, leading to greater variability over short length scales, up to the transition into San Pablo Bay proper, at 25 km from Golden Gate. The greatest amplification occurs in the middle of San Pablo Bay. In a sense, San Pablo Bay can be considered a "leaky" reflective basin, particularly at present day sea level. The landward outlet for tidal energy, Carquinez Strait, is relatively small; M2 velocity phase lead within the strait itself dips below 30° (figure 4.6), but this depression of ϕ_{M2} extends only partially into San Pablo Bay itself, approximately to the point of maximum amplification at 25 km from Golden Gate. In fact the northern shoal areas are far enough from the main channel to see additional, though minor, amplification relative to the main channel (not shown). This effect can be traced to the wide aspect ratio of San Pablo Bay and how the northern shoals behave somewhat like an off-axis standing wave basin with the accompanying amplification of a standing wave basin.

The progression of tidal amplification from scenario s0 to s100 in figure 4.9 demonstrates the combined, attenuating effect of inundation and sea level rise. Beyond Central Bay, i.e. -20 km – 15 km, the attenuation due to inundation more than offsets the amplification expected from a deeper basin. The locations at which the scenarios begin to diverge roughly correspond to where the inundatable regions occur, notably south of -40 km and north of 25 km. Figure 4.10 shows the incremental extent of inundation for each soft shoreline scenario. The most marked change in inundation in South San Francisco Bay occurs when sea level rise approaches 1.0 m, compared to a relatively small change in inundated area between the s0 and s60 scenarios. Consistent with the inundation distribution in the south, the greatest change in amplification is between the s60 and s100 scenarios, where inundation leads to a significantly attenuated $\Delta\eta$ south of -40 km. San Pablo Bay has a more even distribution of inundated area, both in terms of where these areas are located and at what rise in sea level they become inundated. In San Pablo Bay the incremental difference in attenuation between s0 and s60 is similar to the difference between s60 and s100. The longitudinal variation in the attenuation is also less pronounced than in South San Francisco Bay.

The comparisons between soft shorelines and hard shorelines approximately separate the effects of deepening from the effects of inundation, but also allow a comparison of local effects versus remote effects, by selectively hardening only a portion of the domain. This demonstrates the dynamic interactions of the basins, and at the same time informs practical management decisions regarding the degree to which mitigation efforts must be coordinated throughout a basin. The local and remote effects of shoreline hardening are quantified in figure 4.11, where the change in relative high water is shown for the four shoreline configurations. The baseline amplification of scenario s0 has been subtracted out, as the changes are small relative to the baseline tidal amplification (i.e. the solid line of figure 4.9). When all shorelines are allowed to inundate the model shows that a small portion of Central Bay is essentially unchanged, but throughout the regions with significant inundatable shoreline

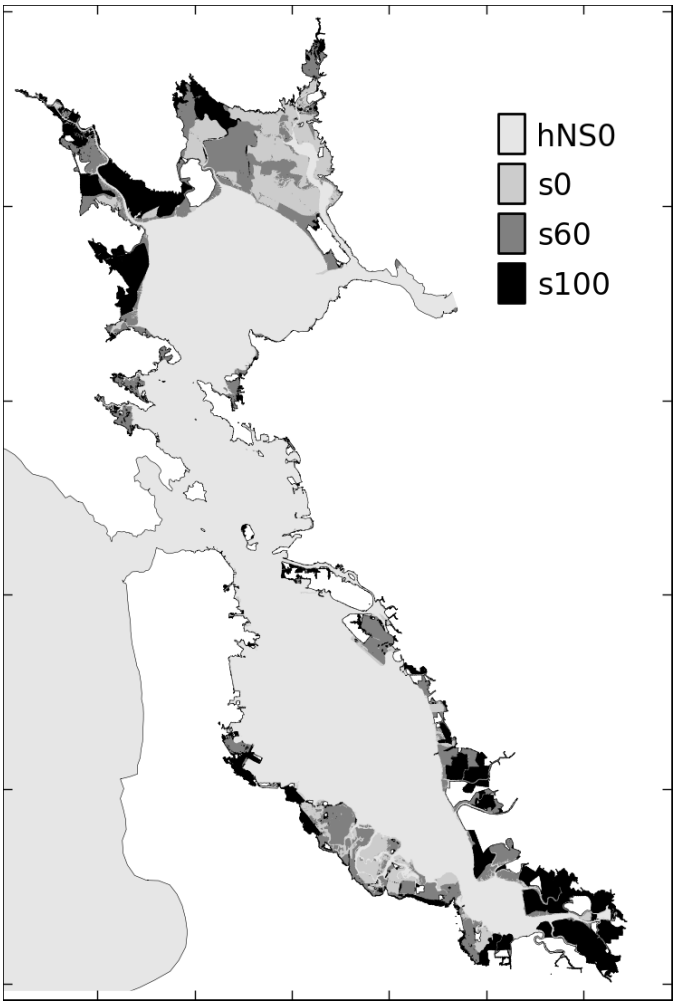

Figure 4.10: Extent of inundation across scenarios, progressing from least inundation to most.

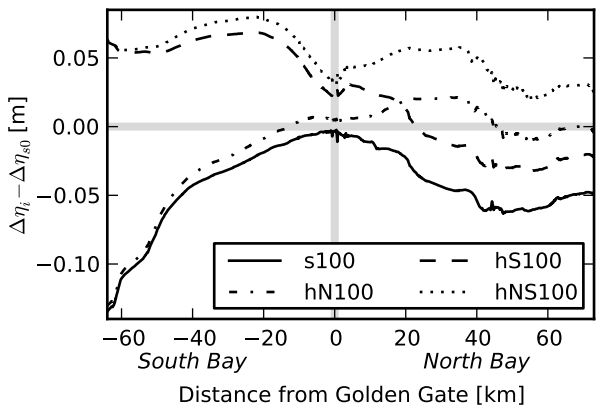

Figure 4.11: Longitudinal transect showing change in $\Delta\eta$ for s100, hN100, hS100 and hNS100, relative to the baseline $\Delta\eta$ of scenario s0. Negative longitudinal distance indicates South San Francisco Bay, positive is San Pablo Bay and beyond.

area tides are attenuated. With a maximum change in $\Delta\eta$ of approximately -0.13 m, the attenuation is significant, though small compared to baseline tidal amplification. Hardening the shorelines of South San Francisco Bay has a more pronounced effect, where the already large tidal amplification of 0.60 m in the far southern reach sees an additional 0.06 m increase. Hardening South San Francisco Bay also affects the high water level in Central Bay. Hardening and sea level rise both shift South San Francisco Bay towards a more reflective, standing wave environment, and it is apparent in figure 4.11 that the reflected tidal wave couples back into Central Bay, which also changes the seaward boundary condition for the northern reach of the bay.

Modifications to San Pablo Bay have similar local effects as in South San Francisco Bay. With soft shorelines, the broad inundatable regions of San Pablo Bay and its adjacent marshlands become a greater sink of tidal energy and tidal amplitudes decrease. Hardening the San Pablo Bay shorelines reverses that trend, leading to some amplification. Contrary to South San Francisco Bay, though, hardening the shorelines of San Pablo Bay has a negligible effect on South San Francisco Bay, as the progressive wave dynamics of the northern reach reflect little energy back to Central Bay.

Analytic approach for converging basin

Analyzing the whole of San Francisco Bay through the analytic lens of converging estuary hydraulics such as [66, 79, 8] is frustrated by the various branching, diverging, and reconverging features. Nonetheless, analytic approaches aid in identifying the dominant factors controlling the tidal response and can quickly predict the general response of a system without detailed observation or involved numerical approaches. Although the complex geometry

of much of San Francisco Bay makes a large-scale application of analytic theory difficult, the central portion of South San Francisco Bay has a smoothly convergent geometry. This section applies the methods of [8] (CST) to this reach, between 30 km and 55 km south of the Golden Gate (roughly the widest point of South San Francisco Bay to the point at which the bay transitions to a broad slough at the far south end). The goal is to understand the predictive skill of the analytic model and its capacity to include inundation effects.

The formulations of CST include the following parameters in predicting the behavior of a basin: a, the length scale of the basin convergence, $\bar{h}$, the mean depth, $K = n^{-1}$, the inverse of the Manning-Strickler friction coefficient, and r_s, the relative width of off-channel storage. The ocean boundary condition is described by the tidal amplitude η_0 and angular frequency ω. These parameters are combined as:

$$c_0 = \sqrt{\frac{g\bar{h}}{r_s}} \tag{4.23}$$

$$\zeta = \frac{\eta_0}{\bar{h}} \tag{4.24}$$

$$\gamma = \frac{c_0}{\omega a} \tag{4.25}$$

$$\tag{4.26}$$

such that c_0 is the effective celerity including off-channel storage, ζ a non-dimensional tidal amplitude, and γ an estuary shape factor relating convergence and tidal wavelength. The tidal response of the system is described in terms of the non-dimensional numbers

$$\delta = \frac{1}{\eta}\frac{d\eta}{dx}\frac{c_0}{\omega} \tag{4.27}$$

$$\chi = \frac{r_s g c_0 \zeta}{K^2 \omega \bar{h}^{4/3}\left(1-(4\zeta/3)^2\right)} \tag{4.28}$$

$$\mu = \frac{1}{r_s}\frac{v\bar{h}}{\eta c_0} \tag{4.29}$$

where δ is the amplification factor, χ the friction number, v the velocity scale, μ the velocity number, and additionally $\epsilon = \pi/2 - \phi$ describes the velocity phase lead. CST then derives the system of equations

$$\mu = \frac{\cos\epsilon}{\gamma - \delta}, \tag{4.30}$$

$$\tan\epsilon = \frac{\lambda}{\gamma - \delta}, \tag{4.31}$$

$$\lambda^2 = 1 - \delta\left(\gamma - \delta\right), and \tag{4.32}$$

$$\delta = \frac{\gamma}{2} - \frac{4}{9\pi}\chi\frac{\mu}{\lambda} - \frac{1}{3}\chi\mu^2. \tag{4.33}$$

The solution to this system of equations is a prediction of the tidal response. Application of CST to real bathymetry and shorelines, even with idealized tidal forcing, leaves considerable room for interpretation. While CST focuses on channels with a constant depth and converging width, the basis for the exponential convergence is in terms of the cross-sectional area, such as in [65]. In the present case the cross-sectional area demonstrates a significantly smoother trend, and base the length scale a on this dimension rather than the the width. In all cases K was taken to be 36.9, based on the range of $\bar{h}$ and corresponding roughness values in table 4.1. The in-channel region is defined as a contiguous set of cells on either side of the thalweg with a mean depth of at least 1.0 m and a tidal amplitude of at least 0.2 m. The across-channel distance between the bounds of this region then define the channel width $\bar{B}$ (as a function of the along-channel distance), and similarly the integration of the cross-sectional area between these bounds gives $\bar{A}$. The storage ratio r_s was calculated as the ratio of the planform area of all tidally active cells (tidal amplitude greater than 0.2 m) to the planform area of all in-channel cells. Both the shallow intertidal margins of the main basin and all slough or pond features off the main channel are counted as off-axis storage, contributing to an increase in r_s. The open boundary tidal amplitude η_0 is taken from the model at a single cell along the thalweg. The mean depth at each cross-section, $\bar{h}\,(x)$, is the equivalent depth of a rectangular channel with area $\bar{A}$ and width $\bar{B}$. While CST assumes a constant depth, South San Francisco Bay has a trend of decreasing depth towards the head of the estuary, and $\bar{h}$ is taken as the mean over all cross-sections.

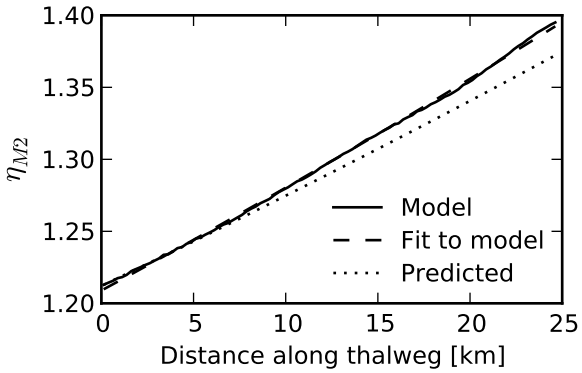

Figure 4.12: Comparison of amplification factor δ in scenario hNS100, as predicted by the model and by the analytic approach of [8]. The longitudinal transect is limited to the convergent portion of South San Francisco Bay.

The input parameters for each scenario and a comparison of phase and amplification between the numerical and analytic models are shown in table 4.3. The range of η_0 is fairly small, showing that, according to the model, dynamics seaward of this section account for about 0.05 m of change in tidal amplitude. The greatest differences in the inputs parameters are the variation of $\bar{h}$ with sea level and the variation of r_s with inundated area. All scenarios

Table 4.3: Application of method of Cai to numerical scenarios

	$\bar{h}$	η_0	a	r_s	model		analytic	
scenario	[m]	[m]	[km]		ϕ [°]	δ	ϕ [°]	δ
hNS0	3.88	1.19	11.5	1.06	83.1	0.237	79.7	0.165
hNS100	4.81	1.21	12.1	1.05	85.2	0.273	84.7	0.240
s100	4.74	1.16	12.2	1.33	79.4	0.236	80.2	0.214

fall within the *amplified* classification of CST, meaning that the net amplification is positive and incremental increases in the depth would result in additional amplification. While the velocity phase lead is predicted well by CST, the amplification factor is uniformly under-predicted, visible in figure 4.12 for the hNS100 scenario. The results from the analytic model reinforce the idea of competing effects of deepening and inundation, where an increase in $\bar{h}$ due to sea level rise is partially offset by an increase in r_s due to inundation. In terms of the amplification, these competing effects are nearly in balance in the numerical model, while the analytic approach is more sensitive to deepening than inundation, suggesting that inundation is not fully captured by considering storage alone, but rather frictional effects within the storage areas must also be considered.

4.6 Overtides

The total change in high water between hNS100 and s100, at -60 km from Golden Gate, is approximately 0.16 m (figure 4.11), but the M2 amplitude explains only about 0.07 m of this difference. While the ocean boundary is forced only with an M2 tide, local generation of overtides leads to significant M4 amplitudes within the domain, shown in figure 4.13. Previous analysis of the nonlinearities in the shallow water equations [55] has shown that M4 overtides are predominantly generated by the depth dependence of the friction term, depth dependence in continuity, and the nonlinear advection term. The depth dependent generation mechanisms are likely significant throughout much of San Francisco Bay, given the O(1 m) tides and O(2 m) depths prevalent in shoals throughout the domain. Furthermore, though the mean M4 amplitude is small (up to about 0.1 m), the differences across scenarios of the M4 amplitude is of the same order as the differences in M2 amplitudes. In addition to varying amplitudes, the distribution of M4 generation and the resulting phase relationships between the M2 and M4 vary greatly between scenarios.

The panels of figure 4.14 show that in all cases Central Bay is a significant source of M4, but portions of South San Francisco Bay may be sources or sinks of M4 depending on the scenario. In all cases, M4 appears to be generated in shallow, off-axis portions of the domain, and propagates seaward.

The differences in M4 generation appear to be driven by three factors: change in mean

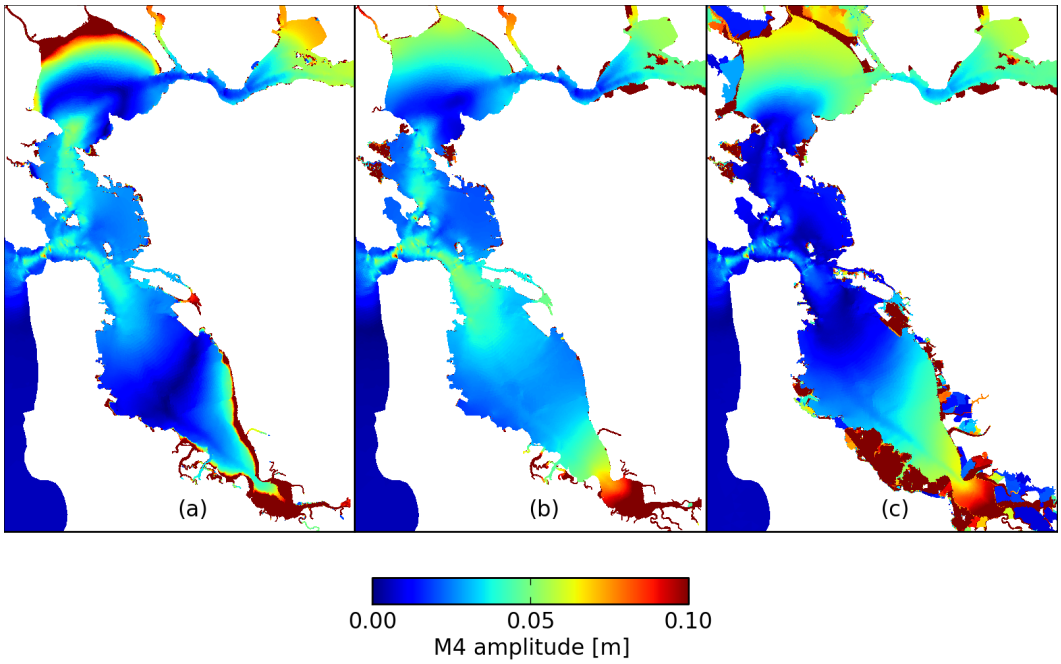

Figure 4.13: Amplitude of M4 overtide for (a) hNS0, (b) hNS100 and (c) s100

depth (i.e. hNS0 versus hNS100 or s100), change in M2 amplitude from which M4 can be extracted, and the potential for locally generated M4 to be locally dissipated when inundation is allowed.

Another complicating factor in the M4 dynamics is the fact that the M4 wavelength is short enough that a node of the standing wave can be present within the basin. In particular, the hNS0 scenario displays M4 minima consistent with nodes both at the widest point of South San Francisco Bay and in the middle of San Pablo Bay. The hNS100 scenario shows a seaward displacement of the node in San Pablo Bay, consistent with deepening and a longer wavelength.

It is clear that the M4 is significant, and taking into account the amplitude as well as the phase relative to the M2 phase, the M4 in hNS100 adds roughly 0.04 m to high water in most of South San Francisco Bay, compared to the M4 in s100 which actually decreases high water by up to 0.04 m. Of the original hNS100-s100 difference of 0.16 m (at -60 km), the combined M2/M4 wave then accounts for roughly 0.15 m. With the M2 wave varying between standing and landward progressing, and the M4 wave primarily seaward progressing, the phase relationships and resulting effect on high water is complex, and highly variable in space. When the M4 generation is plotted per-cell as in figure 4.15, as opposed to regionally aggregated as in figure 4.14, the field is both noisy and shows large variation over very short distances. Further analysis will require a more nuanced consideration of wetting/drying, which may be contaminating the results, and possibly aggregating generation and dissipation

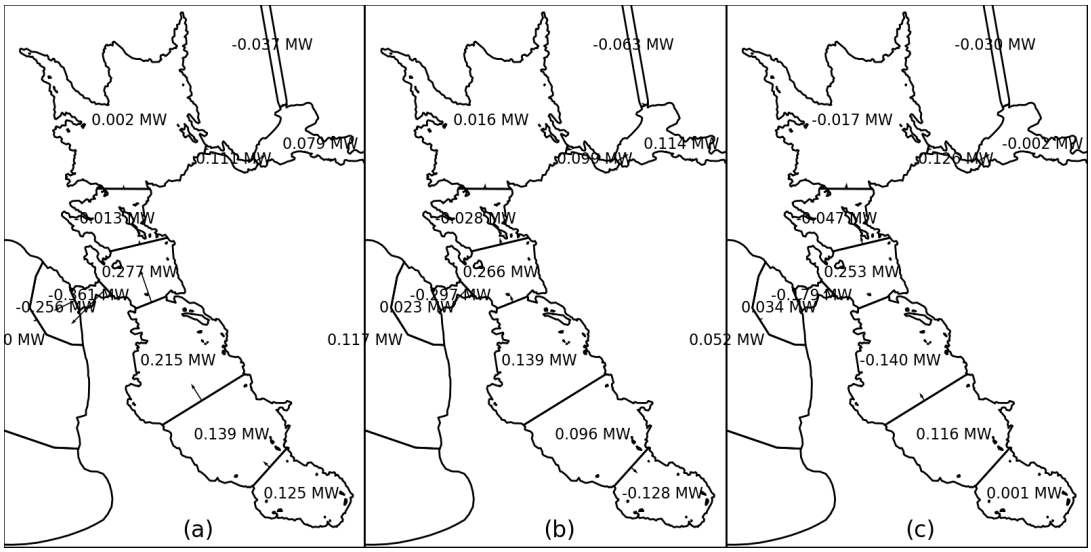

Figure 4.14: Region-integrated M4 generation and dissipation for (a) hNS0, (b) hNS100 and (c) s100. Arrows indicate integrated M4 energy flux between regions.

of the overtide by depth rather than region.

4.7 Discussion

The tidal and inundation response of an estuary to sea level rise is a complex interaction between a number of factors as well as interaction with adjacent tidal basins and the coastal ocean. Within a particular estuary or bay the dominant factors controlling the response to sea level rise include geometric factors, such as the aspect ratio of the basin, the baseline phasing of the tidal wave, and the spatial distribution of inundated areas

The aspect ratio determines the relative importance of longitudinal variation versus lateral variation. In the longer, high aspect ratio, South San Francisco Bay, changes due to sea level rise were relatively consistent across lateral transects, and lateral dynamics appear secondary. In contrast, San Pablo Bay, with a round, low-aspect ratio footprint, showed significant shifts of tidal propagation from the channel to shoals, and changes in the direction of tidal energy propagation.

Tidal phasing is also a major factor in determining the response to sea level rise, both in terms of local tidal amplification and how much the tides in one part of a basin feed back to other parts of the system. A standing wave system such as South San Francisco Bay appears more sensitive to sea level rise, both in the case of deepening only and deepening with inundation. Additionally, standing wave systems tend to have greater tidal range such that even small changes in phasing or dissipation lead to large changes in energy flux and

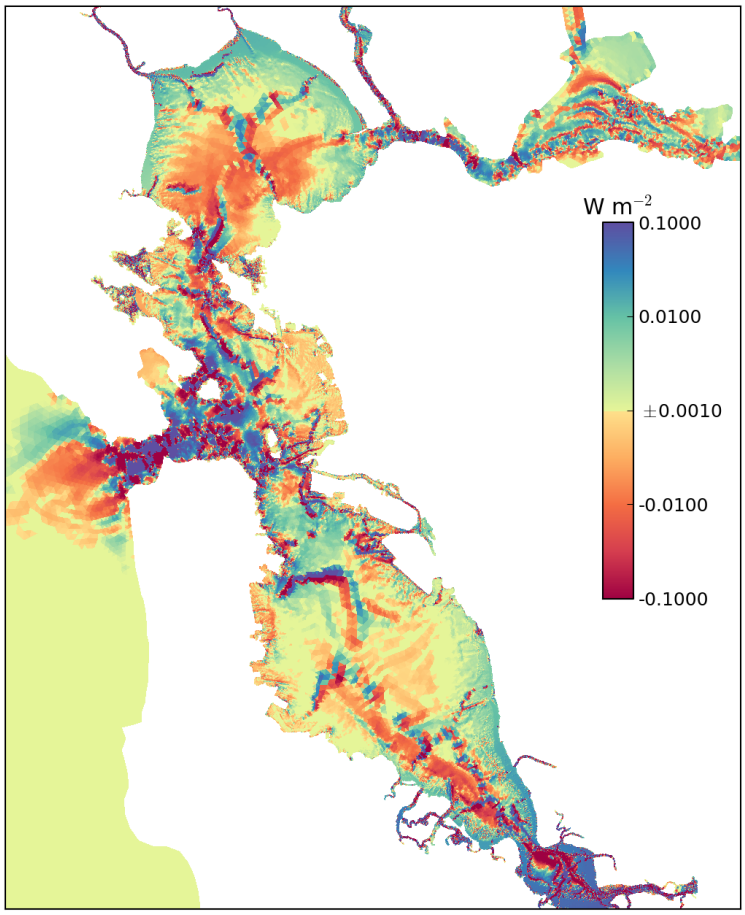

Figure 4.15: Non-aggregated M4 generation and dissipation for hNS100.

net amplification.

The quantity and relative location of inundatable areas also affects the tidal response. Greater expanses of inundatable areas relative to the subtidal area lead to greater attenuation of the incident tidal wave. The location of inundatable areas, along with the tidal phasing within a bay, affect the spatial extent of the attenuation due to inundation. In a purely progressive wave system, any attenuation due to inundation can only alter the tidal amplitudes landward of the attenuation. In a reflective, standing wave system, though, inundation even at the head of the estuary can attenuate the tidal range throughout the bay and even in adjacent tidal basins.

In addition to the incoming tidal wave constituents, overtides generated within a basin may add to or subtract from the high water elevation and appear to be very sensitive to shoreline conditions and incident tidal wave amplitudes. The importance of overtide generation depends on the tidal amplitude relative to water depth, at least for the depth-dependent

M4 generation mechanisms. In an amplifying basin the greatest tidal ranges will be near the head of the basin, coinciding with shallow mean depths. One would then expect that a seaward propagating M4 wave is not at all unusual in such basins. With the complexity of a seaward propagating M4 and landward propagating M2, along with the potential for M4 resonance, modulation of overtides by sea level rise is a non-obvious but important aspect of predicting peak sea level within tidal basins.

The net physical response to coastal sea level rise clearly depends on a broad set of factors. This chapter focuses on M2 tidal forcing, but diurnal tides and interactions between diurnal and semidiurnal tides are likely significant. The long wavelength of diurnal tides leads to phasing closer to a standing wave, though net amplification is typically smaller at longer wave lengths (for example, the analysis of section 4.54.5 when applied to the s0 scenario yields 24% less amplification when the tidal period is doubled). Similarly, the dissipative effects of inundation are likely to be less important for diurnal tides.

Perhaps the largest uncertainty in predicting what will happen in a particular estuary is the unknown evolution of morphology, whether by natural or managed actions. Understanding how basins respond to sea level rise when morphology is kept static is the first step towards understanding what natural changes are likely to occur and what management decisions may be deemed necessary.

4.8 Conclusions

The effects of basins becoming deeper and the effects of increased inundation under sea level rise are approximately separated with a collection of numerical experiments with varying coastal sea level rise and shoreline configurations. Comparisons of phase and energy flux of the M2 tidal wave show that deepening decreases the influence of friction, while inundation adds considerable dissipation in the margins of the basin.

Deepening allows additional tidal amplification (consistent with an amplified estuary in the parlance of [8]), which was observed in all hard shoreline cases with sea level rise. Changes in energy flux and tidal phase are spatially variable and depend on details of each basin. The long, convergent southern arm becomes more reflective when deepened, while the shorter, transmissive northern arm primarily shows a lateral redistribution of the landward energy flux towards the shoals.

In both the progressive wave northern arm of San Francisco Bay and the standing wave southern arm, inundation introduced large energy sinks at the margins of the bays. Most inundation occurred off perimeter sloughs and rivers, causing the most drastic changes in tidal phasing and tidal prism at points where these features join the larger bays. Energy sinks in newly inundated regions caused a progressive shift in tidal phasing, a decrease in tidal amplification, and an increase in the landward tidal energy flux.

In the case of South San Francisco Bay, local changes in the shoreline altered both the tidal range within the basin, and also the magnitude of the reflected wave. The reflected wave subsequently affected tidal range in other parts of the domain. In contrast, local changes

in the shoreline of San Pablo Bay have limited effects on tidal range downstream of San Pablo Bay, as a much smaller fraction of tidal energy is reflected back compared to South San Francisco Bay.

Another key distinction between the two basins is the role of lateral dynamics and the degree to which the system is better described as a one-dimensional or two-dimensional tidal basin. South San Francisco Bay, with its higher aspect ratio and the inundating areas concentrated near the head of the estuary, is amenable to a one-dimensional treatment, and tidal amplification in portions of this bay can be roughly described by one-dimensional, analytic estuary models. San Pablo Bay exhibits greater lateral variation, due to its low aspect ratio and the fact that majority of the inundating areas are displaced laterally from the thalweg, rather than being situated near the head of the channel.

The present study focuses on the semidiurnal tide, and while the baseline M2 amplitude is much larger than the amplitudes of overtides, the variation in M4 amplitude across numerical experiments is comparable to the variation in M2 amplitude between experiments. Together with significant variation in the phase of the M4 wave relative to the M2 wave, overtides are a significant component of the variation in high water. Depth-dependent nonlinearities in the shallow water equations are the most likely sources of M4 generation, consistent with the large expanses of shoals in San Francisco Bay with a depth on the same order as the tidal range. Further work is needed to better understand the generation and dissipation of the M4 wave in intertidal bay margins.

Overall, the coupling between sea level rise, tidal amplification, and inundation is important and must be taken into account for accurate assessment of future restoration and mitigation questions. Quantitatively approaching these problems in a manner suitable for aiding management decisions still presents many challenges. Realistic boundary forcing will be required for the results to accurately predict future conditions. Some amendments to the boundary forcing are simple, such as including the full range of tidal harmonics, but in the grander scheme of climate change, many aspects of the forcing are difficult to quantify. Establishing future forcing conditions was a central aspect of [35], including measures of increased storm surge, trends in tidal constituents, El Niño, and of course sea level rise. In regions where mean sea surface height is also affected by river discharge, such as Suisun Bay in the current study area, river flow predictions will also be part of making model scenarios as accurate as possible. The biggest challenge, though, for coupled tidal-inundation modeling is geomorphology. Each scenario in this chapter assumes a static landscape, but in reality the landscape will evolve in step with sea level rise. Studies of past sea level variation have found that marsh elevation generally keeps up with rising sea levels [76], but recent studies such as [70] have found that accretion will not be sufficient relative to the highest anticipated rates of sea level rise. Predicting the elevation response of individual marshland areas will be difficult and likely variable between sites. The human element is also difficult to predict, as the fate of mitigation efforts is tied to funding and public sentiment, no longer a question of environmental fluid mechanics but public policy. A best-effort approach will probably require ensembles of model scenarios covering many variations in future topography and forcing, not unlike the current state of global climate models.

In many estuaries and bays rising sea level in the coastal ocean will lead to newly inundated areas. To a degree this inundation can mitigate sea level rise by decreasing tidal amplification within the basins. Reinforcing and hardening impacted shorelines can actually increase flooding risk in adjacent areas, and in highly reflective basins the effects can be far-reaching. Restoration of tidal marshland and construction of new low-lying tidal areas offer significant protection from rising tides by dissipating incident tidal energy, and these benefits may extend well beyond the areas directly sheltered by marshland.

Chapter 5

Conclusion

Understanding the physical processes at work in estuaries, and how those processes interact with human action, is a broad and ongoing field of study. This dissertation has explored fundamental dispersion processes, the artifacts of numerical models when applied to mixing and transport problems, and the application of numerical models to questions of sea level rise and flooding. Each of the preceding chapters ends with concluding remarks specific to the respective material; rather than repeating and further summarizing those remarks, the following section instead considers potential research directions for the future, namely improvements to the architecture of numerical models and eddy-scale investigations into mechanics of estuarine mixing.

In the physical context, there are opportunities for exploration into smaller scale dynamics, where eddy-resolving models and high resolution spatiotemporal observations open the door to answering questions related to the mechanics of mixing. The foray into sea level rise in chapter 4 brings up an additional, open-ended question: how can we make progress in understanding the future of complex systems which combine scientific questions of fluid mechanics, geomorphology and marsh ecology, engineering questions of levee construction and wetland restoration, all with the social and political questions of what actions we are willing to support. New approaches to interdisciplinary studies and the ways we communicate our findings may both have to adapt as we study the complex intersection of science, engineering and policy. These new approaches may range from visualization to error analysis and stochastic approaches.

A common remark from scientists in physical oceanography is that essentially every proposal has a modeling component now, increasingly requested and sometimes even required by funding agencies. Studies which may not have involved any modeling ten years ago now have some model-based aspect, and pure-modeling studies involve increasingly complex simulations which combine multiple models across sub-disciplines to characterize the coupling between an array of processes. Observational powers are also increasing, creating loads of scientific data which rival the volume of model data created by modern simulations. As the observational data streams grow, the same technologies which support large-scale modeling will also be indispensable for analyzing large observational datasets. There are many oppor-

tunities to improve not just the numerics and parameterizations in these models but also to improve the interface between the scientist and the model. To match the increasing reliance on and ubiquity of these models there is a need to move modeling from the domain of the specialist to a commodity tool accessible for all scientists. At the same time, the fast pace of hardware development and the increasing power of compiler and language technologies will necessitate strong and rich abstractions between the scientist and the numerical kernel.

As models are developed and utilized by a broader audience and for a wider range of applications, there is increasing value in looking into ways of commoditizing modeling technology. More simply expressed, setting up a numerical model should not solely be the purview of modeling specialists. Standardization of steps in the modeling process will help, but there is also a need to make the models and associated tools smarter. In some ways, creating a numerical model is like typesetting: it was once a laborious process left to the specialist, but algorithms improved, new systems were developed, and today anyone can publish a high quality typeset article, concerning themselves only with content and leaving the low-level details to the codified knowledge of the algorithmist. A similar trend in model development is necessary to make numerical modeling usable for the broader community.

There are plenty of potential steps in this direction, some general and some specific to the physical oceanography community. For many estuarine models, the grid generation step requires a significant time investment. Advances in adaptive methods like Gerris [57] and Chombo [16] are perhaps leading the way, but the efficiency of these adaptive methods still lags far behind curvilinear codes like ROMS and even static unstructured codes like SUNTANS. Mixed triangle/quadrangle grids have seen great success in combining advantages of unstructured and structured models, but require extensive manual input during the generation step.

The software architecture behind the scenes of numerical models is ripe for innovation. Quantifying model error should also come to the forefront and be supported as a first-class objective of models. The mathematical machinery for error estimation is well established but without strong software architecture which supports extensible numerics, such modifications are time-consuming and lead to brittle implementations. By improving abstractions within the modeling frameworks we can improve reliability and flexibility. Separating the description of the computation from the implementation is valuable, perhaps exploiting the notion of domain-specific languages. This will allow scientists and engineers more freedom to choose the most appropriate implementation, and cleanly separate the roles of physical oceanographer and applied mathematician.

Observational and computational capabilities are continually enabling investigations into broader ranges of spatial and temporal scales. The small-scale portions of these data are particularly rich, where turbulent and other transient flow features evolve. In a sense, there appears to be a gap in the scales on which we focus. Basin-scale circulations are well-characterized by numerical models and sparse observational arrays. Inertial and viscous range measurements from acoustic doppler velocimetry can say a lot about local turbulent quantities. Topographically-induced horizontal eddies are likely important, but it is not clear how these transient, intermediate-scale features affect mixing, dispersion or smaller-

scale turbulence. Autonomous underwater vehicles, passive drifters, high resolution CODAR, and drone-based remote sensing are all under active development. These observational platforms can be expected to provide increasingly detailed snapshots of environmental flows at intermediate and small scales.

At the same time as observational capabilities are extending into this range of spatial scales, numerical models are also reaching a point where eddy-resolving estuarine simulations will be possible. The LES approaches typical of atmospheric and engineering CFD codes have established capabilities in resolving portions of the turbulence spectrum. A major challenge along the way to fully three-dimensional large eddy simulation in estuaries, though, will be confronting the highly anisotropic spatial scales. Horizontal length scales range from a kilometer to a hundred kilometers while vertical motions are constrained to scales on the order of a meter by the sea surface, the bed, and stratification. Existing LES methods will have to be adapted and perhaps hybridized with the RANS models common in modern estuarine codes.

One of the defining features of environmental fluid mechanics, under the umbrella of environmental engineering, is the acknowledgment that society depends on and has an impact on the environment. In cases like the construction of new wetlands, the impacts are intentional. In other cases the impacts are not only unintentional but may be so subtle as to elude detection for many years, such as anthropogenic climate change. Questions of climate change also highlight the importance of multi-scale modeling and data analyis. While the root causes of climate change exist primarily at the global scale, the effects are observable at many scales and it is not unlikely that most dramatic effects will be felt at the scales of sloughs and intertidal shorelines. In all cases, though, the systems in question are complex and represent the combined influences of many factors, spanning multiple disciplines. Adapting modeling tools to support the integration of processes from various sub-fields will be essential for informing management decisions. The complexity of these systems and nuance of the results also dictate a need for improved communication of findings. Visualization is likely to become more important in a world of data-driven discovery.

Estuaries are home to some of the richest ecosystems but also experience the impacts of human action in myriad ways. The underlying theme of the research presented here and the potential trajectories for future work is to improve the tools with which we can predict future scenarios and to further the understanding of how human action affects these water bodies.

References

[1] F. Alouges, G. Coq, and E. Lorin. Two-dimensional extension of the reservoir technique for some linear advection systems. *Journal of Scientific Computing*, 31(3):419, 2007.

[2] R. Aris. On the dispersion of a solute in a fluid flowing through a tube. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 235:67–77, 1956.

[3] N.G. Barton. On the method of moments for solute dispersion. *Journal of Fluid Mechanics*, 126:205–218, 1983.

[4] P. D. Bates, K. J. Marks, and M. S. Horritt. Optimal use of high-resolution topographic data in flood inundation models. *Hydrological Processes*, 17(3):537, 2003.

[5] D. Bouche, J.-M. Ghidaglia, and F. Pascal. Error estimate and the geometric corrector for the upwind finite volume method applied to the linear advection equation. *SIAM Journal of Numerical Analysis*, 43(2):578–603, 2005.

[6] P. D. Bromirski, R. E. Flick, and D. R. Cayan. Storminess variability along the California coast: 1858–2000. *Journal of Climate*, 16:982–993, 2003.

[7] H. Burchard and H. Rennau. Comparative quantification of physically and numerically induced mixing in ocean models. *Ocean Modelling*, 20(3):293, 2008.

[8] H. Cai, H. H. G. Savenije, and M. Toffolon. A new analytical framework for assessing the effect of sea-level rise and dredging on tidal damping in estuaries. *Journal of Geophysical Research*, 117(C9), 2012.

[9] K. S. Carignan, L. A. Taylor, B. W. Eakins, R. J. Caldwell, D. Z. Friday, P. R. Grothe, and E. Lim. Digital Elevation Models of Central California and San Francisco Bay: Procedures, Data Sources and Analysis. Technical Memorandum, NOAA, Boulder, CO, 2011.

[10] V. Casulli. Semi-implicit finite difference methods for the two-dimensional shallow water equations. *Journal of Computational Physics*, 86:56–74, 1990.

[11] V. Casulli and R. Walters. An unstructured grid, three-dimensional model based on the shallow water equations. *International Journal for Numerical Methods in Fluids*, 32:331–348, 2000.

[12] C. Chen, R. C. Beardsley, and G. Cowles. An unstructured-grid, finite-volume coastal ocean model system. *Advances in Computational Oceanography*, 19(1):78–89, 2006.

[13] H. Childs, E. S. Brugger, K. S. Bonnell, J. S. Meredith, M. Miller, B. J. Whitlock, and N. Max. A contract-based system for large data visualization. In *Proceedings of IEEE Visualization 2005*, pages 190–198, 2005.

[14] Y.-J. Chou, R. C. Holleman, S. Lee, O. B. Fringer, M. T. Stacey, J. R. Koseff, and S. G. Monismith. Three-dimensional coupled wind-wave and cohesive sediment transport modeling in san francisco bay. in preparation.

[15] V. P. Chua and O. B. Fringer. Sensitivity analysis of three-dimensional salinity simulations in North San Francisco Bay using the unstructured-grid SUNTANS model. *Ocean Modelling*, 2011.

[16] P. Colella, M. Dorr, J. Hittinger, D. F. Martin, and P. McCorquodale. High-order finite-volume adaptive methods on locally rectangular grids. *Journal of Physics: Conference Series*, 180(1):012010, 2009.

[17] J. N. Collins and R. M. Grossinger. Synthesis of scientific knowledge concerning estuarine landscapes and related habitats of the South Bay Ecosystem. Technical report of the South Bay Salt Pond Restoration Project, San Francisco Estuary Institute, Oakland, CA, 2004.

[18] M. S. Darwish and F. Moukalled. TVD schemes for unstructured grids. *International Journal of Heat and Mass Transfer*, 46:599–611, 2003.

[19] B. C. Douglas. Global sea rise: a redetermination. *Surveys in Geophysics*, 18(2-3):279–292, 1997.

[20] D. Durran. *Numerical Methods for Fluid Dynamics*. Texts in Applied Mathematics. Springer, 2nd edition, 2010.

[21] California DWR. An estimate of daily average Delta outflow. Website, 2011. Accessed May 1, 2011.

[22] S. Feng, R. T. Cheng, and P. Xi. On tide-induced Lagrangian residual current and residual transport 2. Residual transport with application in South San Francisco Bay, California. *Water Resources Research*, 22(12):1635–1646, 1986.

[23] H. B. Fischer, E. J. List, R. C. Y. Koh, J. Imberger, and N. H. Brooks. *Mixing in Inland and Coastal Waters*. Academic Press, Inc., 1979.

[24] R. E. Flick, J. F. Murray, and L. C. Ewing. Trends in United States tidal datum statistics and tide range. *Journal of Waterway, Port, Coastal and Ocean Engineering*, 129(4):155–164, 2003.

[25] A. Foxgrover, R. E. Smith, and B. E. Jaffe. Suisun Bay and Delta Bathymetry. Bathymetric Dataset, USGS, 2005.

[26] A. C. Foxgrover, D. P. Finlayson, and B. E. Jaffe. 2010 Bathymetry and digital elevation model of Coyote Creek and Alviso Slough, south San Francisco Bay, California. Open File Report, USGS, 2011.

[27] A. C. Foxgrover and B. E. Jaffe. South San Francisco Bay 2004 Topographic Lidar Survey: Data Overview and Preliminary Quality Assessment. Open-File Report, USGS, 2005.

[28] O. B. Fringer, M. Gerritsen, and R. L. Street. An unstructured-grid, finite-volume, nonhydrostatic, parallel coastal ocean simulator. *Ocean Modelling*, 14:139–173, 2006.

[29] P. L. George and E. Seveno. The advancing-front mesh generation method revisited. *International Journal for Numerical Methods in Engineering*, 37:3605–3619, 1994.

[30] S. K. Godunov. A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Mat. Sb. (N.S.)*, 47(89)(3):271–306, 1959.

[31] E. S. Gross. *Numerical modeling of hydrodynamics and scalar transport in an estuary*. PhD thesis, Stanford University, Palo Alto, California, December 1997.

[32] E. S. Gross, M. L. MacWilliams, R. C. Holleman, and T. A. Hervier. Particle tracking model: testing and applications report. Study, Interagency Ecological Program, Berkeley, CA, May 2011.

[33] C. W. Hirt. Heuristic stability theory for finite-difference equations. *Journal of Computational Physics*, 2:339–355, 1968.

[34] H. Jasak. *Error Analysis and Estimation for the Finite Volume Method with Applications to Fluid Flows*. PhD thesis, Imperial College of Science, Technology and Medicine, June 1996.

[35] N. Knowles. Potential inundation due to rising sea levels in the San Francisco Bay Region. *San Francisco Estuary and Watershed Sceince*, 8(1), 2010.

[36] P. Kundu and I. Cohen. *Fluid Mechanics*. Elsevier, San Diego , California, 3rd edition, 2004.

[37] E. M. LaBolle, J. Quastel, G. E. Fogg, and J. Gravner. Diffusion processes in composite porous media and their numerical integration by random walks: Generalized stochastic differential equations with discontinuous coefficients. *Water Resources Research*, 36(3):651–662, 2000.

[38] J. R. Lacy, S. Gladding, A. Brand, A. Collignon, and M.T. Stacey. Lateral baroclinic forcing enhances transport of sediment from shallows to channel in an estuary. *Estuaries and Coasts*, in review, 2013.

[39] F. Lagoutière and B. Després. Genuinely multi-dimensional non-dissipative finite-volume schemes for transport. *International Journal of Applied Mathematics and Computer Science*, 17(3):321, 2007.

[40] S. Lanzoni and G. Seminara. On tide propagation in convergent estuaries. *Journal of Geophysical Research*, 103(C13):30,793–30,812, 1998.

[41] W. G. Large and S. Pond. Open ocean momentum flux measurements in moderate to strong winds. *Journal of Physical Oceanography*, 11:324–336, 1981.

[42] R. J. LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations*. Society for Industrial and Applied Mathematics, Philadelphia, 2007.

[43] C. Li and J. O'Donnell. The Effect of Channel Length on the Residual Circulation in Tidally Dominated Channels . *Journal of Physical Oceanography*, 35:1826–1840, 2005.

[44] C. Li and A. Valle-Levinson. A two-dimensional analytic tidal model for a narrow estuary of arbitrary lateral depth variation: the intratidal motion. *Journal of Geophysical Research*, 104(C10):23,525–23,543, 1999.

[45] M. MacWilliams and E. Gross. UnTRIM San Francisco Bay-Delta Model Calibration Report. Technical report, California Dept. of Water Resources, March 2007.

[46] F. Marcuzzi., M. M. Cecchi, and M. Venturin. An anisotropic unstructured triangular adaptive mesh algorithm based on error and error gradient information. *Mathematics and Computers in Simulation*, 78(5-6):645, 2008.

[47] G. L. Mellor and T. Yamada. Development of a turbulence closure-model for geophysical fluid problems. *Reviews of Geophysics*, 20(4):851–875, 1982.

[48] M. Namin, B. Lin, and R. A. Falconer. Modelling estuarine and coastal flows using an unstructured triangular finite volume method. *Advances in Water Resources*, 27(12):1179–1197, 2004.

[49] NCALM. LiDAR data acquired by the National Center for Airborne Laser Mapping. online, 2003.

[50] NOAA. 2009 - 2011 CA Coastal Conservancy Coastal Lidar Project: Hydro-flattened Bare Earth DEM . http://www.csc.noaa.gov/dataviewer/, 2012.

[51] NOAA. NOAA Tides and Currents (online). http://tidesandcurrents.noaa.gov/, 2013. Accessed March 20, 2013.

[52] L.-Y. Oey, T. Ezer, C. Hu, and F. E. Muller-Karger. Baroclinic tidal flows and inundation processes in Cook Inlet, Alaska: numerical modeling and satellite observations. *Ocean Dynamics*, 57(3):205, 2007.

[53] A. Okubo. Effect of shoreline irregularities on streamwise dispersion in estuaries and other embayments. *Netherlands Journal of Sea Research*, 6(1-2):213–224, 1973.

[54] S. J. Owen. A survey of unstructured mesh generation technology. In *Proceedings of the 7th International Meshing Roundtable*, pages 239–267, 1998.

[55] B. Parker. *The relative importance of the various nonlinear mechanisms in a wide range of tidal interactions (review).* Wiley, 1991.

[56] F. W. Platzek, G. S. Stelling, J. A. Jankowski, and R. Patzwahl. On the representation of bottom shear stress in z-layer models. In *10th International Conference on Hydroinformatics*, 2012.

[57] S. Popinet and G. Rickard. A tree-based solver for adaptive ocean modelling. *Ocean Modelling*, 16(3-4):224–249, 2007.

[58] S. Rahmstorf. A semi-empirical approach to projecting future sea-level rise. *Science*, 315(5810):368–70, 2007.

[59] S. Ramp, F. Chavez, and L. Breaker. Sea level off California: rising or falling? `http://www.cencoos.org/sections/news/sea_level.shtml`, July 2013.

[60] O. N. Ross and J. Sharples. Recipe for 1-D Lagrangian particle tracking models in space-varying diffusivity. *Limnology and Oceanography: Methods*, 2:289–302, 2004.

[61] P. G. Saffman. The effect of wind shear on horizontal spread from an instantaneous ground source. *Quarterly Journal of the Royal Meteorological Society*, 88(378):382–393, 1962.

[62] S. Salah-Mars and M. W. McCann Jr. Delta Risk Management Strategy: Water Analysis Module. Technical memorandum, URS Corporation, Oakland, CA, 2007.

[63] L. P. Sanford. Modeling a dynamically varying mixed sediment bed with erosion, deposition, bioturbation, consolidation and armoring. *Computers and Geosciences*, 34:1263–1283, 2008.

[64] S. Saramul and T. Ezer. Tidal-driven dynamics and mixing processes in a coastal ocean model with wetting and drying. *Ocean Dynamics*, 60(2):461, 2010.

[65] H. H. G. Savenije. Lagrangian Solution of St Venants equations for alluvial estuary. *Journal of Hydraulic Engineering*, 118(8):1153–1163, 1992.

[66] H. H. G. Savenije, M. Toffolon, J. Haas, and E. J. M. Veling. Analytical description of tidal dynamics in convergent estuaries. *Journal of Geophysical Research*, 113(C10), 2008.

[67] J. R. Shewchuk. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In Ming C. Lin and Dinesh Manocha, editors, *Applied Computational Geometry: Towards Geometric Engineering*, volume 1148 of *Lecture Notes in Computer Science*, pages 203–222. Springer-Verlag, May 1996. From the First ACM Workshop on Applied Computational Geometry.

[68] M. S. Spydell and F. Feddersen. The effect of a non-zero Lagrangian time scale on bounded shear dispersion. *Journal of Fluid Mechanics*, 691:69–94, 2012.

[69] H. Stommel and H. G. Farmer. On the nature of estuarine circulation. Technical Report, Woods Hole Oceanographic Institution, 1952.

[70] K. Swanson, J. Z. Drexler, D. H. Schoellhamer, K. Thorne, K. Spragens, and J. Takekawa. Development of a data-driven numerical model for San Francisco Bay marsh habitat sustainability. In *American Geophysical Union, Fall Meeting*, 2011.

[71] P. K. Sweby. High resolution schemes using flux limiters for hyperbolic conservation laws. *SIAM Journal on Numerical Analysis*, 21(5):995–1011, 1984.

[72] S. A. Talke and D. A. Jay. Nineteenth century North American and Pacific tidal data: lost or just forgotten? *Journal of Coastal Research*, 2013.

[73] A. Tam, D. Ait-Ali-Yahia, M. P. Robichaud, M. Moore, V. Kozel, and W. G. Habashi. Anisotropic mesh adaptation for 3D flows on structured and unstructured grids. *Computer Methods in Applied Mech. Engr.*, 189:1205–1230, 2000.

[74] G. I. Taylor. Dispersion of soluble matter in solvent flowing slowly through a tube. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 219:186, 1953.

[75] G. I. Taylor. The Dispersion of matter in turbulent flow through a pipe. *Proceedings of the Royal Society of London Series A-Mathematical and Physical Sciences*, 223(1155):446–468, 1954.

[76] J. G. Titus, T. W. Kana, B. J. Baca, W. C. Eiser, M. L. Williams, T. V. Armentano, R. A. Park, and C. L. Cloonan. Greenhouse effect, sea level rise and coastal wetlands. Technical report, Environmental Protection Agency, Washington, DC, 1988.

[77] E. Turkel. Accuracy of schemes with nonuniform meshes for compressible fluid flows. *Applied Numerical Mathematics*, 2:529–550, 1986.

[78] USGS. USGS water quality monitoring. Website, 2011. Accessed May 1, 2011.

[79] L. C. van Rijn. Analytical and numerical analysis of tides and salinities in estuaries; part I: tidal wave propagation in convergent estuaries. *Ocean Dynamics*, 61(11):1719, 2011.

[80] B. Wang, O. B. Fringer, S. N. Giddings, and D. A. Fong. High-resolution simulations of a macrotidal estuary using SUNTANS. *Ocean Modelling*, 28(1-3):167, 2009.

[81] B. Wang, G. Zhao, and O. B. Fringer. Reconstruction of vector fields for semi-lagrangian advection on unstructured, staggered grids. *Ocean Modelling*, 40(1):52 – 71, 2011.

[82] R. F. Warming and B. J. Hyett. The modified equation approach to the stability and accuracy analysis of finite-difference methods. *Journal of Computational Physics*, 14(2):159–179, 1974.

[83] C. J. Willmott, S. G. Ackleson, R. E. Davis, J. J. Feddema, K. M. Klink, D. R. Legates, J. O'Donnell, and C. M. Rowe. Statistics for the evaluation and comparison of models. *Journal of Geophysical Research*, 90(C5):8995–9005, 1985.

[84] W. R. Young and S. Jones. Shear dispersion. *Physics of Fluids A: Fluid Dynamics*, 3(5):1087–1101, 1991.

[85] W. R. Young, P. B. Rhines, and C. J. R. Garrett. Shear-flow dispersion, internal waves and horizontal mixing in the ocean. *Journal of Physical Oceanography*, 12:515–527, 1982.

Appendix A

Calculation of Diffusion Coefficients

Starting from the closed-form recurrence expression for B_1^n:

$$\begin{aligned} B_1^n &= \frac{1}{1-\lambda}\left[C_\alpha A_1^n + C_\gamma A_2^n\right] - \frac{\Delta t}{(1-\lambda)^2}\left[C_\alpha A_{1,t}^n + C_\gamma A_{2,t}^n\right] \\ &\quad + \frac{(1+\lambda)\,\Delta t^2}{2\,(1-\lambda)^3}\left[C_\alpha A_{1,tt}^n + C_\gamma A_{2,tt}^n\right] + O(l^3) \end{aligned} \tag{A.1}$$

And the update equation for A_3^{n+1}:

$$A_3^{n+1} = (1 - C_\beta)\, A_3^n + C_\beta B_1^n \tag{A.2}$$

Expressions for the volumetric Courant numbers are:

$$C_\alpha = \frac{4U\Delta t \sin\theta}{\sqrt{3}l} \tag{A.3}$$

$$C_\beta = \frac{4U\Delta t \sin\left(\theta + \frac{\pi}{3}\right)}{\sqrt{3}l} \tag{A.4}$$

$$C_\gamma = \frac{4U\Delta t \sin\left(\theta + \frac{2\pi}{3}\right)}{\sqrt{3}l} \tag{A.5}$$

Taylor expansions of A_1, A_2 and respective derivatives, where a lack of subscript on A implies evaluation at A_3:

$$A_i^n = A^n + \delta_{xi}A_x + \delta_{yi}A_y + \frac{\delta_{xi}^2}{2}A_{xx} + \frac{\delta_{yi}^2}{2}A_{yy} + \delta_{xi}\delta_{yi}A_{xy} + O\left(l^3\right) \tag{A.6}$$

$$(A_i^n)_t = -U\left(A_x + \delta_{xi}A_{xx} + \delta_{yi}A_{xy}\right) + O\left(l^2\right) \tag{A.7}$$

$$(A_i^n)_{tt} = U^2 A_{xx} + O\left(l\right) \tag{A.8}$$

While the expansions for first and second time derivatives have been truncated at $O\left(l^2\right)$ and $O\left(l\right)$, the coefficients on these terms after substitution below retain the overall $O\left(l^3\right)$ bound

on the modified equation. Distances between cell centers are defined:

$$\delta_{x1} = -l\cos\left(\theta - \frac{\pi}{3}\right) \tag{A.9}$$

$$\delta_{y1} = -l\sin\left(\theta - \frac{\pi}{3}\right) \tag{A.10}$$

$$\delta_{x2} = -l\cos\theta \tag{A.11}$$

$$\delta_{y2} = -l\sin\theta \tag{A.12}$$

Substituting (A.3)-(A.12) into (A.1), B_1^n can be written in the form

$$B_1^n = DA + D^x A_x + D^{xx} A_{xx} + D^{xy} A_{xy} + D^{yy} A_{yy} + O\left(l^3\right) \tag{A.13}$$

Where the coefficients are identified by superscripts, and subscripts continue to denote partial derivatives. Expressions for each of the coefficients simplify to

$$D = 1 \tag{A.14}$$

$$D^x = -\frac{\sqrt{3}}{4\sin\left(\theta + \frac{\pi}{3}\right)} l \tag{A.15}$$

$$D^{xx} = \frac{\sin\left(\theta + \frac{\pi}{3}\right)\left[\cos^2\theta\cos\left(\theta + \frac{\pi}{6}\right) + \sin\theta\cos^2\left(\theta - \frac{\pi}{3}\right)\right] - \frac{3}{8}}{2\sin^2\left(\theta + \frac{\pi}{3}\right)} l^2 - \frac{\sqrt{3}}{8\sin\left(\theta + \frac{\pi}{3}\right)} Ul\Delta t \tag{A.16}$$

$$D^{yy} = \frac{\sin(3\theta)}{8\sin\left(\theta + \frac{\pi}{3}\right)} l^2 \tag{A.17}$$

$$D^{xy} = \frac{\frac{\sqrt{3}}{4}\sin\theta - \cos\theta\sin^2\theta}{\sin\left(\theta + \frac{\pi}{3}\right)} l^2 \tag{A.18}$$

Now considering (3.4), the update equation for A_3, it remains to expand A_3^{n+1} in time and substitute for B_1^n to get

$$\begin{aligned} &A + \Delta t A_t + \frac{\Delta t^2}{2} A_{tt} + O\left(\Delta t^3\right) \\ &\quad = (1 - C_\beta) A + C_\beta \left(DA + D^x A_x + D^{xx} A_{xx} + D^{xy} A_{xy} + D^{yy} A_{yy} + O\left(l^3\right)\right) \end{aligned} \tag{A.19}$$

Canceling A and replacing second-order time derivatives with spatial derivatives and omitting all higher-order terms for brevity, we obtain the modified equation

$$A_t - \frac{C_\beta D_x}{\Delta t} A_x = \left[\frac{C_\beta D_{xx}}{\Delta t} - \frac{U^2 \Delta t}{2}\right] A_{xx} + \frac{C_\beta D_{xy}}{\Delta t} A_{xy} + \frac{C_\beta D_{yy}}{\Delta t} A_{yy} \tag{A.20}$$

Considering each term separately, the effective advective speed

$$\begin{aligned} U_{eff} &= -\frac{C_\beta D_x}{\Delta t} \\ &= -\frac{4U\sin\left(\theta+\frac{\pi}{3}\right)}{\sqrt{3}l}\left[-\frac{\sqrt{3}l}{4\sin\left(\theta+\frac{\pi}{3}\right)}\right] \\ &= U \end{aligned} \tag{A.21}$$

Longitudinal diffusion:

$$\begin{aligned} 2K^x &= \frac{C_\beta D_{xx}}{\Delta t} - \frac{U^2\Delta t}{2} \\ &= \left[\frac{2}{\sqrt{3}}\left(\cos^2\theta\cos\left(\theta+\frac{\pi}{6}\right)+\sin\theta\cos^2\left(\theta-\frac{\pi}{3}\right)\right)-\frac{\sqrt{3}}{4\sin\left(\theta+\frac{\pi}{3}\right)}\right]Ul - U^2\Delta t \end{aligned} \tag{A.22}$$

Lateral diffusion:

$$\begin{aligned} 2K^y &= \frac{C_\beta D_{yy}}{\Delta t} \\ &= \frac{Ul\sin(3\theta)}{2\sqrt{3}} \end{aligned} \tag{A.23}$$

And the cross-diffusion coefficient:

$$\begin{aligned} 2K^{cross} &= \frac{C_\beta D_{xy}}{\Delta t} \\ &= Ul\left[\frac{4}{\sqrt{3}}\sin\theta\sin\left(\theta-\frac{\pi}{6}\right)\sin\left(\theta-\frac{\pi}{3}\right)\right] \end{aligned} \tag{A.24}$$

Appendix B

Connectivity-based Model Bathymetry

Algorithm

In potentially flooded and intertidal regions, it is more important to capture the connectivity (or lack thereof) between hydrologic features than to preserve cross-sectional area or mean depth. This appendix describes the method used for assigning bathymetry to each edge of the numerical model in a manner which explicitly includes connectivity. The base input to the algorithm is a high resolution DEM, typically a 1 m LiDaR dataset including intertidal and supertidal regions). Where available, additional datasets of the present-day extent of tidal action are used to correct for erroneous LiDaR "dams." These dams are a consequence of overpasses, bridges, and aerial clutter which appear as an obstruction obscuring the underlying channel.

Bathymetry in the numerical model is stored at both cell edges and cell centers. For the present application edge bathymetry is the primary concern, and cell depths follow from the deepest adjacent edge. For each edge, a graph search algorithm is executed to determine the minimum water surface elevation at which the two cells adjacent to the edge are connected. For a given water surface elevation η_w, cells are considered connected via their shared edge if there exists a path along pixels of the DEM which connects representative points within the two cells, subject to the conditions that all DEM pixels along the path are below the elevation η_w, and that the path lies entirely within the footprint of the two cells. Depending on prior knowledge of the region, one of two approaches is taken for choosing a representative point from each cell. In the general case where no additional information is available on the tidal regions of the domain, the lowest point of the cell is chosen as the representative point. When a single deep channel is flanked by areas of higher elevation, this approach is effective at maintaining connectivity along the channel and enforcing a lack of connectivity through the channel banks. For portions of the domain, however, additional information was available in the form of a polygon which overlays the extent of areas known to be tidally active. Incorporating this additional information into the edge-centered bathymetry was a two part

process. First, the medial axis of the tidal regions was extracted, yielding a branching set of polylines along the "skeleton" of the tidal regions, roughly similar to thalwegs. Pixels of the DEM along the medial axis lines were lowered as necessary to force the pixels to be intertidal. This corrects for missing connections in the LiDaR data due to bridges, culverts and other unresolved channel features. Second, the choice of representative point when evaluating cell-cell connections, when a cell overlaps the medial axis, is chosen to be on the medial axis. In cases where a cell overlaps multiple channels, this favors channels which are known to be tidal.

The results of a simple averaging approach are shown in figure B.1. A narrow, deep channel is flanked by levees of a similar width, and broad ponds lie beyond the pair of levees. The widths of the features are small compared to the 50 m grid spacing, and the averaging yields model bathymetry which fails to resolve the channel and significantly under-estimates the height of the levees preventing flow from one pond to the other. Figure B.2 shows the same region, now with the tidal area, medial axis of the tidal area, and resulting edge bathymetry. Channel connectivity is present at a more representative elevation, while levees flanking the channel on either side are resolved with much higher elevations. While the resolved channel is much wider than the true channel, the grid properly connects cells along the channel and prevents unrealistic overtopping flows from connecting the channel and the adjacent ponds.

This graph search approach is considerably more expensive than a typical averaging approach, but with the help of accelerated image processing methods (primarily a binary image connected components algorithm), the algorithm is not prohibitively slow. For a model grid with 937,759 cells and 1,413,110 edges, the bathymetry processing takes approximately 4 h running serially on a 2.4 GHz desktop workstation.

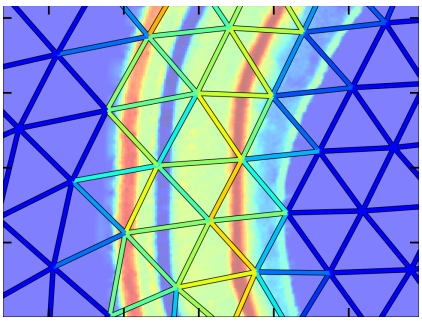

Figure B.1: Edge-centered depths taken from averaged bathymetry along edges. Background bathymetry is shown desaturated but otherwise with the same color scale as the edge bathymetry.

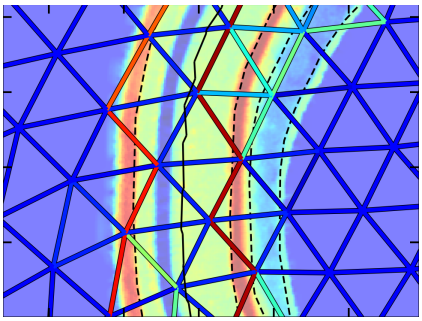

Figure B.2: Edge-centered depths by levee and channel embedding method. Dashed black lines show the boundary of the tidal polygon, and the solid black line the medial axis of that polygon.

Static analysis

A non-dynamic inundation analysis is applied to the computational grid as a simple means of understanding the efficacy of the above bathymetry processing to embed channels and levees. The analysis is similar to simply considering all areas below a given elevation contour to be inundated, but takes into account the connectivity between the bay, channels and ponds at a given sea surface elevation.

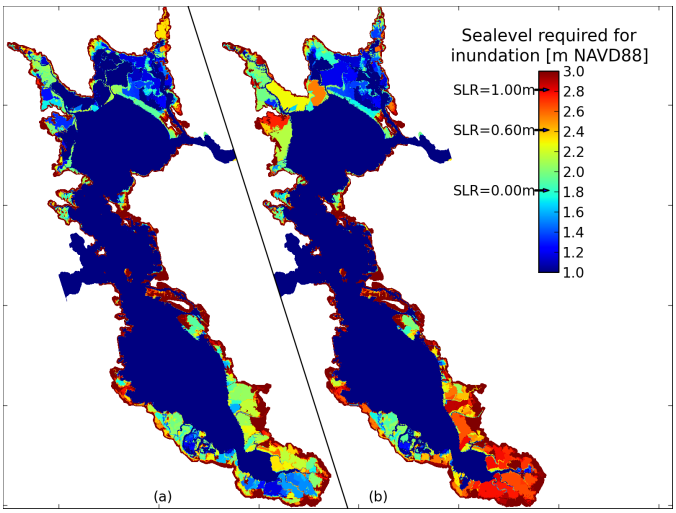

Figure B.3: Sea level of inundation for (a) edge-averaged bathymetry and (b) connectivity-based bathymetry. Arrows on the colorbar indicate the NAVD88 heights of MHHW in the coastal ocean under 0,0.6 and 1.0 m sea level rise.

Each grid cell is assigned a value equal to the sea surface elevation at which that cell

becomes connected to the open ocean boundary. The result is a spatial map of "sea level of inundation", though with no consideration of tidal dynamics. The method for this calculation stems directly from graph theory and connected components. Given the primary interest in ranges of sea level rise, a starting sea level elevation η_0=1 m NAVD88 is chosen, below present day mean high water throughout the domain. Considering the cells of the model grid as nodes in a graph, and edges in the model grid with elevation below η_0 as edges of a graph connecting those nodes, the problem is equivalent to a connected components analysis. The ocean component is defined as the component containing a reference point in the coastal ocean, and cells which are initially members of this component are labeled with a sea level of inundation of η_0. The remaining edges with elevations above η_0 are then sorted in ascending order of elevation, and incrementally added to the graph. At each step, the next lowest edge, with some elevation η_i, is added to the graph, and the connected components are updated with this new connection. Any cells which become part of the ocean component due to this connection are then given a sea level of inundation equal to η_i.

Figure B.3 compares this static inundation analysis as applied to the averaged bathymetry and connectivity-based bathymetry. From these plots it is clear that an abundance of connectivity information is lost when topography is averaged across 50 m edges. Panel B.3b also highlights the extensive preexisting tidal marshes adjacent to the Napa River in the northeast, and the extensive levees in the south which isolate low-lying salt ponds from the high tidal ranges of South San Francisco Bay.

Appendix C

Model Calibration and Validation

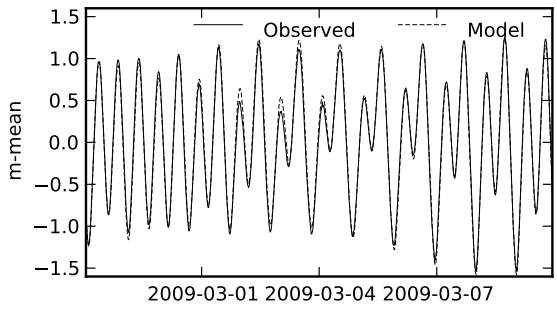

Figure C.1: Observed and predicted sea level near the San Mateo Bridge. Heights are relative to local MSL.

The tidal boundary condition was calibrated to match phase and amplitude of the sea surface height at the Golden Gate over the period February 24, 2009 to March 15, 2009, which required scaling measured tidal amplitude by 0.931, and compensating for a 120s phase lag. The model has been validated against observed tidal stage at two locations and depth-averaged velocity at two locations, over the period February 26, 2009 to March 9, 2009 (with the exception of the velocity validation in South Bay, for which observations are truncated at March 6, 2009). Model forcing for the validation run was taken from observed coastal ocean sea level as measured at Point Reyes, and observed winds from Point Reyes, Port Chicago, Alameda, Redwood City, Richmond, and Union City. River flows were included for the San Joaquin and Sacramento Rivers, where the Net Delta Outflow Index [21] was apportioned 25% and 75% respectively. Table C.1 summarizes the comparison between observations and model predictions.

Figure C.1 shows time series comparisons in South San Francisco Bay, near the 40 km mark in figure 4.1 and laterally situated at the eastern edge of the channel at the foot of

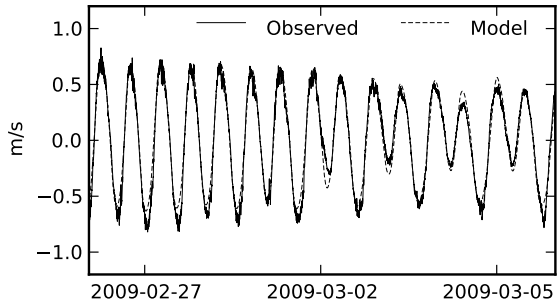

Figure C.2: Observed and predicted along-channel, depth-averaged velocity near the San Mateo Bridge, edge of the main channel. Velocities are rotated to respective principal directions, with positive up-estuary.

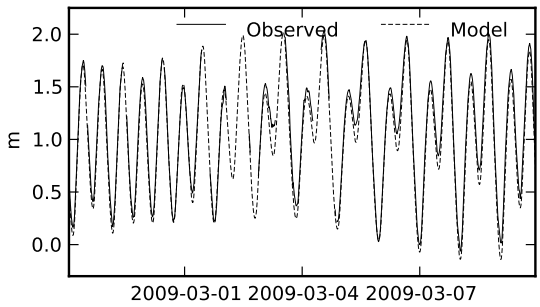

Figure C.3: Observed and predicted sea level at Mare Island. Heights are relative to NAVD88.

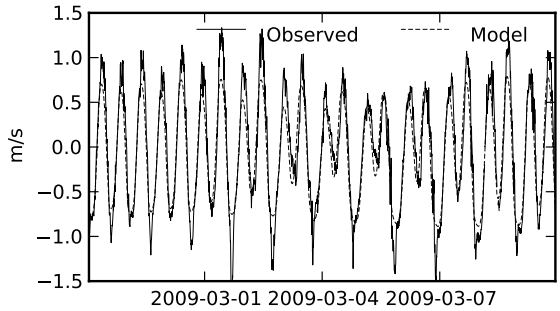

Figure C.4: Observed and predicted along-channel, depth-averaged velocity at Martinez, each rotated to respective principal direction with positive up-estuary.

Table C.1: Velocity and stage comparisons, forced with observed tides and winds. r denotes the Pearson correlation coefficient. Lags are computed as the time offset which maximizes the correlation coefficient. rms ratio is the ratio of model rms amplitude to observed rms amplitude.

	r	bias	lag	rms ratio
η @ -40 km	0.997	-0.003 m	92 s	1.009
u @ -40 km	0.983	0.02 m s^{-}1	-151 s	0.985
η @ 40 km	0.996	-0.076 m	-146 s	1.016
u @ 53 km	0.948	-0.06 m s^{-}1	-1517 s	0.826

the slope leading into the shoal. A storm system passed through between March 2, 2009 and March 4, 2009. Uncertainty in the exact distribution of wind-based surface stress is the likely cause for the trend of over-predicted sea surface height during this period. Depth averaged currents at the same location are shown in figure C.2.

Validation for sea surface height in San Pablo Bay is taken at the tide gauge at Mare Island, immediately west of the mouth of the Napa River. The comparison is shown in figure C.3. Long-term measurements of velocity in San Pablo Bay during the validation period were not available. Velocity measurements at the other end of Carquinez Strait are available for a site near the southern shore (near 53 km along the thalweg shown in figure 4.1). Comparison at this location is shown in figure C.4. Note that this site is beyond the intended study area, and validation here is adversely affected by proximity to the false deltas, decreased grid resolution outside the study area, and the highly energetic and spatially variable flows in this constricted tidal strait. Nonetheless, velocity phase and temporal patterns of variation in current magnitude are reasonably captured by the model.

Appendix D

Tools for Estuarine Modeling

This appendix outlines the collection of modeling tools developed in the course of the research for the preceding chapters. This includes a range of modifications to the SUNTANS hydrodynamic code, tools for preparing inputs to the model, and tools for processing model outputs.

D.1 Changes to the Numerical Kernel

The majority of the changes to the internal core of SUNTANS have been motivated by the specific applications covered in this dissertation. While these modifications may not represent novel advances in numerical methods, they have hopefully resulted in a more practical, usable modeling ecosystem. The SUNTANS code base utilized throughout this dissertation was forked from the trunk SUNTANS version in the spring of 2008. More recently, the trunk SUNTANS version has been made publicly available, and has seen significant enhancements, especially in terms of core numerics for advection of momentum. The features described below have so far been introduced only into the forked code base.

Scalar substepping

Many hydrodynamic models advance the freesurface via an explicit time-stepping method, and are therefore subject to a Courant number limitation based on the barotropic wave speed. SUNTANS, to the contrary, follows the formulation of [10] in which the freesurface perturbation is advanced semi-implicitly with the θ method. By selecting the unconditionally stable Eulerian-Lagrangian method for momentum of advection, the model is generally stable even for Courant numbers greater than unity.

In many cases the most restrictive time step limitation in SUNTANS is then due to scalar advection. Where advection of momentum in shallow estuaries is often captured sufficiently by the non-conserving ELM, in the typical estuarine modeling context there is a fundamental need to conserve scalar mass. Even slight discrepancies in the conservation of salt mass may

have significant impacts on the structure of density driven flows or the analysis of the salt balance. The basic scalar transport discretization in SUNTANS is first-order upwind; a less diffusive total-variation diminishing method is also available. Either method, though, requires that the timestep be sufficiently small to control the maximum Courant number in the domain (for first order upwind, this maximum is 1.0, while for TVD the maximum Courant number is 0.5). Note that this Courant number is formally the volumetric Courant number, utilizing the θ method semi-implicit velocity field, and corresponds to the fraction of a cell's volume which is flushed out of the cell in the passing of a single time step of the simulation.

In order to make the most of the implicit free surface solver and the unconditionally stable ELM for advection of momentum, a scalar substepping modification has been implemented for all scalar advection in SUNTANS. At each time step, once the new velocity field has been calculated, the global maximum volumetric Courant number is calculated. A number of substeps is chosen such that within this smaller time substep the maximum Courant number is within the limits defined by the base scalar advection scheme. In this manner a long time step is achievable for the hydrodynamics, and the scalar fields will be advanced with a shorter time step chosen as needed to ensure stability.

This feature is always active in the present version, and is applied for all scalar species including turbulent quantities and sediment species.

Manipulation of interprocessor boundaries

SUNTANS is implemented as a parallel code, based on MPI and a spatial partitioning of the computational grid. This partitioning is optimized by the PARMETIS library, and is configured to minimize the amount of interprocessor communication required at each step. While in general this is exactly the optimization one would desire, there are some hidden implications for hydrodynamics when ELM is used for advection of momentum. The problem stems from the fact that in the current implementation, the tracebacks for ELM are not allowed to span multiple processors. Essentially, processor boundaries function as blockades for advected momentum. This is not the case for conservative advection schemes such as upwind, since in that case the ghost cells provide sufficient information to advect momentum between processors. The crux of the issue between optimizing the partitions and ELM is that graph cuts for the partitioning tend to coincide with headlands, peninsulas and other features near which the advection of momentum is a significant term. While the weakness of ELM at partition boundaries is negligible in regions of low velocity gradients, it can lead to large errors and instabilities when velocity gradients coincide with interprocessor boundaries.

As a workaround for this issue, an option has been added to bundle groups of cells together for the purpose of partitioning. A set of axis-aligned rectangular regions are specified during the partitioning step in a file "contiguous.dat," containing lines such as

```
560400 4160700 560800 4161100
```

which each specify a region by the coordinates $x_{min}, y_{min}, x_{max}, y_{max}$. For the purposes of partitioning, all cells within that region will be lumped together and partitioned as a single entity (with a computational weight equal to the sum of the computational weights of the original cells). In most applications this step is unnecessary, but in locally advection-dominated regions this can be a useful tool for combatting numerically unstable eddies at headlands or other points of land.

Sediment transport

While there has been significant previous development of a SUNTANS-integrated suspended sediment, bed sediment and wind-wave model, described in [14], this development has tended towards the ideal of including as many processes as possible. The downside of such a broad formulation is that many processes require parameterizations which in turn require the user to have detailed knowledge of tunable coefficients that are obscure or difficult to determine. As a simpler method of modeling basic sediment transport processes in a more "turn-key" approach, the sediment model of [63] has been implemented. This model is both well-regarded in the field and relatively direct in its parameterizations. For many applications this approach is sufficient and avoids the complexities of a much larger set of parameterizations.

The translation of the original Sanford implementation to SUNTANS is relatively straightforward. Vertical diffusion of sediment properties has not been included in the present implementation, and no attempt was made to apply smoothing or diffusion of bed properties in the horizontal.

Perhaps the most important link between the hydrodynamics and the sediment model is the bed stress, and special care is required to avoid numerical artifacts in the bed stress field. Bed stress in the sediment model is used at cell centers to calculate resuspension and deposition fluxes, but the hydrodynamic kernel evaluates the momentum balance, including bed stress, at edges, not cell centers. Further complicating the bed stress calculation is the z-level vertical grid utilized in SUNTANS.

Using z-levels for the vertical grid avoids pressure errors common in sigma-coordinate models, and typically allows wetting and drying of cells with tighter tolerances. Unfortunately, z-levels lead to two specific issues which can introduce artifacts into bed stress calcuations. The first issue is a discrepancy in the number of vertical levels between nearby cells or edges. In areas of sloping bathymetry, it is unavoidable that there will be adjacent pairs of water columns with different numbers of z-levels, and likewise there is no guarantee that the three edges of a single cell have the same number of vertical levels. Blindly interpolating velocities from the same z-level may lead to artifacts where one velocity is from the bottom of a water column and and another is from the middle of a water column. Choosing an edge velocity from each edge's respective bed z-level still has the problem that the height of that edge may be different between the cells, and reflect different heights above bed. To decrease the sensitivity to differing number of z-levels, the bed stress is calculated based on a per-edge velocity interpolated to a reference height. Obtaining the vertically interpolated velocity on each edge starts with averaging the edge-normal velocities from the bed up to a

reference height of $O(1m)$. The velocity, averaged over the lowest portion of the watercolumn, is then converted to a velocity at the reference height above the bed by assuming a log-law profile and using the prescribed bed roughness. These edge velocities, now interpolated to a consistent height above the bed, can be combined to a single, cell-centered velocity vector at the reference height, from which the bed stress is calculated. This approach does mean that the bed stress for sediment departs from that used for the hydrodynamics. However, tests in idealized domains show that the actual bed stress from hydrodynamics is highly variable even between adjacent cells. This is not problematic for hydrodynamics, but with resuspension fluxes being a nonlinear function of the excess bed stress, this variation can lead to adjacent cells experiencing wildly different sediment fluxes at the bed.

A more subtle issue with z-levels is that the vertical spacing near the bed is typically non-uniform. Even if the prescribed vertical spacing of z-levels is uniform, the bed cell is typically "shaved down" such that the bed is at the correct local elevation. When such shaved cells are allowed to be very thin, vertical finite differences are poorly behaved due to the nonuniform Δz. This is especially evident when evaluating the turbulence closure in bed cells. The near bed region typically has high shear, high gradients in TKE production and dissipation, and significant vertical fluxes of TKE. Poorly conditioned finite difference operators in this portion of the watercolumn can lead to oscillatory behavior in the turbulence closure. Even when temporal oscillations are not triggered, the depth-averaged velocity can differ significantly between adjacent water columns where the number of z-levels differs but the total depth is nearly the same. As a rough measure to counter the issues of thin cells, the vertical discretization has been updated to allow more aggressive lumping of vertical layers. Assuming that one wishes to keep the ratio of successive layer thicknesses as close as possble to unity, then a simple geometric argument leads to the conclusion that layers should be combined when the ratio of their thicknesses reaches the inverse of the golden ratio, i.e. when

$$\frac{\Delta z_b}{\Delta z} < \frac{\Delta z}{\Delta z_b + \Delta z} \tag{D.1}$$

where Δz_b is the original thickness of the bed layer and Δz the thickness of layers not truncated by the bathymetry. This adjustment does not entirely remove the artifacts from changes in the number of layers, but tests show that the artifacts are greatly reduced compared to a simulation without lumping. Recent work in [56] considers this problem in more depth, with an even more aggressive lumping scheme extending higher into the water column. Future work may include adapting more specialized boundary discretizations for the turbulence closure to reduce sensitivity to thin z-levels without affecting near-bed transport.

Despite the above measures intended to reduce the artifacts created by z-levels and variable bathymetry, the sediment resuspension fluxes typically still have a high degree of non-physical spatial variation, often showing up as a sawtooth or $2\Delta x$ pattern. Typical methods of assigning bathymetry to a grid either evaluate the input DEM at cell centers or edge centers. Taking the example of a regular grid aligned with a channel, the cells along a particular isobath can be viewed as a "strip" of triangles with alternating left-pointing and right-pointing triangles. The respective cell-centers also alternate, offset to either side by

a small amount. When the isobath lies on a slope, this lateral offset between adjacent cell centers leads to an alternating pattern of cell depths. Similarly, if bathymetry is evaluated at edges, half of the cells in the strip will have two edges straddling the isobath and one shallower. The other half of the cells in the strip will have two edges straddling the isobath and one deeper edge. Again, the resulting grid will have cell depths alternating shallower and deeper. Conceptually this leads to along-channel flow encountering a series of bedforms at the scale of the grid, creating alternating ridges of high bed stress and troughs of low bed stress.

Methods for reducing this artifact are commonly termed "terracing," and amount to smoothing the bathymetry along isobaths. Rather than explicitly extracting isobaths and smoothing grid bathymetry along those contours, a more general and simpler approach is to compute depths from the DEM at edge centers and then define the cell depths as the median of the edge depths. This yields a depth for each cell, and the final edge depths are then dictated by the shallower of the two adjacent cells for each edge. In the simple case of a regular grid aligned with a channel with an arbitrary lateral bathymetry profile, this approach yields a bathymetry field which is constant along isobaths. As the grid departs from regularity or being aligned with the channel the method will tend to create small patches of smooth bathymetry, at least reducing the number of sites where step changes in bathymetry introduce bed stress artifacts.

D.2 Grid Generation

The discretization used in SUNTANS and other Arakawa C grid models places several constraints on the geometry of the computational grid. These models calculate gradients of cell-centered quantities under the assumption that the line between adjacent cell centers intersects and is perpendicular to the common edge. A consequence of this requirement is that cell centers must be the circumcenters of the triangles. In order for the circumcenter to lie within the cell itself the cell must be an acute triangle. The primary motivation for development of TOM (Triangular grid Orthogonal Mesh generator), is to provide automated creation of triangular meshes with all cells obeying the acute angle constraints. The base algorithm is a variant of the advancing front method [29], with extensions for preserving orthogonality and embedding linear features. Furthermore, as noted in [81], truncation errors when interpolating on a triangular grid are minimized when the cells are regular and equilateral. Where the requirement of acute triangles is relatively strict, there is an additional soft requirement that cells be as close to equilateral as possible. Grid generation approaches such as that of Triangle [67] provide a mechanism for accepting or rejecting a triangle based on parameters like area and internal angle, but do not provide a way to heuristically push the triangulation towards an equilateral tiling.

Inputs and preprocessing

Inputs to the mesh generation process are (a) a polygon (possibly with holes) representing the shoreline and computational boundaries, (b) a set of points and associated length-scales giving the *requested* grid scales, and optionally (c) a collection of polylines in the interior of the domain for controlling cell alignment. The grid scale refers to the nominal edge length, defined as a continuous function across the entire domain. In the simplest usage of TOM this scale function is based on a natural neighbors interpolation from the set of input points. More complicated approaches to specifying the scale have been implemented and are discussed below. Before grid creation can begin the shoreline must be made consistent with the requested scale. For example, if a region contains a slough with a width of 50 m, but the requested resolution for the area is 100 m, one of three options must be taken: widen the representation of the slough to 100 m, remove the slough, or locally decrease the grid scale down to 50 m. The most common approach is to remove the features which are too small (an option for instead modifying the requested scale is described later). Starting with a constrained Delaunay triangulation of the input shoreline (Figure D.1a), Steiner points are added until the circumcenters lie near the medial axis of the shoreline. Specifically, for each constrained edge in the triangulation, the circumradius of the adjacent triangle internal to the domain is compared to the point-line distance from the triangle's center (the Voronoi point) to the constrained edge. If ratio between these distances is greater than a prescribed threshold (empirically chosen to be 1.2), the edge is subdivided with a Steiner point inserted at its midpoint (Figure D.1b). The process is repeated until no more edges require subdividing. This resampling ensures that the circumradius reflects the length scales between features and not the length scale between successive vertices on the shoreline. Any triangles with a circumradius smaller than half the requested grid scale are removed, and a new shoreline polygon is constructed from the union of the remaining triangles (Figure D.1c).

In applications where it is desirable to control the alignment of cells, additional polylines internal to the domain may be specified. These lines are treated identically to shorelines, with the caveat that both sides of each line are internal to the domain. In addition to controlling alignment, such polylines can also be used to efficiently resolve linear bathymetric features such as underwater ridges or "zero-width" dams.

Paving

In an advancing front method, the *front* is the boundary of the region in which no cells have been created, initially identical to the shoreline. Starting with the preprocessed shoreline polygon and any interior lines, the paving process iteratively applies the following steps, illustrated in Figure D.2:

1. Find the point along the front with the smallest internal angle, label it B, and keep a reference to the points A and C that are immediately clockwise and counterclockwise

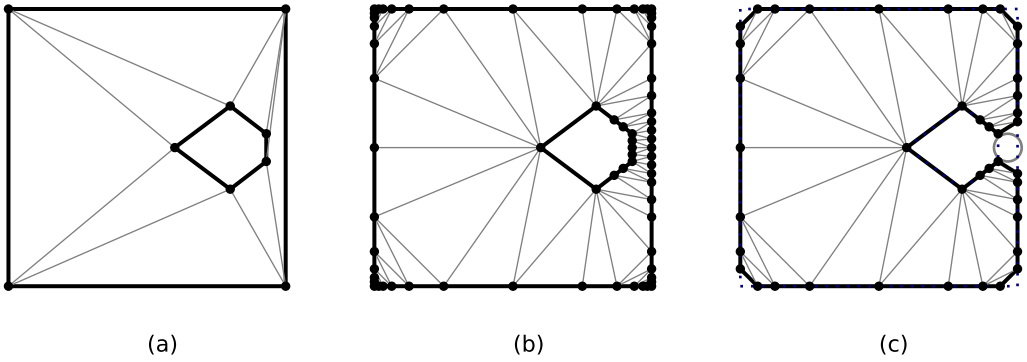

Figure D.1: Removal of features smaller than local grid scale. (a) Input shoreline and constrained Delaunay triangulation. (b) Constrained Delaunay triangulation after adding Steiner points. (c) Smoothed shoreline after removal of triangles with circumradius smaller than $l/2$ (illustrated by minimum circumcircle shown on the right).

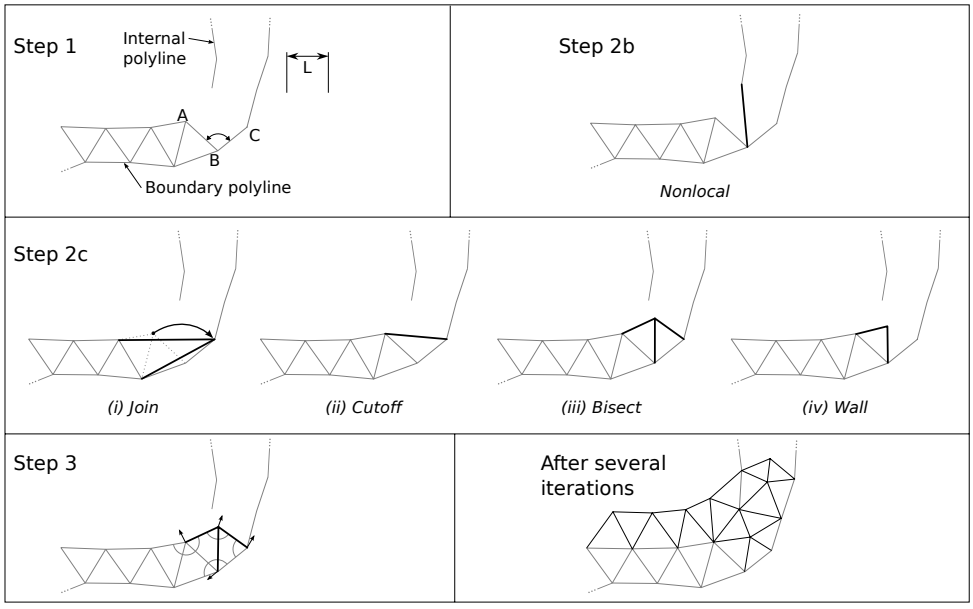

Figure D.2: One iteration of the paving algorithm. The node B is chosen based on the smallest internal angle. The strategy for how to add edges and cells at B is chosen based on the internal angle and whether the local scale is smaller than, equal to, or greater than the desired scale. After updating the topology, node locations are refined by minimizing a cost function involving cell angles and edge lengths.

respectively from B along the front. This will be the site of the cell or cells created during this iteration.

2. There are multiple choices or strategies for how new cells will be added at B. Topologically these strategies correspond to how the degree of node B, the number of edges adjacent to B, is modified. For example, adding a line AC to create the cell ABC preserves the degree of B, while adding a new point D to form two new cells increases the degree by one. In a perfectly regular triangular grid all nodes have a degree of six. Strategies that result in a greater degree will tend to create shorter edges, while the edge length can be increased by limiting the degree of B. In this way, the choice of strategy is the key mechanism by which a specific requested scale is achieved. In many situations, though, local geometric constraints rule out some strategies, so it becomes necessary to construct a prioritized list of strategies based on the local scales and geometry. The steps in constructing the prioritized list are:

   a) Calculate a *scale factor* s as the ratio of the *current scale* (the average of edge lengths, $\frac{1}{2}(|AB| + |BC|)$) to the requested grid scale L.

   b) Scan the vicinity around B for the possibility of non-local connections, i.e. when the advancing front encounters an island, an opposing shoreline, or an internal polyline. A constantly updated constrained Delaunay triangulation of all the grid features (shorelines, internal polylines and completed triangular cells) enables querying of nearby features in $O(1)$ time. If the non-local feature is within $1.25L$, the highest priority strategy will be creating an edge from B to this feature. In the event that the nearest point on the non-local feature is not a preexisting vertex, a new vertex will be inserted in the feature at the point closest to B.

   c) Based on both the interior angle $\theta = \angle ABC$ and the scale factor s, the remaining local strategies are prioritized. For example, if the angle were 60°, a *cutoff* will add the edge AC, making a roughly equilateral triangle ABC. The scale factor changes the thresholds between choosing one strategy over another. If the angle were instead 80° but the current scale were large compared to the requested scale, the *bisect* strategy would attempt to bisect the angle $\angle ABC$ by adding a new point D and creating two triangles ABD and BCD (and increasing the degree of B in an attempt to reduce the edge lengths). If the current scale were small compared to the requested scale, the cutoff strategy would be tried first, creating a single triangle ABC with longer edge lengths. Aside from non-local connections, the strategies (and the range of θ for which each is the highest priority strategy) are:

      i. join - merge AB and CB into a single edge $[\theta < 30°]$
      ii. cutoff - add the edge AC $[30° \leq \theta < 85° s^{-3/2}]$
      iii. bisect - add a new node D, and edges AD, BD, and CD $[85° s^{-3/2} \leq \theta < 160° - 50°\,(s-1)]$

iv. wall - add a new node D, but only create the edges AD and BD $[\theta \geq 160^\circ - 50^\circ\,(s-1)]$

3. For each attempted strategy, the mesh topology is updated and a local optimization step adjusts vertex locations to restore the acute angle constraints. The cost function minimized during this optimization is

$$C(p) = 10\left[\frac{\max_{i\in I}\left(|\phi_i - 60^\circ|\right)}{\phi_{max} - 60^\circ}\right]^5 + \exp\left(\frac{\max_{i\in I}\left(\phi_i\right) - \phi_{max}}{3^\circ} - 1\right)\cdots \quad \text{(D.2)}$$
$$\cdots + 2\left(\left[\frac{\max_{e\in E}\left(l_e\right)}{L}\right]^2 + \left[\frac{L}{\min_{e\in E}\left(l_e\right)}\right]^2\right)$$

where I denotes the set of internal angles ϕ_i of triangles that include the vertex p, ϕ_{max} is the target maximum internal angle, chosen to be 85°, E the set of edges with p as an endpoint, and l_e the respective length of each of those edges. The first term favors equilateral triangles, the second term penalizes angles approaching ϕ_{max}, and the third term favors edge lengths close to L. If the optimization is not able to satisfy the angle constraints or the optimization causes a self-intersection, the mesh is reverted to its original state and the next strategy is attempted.

Post-processing

While the paving process creates only cells which respect the interior angle criterion, poor quality cells that only barely satisfy the angle requirements may still exist. Two post-processing methods are applied to improve overall grid quality: relaxation and repaving. The relaxation steps apply the same optimization method used during paving to regions containing low quality cells. By applying the optimization over a larger group of vertices it is sometimes possible to nudge the vertices toward a better layout without modifying the grid topology. Where this approach is unable to improve the cells, the offending cell and its neighbors are removed and the paving method is re-applied locally. While repaving usually creates cells with better geometry, it may upset the alignment of cells that are not constrained by an adjacent shoreline or interior line. For most cases, only the most borderline low quality cells are selected for repaving in an effort to preserve as much of the alignment and regularity as possible.

Implementation and complexity

The current implementation of TOM is primarily in Python, with broad utilization of the numerical libraries *numpy*, *CGAL* and *scipy*. While the asymptotic complexity is $O\left(n \log n\right)$, where n is the total number of cells, in practice the optimization step is the most expensive and for grids at least up to 5×10^5 cells the run time is very close to linear in n. On a typical 64-bit desktop computer the speed ranges from approximately 40 cells/second in tightly constrained areas to 110 cells/second in broad, unconstrained areas.

Additional methods for defining scale

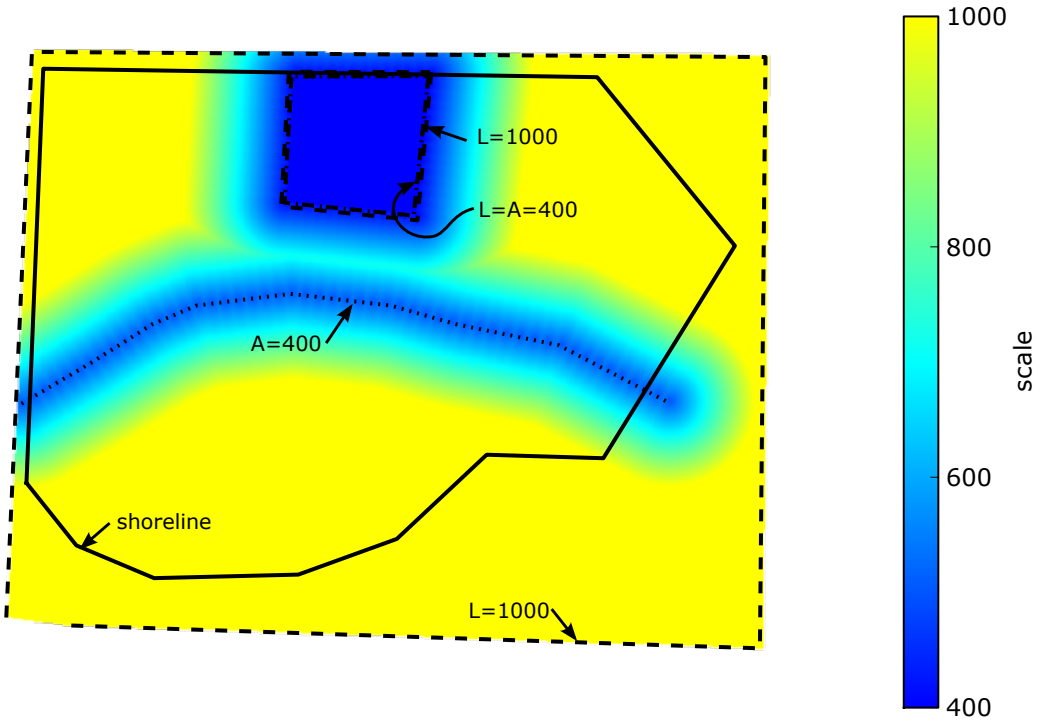

Figure D.3: Inputs to the grid generation process. Solid black denotes the intended shoreline. Dashed lines define the constrained Delaunay triangulation and interpolated scale values. Dotted lines denote auto-telescoping scale constraints (with the telescoping rate set to 8%). Background color denotes the requested scale based resulting from these scale inputs.

The simplest approach to defining the requested scale is a collection of points with a scale value assigned to each. A Delaunay triangulation is constructed over this set of points, and a natural neighbors interpolation method is used to interpolate the point values to a continuous field covering the domain.

In many cases, though, the exact structure of the Delaunay triangulation does not conform to the expectations of the user. By associating scale values with polylines rather than individual points, the user can control the structure of the triangulation (now technically a constrained Delaunay triangulation). This is helpful in making sure that a small scale value in one location does not "bleed over" into an adjacent area.

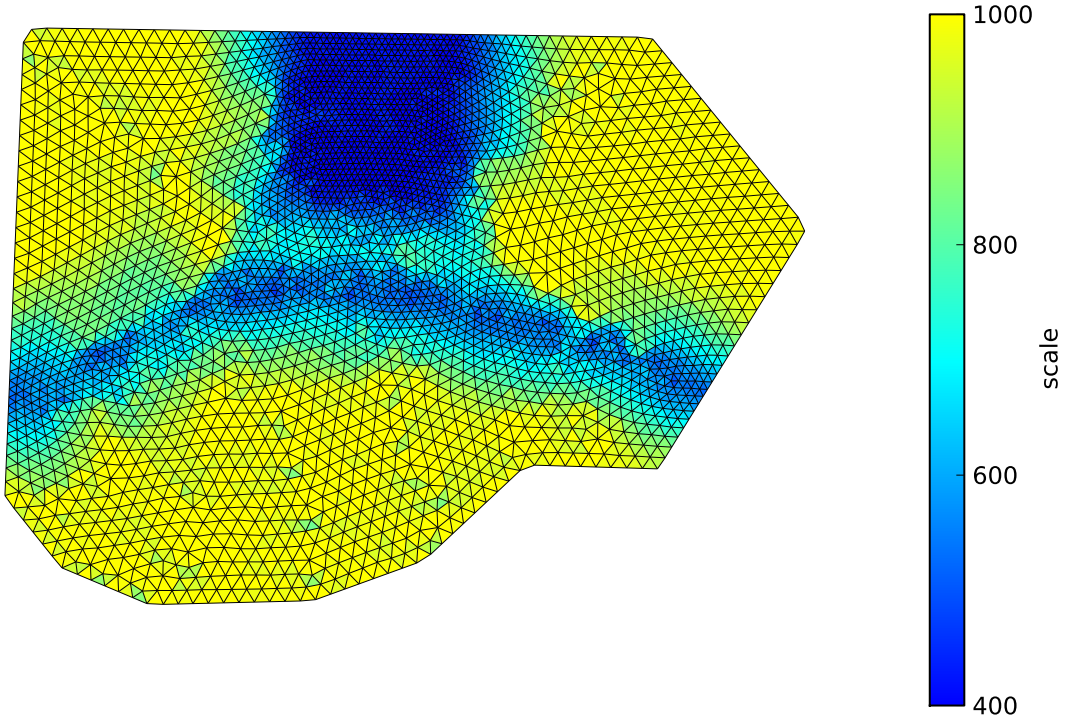

Figure D.4: The triangulation resulting from the inputs in figure D.3. Cells are colored by area-equivalent equilateral side-length, $l = \left(\frac{4}{\sqrt{3}}A\right)^{\frac{1}{2}}$.

There is a common rule of thumb in unstructured grids that the change in linear scale between adjacent cells should not exceed 10%. The modeler is often presented with cases where a particular discrete feature needs to be resolved at a particular scale (for example, a slough or breach), but this increased resolution is unnecessary in neighboring areas. In this case, the nearby resolution should be dictated only by the telescoping rate.

TOM supports specifying such "auto-telescoping" scales through the use of a novel datastructure known as the Apollonius graph. In the non-telescoping, purely interpolating method from above, it is simple to compute which input points are relevant for a particular scale query: they are the points forming the Delaunay triangle containing the query point. In the telescoping method, though, this test becomes significantly more complicated. A nearby point with a large scale may be superseded by a far away point with a very small scale. A naïve implementation of the scale query would have to check all input points, considering their respective scales and distances from the query point, in order to find the correct tele-

scoped scale. The Apollonius graph, an algorithm borrowed from computational geometry, allows efficient queries of this sort. This datastructure is a generalization of the Delaunay triangulation in which each point is allowed a finite radius. By assigning points with small scales larger radii, and points with large scales small radii, graph queries can efficiently account for both scale and distance. Global graph search is still slow, but the spatial coherence between successive queries can be exploited for significant performance gains.

Most uses of the telescoping scale functionality rely on a user-provided set of points and respective scales. As described above, the resulting scale distribution is then used to simplify the shorelines to the point that they are consistent with the requested scale. A novel use of the telescoping scale functionality, though, is to reverse this relationship. The shoreline can be triangulated as above, but rather than using the local feature scale to determine which portions of the domain to remove, this scale can be used instead to dictate the requested scale. Telescoping scale constraints then enforce that the scale does not change too quickly. This was in fact the original approach for the slough and levy-resolving grids of chapter 4, but difficulties in robustly defining levee features from LiDaR made this ultimately untenable.

Mixing interpolated and telescoped scales is as simple as querying both and returning the smaller of the two values. The end result is that complex distributions of grid resolution can be specified with minimal inputs from the modeler. A sample input demonstrating mixed scale specification is shown in figure D.3. This setup might be relevant for a basin in which a narrow channel runs east to west through a broad shoal, with a study site in the center-northern portion of the domain. Telescoping 400 m resolution around the channel ensures that it is captured in the grid, without costing too many unnecessary grid cells in the shoals. Resolution throughout the study area (perhaps a wastewater outfall) is pinned at 400 m, and allowed to safely telescope up to the background scale of 1000 m outside the study area. The resulting grid is shown in figure D.4. Here the resulting edge length, for plotting purposes, is calculated as the edge length of an equilateral triangle with equivalent area.

D.3 Boundary Forcing

The stock SUNTANS code base has historically defined boundary conditions from a customized C source file. This provided the ultimate in flexibility as all aspects of the model are accessible and the complexity of a boundary forcing scheme was limited only by the savviness of the programmer. However, there are several potential issues with an open-ended and low-level description of boundary conditions. First, the model setup is tied to model internals. If a future version of SUNTANS changes the way that wind stress is evaluated or the name of a boundary velocity variable, all models developed to the older standard become unusable. Second, it discourages code reuse, and instead puts each model developer in the position of re-inventing the wheel when it comes to reading in time series, evaluating harmonic constituents, or coding up spatial interpolation methods. There is also the issue that data and code become intertwined, where boundary condition code must know something

about the input grid in order to apply the correct boundary condition to the correct portion of the boundary.

In an effort to streamline model development, reduce the chance of boundary condition errors, and improve forward-compatibility, a framework for defining and implementing a range of boundary conditions common to a hydrodynamic model has been developed. Instantiating boundary conditions requires detailed knowledge of the computational grid and simulation parameters. Consequently, the codebase has evolved into a convenient framework for handling a range of model deployment tasks from automated grid generation to post-processing of model data for validation.

The framework is primarily in the form of an object-oriented class library, with essential top-level classes:

Domain the procedural glue, typically common to multiple simulations spanning multiple time periods, possibly run on multiple architectures, but all using the same grid (or the same recipe for creating a grid), and similar boundary conditions.

TriGrid a geometric description of the triangular 2-D grid which supports many geometric and topological operations and queries.

SunReader a specific model run, whether being prepared for execution or already executed. If the run has been completed, this class also provides access to the model output.

Field A generic class for describing a 2-D field, i.e. a mapping $\mathbb{R}^2 \rightarrow \mathbb{R}$. Useful for describing variation in grid scale, bathymetry data, wind fields, spatiall varying roughness values, etc.

GlobalForcing a collection of boundary conditions for a particular grid. Subsets of the grid are chosen based on indices or geographic filters, and multiple variables may be forced within each subset of the grid.

DataSource For each variable which is forced, the relevant scalar or vector data is defined in a DataSource, which may represent a constant value, a time series, or a harmonic decomposition. Subclasses determine how the data is populated, including on-demand downloading from internet sources.

The intended usage pattern is to derive a project-specific subclass of Domain, with grid specifics, boundary conditions and custom "recipes" (methods prefixed with `cmd_` and accessible from the command line) defined in the subclass. This allows building up complex recipes using smaller recipes as steps. One example is the `batch` recipe, which roughly breaks down into the following pseudocode:

Data: folder,start,stop
folder: path prefix for the series of runs
[start,stop]: time interval for the first run in the series.

begin initialize run in folder
  if *no grid present* **then**
    generate grid
  end
  if *grid has changed or is missing bathymetry data* **then**
    interpolate bathymetry onto grid
  end
  if *grid has changed* **then**
    update marks in grid to reflect boundary conditions
    partition grid into subdomains
  end
end
repeat
  download data for the period [start,stop]
  write forcing data to per-processor data files
  queue the simulation to the local cluster
  // C code within SUNTANS parses and applies specified boundary conditions
  wait for simulation to complete
  start, stop ← stop,stop + (stop - start)
until *interrupted by user*

Algorithm 1: process for batch recipe

The full power of an object-oriented framework becomes evident when suites of simulations are desired. A minimal subclass can substitute a different value for a parameter or override an arbitrary step of the framework. Multiple runs differing only in the simulation period can be invoked from the same subclass, just specifying different time periods on the command line.

D.4 Visualization

As computational power and numerical models progress, both the volume of model data and its complexity increase. Low resolution, two-dimensional model data can be plotted directly and quickly allowing efficient inspection, analysis and debugging by the modeler. High resolution, three-dimensional models, with perhaps ten or more different scalar and vector quantities to examine, spread across dozens of subdomains, present a challenge to the modeler who wants a glimpse of the model results.

The need to efficiently visualize and explore large volumes of model data is common among modern numerical modelers, and high performance tools are increasingly available to the modeler. One particular package, VisIt [13], is particularly well-suited to CFD and

cluster computing. Features of VisIt include support for scalar, vector and tensor fields, two- and three-dimensional domains, and a client-server architecture which takes advantage of local graphics acceleration while minimizing network traffic to remote compute resources.

To facilitate debugging, exploration, and presentation of simulation results, a plugin for VisIt has been developed which allows reading of the native SUNTANS output binaries. The plugin currently parses grid geometry, bathymetry, freesurface elevation, ν_T, salinity, temperature, and cell-center velocities. The full-state output files (typically `store.dat.`N) can also be parsed to obtain instantaneous edge-velocity and freesurface fields, helpful for post-mortem diagnosis of crashes.

An additional feature, when the Berkeley SUNTANS codebase is used, is the automatic recognition of "raw" files. `logger.c` provides a set of utility functions for outputting static or time-varying fields to a basic binary format. A simple naming scheme describes the data layout, such that the VisIt plugin can parse the filename and associate the binary data with the correct elements of the grid. This provides a quick and flexible mechanism for debugging, for example by outputting intermediate values in the momentum calculation, or to add nonstandard output of values like bed stress. The sediment module also uses these output methods for bed layer mass and suspended sediment concentration output.

In the long term, this plugin is expected to become obsolete as SUNTANS output moves to CF-compliant NetCDF and CF-compliant VisIt plugins become available. In the meantime, though, this is likely the most complete visualization approach for SUNTANS output.