

http://doi.org/10.25923/2sfh-7827

Collected Notes on the Basics of Pressure-Equipped Inverted Echo Sounder Analysis

C.S. Meinen

NOAA-Atlantic Oceanographic and Meteorological Laboratory Miami, Florida

noaa

February 2019

U.S. DEPARTMENT OF

NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION OFFICE OF OCEANIC AND ATMOSPHERIC RESEARCH

Suggested Citation

Meinen, C.S., 2019: Collected notes on the basics of pressure-equipped inverted echo sounder (PIES) analysis. NOAA Technical Report, OAR-AOML-51 (doi:10.25923/2sfh-7827), 18 pp.

Acknowledgments

Comments from Daniel Valla, Marion Kersalé, Ryan Smith, Renellys Perez, and Rigoberto Garcia helped to improve earlier drafts of this report, and their assistance is greatly appreciated. Gail Derr graciously helped with the formatting and preparation of this report. The report was supported by the NOAA-Climate Program Office/Ocean Observing and Monitoring Division-funded Southwest Atlantic Meridional Overturning Circulation project and by NOAA's Atlantic Oceanographic and Meteorological Laboratory.

Disclaimer

NOAA does not approve, recommend, or endorse any proprietary product or material mentioned in this document. No reference shall be made to NOAA or to this document in any advertising or sales promotion which would indicate or imply that NOAA approves, recommends, or endorses any proprietary product or proprietary material herein or which has as its purpose any intent to cause directly or indirectly the advertised product to be used or purchased because of this document. The findings and conclusions in this report are those of the authors and do not necessarily represent the views of the funding agency.

NOAA Technical Report, OAR-AOML-51

http://doi.org/10.25923/2sfh-7827

Collected Notes on the Basics of Pressure-Equipped Inverted Echo Sounder (PIES) Analysis

Christopher S. Meinen

NOAA-Atlantic Oceanographic and Meteorological Laboratory Miami, Florida

February 2019

UNITED STATES DEPARTMENT OF COMMERCE Mr. Wilbur L. Ross, Jr., Secretary

NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION RDML Timothy Gallaudet, Acting Under Secretary for Oceans and Atmosphere and NOAA Administrator

OFFICE OF OCEANIC AND ATMOSPHERIC RESEARCH Mr. Craig N. McLean, Assistant Administrator



This page intentionally left blank.

Table of Contents

Fig	uresiii		
Acr	ronymsv		
Ab	stract vii		
1.	Brief History of the IES/PIES1		
2.	Travel Time Measurement.12.1Initial processing to obtain hourly data values.12.2Seasonal correction of travel times32.3Calibrating travel time to a fixed pressure surface.52.4Applying the GEM technique to obtain profile data.72.5Seasonal GEM correction.9		
3.	Bottom Pressure Measurement103.1 Initial processing to remove tides103.2 Initial processing to remove drift11		
4.	Merging Travel Time and Bottom Pressure Records from Subsequent Deployments12		
5.	Estimating Relative Velocity Profiles135.1Use of PIES-GEM profiles135.2Use of PIES bottom pressures14		
6.	Combining Relative and Reference Velocity for Absolute Velocities15		
7	Final Thoughts and Recommendations16		
8.	Papers for Further Reading17		
9.	References		

This page intentionally left blank.

Figures

1.	Photo of a PIES mooring being prepared for deployment in December 20121
2.	Example of a travel time record from a PIES deployed in the South Atlantic at 34.5°S as part of the Southwest Atlantic Meridional Overturning Circulation project
3.	Plots showing the seasonal variations in temperature as a function of year-day from a collection of 565 CTD and Argo profiles
4.	Plot showing the seasonal correction in travel time determined between the surface and 200 dbar using a collection of 565 CTD and Argo profiles
5.	Plots illustrating how PIES travel time data were interpreted for much of the 1960s through the 1990s
6.	Plots illustrating the steps for building the GEM look-up table of temperature for the Southwest Atlantic Meridional Overturning Circulation study region at 34.5°S7
7.	Example of a GEM look-up table of temperature from the Southwest Atlantic Meridional Overturning Circulation project region at 34.5°S in the Atlantic
8.	Example of a bottom pressure record from a PIES11
9.	Example of drift removal from a PIES record from the Western Boundary Time Series Site E with the x-axis time in days

This page intentionally left blank.

Acronyms

CTD	Conductivity-temperature-depth
CPIES	Current-and pressure-equipped inverted echo sounder
GEM	Gravest empirical mode
GMT	Greenwich mean time
IES	Inverted echo sounder
MOC	Meridional overturning circulation
MPIES	MicroCAT-and pressure-equipped inverted echo sounder
OFES	Ocean General Circulation Model For the Earth Simulator
PIES	Pressure-equipped inverted echo sounder
SAM	Southwest Atlantic MOC
URI	University of Rhode Island
WBTS	Western Boundary Time Series

This page intentionally left blank.

Abstract

The pressure-equipped inverted echo sounder (PIES) is a powerful oceanographic tool that can provide, when combined with historical hydrographic data, full-water-column estimates of temperature, salinity, density, and dynamic height anomaly. Arrays of PIES can provide full-water-column estimates of geostrophic velocity as well. As with any measurement system, the PIES has limitations, but it also has significant strengths, including: relatively low equipment costs; simple deployment/recovery requirements; long deployment lengths (up to 5 years); and the ability to acoustically transmit daily-averaged data to a nearby research ship without recovering the instrument. These strengths make the PIES a good tool for the study of ocean currents in many deep ocean regions.

This technical report is essentially a somewhat informal set of notes I've put together to aid newcomers to PIES work, and/or those considering future work with PIES. The notes are based on what I've learned over the years to teach/tutor folks in PIES analysis. This document is in no way complete or definitive—it is intended to be a simple overview of the type of work I personally do/have done with PIES. I highly recommend that folks interested in working with PIES obtain and read both the instrument manual and the processing manual (Kennelly *et al.*, 2007) that are provided by the manufacturers (Randy Watts' group at the University of Rhode Island). The peer-reviewed literature cited herein also presents a more thorough and complete description of PIES analysis (e.g., Meinen *et al.*, 2004, 2006, 2012; Meinen and Garzoli, 2014; Meinen and Luther, 2016).

This page intentionally left blank.

1. Brief History of the IES/PIES

The inverted echo sounder (IES) was invented by Tom Rossby, who was looking for a way to study thermocline changes and heat content changes in the Gulf Stream in the late 1960s (Rossby, 1969). After a subsequent move to the University of Rhode Island (URI), Tom continued to work on the idea together with Randy Watts, and they produced what was perhaps the first science-ready version of the IES in the mid-1970s (e.g., Watts and Rossby, 1977). Randy took over the development of the IES in these years and essentially became the "father of the IES." Randy continued to develop the IES over the 1980s and 1990s, along the way creating a version that incorporated an additional measurement device-a bottom pressure sensor-to create the pressure-equipped inverted echo sounder, i.e., the PIES (Figure 1). There were other versions that were created over the years, and this history is far from all-inclusive.

I'll jump forward to the modern era now. The modern PIES dates from the late 1990s and early 2000s when Randy and his team at URI designed a new version of the PIES that replaced the older technology inside the PIES with modern solid state memory and more. There have been a few subsequent versions, such as the current-andpressure-equipped inverted echo sounder (CPIES) that adds a single-depth acoustic current meter 50 m above the bottom, as well as the MicroCAT-and-pressure-equipped inverted echo sounder (MPIES) that adds a Sea-Bird Instruments MicroCAT computed tomography recorder attached on the outside of the instrument to measure salinity, a more accurate temperature, and optionally, a secondary pressure record.

In my opinion, the CPIES can be a great upgrade, but only in very specific circumstances, i.e., when the instruments will be deployed in a tightly-spaced array located within a horizontal correlation length scale to one another. This is so that one can reasonably argue that the average of the velocity measurements from the current meters on two neighboring CPIES is representative of the true average velocity at that depth between those locations. If the instruments are not going to be deployed close together, I do not think the roughly 50 percent increase in cost for a CPIES over a PIES is worthwhile.



Figure 1: Photo of a PIES mooring being prepared for deployment in December 2012.

I personally haven't seen much benefit from the MPIES upgrade yet—it's a new option that may show more utility in the future if adopted by the community. The MPIES option increases the cost by roughly \$1000 (for the modifications implemented at URI), plus the cost of the MicroCAT. Note that while it is still possible to purchase a simple IES from Randy's group at URI (i.e., without the pressure sensor or any other sensors), I cannot think of a reasonable scenario why one would do so. In these notes I focus primarily on the PIES.

2. Travel Time Measurement

2.1 Initial Processing to Obtain Hourly Data Values

The PIES travel time measurement is made as follows. The PIES transmits a sound pulse and simultaneously starts a timer. The sound pulse travels vertically up from the instrument anchored on the sea bottom and, when the sound pulse reaches the sea surface, it reflects off of the strong density interface between the water and the air. The pulse then travels downward, and when the PIES detects the pulse return, it stops the timer and records the round-trip travel time. Because the speed of sound in seawater is dependent on the temperature and salinity of the water, this round-trip travel time tells us something about the full-depth temperature and salinity structure. Later in these notes, I'll explain a bit more about the scientific interpretation of the travel time measurement. The PIES, as it is normally configured, makes a travel time measurement every hour. It does this by sending out a series of 24 acoustic pings at 12 kHz (note: in the earlier IES/PIES built through much of the 1960s-1990s, 10 kHz was used). The choice of 24 pings per hour is hard-coded into the instrument; the user does not have a choice to have more or less pings in an hour (without getting a special model of the instrument built by the manufacturer). The PIES sends out each ping and measures the amount of time required for that ping to reach the sea surface, reflect, and return and be heard by the PIES.

As each acoustic pulse radiates upward from the PIES, the acoustic footprint of the sound pulse grows because the sound spreads upward in a cone shape, i.e., the deeper the instrument, the larger the footprint. Because of this, I would not recommend deploying a PIES at any depth shallower than about 500 m as the acoustic footprint would become, for example, small enough that it could be blocked by a school of fish. For a PIES at greater depths, the school of fish would have to be much larger in horizontal size to cause a problem. This limitation only applies to the travel time measurement—a PIES would still collect good bottom pressure data if deployed at 200 m depth, just poor quality travel time data.

There are a number of different settings for the PIES that allow the user to distribute the 24 acoustic travel time pings through the hour. I have usually recommended that our instruments be set to do all 24 pings at the start of the hour, but Randy and his folks usually set their instruments to send out groups of four pings every 10 minutes throughout the hour. Regardless of this setting, the pulses in each group are sent out on an alternating schedule every 16 and 18 seconds. For example, pulse one goes out, then 16 seconds later pulse two goes out, then 18 seconds later pulse three goes out, then 16 seconds later pulse four goes out, etc. This alternating 16 and 18 second gap helps avoid interpretation issues from echoes. I doubt it makes much difference whether one selects the all-in-one group or the six-groups-of-four pulses option. I made my choice based on the idea that the instrument wakes up less frequently with the once-an-hour setup, and thus we may be saving battery life. I also anticipate that the ocean can change over the course of an hour, and I'd like to have more samples in a single "burst" for better statistics. Again, I'm dubious about how important this setting is in practice.

Some notes regarding the travel time data file contained within the PIES may be helpful. The time (clock) units used in the PIES are hours elapsed since zero hours on January 1, 1970. The travel time units are in units of seconds but the decimal point is missing, so a travel time value of 671123 is 6.71123 seconds. The travel time file will have a header line that begins with "T," and after that it will have a regular array with 25 columns and one row for every hour the PIES was operating. The first column is the time (clock) for that hourly measurement, and columns two through 25 are the 24 measurements of travel time for that hour. For more details, check the instrument manual from URI.

When a PIES sends out a sound pulse, it does not listen for a reply immediately. Instead, it waits a few moments to start listening to avoid recording returns that have bounced off close-by objects such as seamounts. This lock-out window is calculated by the instrument as a function of the planned deployment depth of the instrument. This is why it is crucial to make sure the depth you use when you program the instrument is fairly accurate (must be accurate to better than 10 percent to avoid potentially losing data).

Once the lock-out window passes, the PIES begins listening, and it will listen for up to 9.99999 seconds. If the PIES does not hear a return ping before 9.99999 seconds after it sends out a ping, it will record 9.99999 seconds in the data file for that ping and then it will prepare to send out the next ping in the group. An example travel time record from a multi-year deployment is presented as **Figure 2.**

The travel time data are processed by first windowing the hourly measurements to remove the obvious reflections off schools of fish (i.e., early returns) and indirect path returns (where the direct path to and from the surface was blocked, but a reflection from the surface some distance to the side was received, i.e., a late return). The IES processing manual (Kennelly *et al.*, 2007) provides a good description of this and illustrates how the URI-provided code completes this step. Once the obvious outliers are



Figure 2: Example of a travel time record from a PIES deployed in the South Atlantic at 34.5°S as part of the Southwest Atlantic Meridional Overturning Circulation project. Black dots in both panels show the raw individual ping values. The lower panel is a zoom near the actual surface reflection signals, while the red line shows the processed hourly values.

removed, the hourly values are sorted into quartiles, and the value for each hour is set to be at the end of the first quartile (red line in **Figure 2**). For more details on this process of removing outliers and selecting the hourly value, see the URI processing manual (Kennelly *et al.*, 2007).

Perhaps a comment on the tides and how they impact the travel time measurement is appropriate here. Open ocean tides are typically of the order of 1 m. Doubling this (because the travel time measurement is round-trip), one would expect a tide impact on the travel time of almost 2 m divided by 1500 m per second (a rough estimate of the speed of sound in seawater), which equals 0.0013 seconds, or about 1 millisecond. Because this value is not much bigger than the 0.5 millisecond resolution of the PIES, and because a low-pass filtering of the data with either a 40-hour or 72-hour cutoff period is standard in the PIES processing, the tidal signal is essentially negligible for the travel time measurement. A comment on calculating the speed of sound in seawater is also important here. The equation of seawater I use and recommend is that of Del Grosso (1974). Our tests (Meinen and Watts, 1997) and the work of others have indicated the Del Grosso (1974) equation is more accurate than that of the United Nations Educational, Scientific, and Cultural Organization (UNESCO)-endorsed equation of Chen and Millero (1977), even after the application of the later Millero and Li (1994) update/correction. Therefore, I recommend using the Del Grosso (1974) equation for simulating a travel time measurement using hydrographic temperature and salinity profile data.

Once the PIES travel time data have been processed to produce a single value for each hour, the data are typically low-pass filtered (see note above regarding tides). In my work, we have normally used a second-order Butterworth filter with a 72-hour cutoff period, passed both forward and back to avoid phase shifting. After filtering the data, we generally subsample the travel time data to one value per day at noon Greenwich Mean Time (GMT).

2.2 Seasonal Correction of Travel Times

Seasonal variations in temperature and salinity in the upper few hundred meters of the water column cause small changes in the measured travel time, typically around 1 millisecond. These signals are considered noise and, while they are quite small and close to the ~0.5 millisecond resolution of the PIES travel time measurement itself, the seasonal variations are usually removed since they can be estimated. As an aside, this step does not impact the ability of the PIES to measure seasonal velocity variations. The seasonal variations of temperature and salinity discussed here have very large spatial scales-whole basins and geostrophic velocities are based on density gradients-so this correction does not impact the calculated velocities, which I will discuss in greater detail later in these notes. Hydrography-derived seasonal corrections are typically applied to the PIES travel time records after they have been processed and subsampled to a single value per day. The first step in making this correction is determining how deep into the ocean the seasonal signals extend for a given location. In my experience, this can vary from as little as 100 dbar to as deep as 300 dbar, with the deeper

penetration depths usually occurring at higher latitudes, which is perhaps not surprising as seasonal variations in temperature and salinity are higher at the higher latitudes as compared to the tropical regions.

To evaluate the depth to which the seasonal signals penetrate at a given location, simply plot temperature as a function of year-day at a variety of depths using hydrographic observations from the region where you're working. In **Figure 3**, an example is shown for the southwest Atlantic Ocean (from our Southwest Atlantic Meridional Overturning Circulation [MOC] project region, or SAM) near 34.5°S with the temperature plotted at the surface, 100 dbar, 200 dbar, and 300 dbar as a function of year-day. The hydrography used here is a collection of 565 conductivity-temperature-depth (CTD) and Argo profiles from the region.

In the example presented in **Figure 3**, an evaluation of the data suggests that there is no significant seasonality below

200 dbar, so this is defined as the base of the seasonal layer for this purpose at this location in the southwest Atlantic Ocean near 34.5°S. Once it has been determined how deep the seasonal variations extend into the water column, a correction can be developed by calculating the simulated travel time between the surface and the base of the seasonal layer (200 dbar in the example) using the hydrographic data (CTD and/or Argo profiles). The overall bulk average of all of the simulated travel times for the set of hydrographic observations should be removed first, as the goal here is to understand the seasonal anomalies, not the time-mean. To avoid edge effects in the curve fitting, the year should be tripled, with year minus one and year plus one being identical to the center year. A polynomial can then be fit to the simulated travel time values, and only the center year is kept. Figure 4 presents an example. Once this seasonal signal is removed, the PIES travel time data are ready for the next step in their processing.



Figure 3: Plots showing the seasonal variations in temperature as a function of year-day from a collection of 565 CTD and Argo profiles. The relationship is shown at four pressure levels.



Figure 4: Plots showing the seasonal correction in travel time determined between the sea surface and 200 dbar using a collection of 565 CTD and Argo profiles. The x symbols show the original hydrographic values, and the red line is the polynomial fit to create the correction.

2.3 Calibrating Travel Time to a Fixed Pressure Surface

The travel time measurement of the PIES is not particularly valuable on its own—it only becomes powerful when combined with hydrographic information to tell us something about the water column. Early studies with IES/PIES would seek to estimate things like the depth of the main thermocline, or the dynamic height anomaly integrated through the upper layer, with their travel time data (e.g., **Figure 5**).

The relationships needed for these interpretations were derived from hydrographic data from the region. In the examples shown in **Figure 5**, the relationships were determined using CTD data. Note that the x-axis in these relationships is the travel time simulated at 1000 dbar, where the travel time is calculated using CTD data and the sound speed equation. However, the PIES is never actually deployed at exactly 1000 dbar. The next essential step in processing PIES travel time data is to calibrate the PIES measured travel times, which are observed at some semiunknown pressure "p," into the corresponding travel time values at a known pressure level such as 1000 dbar.

Many of the older PIES studies did not use a travel time at 1000 dbar as their x-axis in these hydrography-derived relationships; in some cases, they did not use travel time at all, but instead used things like the depth of the 12°C isotherm. Because there are tight relationships like those shown in **Figure 5** between many of these quantities, what I'm discussing here doesn't change much for those alternative options for the x axis. The Meinen and Watts (1998) paper listed in the bibliography goes into this issue in some detail. In brief, we calibrate the travel time



Figure 5: Plots illustrating how PIES travel time data were interpreted for much of the 1960s through the 1990s. Left: Estimating the vertical location of the 12°C isotherm, which was an approximation for the main thermocline depth in many Gulf Stream studies. Right: Estimating the dynamic height anomaly, gradients of which gave geostrophic velocities and transports. Examples here are based on a dataset of 333 CTD profiles in the Mid-Atlantic Bight near the region of the Synoptic Ocean Prediction study.

measured at the depth of the PIES into the corresponding travel time at a fixed level (e.g., 1000 dbar) through what we refer to as calibration CTDs. As noted earlier, interpretation of the PIES travel time measurement is performed via comparison with hydrographic data (such as CTD profiles); however, these hydrographic data are not necessarily collected at the PIES site or during the time period when the PIES is in the water. The calibration CTD data are different, i.e., they are explicitly collected at the site of the PIES during the period when the PIES is on the bottom making its measurements.

Because you want the PIES to produce good measurements during the CTD cast, this means you always want to do a CTD before you recover a PIES, and you always want to deploy a CTD after you deploy a PIES. Of course, this is in a perfect world; in practice, keep in mind you'll be doing a 72-hour low pass filter of the travel time data so a less than perfect scenario can still work even if you deploy your CTD a bit too early or late. Ideally, you'd have at least two calibration CTD casts for each PIES deployment one just after the PIES is deployed and has settled on the bottom and one just before the release command is sent and the PIES is recovered.

More CTDs at the site during the ~4 years of a PIES deployment are, of course, even better. Essentially, the simulated travel time calculated with the calibration CTD data is compared to the concurrent measurement of travel time made by the PIES. The difference between the two travel time values is one realization of the calibration offset; when multiple CTD casts have been collected at the PIES site over the course of a single PIES deployment, the various calibration offsets can be averaged to obtain a more accurate offset value. This offset value is then subtracted from the PIES-measured travel time record to calibrate it into the corresponding travel time record at the fixed pressure surface (for more details see Meinen and Watts, 1998). Note that the Meinen and Watts (1998) paper presents a method for calibrating the PIES without the use of calibration CTDs based on the pressure measurement. In practice, we rarely apply this method, but it is available as a backup method if no calibration CTD casts are collected for a particular PIES deployment.

As a side note, it is important that the calibration CTDs (and all CTD and Argo profiles used to build the analysis look-up tables discussed in the next section) have data values near the surface. If, due to a failure of the CTD system, for example, the shallowest data values are at 100 dbar, that cast is not useful. The cast should have data within the mixed layer for it to be used. If the first value is at a depth of 10 m, that is generally fine, and you should simply extrapolate the cast up to the surface assuming a constant mixed layer. Obviously, if there are no data above 100 m, such an extrapolation would not be a reasonable thing to do. Where the cutoff is between reasonable and unreasonable probably varies by region (i.e., the depth of the mixed layer where you're working) but, in general, a good rule of thumb is if you have no data in the upper 50 m, you can probably not use that cast.

It is important to note that inherent in the calibration CTD method is the idea that the bulk of the baroclinicity of the water column is above whatever fixed pressure surface you are using-essentially you're saying there is no independent baroclinicity vertically between that fixed level and the level of your PIES. That assumption is, of course, not perfectly true, so it is always best to use a fixed level that is as deep as is possible. At the same time you cannot, of course, use a fixed level that is deeper than the shallowest instrument in your array (for the Western Boundary Time Series [WBTS] that's around 1100 dbar and for SAM it's about 1300 dbar). Finally, if you select a fixed level that is too deep, not all of the hydrographic data you collected will have observations that deep, so a very deep fixed level can reduce the available hydrographic data. For example, the modern Argo float only collects data to a depth of 2000 dbar. If you're using Argo profile data for your analysis and interpretation of the PIES travel times, you cannot use a fixed pressure level deeper than the 2000 dbar. I have used 1000, 2000, and 3000 dbar fixed levels when calibrating my PIES travel times, depending on the project.

2.4 Applying the GEM Technique to Obtain Profile Data

The travel time interpretation methods idealized in Figure 5 are no longer generally used, although the calibration into travel time on a fixed level is still required. For essentially all PIES travel time analysis that is performed now, the Gravest Empirical Mode (GEM) technique pioneered by Meinen and Watts (2000) is used. The word "gravest" is not often used in modern English (Randy Watts likes the word). In the context here, it essentially means lowest or most basic. An important point to keep in mind with the GEM technique is that it involves an empirical mode, not a theoretical mode. No assumptions about the vertical structure are made in creating the smoothed two-dimensional look-up tables that result from the GEM technique (e.g., no assumption is made about the buoyancy profile, which would be required for the normal mode method).

Details of the GEM technique and the methods of building the two-dimensional look-up tables of temperature, salinity, and density are presented in the Meinen and Watts (2000) paper, with an application to the North Atlantic Current. (Note: there is no salinity GEM lookup table presented in that paper—not because I didn't think we could do it, in fact I had a salinity table created, but because my adviser Randy didn't think anyone would believe us if we tried to publish it. He'd had a lot of years of trying to convince folks that IES and PIES were good tools for oceanography already at that point. A year or so later, he'd convinced himself otherwise and published a salinity GEM table for a different region.) Regardless, the detailed methods for building the GEM look-up tables are presented in the Meinen and Watts (2000) paper. Here I'll only review them briefly.

In essence, to build a GEM two-dimensional look-up table of temperature, for example, one looks at the temperature at each pressure surface (**Figure 6**) and fits a smooth curve to the temperature as a function of the simulated travel time. (As with the earlier PIES travel time interpretation methods, the GEM fields are built using hydrographic data



Figure 6: Plots illustrating the steps for building the GEM look-up table of temperature for the SAM study region at 34.5°S.

from the region. These data do not need to be concurrent with the actual PIES deployment, but long-term trends and such should be considered.) The method I use for building the GEM fields uses cubic smoothing splines to fit to the data points (see the red lines in **Figure 6**), and then values from the splines are extracted at a regular grid of travel time points (e.g., every 1 millisecond). There is nothing magical about these splines. If you prefer, you can use optimal interpolation or any other method that you like. The key is to get reasonably smooth curves going through the temperatures at each depth.

For the vertical resolution of this fitting, in the early days of doing GEM fields when computer power wasn't as good as it is today, I used 20 dbar vertical resolution in the upper 1000 dbar, and 50 dbar vertical resolution below 1000 dbar. Nowadays, you could easily use uniform 20 dbar or even 10 dbar vertical resolution for the full water column, and the computations still won't be prohibitively intensive in any way. After the first smoothing of the temperature data onto a regular grid of simulated travel times is completed, I perform a secondary smoothing vertically using "splines with knots," again nothing magical about this method. This second smoothing removes the small wiggles that can be observed from one pressure surface to the next. In the case of the SAM project example in Figure 6, you can't even see the difference between the lines I plotted based on the first splines and the lines I plotted after the second smoothing, as this particular region wasn't too problematic in this regard.

You will note that the red lines in **Figure 6** extend a bit farther to the right and left (longer and shorter travel times) than the original hydrographic observations (black x symbols). This is intentional. With experience, what we have found is that even when we have a good hydrographic database for the region of the PIES, several years of PIES observations inevitably observe more extreme conditions than have ever been observed by hydrography. This isn't too surprising if you think about it: each PIES observes 365 daily profiles annually. It doesn't take much imagination to understand that the PIES might observe more conditions than the limited number of observations from ships, or even the 10-day Argo float profiles, may ever encounter. An example of a GEM look-up table, or GEM field, of temperature is shown in **Figure 7**.



Figure 7: Example of a GEM look-up table of temperature from the SAM project region at 34.5°S in the Atlantic. Top panel is the GEM field. Middle panel is the root-mean-squared differences between the original CTD/Argo observations and the smoothed field, with the location of the original observations indicated by the gray vertical dotted lines. Bottom panel is the signal-to-noise ratio at each pressure level.

For building my GEM fields, I've typically selected hydrographic observations by collecting all casts within a latitude-longitude box around the PIES array that spans one, two, or three degrees of latitude beyond the PIES locations, and roughly the same extent in longitude beyond the array as well. The goal is to capture casts that represent the breadth of conditions that are likely observed at some point at the PIES sites. Of course, with the finite resources available for observing the ocean, we'll never have a CTD profile that captures every single set of conditions that occurs within the ocean. By using CTD (and Argo) data from a fairly broad region, we hope to capture many of the possibilities. On the other hand, however, we do not wish to include profiles capturing conditions that would never be observed at the location (i.e., we would never want to use a CTD profile from the subpolar gyre in the North Atlantic when building GEM look-up tables for the South Atlantic, as an extreme example). The bounding boxes for the hydrographic observations used in building the GEM fields are shown in most of the GEM-based papers.

Notice that the GEM field is complete at all depths from the surface down to whatever maximum pressure you've selected for your GEM fields. This is, of course, by construction, and as a result one can predict values below the bottom of the ocean in some cases. For example, with the SAM array where the shallowest PIES site is at roughly 1300 dbar, I can still use the calibrated travel time record together with the GEM field (Figure 7) to predict a profile down to 5000 dbar. Of course, the data below the bottom are not something one should analyze. What I typically do after creating my time series of data profiles at a PIES site is go back through and replace all of the data below the bottom with a "NaN" (if you're using Matlab). However you decide to deal with this is up to you of course, but clearly you shouldn't be trying to interpret the empirical estimates from the GEM method below the bottom of the ocean. Interestingly, there is one useful way that the below the bottom GEM estimates can be used. When you are dealing with bottom triangles, the estimates can be useful. I'll return to this topic when I discuss geostrophic velocity estimates.

While the example shown in **Figure 7** is for a GEM field of temperature, similar fields can be made for salinity and density. And, of course, if we can get a density profile at a PIES site, we can then vertically integrate it to get a profile of the dynamic height anomaly (or the geopotential anomaly, if you prefer the more modern terminology). The travel time measurement of the PIES, when combined with the hydrography-derived look-up tables created via the GEM method, can yield full-water-column profiles of temperature, salinity, density, and dynamic height anomaly (see the Meinen and Watts [2000] paper for more details on how this is done).

One of the major strengths of the GEM technique is that, in addition to providing full-water-column estimates of the temperature, for example, the GEM method also provides natural accuracy estimates to go with that profile. It does this by determining the scatter between the original hydrographic observations and the smoothed look-up table values (middle panel, **Figure 7**). These errors can be viewed in raw numbers or as a signal-to-noise ratio (lower panel, **Figure 7**). The signal-to-noise ratio illustrates that the PIES-GEM method is most accurate within the depth range where the main thermocline is observed, for example, from roughly 200 to 1500 dbar in **Figure 7** for the SAM region where the signal-to-noise ratio is ten or better. The signal-to-noise ratio decreases in the uppermost 100-200 dbar (see the next section on seasonal variability), and it also decreases down to single-digit values below roughly 1500 dbar. If you look at the middle panel of **Figure 7**, or the bottom two panels in **Figure 6**, you can quickly see that this low signal-to-noise ratio at depth is occurring just as much because the signal is getting quite weak as because of the scatter around the GEM field.

Nevertheless, a limitation of the GEM method and PIES analysis is that these tools are not well suited to observe the small water property variations that occur in the deep ocean below roughly 2000 dbar. This has proven to be true at all of the locations where I have applied the GEM technique over the years (e.g., North Atlantic Current, Subantarctic Front, Gulf Stream, and the Deep Western Boundary Current in both the North and South Atlantic). Keep in mind, however, that this does not mean the PIES-GEM method cannot capture the velocity structure in the deep ocean. The small water property variations observed below 2000 dbar have, in my experience, never had a major impact on the velocity structure. The PIES-GEM method has always done a good job of describing the velocity structure and its variability at essentially all depths in the aforementioned oceanic currents even if it cannot capture the small water property variations below 2000 dbar or so. I'll discuss the velocity structure a bit more in a later section after the discussion on bottom pressure data processing and analysis.

2.5 Seasonal GEM Correction

You will perhaps have noted that the scatter around the red line in the top left panel of **Figure 6** is much larger than observed in the other panels. This should not come as a surprise since when we build a GEM field we are creating look-up tables that do not consider the time of year. Recall that when we seasonally correct the travel time data (e.g., the discussion around **Figures 3** and **4**), we are not actually changing the hydrographic data. The seasonal variability is still in the CTD/Argo data used to build the GEM fields. As such, there is significant seasonal scatter in the upper few hundred meters of the water column (e.g., **Figure 3**) and, because that scatter is uncorrelated with the "gravest" structure changes that are happening farther down (which one could call the geostrophic scale changes), the seasonal variability shows up as scatter.

Note also the decrease in the signal-to-noise ratio in the bottom panel of Figure 7 near the surface. This is another "limitation" of the GEM technique; however, I put limitation in quotes here because there are two important caveats to this. First, if the primary goal of the analysis is to get geostrophic velocity/transport estimates, then this seasonal variability has little or no impact because the spatial scales of the seasonal heating or cooling are of very large scale, and thus they have no impact once one starts looking at density gradients within an instrumentation array. Second, one can build a seasonal correction GEM field to reduce/eliminate the errors in the PIES-GEM estimated temperature and salinity profiles if one wishes to use exactly the same CTD/Argo data used to build the initial GEM fields. Essentially, one maps the temperature/ salinity/density anomalies between the original hydrographic data and the smooth GEM field at each depth in the seasonally-affected layer to year-day. These variations themselves are then smoothed (via splines or whatever one chooses) onto a regular grid to create a lookup table of seasonal variations. The seasonal correction GEM method was created by Watts et al. (2001), and more details can be found there. Because these seasonal variations have no impact on the geostrophic velocities, I have not generally implemented seasonal correction GEM fields into my analyses.

3. Bottom Pressure Measurement

3.1 Initial Processing to Remove Tides

The other measurement made by a PIES is bottom pressure. As with the travel time, there are some setting choices that the instrument user has for pressure, mainly a choice of how many observations to make during each hour. Again, to conserve battery power I've always recommended that we simply make one pressure measurement each hour. In this way, the PIES stays in a lower-power status longer for each hour. Randy and his group usually configure their instruments to make a measurement every 10 minutes, and the 6-hourly samples are ultimately averaged. I'm not sure how much difference this choice makes, probably not a lot.

The pressure file in the PIES usually starts with a letter "p," and the first row of the file contains a "p" and then the coefficients of the pressure sensor. The data columns start with the same format time word as the travel time file, and each row once again includes all of the data from one hour of measurements. Each pressure sensor contains a temperature sensor that is needed to convert the engineering units into pressure; the second through final columns in the file are alternating columns of pressure and temperature. The pressure values are in decibars with three decimal places; however, the decimal point is not included—a value of 4876738 is thus 4876.738 dbars. The temperature is in units of degrees Celsius with three decimal places, so a temperature of 22878 would be 22.878°C (and would also indicate that the instrument is still on the ship prior to deployment). If the PIES is set as we usually do with only one measurement made each hour, there will be only one column with meaningful pressure values and one column with meaningful temperatures, and all of the other columns will be filled with zeros (see the URI instrument manual for more details on the pressure file format).

The hourly observations of the bottom pressure are strongly influenced by the tides and, unlike the travel time situation, tide signals are an order of magnitude larger than the signals we're interested in for looking at geostrophic velocities. (Typical tide amplitudes are around 1 m in the open ocean, while the signals we're interested in are akin to a few centimeters.) An example of a pressure record from a PIES is shown in **Figure 8**. Notice that the tide signals (included in the full record shown by the gray line) greatly exceed in amplitude the small variations associated with the signals we generally focus on (black line). The initial processing package provided by the URI group uses a Response Analysis method to remove the tide signals (e.g., Munk and Cartwright, 1966; Donohue



Figure 8: Example of a bottom pressure record from a PIES. This instrument was in place at Site B within the WBTS array for about 4 years. The gray line is the raw hourly data, while the red line is the record after the tides were removed via the Response Analysis program. The black line is the record after a first effort at removing the sensor drift.

et al., 2010). Note that the Response Analysis code used in the URI software package does not remove the fortnightly (14-day) tide signals. This is an important point in some areas of the globe and/or when you compare these data with other bottom pressure data sets, as some people remove the fortnightly tides as part of their standard processing (e.g., the UK partners involved in the 26.5°N trans-basin MOC array). The URI package also does not remove the semi-annual or annual tide components. Personally, I think of the fortnightly tides as likely being in quasi-geostrophic balance, so I don't think they should be removed, and the same goes for the semi-annual and annual tides. Certainly they shouldn't be removed thoughtlessly and automatically-the science user should think about their application and should remove them only when they decide to do so based on their problem of interest.

3.2 Initial Processing to Remove Drift

Another significant issue with bottom pressure data is sensor drift. This has been a well-known problem with all bottom pressure sensors for quite some years (e.g., Watts and Kontoyiannis, 1990; Donohue *et al.*, 2010). Essentially, all bottom pressure sensors exhibit an exponential drift over the first few weeks/months of the record, a linear drift over the full record and, in some cases, both. Amusingly enough, 25-30 years ago when the quality of the sensors was worse, it was often much easier to identify the exponential drifts than it is today, as the amplitude of the drift was much larger than the amplitude of any of the real signals observed during the record.

Nowadays, the exponential drift is oftentimes (but not always) of a similar order as the observed signals, which makes it difficult to determine if what we observe at the start of a record is a drift or if we just happened to deploy our instrument in the midst of a large event. Removing the exponential at the beginning of the record when the amplitude is small is quite subjective, unless you are lucky and you have overlapping, redundant observations from another sensor. My personal tendency is to remove only those exponential drifts at the start of the record that are larger in amplitude than any of the other events during the record. The linear drift over the whole record is a bit easier to detect, particularly when you have a 3+ year record from the PIES.

An important point to keep in mind with drift removal is the implication for erroneously removing ocean signals. For the exponential signal, since it commonly only spans the first 1-3 months of the record, this is the only period that can be influenced. For the record length linear trend, when we remove it, we can also be removing/altering the long-period variability. The time scale of variability we could be impacting by that is a function of the length of the record. For example, if we have a 4-year long record and we remove a linear drift from it, we're obviously not influencing the seasonal and shorter-period variability. If one imagines a sine wave and then evaluates how long the period of that sine wave would need to be for a 4-year segment of it to appear linear, you can quickly see that the sine wave would need to have a period greater than a decade. For a 2-year long record, perhaps we might be

looking at a sine wave with a period of 9-10 years. Since we typically deploy our PIES for 3-4 years, I argue that we cannot say much about variations at periods longer than a decade due to the linear pressure drift removal. However, year-to-year variations, for example, should still be robust in the bottom pressure sensor data. Some folks in the community kind of throw up their hands in defeat regarding the drift problems in the bottom pressure. Personally, I think that's too pessimistic for the above reasons. It is, however, something that the science users of PIES and other bottom pressure measuring systems need to consider and keep under consideration.

Figure 9 illustrates the steps of the pressure processing and removal of the drift. As you can see in the example, the exact location where the exponential behavior ends and where the drift becomes linear is somewhat subjective. Obviously, if you remove just a linear drift from the full record, the events at the beginning will be much higher in amplitude than any other events during this record. Somewhere after 100-200 days, the exponential component of the drift peters out and the drift becomes linear, but this example is a good one to illustrate that the choice of where that transition occurs is somewhat subjective at times. Most of the time, the exponential part is done faster than is observed in this example. The final drift-removed version of the pressure record is shown in the bottom panel of Figure 9. After drift removal, the final step in processing the bottom pressure is smoothing the time series with a second-order Butterworth filter with a 72-hour cutoff period to remove any remaining residual of the tides, passing the filter both forward and back to avoid phase shifts; the data are then subsampled to one value per day at noon GMT (at least that's how we usually do it).

4. Merging Travel Time and Bottom Pressure Records from Subsequent Deployments

For long-term studies like the WBTS (e.g., Meinen *et al.*, 2013a) and SAM projects (e.g., Meinen *et al.*, 2017), you are likely to have multiple PIES deployments one after another at the same site. To analyze the 5-10+ years of data, you'll want to merge the data from the deployments



Figure 9: Example of drift removal from a PIES record from the WBTS Site E with the x-axis time in days. Top panel illustrates the full hourly pressure record collected. Note the roughly 1 m range that is dominated by the tide signals. Middle panel shows the record after the tides have been removed using the URI Response Analysis script (red line) and the exponential-linear drift that has been fit to the time series (blue line). Bottom panel shows the final hourly pressure record after removing the drift and the tides.

to build long 5-10+ year travel time and bottom pressure records at each site. Generally, what we do on our projects is blend together the subsequent records (both travel time and bottom pressure) using the time series of daily values after the records have been smoothed with the low-pass filter and subsampled to one value per day at noon GMT.

For travel time data, the records that are merged are those that have been calibrated into travel time on a fixed level, e.g., "tau1000" values of travel time at 1000 dbar. In the best case scenario where you have multiple CTD casts that occur during each of the deployments, there should be no abrupt shifts at the change-over points between the end of one PIES calibrated travel time record and the start of the next PIES calibrated travel time record at the same site. At least this is true if there is no time gap between the deployments, e.g., if the instruments were recovered and deployed in quick succession. Obviously, if there was a long gap (i.e., longer than 1-2 days) a major shift would be observed between deployments. If there is no long gap between records, the travel times can be merged into a single record with no other steps needed. If there is a long time gap between the end of one record and the start of the next (e.g., longer than 1-2 days), and/or if there is an abrupt shift right at the change-over point (when the instrument is being turned around), a bit more thought is required.

A spot of bad luck could cause a turn-around to be happening right in the middle of a big event. Therefore, an abrupt shift is not necessarily a problem but if, for example, you see that the mean of the first multi-year deployment is 5 milliseconds lower than the mean of the second multiyear deployment, and most of the signals during those two records are only 2-3 milliseconds in amplitude, there is likely a problem. Obviously, this is somewhat subjective, but if you observe such a big shift in the travel time, you have to use whatever other information you have to determine whether the calibration of one or the other record is in error. Is one set of calibration CTDs of known lower quality than the other? Is there a similar abrupt shift in the neighboring station that was turned around a few days later instead of during the event? Are the high travel time events in the second deployment all several milliseconds higher than the high travel time events in the first deployment? These are the sorts of questions to ask to help make a decision on how to fix or not fix a big shift at the instrument change-over point.

For bottom pressure, the merging of subsequent records is a bit more complex. As shown in **Figure 9**, the process of removing drift from the pressure record also removes the mean, yielding a zero record-length mean for each subsequent record. By default, there is no reason why the record-length means of subsequent records would need to be exactly equal (i.e., interannual/decadal variability). It is important when merging subsequent pressure records to evaluate whether an offset/correction should be applied to subsequent records as part of the merging process. This is, of course, once again somewhat subjective. What I do is look and see once again whether the highest high frequency events in record one all seem to exceed all of the highest high frequency events in record two, and are all of the lowest lows in record one weaker than the lowest lows in record two? If yes, maybe an adjustment is necessary. Keep in mind that when you redeploy a new PIES at the same location where you just recovered one, the new PIES will not be at exactly the same location on the bottom. (Of course you're going to remove the record-length mean pressure value from each instrument as part of the merging process, but there may have been legitimate interannual/ decadal variability that you're removing at the same time.)

It is easy to imagine the second PIES could be 3 centimeters shallower, for example, which would have a very significant (~0.03 dbar) impact on the pressure but a travel time impact (0.04 milliseconds) that could not be observed. As with the travel time, it is, of course, possible that you might have a PIES turn-around that happens right at a period when the bottom pressure is changing rapidly. This could happen. The best thing you can do is use the information from neighboring sites to evaluate whether such an abrupt shift was occurring or if it is just an arbitrary shift associated with removing the mean values from the subsequent pressure records.

5. Estimating Relative Velocity Profiles

5.1 Use of PIES-GEM Profiles

As noted earlier in the discussion of the application of the GEM method, density profiles estimated at a PIES site can be vertically integrated to get dynamic height anomaly profiles. Dynamic height anomaly profiles at neighboring PIES sites can also be differenced via the standard geostrophic method (i.e., the thermal wind equation) to yield profiles of the baroclinic geostrophic velocity relative to an arbitrary level of no motion. As with any application of the geostrophic method, only the horizontal component of the velocity orthogonal to the line between the PIES sites is obtained. The velocity represents a true average across the span between the two PIES sites. Two-dimensional arrays of PIES can provide both horizontal components of velocity (e.g., Meinen et al., 2009). Regardless of whether one is working with a single line of PIES (e.g., the WBTS and/or SAM arrays) or a two-dimensional array of PIES (e.g., the Synoptic Ocean

Prediction array), these methods can provide full-watercolumn profiles of velocity.

Whatever arbitrary level of no motion that you implement is up to your application. I often select the sea surface as an initial level of no motion and then shift it to the deepest common depth between the pair of PIES if I'm looking at a surface trapped current (e.g., the Gulf Stream), where I might instead shift the "level of no motion" to something like 800 dbar if I'm studying the Deep Western Boundary Current. Regardless of what you choose, it doesn't matter much for most boundary current studies because the idea of a time-invariant level of no motion is pretty much dead within the field of physical oceanography these days.

Many papers and analyses (such as those cited herein) have shown that there is no time-invariant level of no motion near strong oceanic currents, and worse yet, the bottom velocity (or reference level velocity) in the real ocean almost always varies quite independently of the baroclinic velocity profile relative to the arbitrary level of no motion. If you only measure baroclinic transport relative to an assumed level of no motion in a strong ocean current, the flow profile often bears little absolute relationship with the actual ocean velocity profile. It isn't that the shape of the profile in the vertical that is determined from the baroclinic shear is incorrect. However, the bottom velocity (reference level velocity) variations in the boundary currents are often strong, variable on a wide range of time scales, and vary independently from the baroclinic variations. To actually say anything about ocean velocity you need to have either direct velocity measurements (e.g., moored current meters or moored acoustic Doppler current profilers) or you need to measure both the baroclinic and the barotropic components of the flow. The latter is where the bottom pressure data becomes important.

5.2 Use of PIES Bottom Pressures

Obtaining a reference velocity from two neighboring PIES is at first glance quite simple and straightforward—you just difference the two pressure records and scale them with the Coriolis parameter and the distance between the two sites via the geostrophic method. Of course, it is a bit more complex than this in application. Keep in mind that neighboring PIES are not at the same level vertically compared to a constant geopotential surface and, in some cases (e.g., Sites A and B in SAM, Sites A and A2 in WBTS), the difference in depth between neighboring sites can be more than 1000 m. If you consider how the geostrophic method is normally discussed, the horizontal gradients of pressure are expected to be calculated along a constant geopotential surface. This is something that must be considered when thinking how to obtain the reference velocity.

Starting with a simpler situation, consider Sites B and C in the WBTS array. The depths of these two instruments are only ~200 m apart, with one at ~4600 m and the other at ~4800 m. If one looks at geostrophic velocity profiles determined between neighboring CTD stations from hydrographic sections along this line, one quickly sees that there is very little vertical shear in the velocity in the bottom few hundred meters between Sites B and C. This means the horizontal pressure gradients will be essentially the same at these depths—using one pressure record from a PIES at 4600 m depth and another pressure record from a PIES at 4800 m depth to calculate the horizontal gradient should be fine. This is what we commonly do, as in most cases with our various PIES arrays (e.g., WBTS and SAM) the depth difference from neighboring sites is usually quite small and is occurring in a depth layer where the vertical gradients in velocity are small.

The counter-example to this, however, is the situation we find ourselves in when we are looking at the pressure data from Site A in either the WBTS or SAM array versus the pressure data at the next site offshore in either array (Site A2 for WBTS or Site B for SAM). In these cases, there is 2000-3000 m of depth difference between the neighboring sites. It is obvious that assuming zero baroclinic vertical shear exists between ~3500 m for WBTS Site A2 and ~1100 m for WBTS Site A is a very different situation than making the same assumption between ~4800 m and ~4600 m for WBTS Sites B and C. What to do in these situations is a question that can only be answered based on the science question(s) being addressed.

One method that I have explored to deal with this situation is to use the PIES-GEM profiles of density at the deeper site in a pair to integrate the baroclinic signal between the nominal depths of the two instruments. In theory, if one integrates the density between ~3500 m and ~1100 m at Site A2, for example, one can subtract these baroclinic signals from the bottom pressure record at Site A2 and get the pressure signals that would have been observed if the Site A2 PIES was at ~1100 m depth. The gradient between Site A and the modified Site A2 pressure records can then be used to obtain a reference velocity between these neighboring sites. I've explored this option but haven't published using it. This is still an area where a science user will need to plan how they want to deal with the problem and explain and justify their method carefully in their publications.

A critically-important issue I haven't touched on yet relates to the time-mean reference velocity determination. Consider two neighboring PIES: if we simply compare their two pressure records and discover that there is a time-mean difference between their pressure data, we have no way of knowing if one PIES is located at a deeper depth than the other or if they are at the same depth. There is, however, a time-mean geostrophic flow orthogonal to the horizontal line between the two sites. This is the well-known leveling problem, which has been discussed many times (e.g. Donohue *et al.*, 2010). Because there is no solution to the leveling problem using only the two bottom pressure records, we normally remove the time-mean bottom pressure from each PIES record (see bottom panel of **Figure 9**).

If the bottom pressure data cannot give us the time-mean reference velocity, how do we get it? Randy Watts invented the CPIES to address this exact problem; however, it only represents a solution if the instruments are located horizontally within a correlation length scale to one another. In that case, one can argue the average of the two current meters from neighboring CPIES is representative of the horizontal average flow between the two sites. Hence, the time mean of those horizontally-averaged currents can be added to the bottom pressure gradient-derived reference velocity variations to yield absolute reference velocities that include the time mean. Unfortunately, nearly all of my PIES work has involved sites that are horizontally much farther apart than a correlation length scale. As such, a current meter measurement at each site cannot reasonably be assumed to be representative of the horizontally-averaged flow between the two sites. For my work, I have generally either obtained the time-mean reference velocity either from historical data (e.g., multiyear current meter observations from the Subtropical Atlantic Climate Studies arrays in the WBTS array region) or from a lengthy run of a numerical model (e.g., the Ocean General Circulation Model for the Earth Simulator [OFES] model for our SAM array work). The time mean reference velocity is one thing that the PIES-GEM method does not provide, at least for widely spaced arrays, so the time-mean bottom/reference flow is generally not well characterized.

6. Combining Relative and Reference Velocity for Absolute Velocities

Once a reference velocity is obtained from the bottom pressure differences, the next step is to combine it with the relative velocity profile to obtain a profile of absolute velocity. The example discussed at the end of the previous section illustrates what is perhaps the biggest complication to this step. Consider Sites A and A2 from the WBTS array, which are at nominal depths of about 1100 m and 3500 m, respectively. The dynamic height anomaly profile at Site A will have values between the surface and 1100 m, while the profile at Site A2 will have values between the surface and 3500 m. When you calculate the difference, you only have meaningful values between the surface at 1100 m. This is the classic bottom triangle problem, which has been discussed many times over the last four-plus decades and which affects all geostrophic velocity methods, not just PIES analyses (e.g., it impacts estimates from CTD sections, dynamic height moorings, etc.).

In addition to the problem of figuring out how to calculate the velocity within the bottom triangle region (i.e., below 1100 m and above 3500 m in the example), this problem also impacts the application of the reference velocity. Nominally, one would apply the reference velocity at the deepest common depth between the two sites, i.e., at 1100 m in the example above. This is one reasonable option, and it is one I have used in the past, but it still leaves the problem of how to deal with the flow within the bottom triangle. There are several ways one can deal with the bottom triangle issue. What is commonly done (by me and many before me) when dealing with the bottom triangle is to extend the velocity profile from the deepest common depth (1100 m in the above example) down to the middepth of the pair of sites ((1100+3500)/2 = 2300 m in the)above example), and then assume the area of the bottom triangle is roughly equal to the rectangular area described by that mid-depth level and the horizontal span between the sites. This is a common method for dealing with a bottom triangle-where you find differences in earlier work is in how one extends the profile downward. One way to deal with it is to make an assumption that the flow is constant from the deepest common depth downward. Another option is to assume that the shear in the layer just above the deepest common depth is constant and can be used to extrapolate the profile downward into the bottom triangle, although one then has the challenge of determining what thickness layer one wishes to calculate the vertical shear over.

A third method I like and often use in my PIES analyses is to take advantage of the hydrography information in the GEM field to extrapolate downward into the bottom triangle on a somewhat more informed basis. If one were sticking with the example above to use the PIES-GEM profiles at Site A down to a depth of 2300 m, instead of replacing all of the data below the bottom with values of NaN, and one then calculated the horizontal gradient between the profiles at Site A and Site A2, the relative velocity profile would extend down to 2300 m. The extrapolation down into the bottom triangle would be informed by all of the hydrography that was used to build the GEM fields. One could then apply the reference velocity at 2300 m depth, and when transport was integrated between Sites A and A2, one would have a reasonable estimate of the flow within the bottom triangle.

At this point you have temperature, salinity, density, and absolute velocity information at and/or between all of your sites. Geostrophic absolute transports can be integrated from the absolute geostrophic velocities, as we've done in many of the WBTS and SAM studies cited herein. One further note: the above is my recommended method for dealing with bottom triangles and for applying the reference velocity. Having said that, there are circumstances where I might do something different.

In our work to estimate the MOC and look at the Deep Western Boundary Current at 34.5°S using the SAM/ South Atlantic MOC Basinwide Array instruments (e.g., Meinen et al., 2013b, 2017, 2018), we decided to not to apply the reference velocity at the bottom because the bottom cell in the OFES model (from which we were getting our time-mean reference flow) was introducing some very odd values. Instead, in that case we chose to apply the reference velocity at 1500 dbar with the assumption that the baroclinic shear between the instrument levels and 1500 dbar was negligible. I mention this to highlight the fact that many of the details of how to implement the PIES-GEM methods presented herein must be reviewed and considered based on the situation and the resources and observations at hand. These techniques are not just simple black boxes that can be applied without thought. It is important to think things through and make sure there isn't a reason to do something differently. The world is made up of special cases, of course.

7. Final Thoughts and Recommendations

The PIES-GEM methods I've discussed herein are powerful, but there has always been a bit of doubt in the broader physical oceanographic community about whether the method has been oversold and/or is promising more than it can possibly deliver. There is some validity to this concern, as there are things that the PIES-GEM method will never be able to do. An example I often use to illustrate what the PIES-GEM method cannot do is monitor for the arrival of freshly-ventilated North Atlantic Deep Water within the Deep Western Boundary Current. The temperature-salinity variations within the North Atlantic Deep Water class are quite small, and they tend to be uncorrelated with variations in the main thermocline/ halocline/pycnocline. The PIES-GEM method is great for capturing the "gravest," what I might call the "dominant," mode of variability. Small variations in temperature and salinity (e.g., internal waves, etc.), which normally have little impact on the large-scale velocity structure, are not well captured by the PIES-GEM method and probably never will be. This must be kept in mind. As the old joke goes: to a simpleton with a hammer, every problem looks like a nail. The PIES-GEM method cannot solve all ocean velocity and transport measuring problems. Keep in mind its limitations.

Having said that, the PIES-GEM method is quite powerful. Direct comparisons of the absolute velocities/transports derived via the PIES-GEM method have been performed against current meters (e.g., Meinen and Watts, 2000) and against the combination of dynamic height moorings and bottom pressure gauges (e.g., Meinen *et al.*, 2013a). These comparisons have shown excellent agreement. Similar comparisons have been made against moored temperature and salinity sensors, for example, and have shown impressive agreement. (A paper I'm presently writing on the Antilles Current shows good agreement with data from sensors on the dynamic height moorings deployed in the WBTS region. In fact, these comparisons were able to point out some problems in the calibration of the salinity sensors on the tall moorings.)

While it is important to bear in mind the limitations of the PIES-GEM method, also be aware of the quality of data that it can provide for the relatively low cost of PIES moorings compared to other methods like tall moorings. Tall moorings may often provide somewhat more accurate values (although those data sets have their own issues, particularly mooring motion); however, the cost differential between a PIES mooring and a tall mooring tends to be a factor of three to five or more, roughly \$45K and \$250K, respectively, now in the year 2018. The PIES can also be deployed for longer periods without recovery and redeployment (3-5 years for a PIES versus 1-2 years for a tall mooring). The cost of a PIES is also essentially the same as the cost of a small bottom-pressurerecorder mooring, such as those used by the University of Miami and UK participants in the 26.5°N MOC array adjacent to each of their tall moorings, with the PIES providing the additional travel time measurement that the small bottom pressure recorder mooring does not provide. The PIES is a powerful tool that can be used to quantify transports and, while I haven't done this myself, one can combine the velocity, temperature, and salinity data to go after temperature and salt transports as well.

8. Papers for Further Reading

I cannot stress enough that there are many fine researchers who are doing or who have done work with IES/PIES/ CPIES, and by highlighting mostly my own papers in these notes I'm simply focusing on the work I know best. There are important and helpful papers on PIES analysis led by Randy Watts, Kathy Donohue, Silvia Garzoli, Mark Wimbush, and many others. I am not in any way attempting to say my work on PIES has been the best or most important. It is just the work for which I am most familiar. Randy Watts' group at URI maintains a fairly complete list of publications that have used inverted echo sounders and the variants including pressure, current meters, and more. These publications that can be accessed at http://www.po.gso.uri.edu/dynamics/ies/iesbibupdated. pdf. I encourage you to read many of these papers to get a flavor of what can be accomplished with PIES.

9. References

- Chen, C.-T., and F.J. Millero, 1977: Speed of sound in seawater at high pressures. *Journal of the Acoustical Society of America*, 62(5):1129–1135 (doi:10.1121/1.381646).
- Del Grosso, V.A., 1974: New equation for the speed of sound in natural waters (with comparisons to other equations). *Journal of the Acoustical Society of America*, 56(4):1084–1091 (doi:10.1121/1.1903388).
- Donohue, K.D., D.R. Watts, K.L. Tracey, A.D. Greene, and M. Kennelly, 2010: Mapping circulation in the Kuroshio Extension with an array of current and pressure recording inverted echo sounders. *Journal of Atmospheric and Oceanic Technology*, 27(3):507–527 (doi: 10.1175/2009JTECHO686.1).
- Kennelly, M., K.L. Tracey, and D.R. Watts, 2007: Inverted echo sounder data processing manual. GSO Technical Report No. 2007-02 (available at: digitalcommons.uri.edu/physical_oceanography_ techrpts/2).
- Meinen, C.S., and D.R. Watts, 1997: Further evidence that the sound speed algorithm of Del Grosso is more accurate than that of Chen and Millero. *Journal of the Acoustical Society of America*, 102(4):2058–2062 (doi:10.1121/1.419655).
- Meinen, C.S., and D.R. Watts, 1998: Calibrating inverted echo sounders equipped with pressure sensors. *Journal of Atmospheric* and Oceanic Technology, 15(12):1339–1345 (doi:10.1175/1520-0426(1998)015<1339:CIESEW>2.0.CO;2).
- Meinen, C.S., and D.R. Watts, 2000: Vertical structure and transport on a transect across the North Atlantic Current near 42°N: Time series and mean. *Journal of Geophysical Research*, 105(C9):21,869– 21,891 (doi:10.1029/2000JC900097).
- Meinen, C.S., and S.L. Garzoli, 2014: Attribution of Deep Western Boundary Current variability at 26.5°N. *Deep-Sea Research, Part I*, 90:81–90 (doi:10.1016/j.dsr.2014.04.016).
- Meinen, C.S., and D.S. Luther, 2016: Structure, transport, and vertical coherence of the Gulf Stream from the Straits of Florida to the Southeast Newfoundland Ridge. *Deep-Sea Research, Part I*, 111:16–33 (doi:10.1016/j.dsr.2016.02.002).
- Meinen, C.S., S.L. Garzoli, W.E. Johns, and M.O. Baringer, 2004: Transport variability of the Deep Western Boundary Current and the Antilles Current off Abaco Island, Bahamas. *Deep-Sea Research, Part I*, 51(11):1397–1415 (doi:10.1016/j.dsr.2004.07.007).
- Meinen, C.S., M.O. Baringer, and S.L. Garzoli, 2006: Variability in Deep Western Boundary Current transports: Preliminary results from 26.5°N in the Atlantic. *Geophysical Research Letters*, 33(17):L17610 (doi:10.1029/2006GL026965).
- Meinen, C.S., D.S. Luther, and M.O. Baringer, 2009: Structure, transport and potential vorticity of the Gulf Stream at 68°W: Revisiting older data sets with new techniques. *Deep-Sea Research, Part I*, 56(1):41–60 (doi:10.1016/j.dsr.2008.07.010).

- Meinen, C.S., A.R. Piola, R.C. Perez, and S.L. Garzoli, 2012: Deep Western Boundary Current transport variability in the South Atlantic: Preliminary results from a pilot array at 34.5°S. *Ocean Science*, 8(6):1041–1054 (doi: 10.5194/os-8-1041-2012).
- Meinen, C.S., W.E. Johns, S.L. Garzoli, E. van Sebille, D. Rayner, T. Kanzow, and M.O. Baringer, 2013a: Variability of the Deep Western Boundary Current at 26.5°N during 2004-2009. Deep-Sea Research, Part II, 85:154–168 (doi:10.1016/j.dsr2.2012.07.036).
- Meinen, C.S., S. Speich, R.C. Perez, S. Dong, A.R. Piola, S.L. Garzoli, M.O. Baringer, S. Gladyshev, and E.J.D. Campos, 2013b: Temporal variability of the Meridional Overturning Circulation at 34.5°S: Results from two pilot boundary arrays in the South Atlantic. *Journal of Geophysical Research*, 118(12):6461–6478 (doi:10.1002/2013JC009228).
- Meinen, C.S., S.L. Garzoli, R.C. Perez, E. Campos, A.R. Piola, M.-P. Chidichimo, S. Dong, and O.T. Sato, 2017: Characteristics and causes of Deep Western Boundary Current transport variability at 34.5°S during 2009-2014. Ocean Science, 13(1):175–194 (doi:10.5194/os-13-175-2017).
- Meinen, C.S., S. Speich, A.R. Piola, I. Ansorge, E. Campos. M. Kersale, T. Terre, M.-P. Chidichimo, T. Lamont, O.T. Sato, R.C. Perez, D. Valla, M. Van den Berg, M. Le Henaff, S. Dong, and S.L. Garzoli, 2018: Meridional Overturning Circulation transport variability at 34.5°S during 2009-2017: Baroclinic and barotropic flows and the dueling influence of the boundaries. *Geophysical Research Letters*, 45(9):4810–4188 (doi:10.1029/2018GL077408).
- Millero, F.J., and X. Li, 1994: Comments on "On equations for the speed of sound in seawater." *Journal of the Acoustical Society of America*, 95(5):2757–2759 (doi:10.1121/1.409844).
- Munk, W., and D. Cartwright, 1966: Tidal spectroscopy and prediction. *Philosophical Transactions of the Royal Society of London, Series A, Mathematical and Physical Sciences* 259(1105):533–581 (doi:10.1098/rsta.1966.0024).
- Rossby, T., 1969: On monitoring depth variations of the main thermocline acoustically. *Journal of Geophysical Research*, 74(23): 5542–5546 (doi:10.1029/JC074i023p05542).
- Watts, D.R., and H.T. Rossby, 1977: Measuring dynamic heights with inverted echo sounders: Results from MODE. *Journal* of Physical Oceanography, 7(5):345–358 (doi:10.1175/1520-0485(1977)007<0345:MDHWIE>2.0CO;2).
- Watts, D.R., and H. Kontoyiannis, 1990: Deep-ocean bottom pressure measurement: Drift removal and performance. *Journal of Atmospheric and Oceanic Technology*, 7(4):296–306 (doi:10.1175/1520-0426(1990)007<0296:DOBPMD>2.0.CO;2).
- Watts, D.R., C. Sun, and S. Rintoul, 2001: A two-dimensional gravest empirical mode determined from hydrographic observations in the Subantarctic Front. *Journal of Physical Oceanography*, 31(8): 2186–2209 (doi:10.1175/1520-0485(2001)031<2186:ATDGEM>2.0. CO;2).



National Oceanic and Atmospheric Administration

OFFICE OF OCEANIC AND ATMOSPHERIC RESEARCH Atlantic Oceanographic and Meteorological Laboratory 4301 Rickenbacker Causeway Miami, FL 33149

www.aoml.noaa.gov