Atlantic Canada's Real-Time Water Level System Observations, Predictions, Forecasts and Datums on the Web

Phillip MacAulay (CHS Atlantic),
Charles O'Reilly (CHS Atlantic), and
Keith Thompson (Dalhousie University)

Abstract

As part of Canada's national initiative to revitalize its water level monitoring capacity, CHS Atlantic Region has upgraded and enhanced its water level monitoring system with: new data collection equipment, new database architecture, and a new real-time, web-based data dissemination system for use by emergency measures organizations (via the Atlantic Storm Prediction Center ASPC and Alaska Tsunami Warning Center ATWC). CHS Atlantic's new water level system now provides: up to the minute water level data from multiple sensors, constituent based predictions, local tidal Datum information, and shortly will include water level forecasts.

A water level forecast differs from simpler constituent based predictions because it seeks to include non-tidal effects on water levels, those of weather, departures of water levels at tidal frequencies from expected values, and the collective result of everything else that produces water level variations (infra-gravity waves, harbour seiches, tsunamis, shelf waves …). CHS Atlantic's forecasts include weather effects derived from the DalcoastII operational oceanographic storm surge model running at Dalhousie University. All other non-tidal effects are included through use of a multiple band-pass Kalman filter based technique.

This paper will provide an overview of Atlantic Region's real-time water level system. It will present and discuss its new forecasting capability, and will look towards future water level system enhancements, including the development of continuous, 3D tidal Datum surfaces and their extension to "anywhere", "anytime" on-Datum water level predictions using WebTide.

Introduction

The paper will provide an overview of the Canadian Hydrographic Service (CHS) Atlantic Region's Real-Time Water Level (RTWL) system. It will present and discuss the system's new water-level forecasting capability, and it will look towards future system enhancement to include anywhere, anytime, on-Datum water level predictions.

Overview of Atlantic Regions Real-Time Water Level Infrastructure

CHS Atlantic's Permanent Water Level Network (PWLN) consists of 16 operational tide gauge stations (see Figure 1). Three have been designated as long term sea level stations of the Global Sea Level Observing System (GLOSS), and six have been designated as storm surge stations (based on the frequency and severity of storm surge activity). Lastly, six others have been assigned duty as tsunami warning stations because the combination of their locations and the bathymetry of the continental shelf conspire to make them first strike points for tsunamis propagating into Atlantic Canadian waters from various deep water directions.

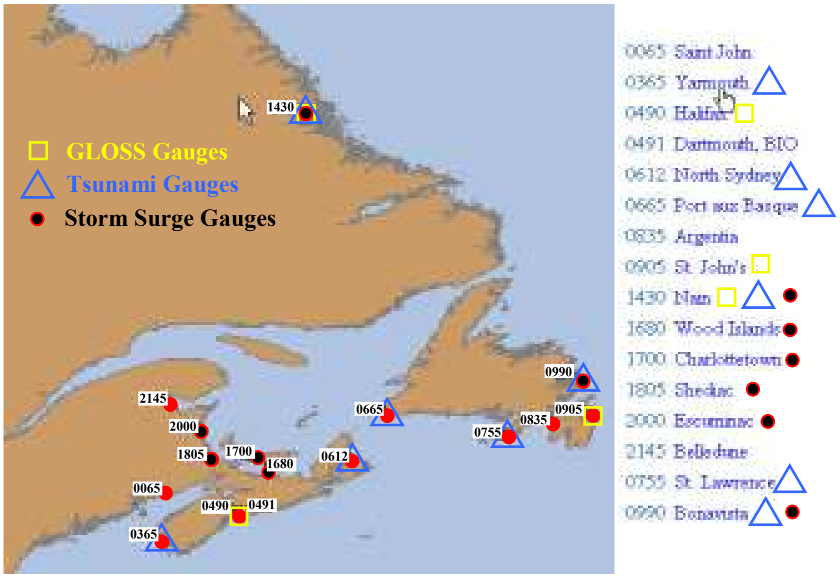

Figure 1: CHS Atlantic's permanent water level stations

Prior to CHS Atlantic's revitalization program each station in the PWLN typically had the following:

- a telephone line connection,
- one or more stilling wells,
- a reference tape drop,
- a Sutron 8210 data logger,
- the primary water level sensor (a float and pulley rotary encoder), and
- if possible a backup pressure or bubbler sensor.

Water level data was logged once every 15 minutes and all stations were polled once a day via modem. In the past, delays of hours, days or even weeks between the collection and dissemination of the water level data had been acceptable and sufficient for basic tidal analysis, traditional hydrographic survey work, and for long term sea level monitoring purposes. However, more contemporary applications like storm surge and tsunami warning systems and the newer navigational aid systems as they come online, for example real-time harbour clearances and dynamic Electronic Navigational Charts (ENCs), require more timely data collection and data dissemination.

To meet these needs CHS Atlantic Region has upgraded and enhanced most aspects of its tidal water level monitoring and data collection infrastructure and CHS Atlantic tide gauge sites are now equipped with:

- versatile Sutron XPert dataloggers
- three water level sensors (see Figure 2)
    - float and pulley optical encoders
    - Sutron Continuous Flow Bubbler sensors
    - Spectre Sensor Submersible pressure sensors
- Sutron Satlink Geostationary Operational Environmental Satellites (GOES) hardware.

In addition, during the 2008 field season highly accurate laser sensors are scheduled to be installed at CHS Atlantic's 3 GLOSS stations. These should be useful additions for accurately tracking sea level change as they provide an absolute distance measurement, eliminating the errors introduced when setting sensors relative to a tape drop. They only require leveling in to existing benchmarks.

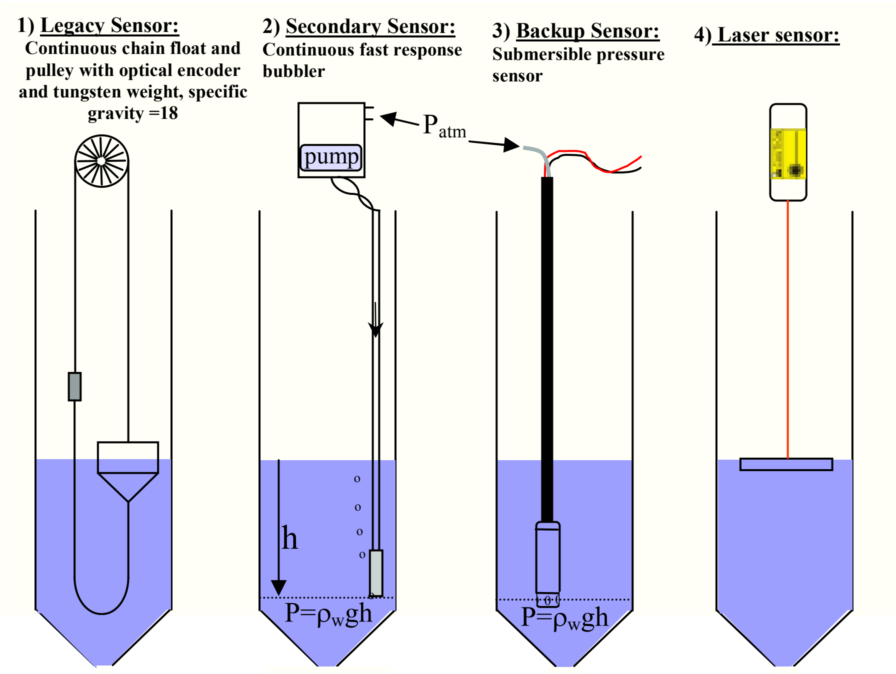

Figure 2: Tide gauge sensors

Time-averaged water level data is now collected from all 3 sensors at all PWLN tide gauges at one minute interval and is downloaded to Atlantic Region's new Real-Time Water Level (RTWL) database once every 10 minute using Sutron's XConnect data collection software. This data is then immediately available to all other aspects of CHS Atlantic's water level system (see Figure 3).

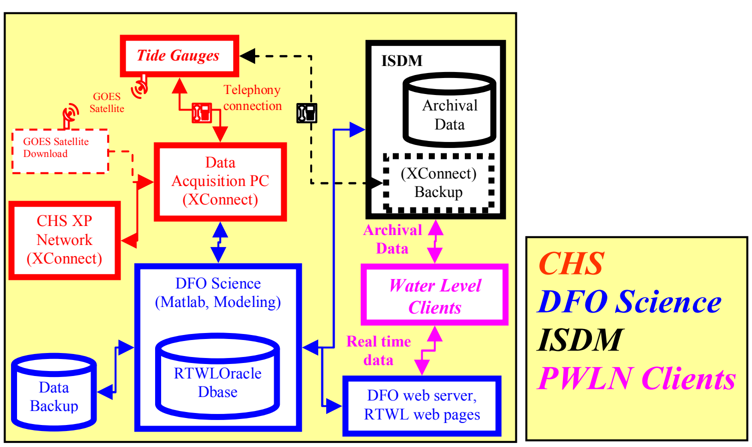

Figure 3: Data flow schematic of CHS Atlantic's water level system

To meet Atlantic tsunami warning system requirements, access to real-time water level data is available to various Emergency Measures Organizations and to the Atlantic Storm Prediction Center (ASPC) and the Alaska Tsunami Warning Center (ATWC) through password protected web pages on a Department of Fisheries and Oceans (DFO) Science Atlantic's web server. Data selection is made through an Environmental Systems Research Institute (ESRI) Map interface and data is displayed using in-house written SQL and Java routines (see Figure 4).

There are similar systems and access to real-time data in use elsewhere. For example, the US has the National Oceanographic and Atmospheric Administration's (NOAA) Tides and Currents web site (Tides and Currents, 2007) and its associated Physical Oceanographic Real-Time System (PORTS) website (Physical Oceanographic Real-Time System (PORTS), 2007). The UK has developed internet access to water level data through the National Tidal and Sea Level Facility (NTSLF) (Proudman Oceanographic Laboratory, Sea level networks, 2008). In Europe, services are available through the European Sea Level Service (ESEAS) (European Sea Level Service, 2006), and the Monitoring Network System for Systematic Sea Level Measurements in the Mediterranean and Black Sea (MedGLOSS) (Monitoring Network System for Systematic Sea Level Measurement in the Mediterranean and Black Sea, 2008).

CHS Atlantic's RTWL web pages, and like systems, provide both observations and constituent-based predictions covering the recent past, up to the very near present, as well as constituent-based predictions into the near future. Display of this combination of data and predictions is useful for numerous reasons. The most significant being that the differences between observations and predictions, the water level residual (observations-predictions), shows all observed non-tidal water level behavior, where tidal behavior is defined strictly as that following constituent based analysis and prediction. The presence of the tidal residual is one of the primary reasons for continued monitoring of coastal water levels at permanent water level stations. Assuming that high frequency water level variations have already been removed either through the use of stilling wells (which act as analog filters), and/or temporal averaging of sensor outputs, and/or through depth of sensors placement (via depth attenuation of the pressure signals of short period surface waves), the residual includes:

- all weather induced affects on water levels (for example due to storm surges, positive and negative),
- all departures of water levels at tidal frequencies from expected values, (for example due to short term variations in the tidal constituent amplitudes and phases due to temporary variations in water column density structure), and
- the collective result of everything else that produces water level variations (infra-gravity waves, harbour seiches, tsunamis, shelf waves …).

Thus, the water level residual shows what we, and the coastal environment in general, are not strictly adapted to and expecting at low frequencies. It is these changes in water levels that lead to the extreme water level events that have the largest short term impacts on coastal environments, populations, and infrastructure, and which often represent (particularly in the case of negative storm surges) one of the least appreciated dangers for navigation.

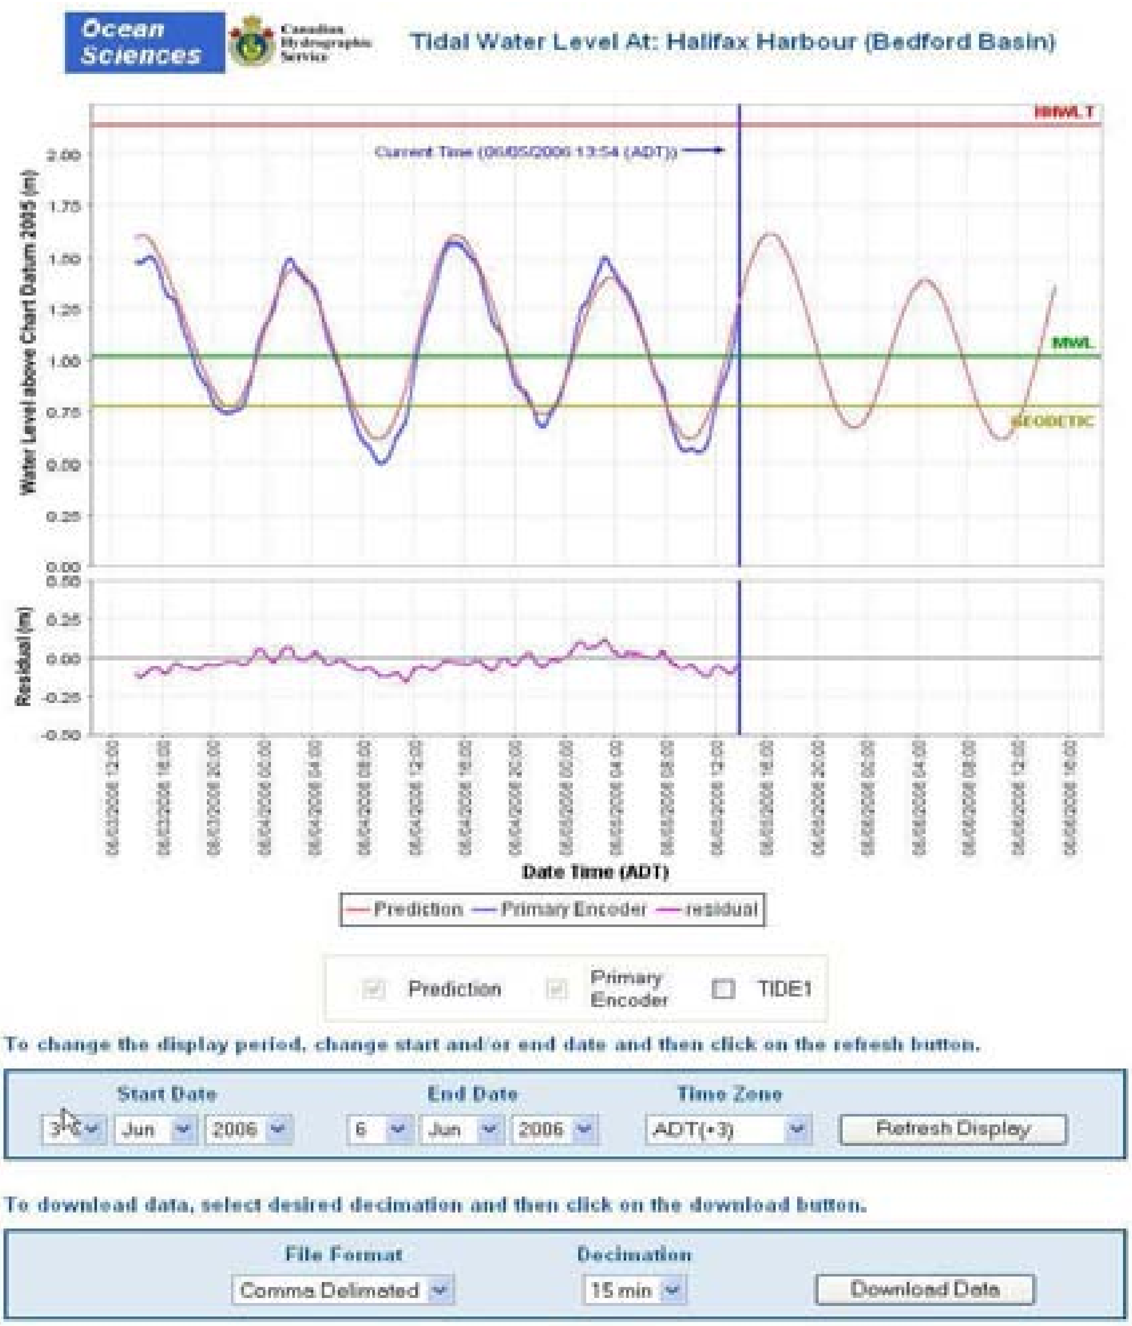

Figure 4: Real Time Water Levels EMO's web page

While real-time observations are useful in their own right, and can be used to mitigate some portion of the undesirable effects of residual water level behavior, the ability to forecast the residual's future response is even more desirable as it can provide advance warning of impending extreme water level events. With this in mind, in partnership with Keith Thompson at Dalhousie University's Oceanography department CHS Atlantic has developed a forecasting component for its real-time water level system.

Water Level Forecasting

A water level forecast differs from a constituent-based prediction because it includes account of the future behaviors of the non-tidal or residual water level components. Forecasting of residual water level behavior can be approached in two ways, either through physically-based dynamical modeling, i.e. using an operational hydrodynamic ocean model driven by forecast atmospheric forcing (Bobanovic and Thompson, 2001), or alternatively, if forecasts are only required for a few individual locations, through a neural-network type of approach (Huang et al, 2003). Each method has its proponents, and advantages and disadvantages. The most obvious being that the neural-network approach only provides estimates at specific locations, while the modeling approach provides forecast information over the whole of the model's geographic extent. For the purposes of CHS's RTWL system, forecasts at 16 tide gauge locations were required, thus both methods could be considered.

Although some studies suggest that with sufficient training the neural-network approach can be effective, its results are not strictly based on physical modeling and understanding of water level behavior (although one might argue that the choice of network inputs and network structure could be). Additionally, although a neural network forecasting system might be simpler to set up from scratch than a dynamical ocean model, the effort required is still significant and it requires specialized skill to choose, implement, and train the model. Without going into further detail, for both ideological and practical reasons, primarily because a hydrodynamic storm-surge model, Dalcoast II, was already running operationally at Dalhousie and ASPC, CHS Atlantic chose the dynamical modeling approach based on the Dalcoast II model.

Dalcoast II is a two-dimensional, barotropic, non-linear model with lateral diffusion with 1/12 of a degree resolution (Bobanovic and Thompson, 2001). It covers the whole of the Atlantic shelf from 38^{o} to 60^{o} N (Cape Cod to Cape Chidley at the northern tip of Labrador). It is driven by the three-hourly forecast winds and atmospheric pressure supplied by Environment Canada's Canadian Meteorological Center. It has been extensively validated, and on average should account for up to 90 % of the synoptic variability in coastal sea level.

At present, the Dalcoast II model is run once a day at Dalhousie. CHS Atlantic receives a daily feed of the model's hourly forecast water levels for the next 48 hrs at all CHS tide gauge locations. These times-series are the surge forecasts. Although Dalcoast II is a powerful tool it will only properly address the first source of residual water level behavior, the synoptic (large scale) variability generated by the weather. To forecast other sources of residual water level behavior requires additional measures. To capture the non-surge components of residual water level behavior a Kalman filter based methodology derived for use as part of the spectral-nudging technique developed by Keith Thompson at Dalhousie University (Thompson et al., 2006) was used.

The spectral-nudging technique employs a fairly simple, zero phase shift, multiple band pass, recursive filter, like a Kalman filter (Priestley, 1981), which is applied to the difference time series between say the observed temperature and salinity climatology and the dynamical ocean model's calculated values. The filter is used to calculate the ongoing nudging used by the model intended to restore the model's output towards a measured reference state, in this case the observed temperature and salinity climatology. The intent of nudging is to suppress model bias and drift. What is unique about the use of the filter to precondition the nudging is that it allows the user to only apply nudging, and thereby only suppress model bias and drift, over discrete, user defined frequency-wavenumber bands, while still allowing the model to evolve freely according to its specified dynamics at all other frequency-wavenumbers, thus the model generally remains eddy permitting.

Fortunately, for dynamical reasons, and as will shortly be demonstrated, the majority of the energy in the residual water level signal not already accounted for by the Dalcoast II storm surge model is also contained in a few select frequency-wavenumber bands. Thus, as will be shown, we can employ the simplicity of the spectral-nudging filter to develop forecasts for this remaining residual energy.

The Forecast method

With the Dalcoast II model output feed, and in the absence of all other residual water level variability, a proto-forecast of water levels for the next 48 hours at any one tide station could be defined as: the constituent-based water level prediction + the surge forecast at that station. This proto-forecast can be tracked as a time-series over the recent past, say 15 days, updating it daily, but only retaining the one-day-ahead (0 to 24 hr) forecast values following each Dalcoast II model update. The result is a sequential time-series of 0 to 24 hr proto-forecasts with a 24 to 48 hour proto-forecast tacked on the end. This time-series becomes the initial forecast model. One limitation of this time-series is that at times it exhibits small jumps at the end of each 24 hr period as better forecast winds and pressures become available for use in successive Dalcoast II model runs. If the Dalcoast II model was run more frequently, 2 or 4 times per day, these errors would be significantly smaller.

Subtraction of the forecast model time-series from the real-time observations (up to their last available value) becomes the forecast model's residual error. Thus, it is the forecast model's residual error that contains all other non-tidal water level behavior. The zero phase shift, multiple band pass, recursive filter developed for spectral nudging is then applied to this residual error time-series, and as will be outlined in the following section it can be used to generate both a time-series of expected residual error nowcast values for all times up to the last recorded real-time observation, and a time-series of residual error forecasts for up to 48 hrs into the future.

The final water level forecast will be calculated as the sum of the proto-forecast, which contains tidal and weather generated water level variability, and the residual error forecast, which should capture most other sources of residual water level behavior.

The zero phase shift, multiple band pass, recursive filter

Let $x_1, x_2, ...,x_t$ denote a sequence of equispaced scalar values to be filtered and let s_t denote an auxiliary state vector that evolves according to

$$s_t = (I - KH)Ms_{t-1} + Kx_t$$

where for with two user defined band pass frequency bands centered on ω_1 and ω_2

$$K = \kappa[1\ 2\ 0\ 2\ 0]^T$$

$$H = [1\ 1\ 0\ 1\ 0]$$

and

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \cos(\omega_1) & -\sin(\omega_1) & 0 & 0 \\ 0 & \sin(\omega_1) & \cos(\omega_1) & 0 & 0 \\ 0 & 0 & 0 & \cos(\omega_2) & -\sin(\omega_2) \\ 0 & 0 & 0 & \sin(\omega_2) & \cos(\omega_2) \end{bmatrix}$$

Using this system of equations, the filtered value of x_t, based on values up to and including time *t*, is given by $\langle x_t \rangle = Hs_t$. It can be shown that the width of the band pass of the filter, defined as the frequency range over which x_t is attenuated by less than one half, is approximately $2\sqrt{3}\kappa$ in radians per modeled time-step, and that the spin up time of the filter is of order κ^{-1} (Thompson et al., 2006). Thus, one may think of κ as the filter gain. Obviously there is a tradeoff as small κ results in narrowly defined prediction frequencies, but results in a long spin-up period. To summarize, the values of frequency and κ selected determine over what frequency bands energy is allowed to pass through the filter and accumulate in the filtered time-series, $\langle x_t \rangle$ and at what rate.

Application of κ through the *K* vector can also be made to be both time and/or frequency dependent. For example for the two frequency case,

$$K = \kappa_t [1 \quad 2f_1(\omega_1) \quad 0 \quad 2f_2(\omega_2) \quad 0]^T$$

where f_1 and f_2 are user defined frequency dependent functions. In the version of the filter used by the RTWL system only frequency dependence was implemented, permitting definition of individual band-pass widths and spin-up periods for each forecast frequency.

A key element of the filter is that in nowcast mode (when real observations are available) it is continuously updating the amplitudes and phases of its user selected frequencies through the auxiliary state vector s_t as new observations come in through the Kx_t term, and as old information decays through the $KHMs_{t-1}$ term. This will not be the case in forecast mode as forecasts will progress 48 hrs ahead of the last observations. We will assume that for forecasts the residual energy previously captured by the filter remains constant, with no decay timescale, thus for forecasts, $s_t = Ms_{t-1}$, and the amplitudes of the filter's selected frequencies remain fixed, although of course their phases proceed with the expected progressions. One might also apply decay terms through $KHMs_{t-1}$, however this was not done for reasons which will shortly be demonstrated.

The RTWL filter application

Recall that the forecast model's residual error contains all the water level residual behavior not previously captured by the Dalcoast II storm surge model. To use the spectral-nudging filter we need only define these frequencies and their frequency dependent filter gains, $\kappa f_i(\omega_i)$, for use in the filter's *K* vector.

Unfortunately, to some degree the optimal choice of frequencies to employ differs from tide station to tide station. Additionally, not all frequencies of all possible sources of residual energy can easily be identified strictly from theoretical arguments. A diurnal and semidiurnal frequency, those of O1 and M2, were chosen automatically from theory, and their gain factors were selected to provide sufficient filter band pass width to also collect the residual energy at other diurnal and semidiurnal constituent frequencies. The choice of the other higher tidal and seiche frequencies to employ in the filter, and their frequency dependent gains, was instead based on spectral analysis of data from each tidal station, although the higher tidal frequencies might just as easily have been chosen based on the station's observed tidal constituents.

Figure 5 presents an example analysis of the residual error for Port aux Basque, Newfoundland. The top right panel shows the full residual error time-series. The top left panel better illustrates its high frequency content. The middle and lower panels present both the cumulative and power spectra for the residual error time-series for frequencies above the semidiurnal frequency band. As expected, the majority of persistent residual energy is contained in several frequency bands in the higher tidal range, and further to the right at a final higher frequency peak corresponding to the Port Aux Basque harbour seiche frequencies. Based on the spectral analysis, four additional frequency bands were selected for the filter at Port Aux Basque, the first three centered on the cyan dots correspond to tidal frequencies, and the last at the magenta dot to seiche frequencies, giving a total of 6 selected filter frequencies.

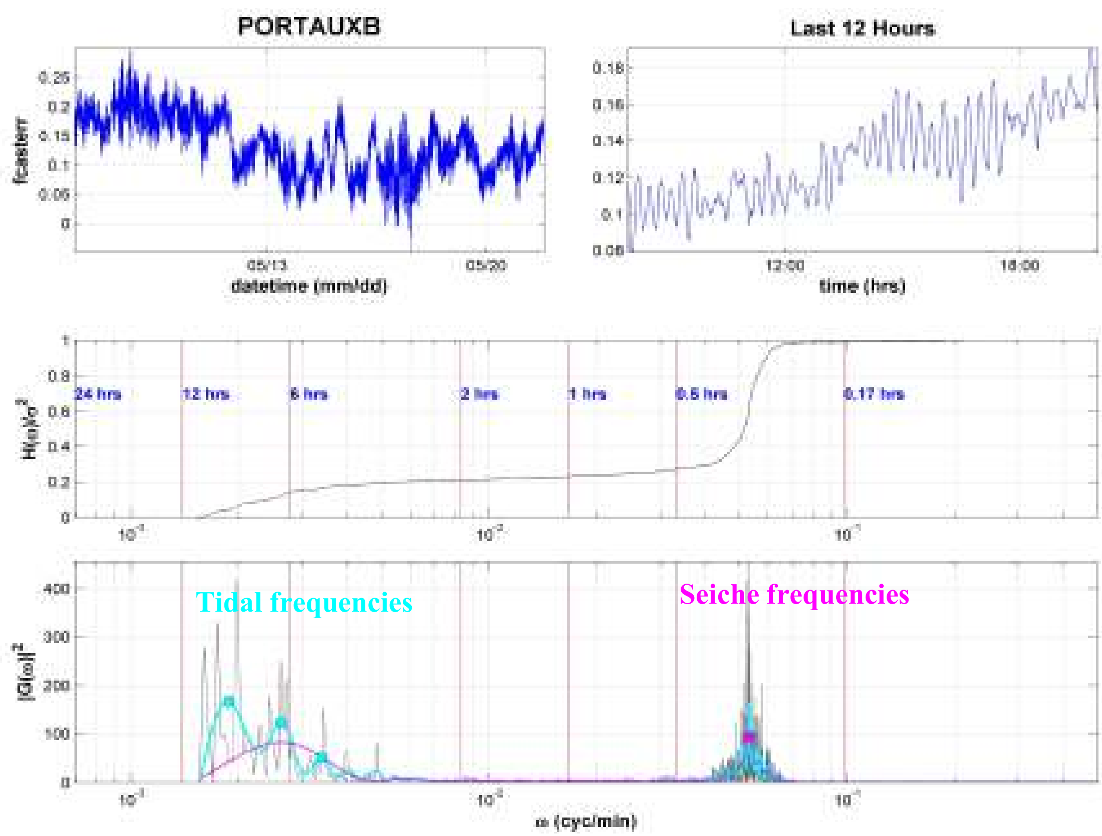

Figure 5: Filter frequency selection at Port Aux Basque, Newfoundland

It should be noted that while there will be residual water level energy at a wide range of high frequencies on an open coast, in Port Aux Basque harbour most of the time the majority of the observed energy will lie at the harbours natural or seiche frequencies. This is because the harbour acts to selectively amplify those frequencies. Thus, although the filter will only collect energy around our selected harbour seiche frequency, most of the time it will still capture the majority of the significant high frequency activity. Of course this will not be true all the time, i.e. specifically in the case of tsunamis and any other large amplitude, high frequency water level disturbances that may not display frequencies close to the seiche frequencies of the harbour in question, but neither the frequencies or occurrences of these events are strictly predictable in the fashion presently under discussion.

The output of the filter is a time series of residual error nowcast values up to the time of the last recorded observation, and a time series of residual error forecast values up to the time of the end of the last Dalcoast II 48 hr forecast. Before the filter is run, the mean is removed from the residual error time series and then later added back to the nowcast and forecast values. This accounts for any persistent bias in the Dalcoast II model not present in the observations as the model is based around in implied mean water level (mwl).

The nowcast and forecast residual error time series are added to the original initial forecast model, predictions + surge forecast, and the result is the stations total nowcast and forecast time series. Although the Dalcoast II model is only run once a day, the filtered residual error portion of the total forecast is recalculated every 2 hours to take advantage of new water level observations as they come in.

Figure 6 presents an example implementation of the full forecast system as it forecasts water levels through a small storm system for the Port Aux Basque tide station in Newfoundland. The top panel presents time series of water level observations, in blue, water level predictions based on constituent analysis, in red, three horizontal lines representing Datum information at the site, and the nowcast/forecast, in green.

The bottom panel shows the normal residual, observations-predictions in magenta, and the full forecast residual, observations-nowcast/forecast in green. The blue vertical line in both panels represents the time that the forecast was generated. Thus, to the left of the blue line are the nowcast values and to the right, the forecast values. Observations collected after the forecast was generated are presented in light blue, and were used to calculate the residuals to the right of the vertical blue line in the bottom panel.

In Figure 6 the data to the left of the vertical blue line illustrate that the combination of the 0 to 24 hr Dalcoast II model forecast, combined with the output of the nowcast portion of the spectral nudging filter, accounts for most of the observed residual energy, indicating that our choices of frequencies for the filter were appropriate. To the left one may also observe the action of the recursive filter as it progresses through time. Note that as the observed high frequency seiche activity waxes and wanes the nowcast values follow, but with a small time lag that represents the frequency dependent spin up time of the filter, in this case set to be equivalent to about 1 to 1.5 wave periods of the frequency in question.

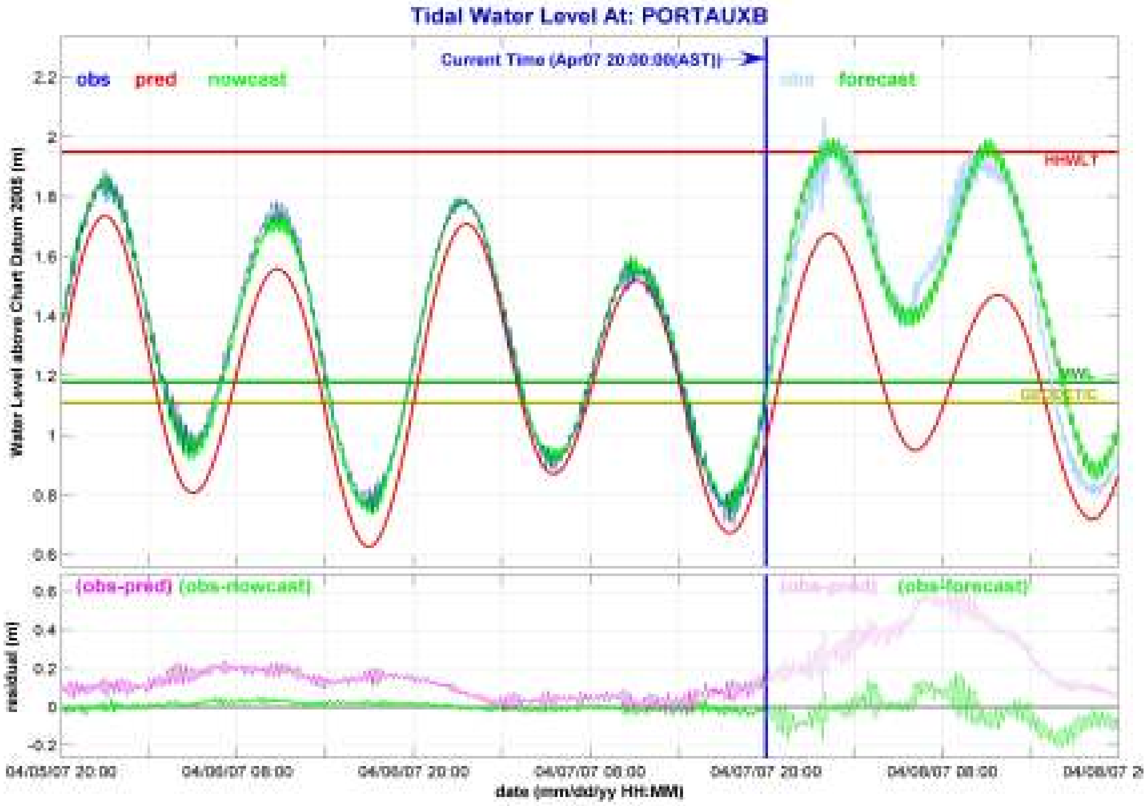

Figure 6: Example water level nowcast and forecast at Port Aux Basque, Newfoundland

The data to the right of the blue line demonstrate that the forecast method as a whole can reasonable be expected to capture the majority of the residual water level activity over the next 24 hours. It also shows that once the observations stop, the amplitudes of the residual error frequencies forecast by the filter remain constant. This is most clearly evident near the first peak in the observed and forecast water level responses where an obvious jump in seiche activity is apparent in the observations, but not in the forecast. The forecast could not capture this detail because the filter was unaware of this future change in seiche activity, it can only estimate future residual responses relative to what has been observed up to the time the forecast was made.

This also illustrates the rational for not including the decay term in the filter's forecast calculations. If the decay term had been present, the seiche energy in the forecast would have decayed to the right of the blue line. Without the decay term, seiche activity in the forecast is maintained and although over a couple of seiche periods the forecast seiche looses its phase relative to the observations, nonetheless in the forecast it still represents an estimate of the water level variability at seiche frequency one might expect to encounter. The same would hold true for the other residual energy at lower frequencies forecast by the filter. However, the lower the frequency the less one might expect the decay term to influence the forecast over the 24 hour forecast period.

The operational forecast portion of CHS Atlantic's real-time water level system is presently undergoing its final testing and will soon be available on the real-time water level web pages.

It was earlier suggested that a primary purpose of continued monitoring of coastal water levels was to observe and record the water level residual. However, the other significant reason for maintaining continued coastal water level monitoring programs is to document the slow changes in water levels arising from relative sea level rise.

From the hydrographic perspective the most significant result of relative sea level rise is the affect it has on the tidal Datums used as references for all hydrographic data collection and charting activities. Tracking these slow changes in water levels through time and space requires a certain minimum number of either continually active monitor sites, or reoccupation, for a year or more, of a minimum number of sites on a regular basis. The argument can be made that because there is a need to monitor residual water levels over the short term, and because reoccupation of mothballed sites for a year or more carries an inherent extra expense over and above the yearly upkeep of a continually active site, it is most appropriate to maintain the required minimum number of active gauge sites for both purposes.

The purpose of the following section is to: i) focus on the primary hydrographic need for continued coastal water level monitoring programs, that of Datum maintenance and development, and ii) to generate discussion between CHS regions and with our American counterparts on expedient methods for development of Canadian 3D tidal Datum surfaces, delivery systems and future Datum related products.

Looking to the Future, Anytime, Anywhere water levels on Datum

Tidal Datums, for example chart Datum and the Datum of elevations, used in Canadian coastal waters are still for the most part treated as discrete point entities, those in the Saint Lawrence Seaway are presently a notable exception. For Hydrographic survey purposes these discrete Datums have often been spatially generalized to govern wide areas based on the assumption that differences between Datums, their transforms, remain constant over said area. In other words, that for example the differences between geodetic, mean water level and chart Datum are spatially constant in the horizontal. Using this assumption it has been the practice to reduce Hydrographic survey data to the discrete Datums associated with one or more tide gauges.

Although this approach has served Hydrography well in the past for near-shore, large-scale survey activities, it is less appropriate for survey projects covering large spatial areas where many tide gauges would be required to achieve a desirable Datum resolution (like route surveys), or where at smaller spatial scale the Datum's target, for example Lowest Astronomical Tide (LAT) changes fairly rapidly in space as in some estuaries and for example in the Bay of Fundy. It also can be problematic when applying high water Datums, for example Highest Astronomical Tide (HAT) to flood mapping of large geographical areas. Finally, it ignores the true spatially and temporally varying nature of the Datum's 'real' target surfaces, the tidal water level thresholds that the Datums are meant to approximate.

It is significant that Hydrographic vessels now commonly employ sophisticated methods of dynamic, motion compensated, vertical positioning using Real Time Kinematic (RTK) GPS, and GPS based modeling solutions like Omnistar and Canada Wide DGPS Correction Service (CDGPS). Using these systems, a vessel's typical Multi-beam system's vertical error budget has shrunk to the point where errors imposed on the final reduced depths by the spatially constant Datum transform ideology become of comparable size, again particularly in areas where tidal amplitudes, and thus the tidal Datum's target surface, change rapidly in space. Finally, anywhere, anytime, water level prediction tools, like WebTide (Dupont et al., 2002), cry out to be referenced to sensible, continuous Datum surfaces.

Awareness and acceptance of the limitations of the spatially constant Datum transform ideology is illustrated by the 3D datum surfaces being generated in the United States, National Oceanographic and Atmospheric Administration (NOAA), National Ocean Service (NOS) ongoing VDATUM projects (Myers, 2005). The approaches already in use by VDATUM for developing and managing 3D Datums serve as concrete examples for how CHS Atlantic might also proceed in the production and management of its own 3D vertical datums.

Vertical Datums

The CHS presently uses several vertical Datums, the most important being Chart Datum and the Datum of Elevations. The target surfaces, the theoretical tidal water level thresholds, for these Datums presently adopted by the CHS are Lower Low Water Large Tide (LLWLT) and Higher High Water Large Tide (HHWLT) respectively. However, Canada intends to migrate its target surfaces for these Datums to Lowest Astronomical Tide (LAT) and Highest Astronomical Tide (HAT) to conform with recent International Hydrographic Organization (IHO) resolutions. In either case, both the low and high water theoretical target surfaces for the Datums are spatially continuous functions and, as previously discussed, for an increasing number of applications it would be appropriate and convenient to establish and use continuous Datum surfaces to better approximate the behavior of the underlying target surfaces.

The problem becomes how to best develop these continuous Datum surfaces such that they most closely approximate the flavor of their underlying target surfaces while still as much as possible locally honoring the existing discrete Datum holdings.

Development of 3D Gridded Datum surfaces

For semi-bounded regions with numerous well established discrete, shore-based Datum reference points, like the Saint Lawrence River, development of continuous Datum surfaces and transforms might reasonably be accomplished by establishing a dense Datum grid, bounded by the existing shore-based reference Datum points, using appropriate contouring, smoothing and gridding methods. To achieve anywhere Datum evaluation relative to this grid one would then later employ appropriate interpolation methods to estimate the Datum surface between given grid points. This is in essence the approach taken by several National Ocean Service VDATUM projects, and by CHS Quebec in the Saint Lawrence River Estuary.

An important point to make is that no matter what the initial steps used to develop the Datum, or more properly the Datum transform grid, once developed the techniques employed for its use, for example the interpolation methods used to estimate the Datum surface between given grid points must be standardized. Additionally, once a grid is developed it should be epoch named, just as the discrete reference Datums should be named.

Even in the simplest of cases there would be challenges to overcome. The process of developing the Datum surfaces is likely to point out inconsistencies and inadequacies in existing Datum reference points and one will be forced to decide whether: to honor the existing, albeit flawed, discrete Datum data, to adjust the existing problem discrete Datums, and/or to what degree one may ignore problem areas in the Datum surface construction. Thus, the process of constructing the Datum will likely force some re-evaluation of the existing discrete Datums or at the least will point to problem areas requiring further attention.

To reasonably develop these Datum surfaces for unbounded areas like open coastlines, covering wide geographical extents, with sparse or potentially poor quality pre-existing discrete Datum data, and to then extend these Datum surfaces out from the coast into deep waters will require a more sophisticated approach than applying basic contouring or gridding based on the existing discrete Datum data. What is required is a theoretical, reasonable estimate of where the Datum's target surface should lie, for example: between established discrete Datum control points, into poorly constrained areas like large gulfs, and offshore onto continental shelves and out into deep water. The required Datum surface estimates can be developed using ocean modeling of varying degrees of complexity. The NOS has taken the same approach in their VDATUM projects.

A co-tidal chart represents a simple model that when combined, or locked down, with existing shore based reference Datums could be used to develop such a theoretical Datum target surface. For the Atlantic Region the tidal prediction software WebTide, whose underlying data sets were developed using a Barotropic hydrodynamic model assimilating Topex-Poseidon altimeter data, is an obvious source from which one might start to generate the Datum target surface estimates. Other similar model solutions might also be employed, for example in the Pacific Mike Foreman's modeling efforts (Foreman et al., 1995) have be employed for this purpose. These models can be employed to supply the necessary target surface estimates for seamless Datums over large geographical extents like the Gulf of Saint Lawrence and out across the continental shelf.

On smaller spatial scales local modeling efforts could be used on a case by case approach to provide target surface estimates where existing Datum information is inadequate, as in some estuaries, or additional observational work could be undertaken to provide the necessary discrete Datum control. These are also approaches adopted by NOS in their VDATUM projects.

An Outline for an Atlantic Approach

Up to this point the details of how to accomplish development of continuous or 3D Gridded Datum surfaces has been left purposely vague. Partly this is because there are options in the approaches one might take. Secondly, the best approach to adopt is likely to become part of the process itself, or in other words, doing the work is likely to suggest the more effective and expedient methods. Nonetheless, for the Atlantic region a basic framework might be suggested.

The challenges for the Atlantic region are that we have an extensive coastline. Our discrete Datum control points are of varying quality, numerous having been established many years ago based on short tidal records. Our tidal systems are fairly complex and in some regions tidal amplitudes and tidal character change rapidly in space. We have the highest tides in the world and because of the nature of our geography in places we experience appreciable technical difficulties measuring them accurately.

Given these limitations, development of continuous Datum surfaces must go hand-in-hand with a reassessment of the existing discrete Datum reference points. Thus, an iterative approach seems appropriate with several passes for both 3D surface and discrete Datum reassessment. A possible scenario is outlined below.

1. Develop an initial theoretical target surface shape based on model data (i.e. for LAT some percentage of the sum, say 95-105%, of the WebTide constituent amplitudes). Figure 7 presents a first look and shows an LAT target surface relative to Mean Water Level (MWL) according to estimation using WebTide. Included in the Figure (the black dots)

are a number of discrete shore based control points for comparison. WebTide is based around a floating MWL, thus a surface computed from it must then be referenced to observed mean water levels. This will raise issues as in some locations mean water level will be better known than others. Also we do not have discrete MWL control away from the coast and thus there we must rely on either Topex-Poseidon data or Geodetic models.

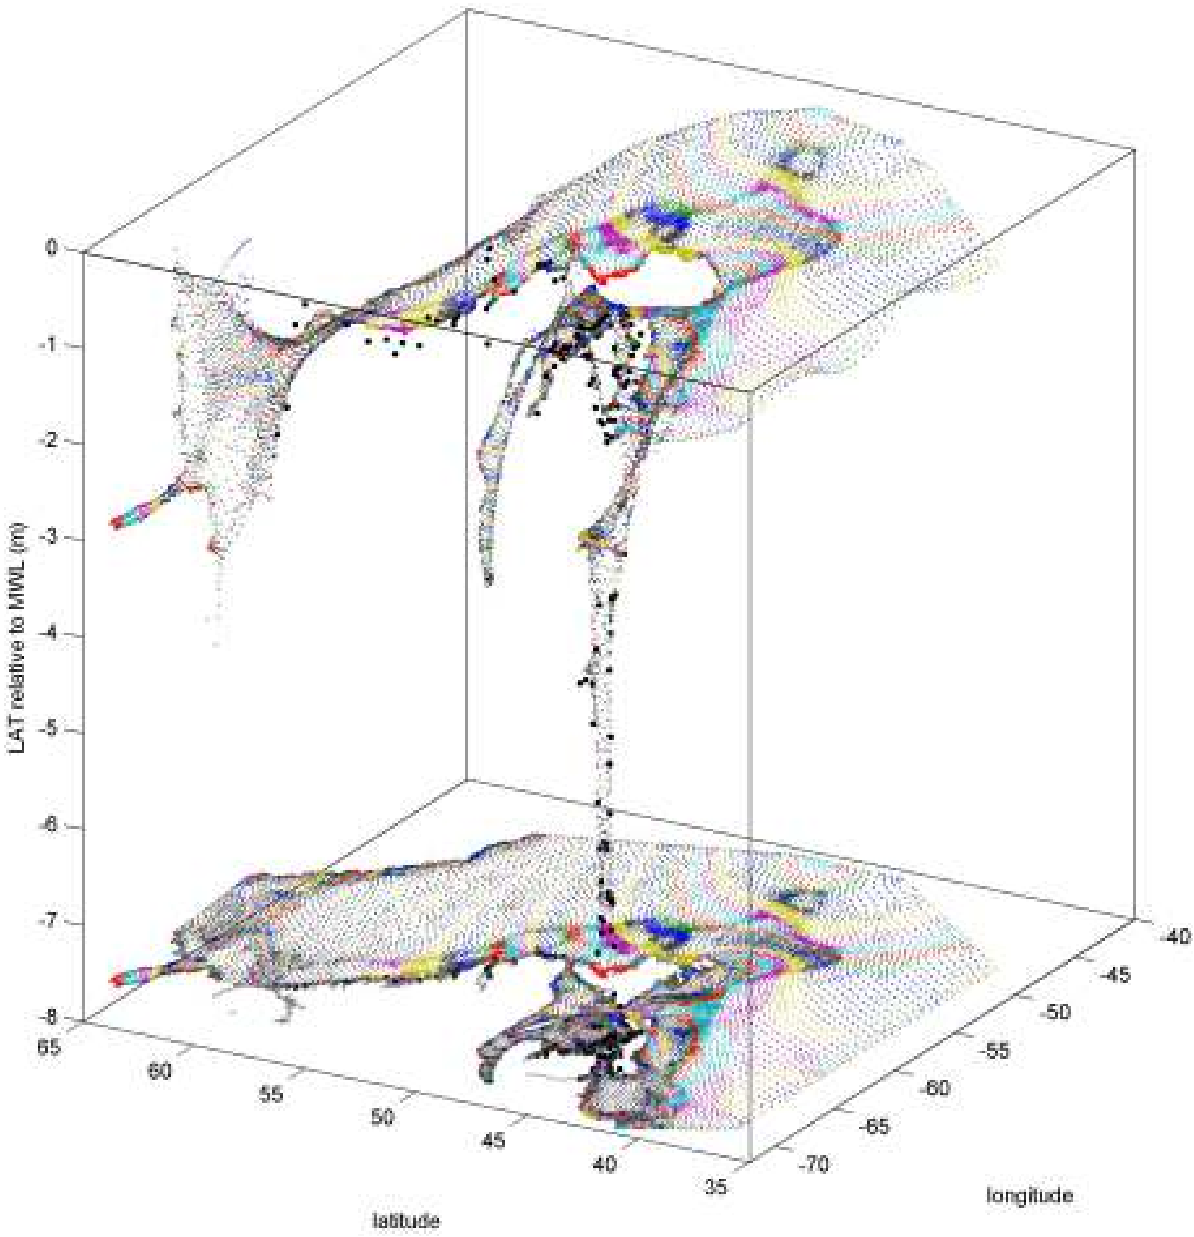

Figure 7: A preliminary LAT target surface with accompanying shore control points

2. Compare existing shore-based Datum reference points to the theoretical target surface. There will be misfit, due to both poor reference data, and theoretical surface inadequacies (offsets). For example, Figure 7 shows that in some regions agreement is good while in others there is significant disagreement between the target surface and the discrete shore

based control. Again, an iterative procedure will be required to find and identify poor reference control, and to iteratively adjust the theoretical target surface to more closely honor high quality shore-based reference control. Iterations between steps 1 and 2 may be necessary.

3. Adjust or supersede the target Datum surface in areas where good surrounding shore-based reference Datums permit independent surface estimation. Sequester inshore areas that will require further data or modeling.
4. Calculate the final seamless Datum surface grid. Final grid spacing might be variable based on the local rate of change of the Datum surface.
5. Agree on standardized methods for interpolation between datum surface grid points and thus calculation of all Datum transformation between Datum surfaces.

Initial definition of 3D Datum surfaces could be accomplished for the Atlantic region as a whole or in a piecemeal fashion, which might then necessitate a final blending or matching of the Datum surfaces between sub-girds. Blending will become necessary in any case when combining the gridded surfaces of one region with those of another, for example between the Atlantic region and Quebec region's existing 3D Datum now extending into the Gulf of Saint Lawrence. It would be highly advantageous to engage in collaborative discussion with the NOS VDatum project administration to ease the transitions between existing or developing 3D NOS Datum surfaces and those contemplated for construction in Canada.

The final combination of the 3D Datum surfaces with WebTide will yield an anytime, anywhere, on Datum water level prediction tool.

Conclusions

This paper has summarized the state of CHS Atlantic's tidal water level network and its developing value-added products. Specific attention was given to its new water level forecasting methods, reflecting the fact that increased access to operational storm surge model forecast data makes water level forecasting an increasingly attractive water level system option. Its final section outlined CHS Atlantic's preliminary plans to develop 3D tidal datum surfaces. In its entirety this document loosely outlines the continued need, from both a hydrographic and socioeconomic perspective, for maintaining a healthy Canadian coastal water level monitoring program.

Acknowledgements

We would like to acknowledge the efforts of Department of Fisheries and Oceans, Science, Coast Guard, and CHS employees for their past and ongoing roles in the maintenance and development of Atlantic PWLN and RTWL systems. Specifically, CHS tidal officers, Chris Collen, and Fred Carmichael, Supervisor, Data Management, Steve Parsons, Coast Guard, Carmen Reid, and Science Informatics Analyst, Kohila Thana without whose continued efforts the Atlantic PWLN and RTWL systems would be severely compromised. Finally we acknowledge Director General, Savithri, Narayanan and the CHS Directorate for their vision and continued support of Canadian coastal water level monitoring programs.

References

Bobanovic, J., and Thompson, K.R., 2001. The influence of Local and Remote Winds on the Sea Level and Circulation of the Gulf of Saint Lawrence. *Continental Shelf Research,* **21***,* pp. 129-144.

Dupont, F., C.G. Hannah, D.A. Greenberg, J.Y. Cherniawsky and C.E. Naimie. 2002. Modelling system for tides for the North-west Atlantic coastal ocean. [online]. [Accessed 21 April, 2008]. Available from World Wide Web: http://www.dfo-mpo.gc.ca/Library/265855.pdf

European Sea Level Service. 2006 [online]. [Accessed 14th April, 2008]. Available from World Wide Web: http://www.eseas.org/products/?page=real_time_data

Forman, M.G.G., Walters, R.A., Henry, R.F., Keller, C.P., and Dolling, A.G. 1995. A tidal model for eastern Juan de Fuca Strait and the southern Strait of Georgia. *Journal of Geophysical Research-Oceans*, **C1(100)**, pp. 721-740.

Huang, W., Murray, C., Kraus, N. and Rosati, J., 2003. Development of a regional neural network for coastal water level predictions. *Ocean Engineering*, 30, pp. 2275-2295.

Monitoring Network System for Systematic Sea Level Measurement in the Mediterranean and Black Sea. 2008 [online]. [Accessed 14th April, 2008]. Available from World Wide Web: http://medgloss.ocean.org.il/

Myers, E. P., 2005. Review of Progress on VDatum, a Vertical Datum Transformation Tool [online]. [Accessed 16th April, 2008]. Available from the World Wide Web: http://nauticalcharts.noaa.gov/csdl/Vdatum_pubs/myersOceans05.pdf

Physical Oceanographic Real-Time System (PORTS). 2007 [online]. [Accessed 14th April, 2008]. Available from World Wide Web: http://tidesandcurrents.noaa.gov/ports.html

Priestley, M. B., 1981. *Spectral Analysis and Time Series*, London, Academic Press

Proudman Oceanographic Laboratory, Sea level networks. 2008 [online]. [Accessed 14th April, 2008]. Available from World Wide Web: http://www.pol.ac.uk/ntslf/networks.html

Thompson, K., Wright, D., Youyu, L., Demirov, E., 2006. A simple method for reducing seasonal bias and drift in eddy resolving ocean models. *Ocean Modeling* **13**, pp. 109-125.

Tides and Currents. 2007 [online]. [Accessed 14th April, 2008]. Available from World Wide Web: http://tidesandcurrents.noaa.gov/

Author Biographies

Phillip MacAulay, macaulayp@mar.dfo-mpo.gc.ca

Canadian Hydrographic Service Atlantic
1 Challenger Dr.
Dartmouth, Nova Scotia
B2Y-4A2
(902)-426-5017

Dr. MacAulay graduated in 2004 with a PhD in Physical Oceanography from Dalhousie University and has been with CHS Atlantic Tidal Section for 2 years. He is presently engaged in redevelopment of CHS Atlantic's water level capabilities and in development of its vertical datum surfaces.

Charles T. O'Reilly, oreillyc@mar.dfo-mpo.gc.ca

Charles O'Reilly graduated in 1975 with a BSc/Honors from Dalhousie University. He worked three years in exploration geophysics for Texaco before joining the Canadian Hydrographic Service/Atlantic. He is presently Chief / Tidal Analysis and Prediction; and has been active researching rising sea levels, storm surges, tsunamis and 3-D vertical datum transforms.

Keith R. Thompson, keith.thompson@dal.ca

Dr. Thompson holds a joint appointment in Oceanography and Statistics in the Faculty of Science of Dalhousie University. His research interests revolve around understanding the dynamics that control the changing physical state of continental shelf seas and the open ocean, and also developing realistic models that can generate useful predictions.