



**National
Oceanography Centre**
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Oceanography Centre

Research & Consultancy Report No. 64

Development of an Automatic Tide Gauge Processing System

Joanne Williams, Andrew Matthews and Svetlana Jevrejeva

2019

National Oceanography Centre
Joseph Proudman Building
Brownlow St Liverpool
L3 5DA

Email: joll@noc.ac.uk

Abstract:

There are many tide gauges around the world for which research quality data is not available for sea-level studies, including gauges that are maintained primarily for tsunami monitoring. In some cases high-frequency data is available for download through the Intergovernmental Oceanographic Commission (IOC) <http://www.iocsealevelmonitoring.org>, but manual QC has been too labour intensive.

In this report we describe Matlab software for automatic quality control (QC) for tide gauges, including generalised comparison of instrument channels, fitting and predicting tides using irregular high-frequency data. We include details of the implementation, testing, and known capability and weaknesses.

3rd April 2019.

Contents

1. Project information.....	6
2. Background.....	6
3. Existing automatic quality control systems.....	6
4. Outline of Matlab automated system.....	7
A. Download of data.....	8
B. Metadata.....	9
C. Primary Quality Control.....	9
D. Channel Comparison.....	12
E. Tidal analysis.....	15
F. QC on non-tidal residual.....	16
G. Single preferred channel.....	16
H. Final tidal analysis.....	16
I. Averaging/Filters to generate hourly/dailies/monthly NTR:.....	16
J. Daily means.....	17
K. Monthly means for PSMSL.....	17
L. Storm surge analysis.....	17
5. Further testing of the automatic quality procedures.....	18
A. Observation of Hurricanes in the record.....	18
B. Retention of tsunamis in the record.....	19
C. Comparison with manual QC.....	21
D. Comparison with PSMSL data.....	22
E. Test cases used for automatic data QC software.....	22
Caribbean.....	22
Senegal, West Africa.....	23
Tanzania.....	24
Pakistan.....	24
Chile.....	24
Indonesia.....	24

UK	24
Denmark	25
6. Technical requirements of the code.....	25
7. Tide gauge data sources.....	25
8. Acknowledgements.....	26

1. Project information

This work is part of the Commonwealth Marine Economies project “Climate Change Impact Assessment: Ocean Modelling and Monitoring for the Caribbean CME states”, work-package 2.8 Development of a coastal data hub for stakeholder access in the Caribbean region.

CME Programme Objective: To develop regional capacity-building for Caribbean SME states in coastal oceanography and climate change resilience.

Deliverable: Scoping and design of automated tide gauge data processing software

2. Background.

Tide gauge installations in the Caribbean region are generating significant amounts of sea level data, which would benefit local communities, businesses, coastal engineers, decision makers if they were processed and easily available. Currently, most of the data are available only in real time for tsunami warning at the IOC website (ioc-sealevelmonitoring.org) and not used for other purposes (e.g. to make climatological data sets, or model validation). This is mainly due to difficulties in processing the data.

A modern tide gauge provides data typically sampled at an interval of several minutes. Many have multiple sensors to provide calibration and redundancy in a range of conditions, for example pressure and radar sensors. Sometimes the timings of tide gauge sensor will require correction because of clock drift in the data logger or occasionally, a record will have gaps due to gauge malfunction or telemetry drop outs, spikes due to false readings, or datum shifts due to a change in the effective reference zero of the tide gauge.

Manual processing of TG data is time-consuming and somewhat subjective, as quality control is based on visual scanning of the data, possibly after an initial tidal prediction fit. Some gauges are managed to include complete processing of the data, for example by UHSLC (University of Hawaii Sea Level Center), BODC (British Oceanographic Data Centre) or SHOM (Service hydrographique et océanographique de la Marine), but this does not necessarily mean that the original-frequency data is available for public download – for example the global repository at UHSLC only includes hourly data. Others gauges, such as Calliaqua on St Vincent are uploaded automatically online at the IOC site but are not included in the UHSLC repository. We also know of gauges, for example those at Georgetown and Grenville in Grenada, which have been installed for local projects or port operations and have not been made openly available in any form. In this last case the gauges have no automatic upload to IOC, instead the data is manually collected every 2 weeks by a Ports Authority representative, and as yet no post-processing is taking place.

Here we describe the automatic quality control processing of the TG data using `tidegaugeQC.m`. This is a new Matlab code developed for the CME project “Climate Change Impact Assessment: Ocean Modelling and Monitoring for the Caribbean CME states”.

3. Existing automatic quality control systems

We take as our starting point the recommendations for quality control for sea-level data from chapter 6 of the EuroGOOS (European Global Ocean Observing System) manual¹.

SELENE 1.0 is an automated tide gauge processing package in Python 3, used by Puertos del Estado, the Spanish Port Authority. We have been able to make use of some elements of this system, particularly the early version of the spike detection algorithm which was kindly provided by Begoña Perez Gomez. However we require a more flexible system that is suitable for a wider range of tide gauge data than is encountered in Spain. We are also able to take advantage of a better tidal processing algorithm than is used in SELENE, the Matlab code NOctidefit.m, which enables us to fit and remove the tide at an earlier stage in the process. This avoids interpolation of missing data which is otherwise necessary to provide a regular timeseries.

4. Outline of Matlab automated system

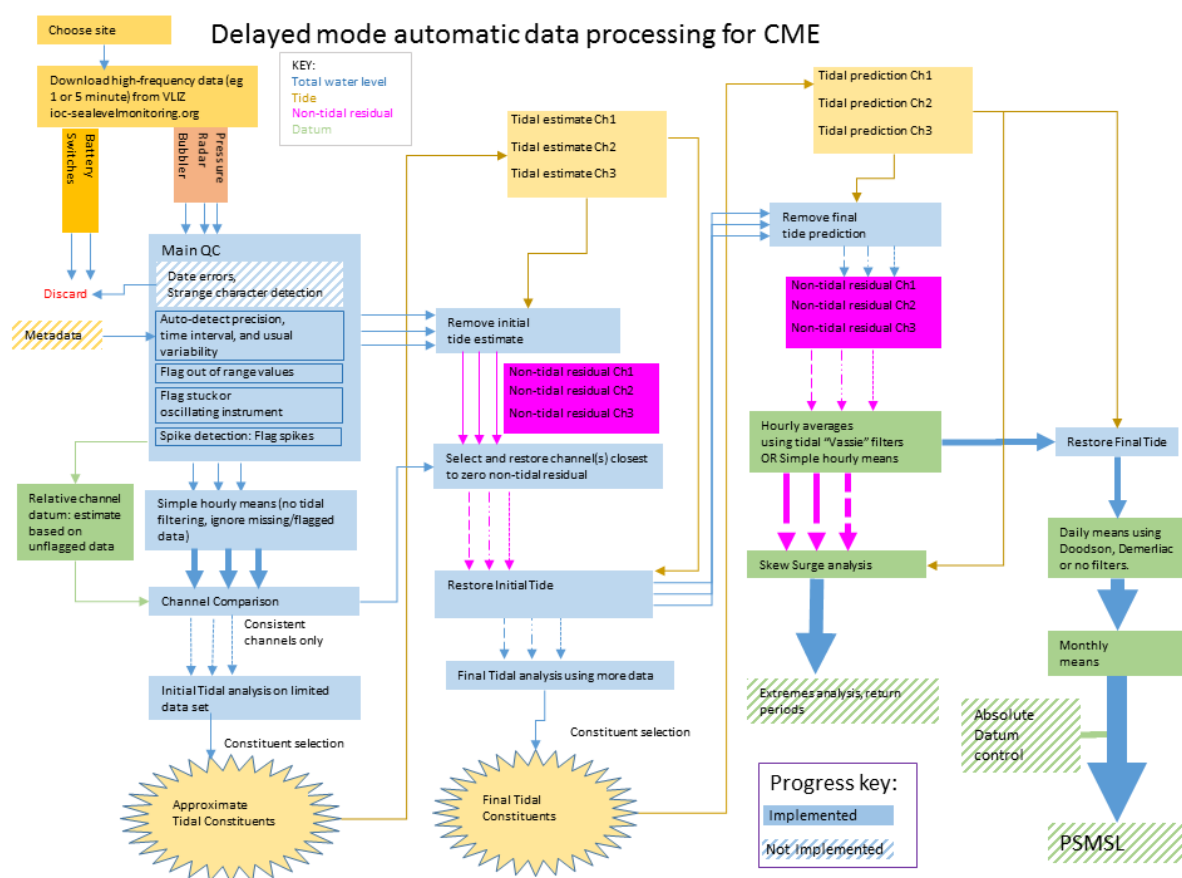


Figure 1: Outline of the TG data-processing code implemented in Matlab. See also QCdataflow.pptx

For convenience, we have initially implemented and tested our new system in Matlab (tidegaugeQC.m, available at <https://psmsl.org/cme/index.php>). However we are aware that Matlab licence costs may prove

¹ "Recommendations for in-situ data Near Real Time Quality Control", Sylvie Pouliquen and the DATA-MEQ working group, EG10.19 December 2010, available at http://eurogoos.eu/download/reference_documents_/rtqc.pdf

a barrier to access for some users. Furthermore there may be a value in providing QC capability as part of an online data portal, for which purposes the code will need to be in a form that can run on a web server, perhaps in Python. If these developments follow they will be linked from the same site.

A. Download of data

Data is available from several sources as described in section 10. For the purposes of the test, data is assumed to come from VLIZ, at ioc-sealevelmonitoring.org (Figure 2) .

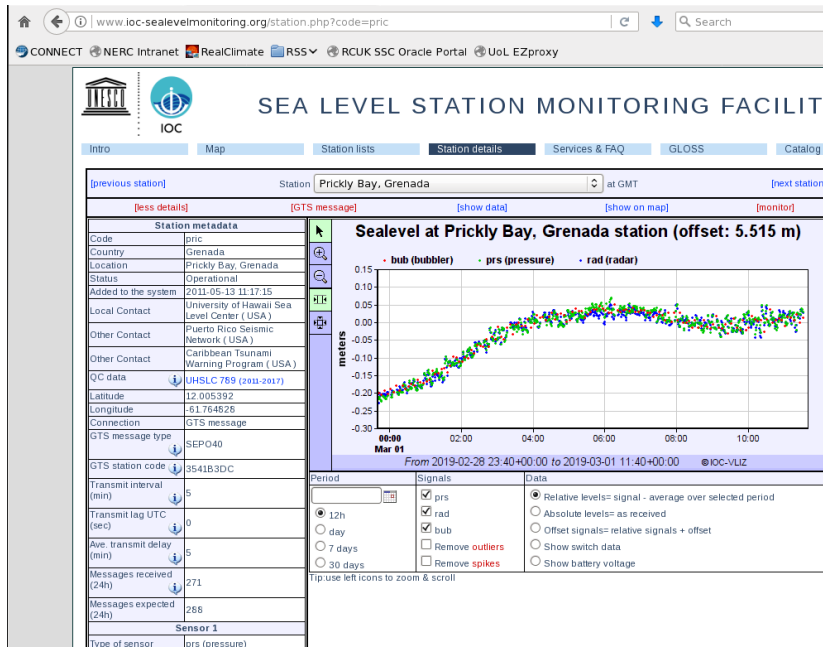


Figure 2: Tide gauge data is available from ioc-sealevelmonitoring.org , operated by VLIZ. This is served in real-time from the gauges and requires quality control.

Data can be downloaded automatically from this site. The website only allows 30 days download in one command, so repeated calls are necessary. This can be automated in two ways:

- using *wget* (Linux commandline). This has been implemented to run from within Matlab as *readIOCxtab.m* . The *wget* method has the problem of leaving repeated headers in the file. These can be removed using Linux commandline *grep* and *cat* .
- or *webread* (entirely within Matlab). This has been implemented to run from within Matlab as *readIOC.m* . It writes separate files for each sensor at the station.

Variable numbers of sensors can be handled, although channel comparison is limited to 3 sensors. Some sites have Battery and/or switch data which is automatically discarded by looking for channels labelled 'bat', 'sw1', 'sw2'. Typical data intervals may be eg 5 minutes for bubbler or pressure sensor, 1 minute for radar. There have been no tests for higher frequency than 1 minute.

To avoid excessive download from IOC during testing, sensor streams are saved locally and can be reloaded.

Within the Matlab code dates are handled as pre-R2014 "datenum" .

The longest test so far is on Prickly Bay, Grenada, 2011-2018, with 3 channels.

B. Metadata.

For the purposes of testing, we assume that the only data available from the gauge is the date/time and water levels, as is available directly from the IOC site. *The datum for the gauge, and also relative positions of instruments, may be obtainable from other sources and should be used if possible.*

C. Primary Quality Control

There are some initial tests that are recommended in the Gloss manual but have not yet been included in this code. They have not proved necessary in the case of data downloaded from the IOC site. In particular:

Strange characters detection (in which case the record is discarded). *Not implemented.*

Checking for timing errors, and that dates are provided in UTC. Not implemented. This can be a problem, especially where the local data provider uploads data in local winter/summertime, and not in UTC. For example, at Hirtshals, Denmark (Figure 3). The records are sent to the IOC site in local time, including Daylight Saving Time when appropriate, but the IOC site assumes that they are in UTC. Hence there is a data gap between 01:50 and 03:00 on 28 March 2016. At the corresponding 30 October there must be ambiguous time stamped data.

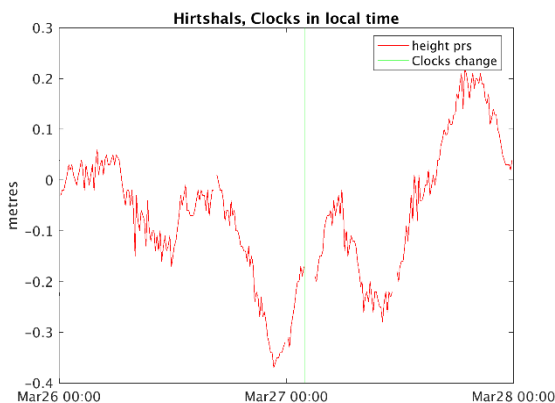


Figure 3: Example of timing error arising from use of local time (including Daylight Saving) instead of UTC.

Flagging of out-of-range values (based on extremes included in the metadata for each station). This has been included in principle, by setting a value of $1e9$ for extremes, but no such metadata is available with the IOC data. Great care must be taken not to exclude genuine surges and other extreme events.

Estimate precision of data (if not provided in metadata). Method: look for the minimum difference between neighbouring values, after allowance for floating point errors.

Estimate the usual sampling interval of data (if not provided in metadata). Method: take the median difference between dates. This will disregard sampling irregularity.

Estimate the typical standard deviation of the data (if not provided in metadata). So long as the channel has more good data than bad, this will help to qualify a likely spike. Method: take the median of a standard deviation windowed over 60 points.

Stuck instrument test. Also known as “stability test” in GLOSS manual. This test is to flag values when there is no change in the water level after a number of time steps. The number of data values or time steps to begin to flag depends obviously on the time interval. A typical value, for example, could be 3 for 5 minute data.

Defaults implemented: If there is no change across 5 or more values, don't flag the first of the stuck ones (unless flagged already) but do flag the rest. This should enable the last recorded point to remain in the record in the event of instrument failure. Tolerance is set to STUCKTOL=0.0001 [ie 0.1mm]. *Defaults numbers of timesteps and tolerance need to be specified for other time intervals and data precision.*

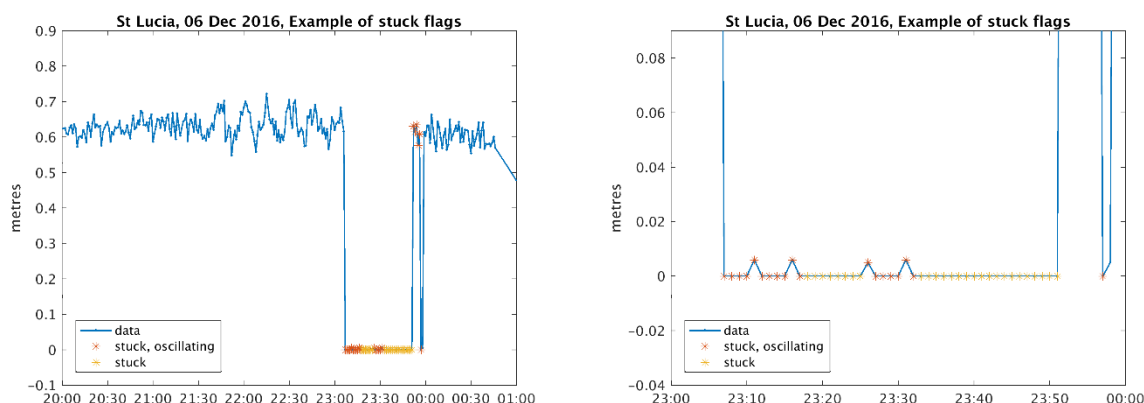


Figure 4: The stuck instrument test may miss faults where the instrument oscillates. The stuck-oscillating test finds these, and flags intermediate points as also suspect.

Stuck oscillating instrument test. Channel pr1 at Ganter’s Bay, St Lucia (Figure 4) has a record that begins with a testing phase. Whilst this example could be completely removed from the record with a start date specified in the metadata, similar errors can occur at other times. It manifests as an irregular oscillation between 2-3 stuck values. Flagging is done based on a repeated value (within a tolerance of 0.1 x estimated precision) occurring at least 20 times within 60 data points, making all 60 suspect. Data gaps are ignored here, as identical values before and after a big data gap are also suspect. The selection of 20 out of 60 is somewhat arbitrary. Very high frequency data at low precision is likely to contain such occasional patterns simply by random chance, requiring a large number of tests, but we want to identify short bursts of measurement error. *Further testing may be required, and dependence on sample interval included.*

Flags from this test are combined with flags from the stuck instrument test.

Spike detection. A spike detection algorithm has been implemented. It was originally based on that used in SELENE and provided by Puertos del Estado, but it has been extensively modified. The key differences are that it allows for irregular data and is more robust at identifying clusters of spikes due to use of median rather than mean.

The spike detection algorithm is an implementation of deviation from a spline (a form of Savitzky-Golay filter). It is a significant improvement on the “spike detection” in TASK, which only considers a point in relation to its immediate neighbours and requires that the jump is manually set.

The polynomial (default quadratic) fit is done using centering and scaling, which allows direct use of large datenumbers and missing data. The polynomial fitting uses an efficient least squares fit based on the Matlab polyfit function. In another change from the SELENE code, the point being tested is omitted from the fit. Figure 5a shows a successful spike detection. In this implementation as in SELENE, the polynomial fit is not robust against outliers (other than the test point), so problems may arise identifying spikes in the neighbourhood of other spikes or datum shifts (Figure 5b, c). However robust polynomial fitting would be significantly slower to implement and run.

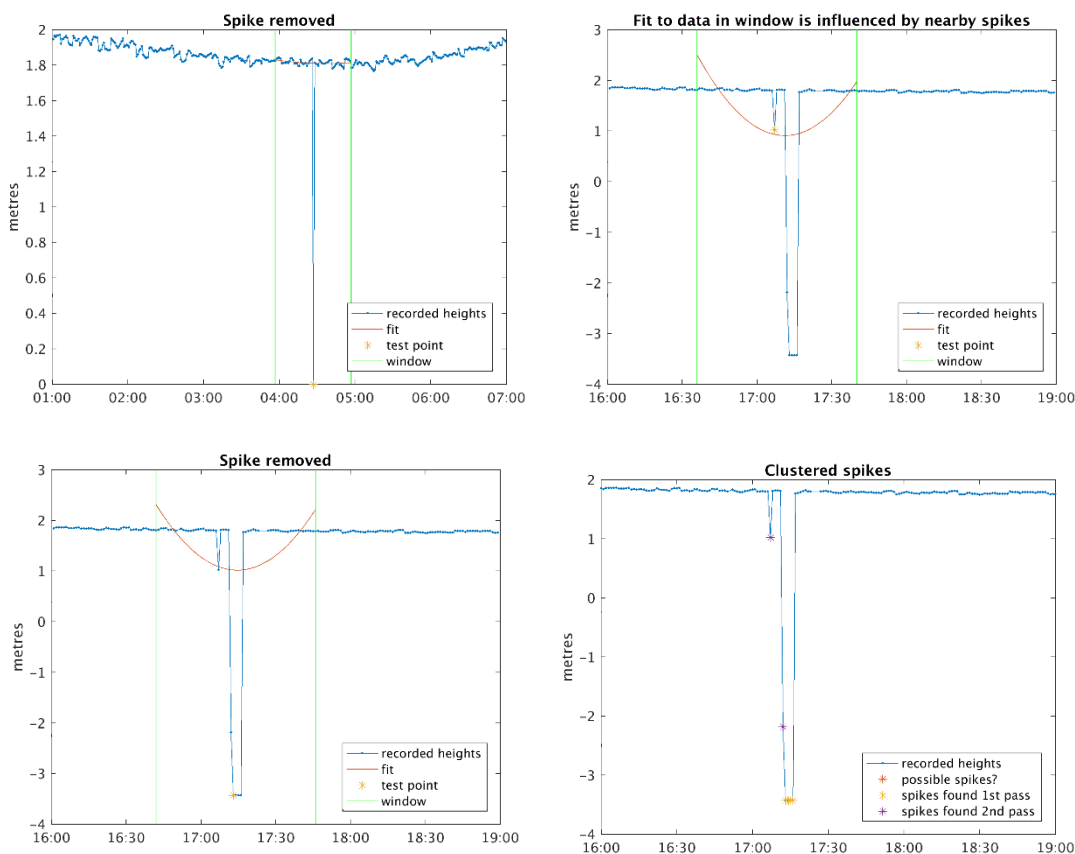


Figure 5: Examples from Prickly Bay, Grenada, May 2013. a) Spike correctly found, b) Nearby spikes prevent robust polynomial fit c) Despite poor fit, spike still eliminated d) After a 2nd pass, all spikes in the cluster are correctly identified.

The deviation from the spline is measured using median absolute deviation (MAD), with a point declared different if it is more than 6 scaled MAD from the spline. The scaled MAD is defined as $c \times \text{median}(\text{abs}(\text{A} - \text{median}(\text{A})))$, where $c=1.4826$, and is derived from $c = -1/(\sqrt{2} \text{erfc}^{-1}(3/2))$.

An option is retained to use RMSE to measure the deviation from the spline (for comparison with the SELENE code).

The default choice of MAD is to improve robustness in identifying spikes in the region of other outliers. Testing using the Prickly Bay data showed that with a window of at least 60 points, and using MAD, 90% of spikes were identified after 1 pass and all spikes were correctly identified after 2 passes of the algorithm, whereas with RMSE half were missed after 2 passes (Eg Figure 5d). The threshold of 6 MAD was chosen to minimise false positives. Most false positives occurred during quiet periods, when a small random deviation from almost flat data appeared out of the ordinary. To eliminate this, a minimum spike size was set at 3 times the typical standard deviation of the complete series.

Spline testing parameter defaults: 2 degree polynomial fitted through 60 points (1 hour at 1 minute-data). Deviation of more than 6 scaled MAD from the polynomial is flagged.

Spike detection is the slowest part of the QC process (typically 3 minutes per channel on the server for 1 year's data at 1 minute frequency). Flag results are therefore saved for reuse during testing.

Note that some spikes could be genuine records for high-frequency waves or tsunamis. Such events should still be discarded for tidal analysis, but in general we attempt to preserve them in the record.

Jump detection Not implemented. The "Standard Normal Homogeneity Test", may be appropriate. Matlab also has a "findchangepts" function which correctly identifies datum shifts, although if no shift exists it finds at least one false positive. It may be tuneable to avoid this. Further testing is required here. If implemented, jump detection may be best used in splitting channels where likely problems arise and treating them as separate.

D. Channel Comparison

Where there are multiple channels available on the TG, we have a powerful data set for cross-comparison. Different instruments may fail or drift at different times, and by comparing neighbouring channels we may be able to automatically judge which instrument is misreading.

The overall method is this:

- Identify the periods for which the multiple channels agree
- Use only those periods for initial tidal analysis
- Assume that during periods they don't agree, the smaller non-tidal residual is more likely to be correct.

At Prickly Bay for example (Figure 6) there is a drift in the pressure sensor that is obvious by eye from around August 2018.

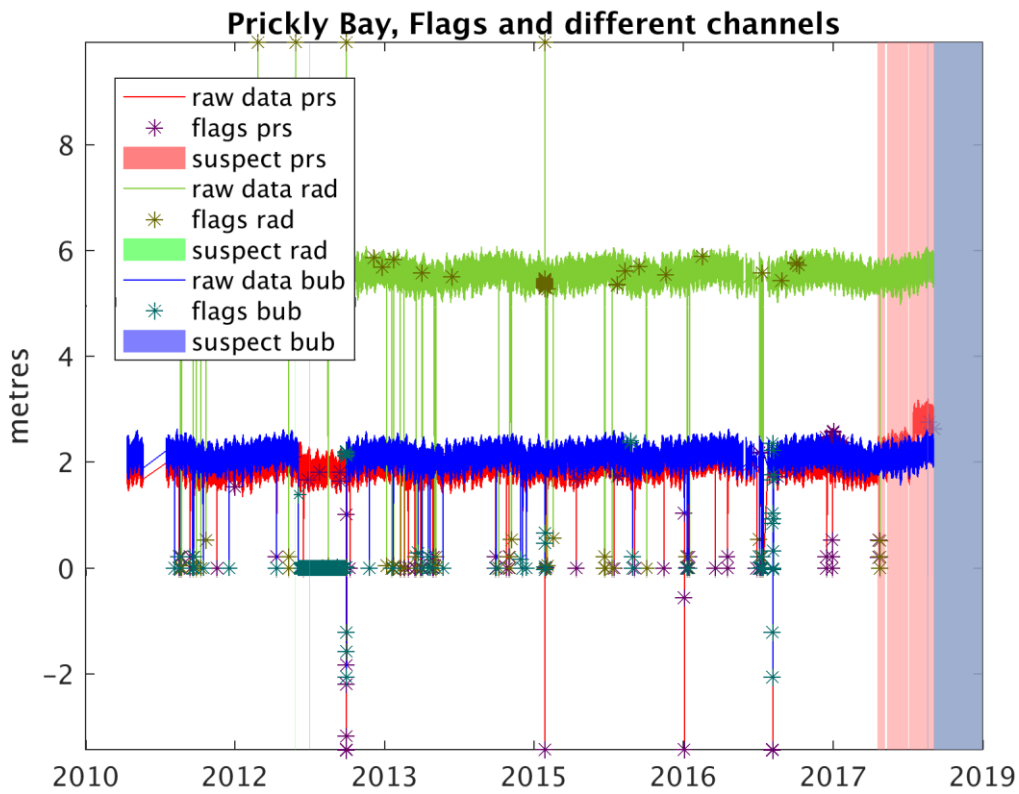


Figure 6: Channel comparison shows a shift in the pressure sensor, clear from around August 2018.

Using the non-tidal residual it is clear that this begins earlier, about the beginning of April (Figure 7). Since the radar and bubble give consistent results over this period, it is very likely that there is a fault on the pressure sensor, and this channel is rejected for this period.

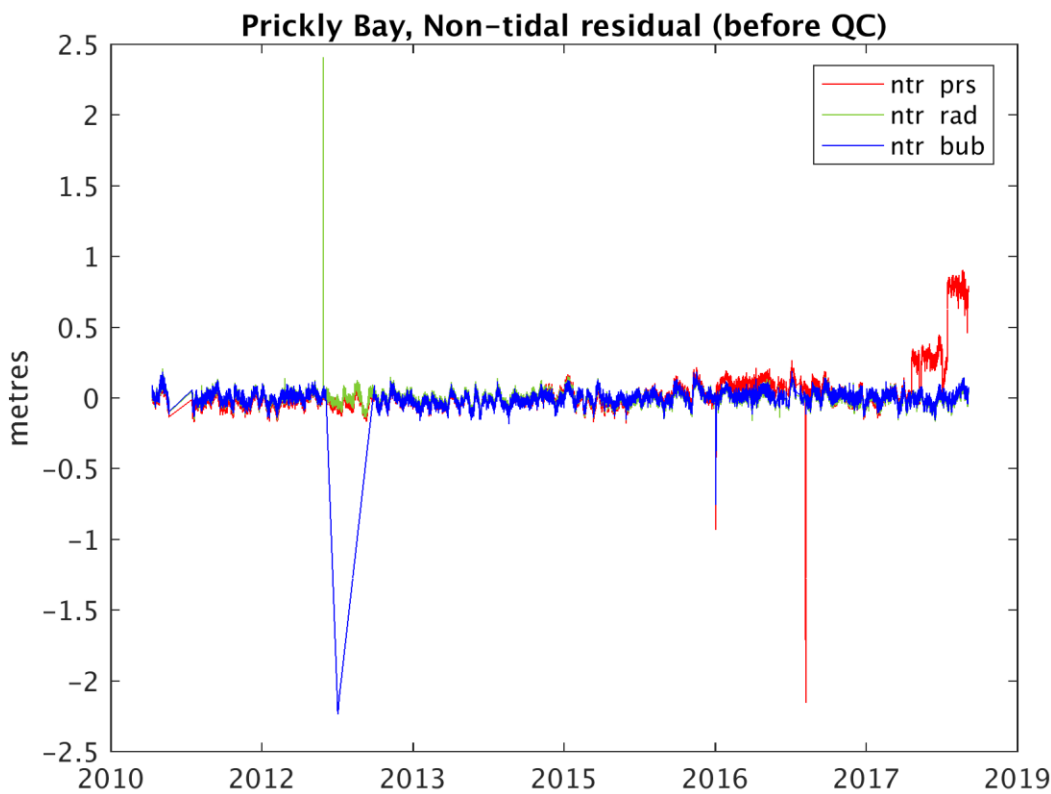


Figure 7: Using the non-tidal residual channel comparison shows up another shift in the pressure sensor around April 2018.

Implementation of generalised multiple channel comparison is quite difficult, as it requires extremely careful handling of cases, especially for the correct handling of flagged data. *The case of more than 3 instruments has not been implemented. We have not yet encountered any gauges for which this is a problem.*

In order to compare channels they must be at the same time interval. Different instrument types may legitimately give different readings (eg high-frequency radar gauges may include wave data that is averaged out by a pressure gauge in a stilling well reporting at 15 minute intervals). We should ignore data that has already been flagged. Therefore the first step is **simple hourly averages ignoring missing data**. (Filters to produce hourly data whilst preserving energy at tidal frequencies (eg the Vassie/Pugh filters) are not necessary at this stage and can be counter-productive, as they tend to propagate data gaps and hence discard useful data. A better filter which permits some gaps may be beneficial.)

Constant offsets between channels are treated as similar. These may be due to datum controls on the instrument(s) that need handling separately. Care needs to be taken to ensure that exactly the same periods are compared, especially if one instrument comes online after another. Where there are 3 channels, any two in agreement are marked as similar. Periods of agreement between channels in hourly data are then applied to the high frequency data.

Long term trends may require special handling to match separate channels. Again, it would be best to specify datum offsets between channels in advance.

Detail of channel comparison

During channel comparison, all channels are temporarily adjusted by a relative datum correction. If not specified, this is estimated as the median of all unflagged data.

1 channel : No comparison possible. Channel labelled as **similar**.

2 channels : Find the dates for which the channels both have unflagged data. Find the absolute difference between the channels for these dates. Find the median of those differences. Label 2 channels as **dissimilar** for a given date if they are more than 4 x median difference apart. Also label them as dissimilar if they are beyond a hardcoded difference. (Defaults to `FlagChannelDiffGreaterThan = 0.1m`). For each channel in turn, assume data between nearby dissimilar points is also dissimilar. (Fill gaps less than 1 day with flags).

Problems arise if one channel is a half-tide sensor – for example at Mtwara, Tanzania, there is a pressure sensor which is dry at low tide. In this case, we would like the channel comparison to discard that channel as dissimilar for a few hours twice a day. This case is not yet handled.

3 channels : Each pair of channels is compared as for 2 channels, with careful handling of the intersecting dates. A channel is marked “similar” for a period if it is similar to either of the others. It is possible for all three to be marked “dissimilar” if they all disagree.

For testing purposes, timeseries plots are produced with shaded sections to indicate disagreement (eg Figure 6).

The information about similar/dissimilar channels is applied to the high-frequency data. If only one channel operates for a period, it is labelled as “similar”.

E. Tidal analysis

Initial tidal analysis is carried out only on the channels in good agreement with each other, with no flagged data or any large extremes. Similar procedures are used again later.

We have `NOctidefit` and `NOctidepred` in Matlab. These were written by Simon Williams, tested by Jo Williams, originally based on David Cartwright's Fortran, via Chris Hughes amongst others! This code is based on the same mathematics as TIRA, which was implemented in the GUI TASK Analyse, but has been independently optimised due to improved memory capability of modern computing.

`NOctidefit` and `NOctidepred` use least-square fitting and work with irregular or gappy data. This is an advantage over the technique in the GLOSS manual, which uses Fourier analysis and so requires interpolation of the data to regular time intervals before tidal analysis.

`NOctidefit` also applies daily nodal corrections, so does not require data to span a particular length.

In `NOctidefit` the harmonic constituents can be automatically selected based on the total data length. However using too many constituents on very gappy data will result in very poor fit, so a wrapper has been added to select constituents based on the total contiguous data length.

A short wrapper is also provided for NOctidepred to convert the date format.

Figure = example of the wrapper and constituent selection

F. QC on non-tidal residual

We linearly remove the tidal estimate to give a high frequency non-tidal residual (NTR). This may still contain tidal frequencies, due to inaccurate tidal prediction and also to short term phase changes due to tide-surge interaction. However most tidal energy should be removed.

For the periods where channels are labelled as dissimilar, the non-tidal residual is used to select a preferred channel, for periods of 10 days at a time. **The assumption is that a channel that is in error is likely to have a larger NTR (averaged over 10 days) than a channel reading correctly.** If there are 3 channels that are all dissimilar, the one with the smallest NTR is preferred.

Further QC could (*and perhaps should*) be carried out at this point – particularly looking for large datum steps or drifts on all channels simultaneously that is difficult before tidal analysis. However a drift on all channels could be considered a genuine signal, and perhaps should be output independently for comparison with vertical land movement rather than labelled “wrong”.

G. Single preferred channel

Where there is relative datum control between the instruments, it is possible to switch between channels, selecting a preferred channel for any given date to provide a continuous timeseries. **Datum correction is essential for this to be done accurately.** If there is no datum available, an estimate is made base on the median of the consistent data in each channel, but the resulting timeseries should be treated with extreme caution.

This is referred to in the code as “stitching” a single channel.

Since there may be more than one channel labelled as “good” data for a given time period (if two or more channels are in reasonable agreement) some criteria is necessary to select one of them. To do this the channels are ranked according to the number of “good” hourly data records in each. (*Allowance should also be made for this to be user-specified if the operator knows one instrument is more reliable*). The channel ranked 1 is then selected wherever possible. If gaps occur in the good data for this channel, then fill with the channel ranked 2 etc.

H. Final tidal analysis

Having selected the best channels, the initial tidal prediction is restored and a new tidal analysis is carried out. The selection of harmonic constituents is allowed to change from the initial set.

This then provides a tidal prediction and linear “non-tidal residuals” (NTR) for each channel.

I. Averaging/Filters to generate hourly/dailies/monthly NTR:

To obtain hourly data we average the non-tidal residual, and then add the hourly tidal predictions. *Strictly, we should add hourly average tide instead of hourly spot values, but the resulting periodic difference is less than 1% of the tidal range, and this method provides for easier handling of irregular dates.*

For the hourly averaging there is optional filtering with the “UsePughFilters” option.

With filtering: The NTR may still contain energy at tidal frequencies, particularly in the case of shallow-water storm surges temporarily affecting tidal phase. Therefore in filtering to **hourly data** an option is provided to use the Vassie & Pugh filters, from 5/6/7.5/10/15/30 minutes to hourly. These are implemented in Matlab as running filters. For higher frequency than 5 minutes there is not a specific tidal filter. Short gaps are interpolated, and high frequency NTR averaged to 5 minutes. (2 or 3 minute data is averaged to 6 minutes).

Without filtering: However a disadvantage of the Pugh/Vassie filters is that they use wide windows and so substantially widen any data gaps. Since we've implemented the tidal processing before moving to hourly data then an option is provided (default) to simply use hourly averages. Data flagged as spike or stuck is not included, and if more than half of the data from an hour is missing it is set to NaN.

The simple hourlymeans is designed to cope with irregular and missing data. However because no account is taken of the position of the data within the hour's window, the phase of the data can end up changing by as much as half the sample interval. Suppose for example data was provided at 09:30, 10:00 and 10:30. The data points for 09:30 and 10:00 are averaged and recorded at 10:00, effectively introducing a 15 minute phase shift. To address this would require windowing beyond the hourly interval. *It may be that a compromise filter can be created, with the benefits of windowing and yet with an allowance for missing data.*

J. Daily means

Daily means are calculated on the heights including tides, rather than NTR. This avoids any confusion about which of the longer tidal frequencies (eg monthly, annual) have been removed. To convert hourly data to **daily** data, it is standard to use a 19-hour Doodson filter, again to suppress the main tidal frequencies. An alternative, the 72 hour Demerliac filter is also included. Both filters have the disadvantage of propagating a one-hour data gap to more than one day – more so for the longer filter. For intermittently gappy data, it might be better to omit these filters and apply simple averages. An option has been implemented to allow a daily mean with no filter, either allowing hourly gaps to propagate or ignoring gaps. *A further option should set a minimum number of hours for a daily mean. Extreme care is required in the case of a half-tide sensor.*

K. Monthly means for PSMSL

To obtain **monthly** data daily means are averaged over calendar months. A minimum of 15 complete days is required for a month, otherwise it is replaced with NaN.

For the data to be entered into PSMSL for long-term sea-level studies, **datum control** is required. This needs further information than is available on the IOC site.

L. Storm surge analysis

Ideally datum control should be applied, as if there is a long term shift in the TG (for example very localised movement of a pier) then surges will apparently change during time. However surges can be found without this. Code to extract skew surges (the difference between peak total water level and nearest predicted high tide) is implemented, using the hourly data.

5. Further testing of the automatic quality procedures.

A. Observation of Hurricanes in the record.

At Calliaqua some effect of Hurricanes Irma and Maria can be seen. (Figure 8a, b) There is little visible in the total water level, but after removal of the tide there is a visible signal in the radar gauge This is most likely the one-minute average of the waves.

At Ganter's Bay, St Lucia (Figure 8c,d) and further north at St Martin (Figure 8e,f), both events were more severe. At St Martin the record was lost, but not before recording non-tidal residuals of over 0.6m. Since

these are one-minute averages they reflect much larger higher frequency waves.

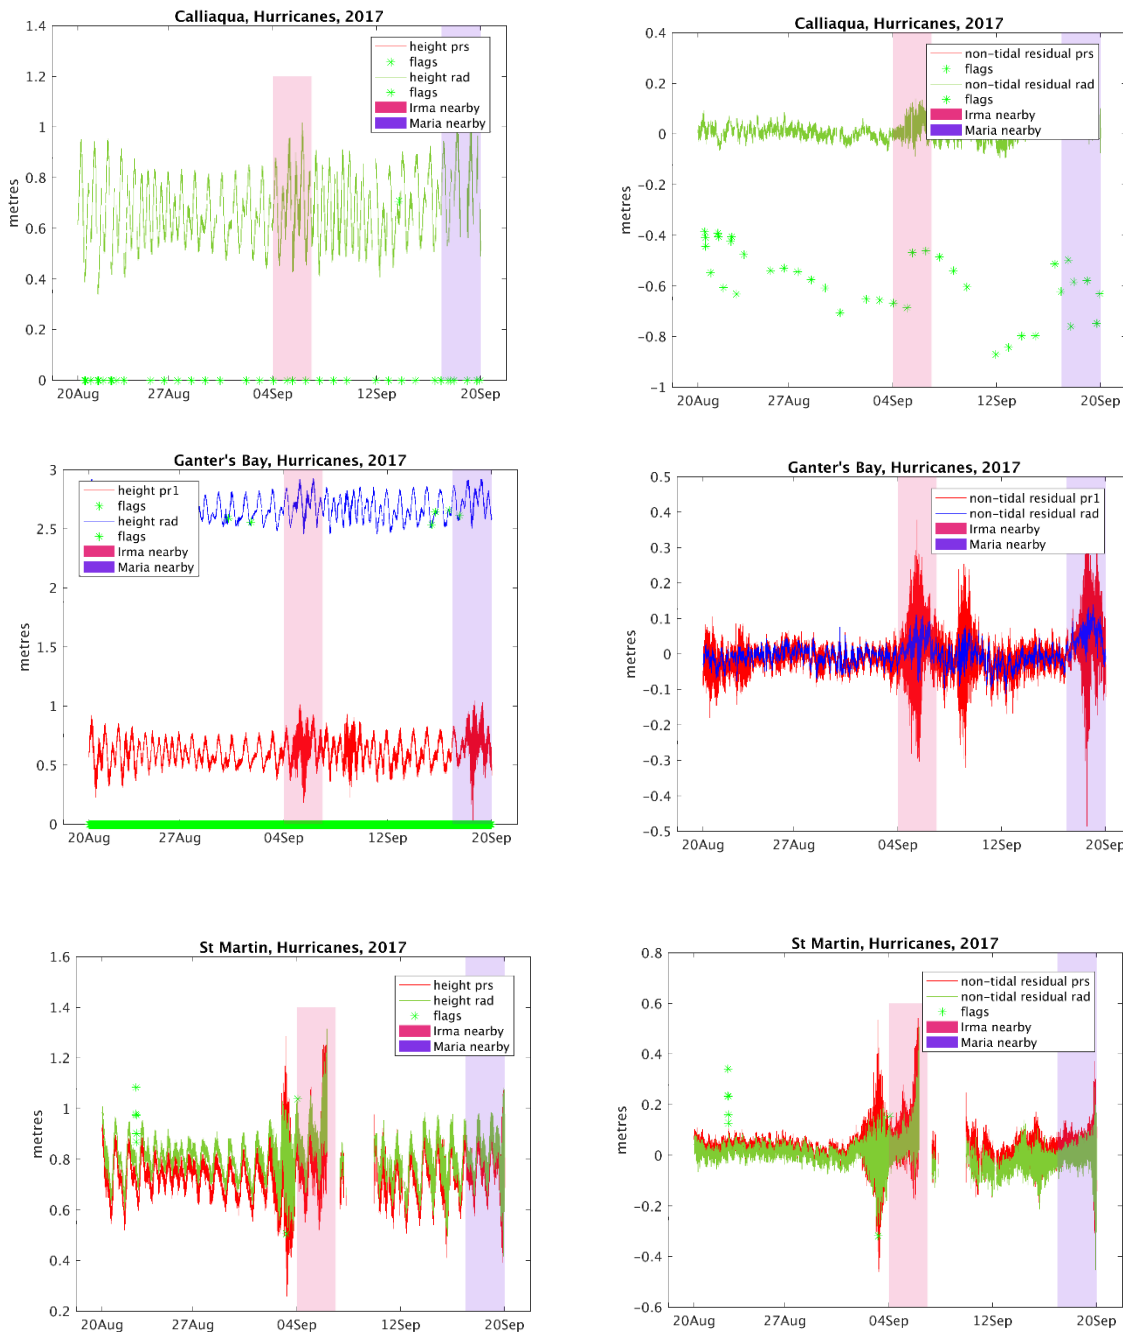


Figure 8: Hurricanes Irma and Maria in 2017, recorded at Calliaqua (St Vincent), Ganter's Bay (St Lucia) and St Martin..

B. Retention of tsunamis in the record.

We do not want the spike detection algorithm to remove tsunamis. We test it using records from Chile from the earthquake of 27th February 2010.

At Corral (Figure 9) there was a low frequency signal in the pressure sensor during the first half of 2010 with a drift of around 4m. Despite this, the algorithm still succeeds in removing tidal frequencies and the tsunami itself is well preserved.

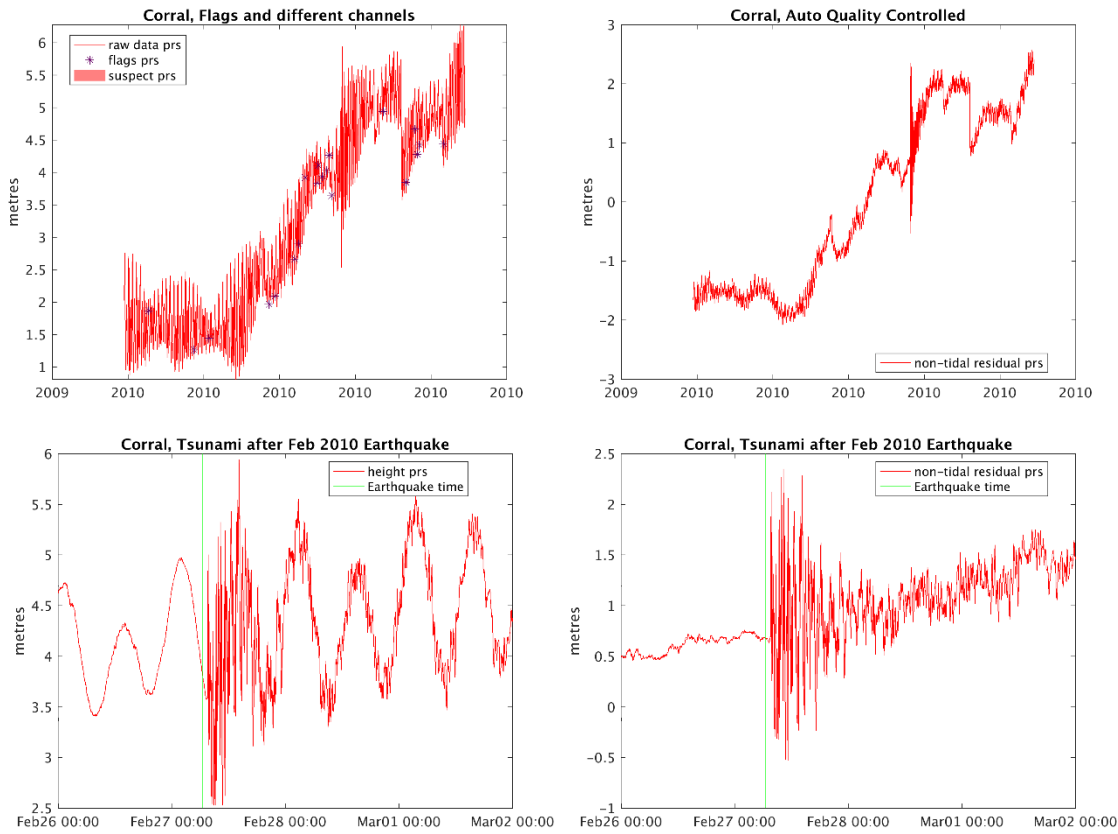


Figure 9: Corral, Chile. Despite the low-frequency drift the tsunami signal is very clear.




Figure 10: Caldera, Chile, showing tsunami after 2010 earthquake.

At Caldera (Figure 10), the tsunami is also well preserved. At Valparaiso however (Figure 11) the first 6 records of the tsunami - not the largest - were incorrectly discarded as spikes. There is another anomaly on the 18th February that was not flagged.

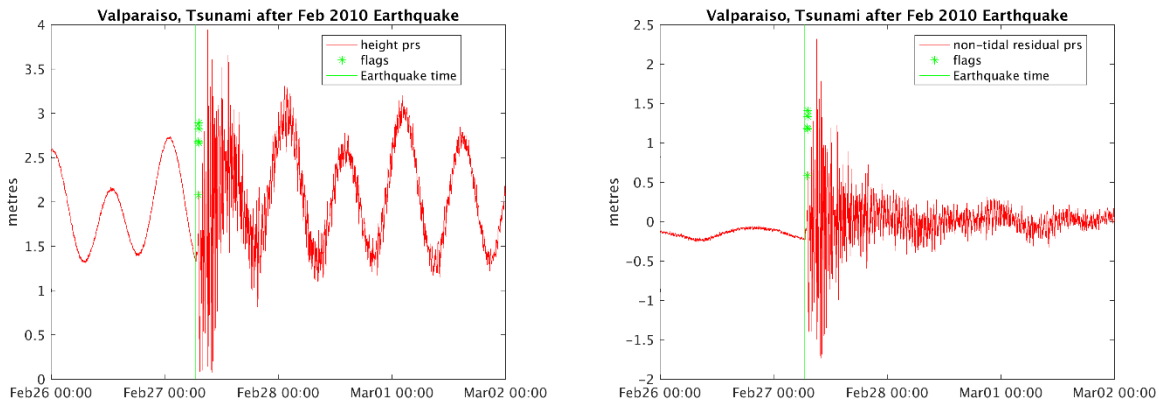


Figure 11: Valparaiso, Chile. The first few records of the tsunami have been discarded by the spike detection algorithm.

C. Comparison with manual QC.

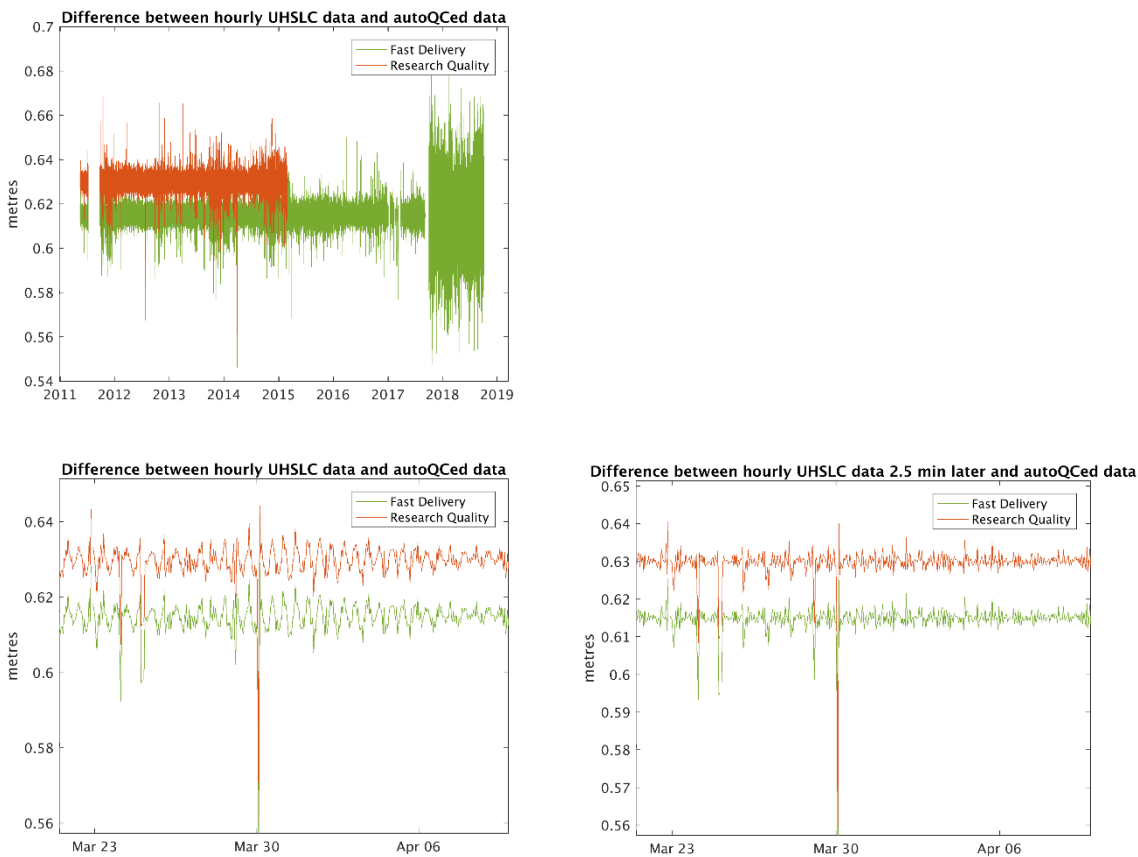


Figure 12: Difference between hourly data at Prickly Bay, Grenada, on UHSLC and from autoQC. (a) All data 2011 to 2018, (b) March 2017, (c) Difference assuming UHSLC data is 2.5 minutes later.

Hourly and daily data for Prickly Bay, Grenada is available at UHSLC (University of Hawaii). The Fast Delivery (FD) service provides data up to recent months, but “Research Quality” (RQ) is only available up to 2015. **Error! Reference source not found.** shows the difference between this manual processing and the single stitched channel from our autoQC. There is a constant difference attributable to datum error - we have not made any datum corrections. This is about 1cm greater on RQ – presumably some manual datum correction from levelling information has been applied at UHSLC. Until 2017 there is also a semi-diurnal periodic difference of about 1.5cm between both the UHSLC streams and the autoQC data. This peaks on

the rising tide and is consistent with a phase difference of 2.5 minutes between the two data sets. This could arise during the averaging from high frequency to hourly data, perhaps from a mis-aligned 5 minute filter, but we have not identified the cause. This would require examination of the high-frequency data at UHSLC system to see whether it is due to an earlier timing difference, and perhaps the averaging code used. The difference remains whether or not we use the Vassie filters or simple hourly averages and averaging on total water levels or non-tidal residuals accounts for less than 3mm.

After October 2017 the difference in the hourly data increases to 10cm and is irregular. Again, the cause is not certain.

D. Comparison with PSMSL data

Data processed using the autoQC software still requires datum control before it is ready for upload to PSMSL, which we have not included at this stage. A simple comparison with the PSMSL monthly data for Prickly Bay (Figure 13) shows some differences in individual months and more gaps – perhaps due to different definitions in what constitutes a complete month.

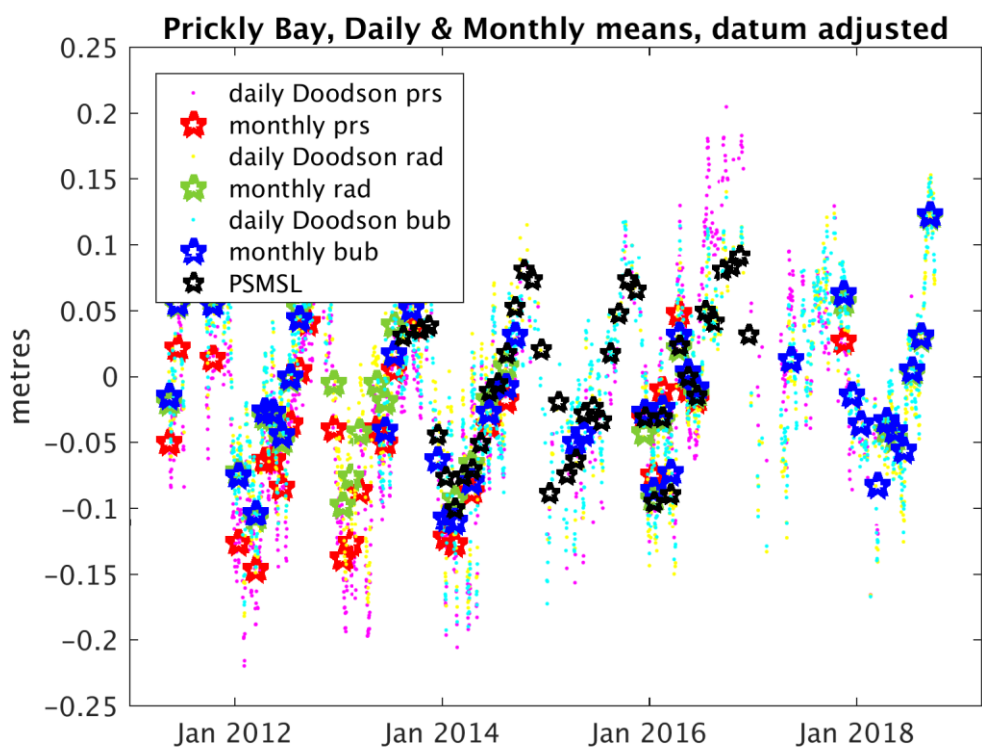


Figure 13: Monthly mean sea level at Prickly Bay, Grenada, from autoQC on IOC data and PSMSL.

E. Test cases used for automatic data QC software

The following sites amongst those that have been used to test the software so far:

Caribbean

Calliaqua (St Vincent). IOC code: calq. Sensors: 'prs','rad'. First record is 2013-12-12

Test: 2013-12-12 - 2014-12-13

Complete record: 2013-12-12 2017-11-23

The separate channels at Calliaqua start out aligned, but drift apart later. To maximise the data it is best to force them into alignment by setting a relative datum of 0.

Prickly Bay (Grenada). IOC code: pric. Sensors: 'prs','rad', 'bub' .

Test (used for workshop): 2017-01-01 - 2017-12-31

Test 2: 2017-01-01 2018-12-20 Includes some tricky problems with the pressure sensor.

Complete: 2011-05-14 2018-10-04

Short test with tricky flags: 2013-02-01 2013-07-01

Ganter's Bay (St Lucia). IOC code: stlu. Sensors: 'pr1','pr2', 'rad' .

2016-08-25 2018-12-31

St Lucia was an example of a site with 3 channels, but channel pr2 did not function well. Tests were done with and without pr2 included.

St Martin. IOC code: stmt. Sensors: 'prs', 'rad'. 2015-12-29 – 2019-02-28

Several other East Caribbean sites have also been processed, in order to examine the impact of hurricanes Irma and Maria.

Senegal, West Africa

St Louis (Senegal) IOC code: stlo2. Sensors: 'rad'. 2016-02-10 2016-12-09

This short and gappy record is a very good example of why you mustn't use too many constituents when fitting the tide. There is actually also some data in 2015, 2017 & March 2018 but it is very poor quality, and because there is only one channel there is no means to automatically reject it. The biggest problem is that the isolated scraps interfere with the tidal fit.

Dakar (Senegal) IOC code: dakar. Sensors: 'prs','rad','ra2'

Dakar is a long record and a good example of what happens when there's a datum shift. There are big shifts around 2011-12-10 and 2014-05-30. With default settings, shifts are not properly handled, so most of the data is wasted. The pressure sensor appears to have a consistent datum, but the radar sensors have jumps of several metres.

One approach to big jumps would be split the time series at an early (or iterative) stage so that they are handled separately. A quick fix is to set a datum correction to align a section of the data, but this requires manual intervention

Complete record: 2007-10-27 2019-02-04

Early record: 2007-10-27 2011-12-09

Late record 2011-12-10 -2019-02-04

Tanzania

Example of a truncated or half-tide sensor. The pressure gauge does not record low water levels, although the radar sensor does. The channel comparison is unable to handle this case, with regular apparent differences between channels. It could be treated using a low water extreme level on that sensor to delete portions of the data.

Mtawa, IOC code: mtwa 2009-2019.

Pakistan

Karachi IOC code: kara. Sensors: 'rad'

Karachi was used as an example of a cyclone event. Cyclone Yemyin – hit Karachi (kara) on 25th / 26th of June 2007. A surge of about 1m was expected.

Again, there was a datum shift in May 2007, which caused some problems with the tidal fit.

Test including datum shift: 2007-01-01 2008-01-01

Test excluding datum shift: 2007-05-09 2008-05-09

Chile

Chilean gauges were used to test the QC routine in the vicinity of a known tsunami, following the earthquake of 27 February 2010. For this purpose only short 3-month tests were done.

Caldera. IOC code: cald. Sensors: 'prs'. 2010-01-01 - 2010-04-01

Coquimbo. IOC code: coqu. Sensors: 'prs'. 2010-01-01 - 2010-04-01

Corral. IOC code: corr. Sensors: 'prs'. 2010-01-01 - 2010-04-01

Valparaiso. IOC code: valp. Sensors: 'prs'. 2010-01-01 - 2010-04-01

Indonesia

Brief examination for evidence of Krakatoa tsunami.

Waikelo. IOC code: waik. 2018-2019

Kolinamil 2. IOC code: koli. 2018-2019.

UK

There is only one channel available on the IOC site for each UK gauge that we examined. This is not the best source for this data, which is available with manual QC from BODC.

Avonmouth Portbury IOC code: avon. Sensors: 'bub'. 2012-01-01 – 2019-01-04

Lowestoft IOC code: lowe. Sensors: 'bub'. 2012-09-04 – 2019-01-24

Newlyn IOC code: newl2. Sensors: 'bub'. 2015-01-01 – 2019-01-24

Liverpool (Gladstone Dock) IOC code: live. Sensors: 'bub'. 2012-09-04 2019-02-04

Denmark

Clock discrepancies.

Hirtshals. IOC code: hirt. 2016-2017.

6. Technical requirements of the code

The Matlab code requires the statistics toolbox, and has been developed on Linux using Matlab R2016b.

Memory appears to be a limiting factor on the demonstration code, particularly if multiple figures are produced. However it was able to run analysis on a 2012-era Windows laptop for one year's data of the 3-channel high-frequency gauges at Prickly Bay.

7. Tide gauge data sources

Data is also available from the following sites. More detail about data availability in the Caribbean is in the accompanying document.

Source	Website	Data Type
PSMSL	https://psmsl.org	Quality controlled monthly and annual data
UHLSC	https://uhslc.soest.hawaii.edu/	Research Quality (JASL): Quality controlled hourly and daily data Fast Delivery: Hourly and daily data with minimal quality control
GESLA	https://gesla.org/	Hourly and higher frequency data – assumed to be quality controlled by supplier
VLIZ	http://www.ioc-sealevelmonitoring.org/	Real time data high frequency data with no quality control
NOAA	https://tidesandcurrents.noaa.gov/	Data from USA, US territories, and some other NOAA-operated sites. Six-minute to monthly quality controlled data.
SHOM	https://data.shom.fr	Data from France, including Overseas territories. One minute and hourly quality controlled data.
SONEL	http://www.sonel.org/	Daily GNSS (GPS) data, and information about land movement at or near tide gauges.

8. Acknowledgements

This work was funded under the Commonwealth Marine Economies Programme, project “Climate Change Impact Assessment: Ocean Modelling and Monitoring for the Caribbean CME states”, workpackage 2.8. It was also supported by the National Environmental Research Council under National Capability Official Development Assistance (NC-ODA), ACCORD programme.

We are grateful for the use of tidal fitting and prediction codes developed over many years at the National Oceanography Centre, most recently by Simon Williams; for the spike detection algorithm from Begona Perez Gomez; and Matlab file exchange entry `parse_pv_pairs` from John D’Errico.

This work would not be possible without a source of high-frequency tide gauge data, and we are very grateful to VLIZ and the IOC for their provision of data at so many gauges; and to the hundreds of teams around the world who fund, install, maintain those gauges, and who make their data available. We hope that our work will assist you in maintaining and developing the global network.