ENACT *IN SITU* DATASET

DATASET VERSION 1.0 DOCUMENT VERION 1.0
EVK-2001-00077 ENACT Deliverable D1.3

Ocean profile processing and quality control

Bruce Ingleby and Matt Huddleston, 29 August 2003

Contents:

Introduction

In 2002/2003 the quality control of oceanic temperature and salinity profiles was rewritten and a comprehensive suite of automated checks introduced. Where appropriate these are based on documented procedures used elsewhere, but often the details differ - in some cases because the checks used elsewhere rely on some human assessment/input and in some cases the detail was not available. Significant effort has been put into coding and tuning an automated track check, because the result of an undetected position error can be quite serious.

This paper describes the system used for ENACT. A version is also being tested for implementation in FOAM in Autumn 2003 - this is similar but does not include the duplicate or stability checks or superobbing. Because the FOAM system will only process one day of data at a time the track check will be less effective. Further differences are noted within the text.

Principles: flag data don't delete them; only correct data if sure and retain original value (the only correction currently applied is to XBT depths). It should be clear from the flags which check(s) a particular value has failed, there is also a final 'reject' flag taking into account all the individual checks. Algorithms and code are shared with the atmospheric processing where possible.

We follow Lorenc and Hammon (1988) and Ingleby and Lorenc (1993) in using Bayesian probability theory to combine information from different checks. In principle this is optimal, but in practice some of the probabilities and pdf's needed are not well known or approximate making it sub-optimal.

System overview

Read observations (Bathy, Tesac and Buoy) from MetDB and concatenate into one array.
ENACT: read observations from NetCDF, thin vertically if necessary, and concatenate Set date/time to missing for any obs outside time window. Set observation type codes Conversions (K to C, background salinity to PSU) Correction for XBT fall rates (depths multiplied by 1.0336 in warm waters, smaller/no correction applied in cooler waters) Track check Internal consistency checks: spike, constant value Superobbing of reports from the same platform: moored buoys (TRITON) only Convert reported temperatures to theta. Stability check for T, S profiles Duplicate check (check for similar position/time) Interpolate background values to obs levels (NB theta not T)
lowest model value used for levels up to half a layer below bottom Convert background potential temperatures to temperature. Interpolate background and observation error estimates to reported levels Perform loose background check - Bayesian check using generic code Average values that have passed qc on to model levels Interpolate background and observation error estimates to model levels Background check the model level values Perform buddy check using generic code Final multi-level check Vertical smoothing (optional, for old style ACobs) *For ENACT map buddy check flags back to reported levels* Listings and statistics. Output observations in ACOBS and/or ENACT format.

The initial checks and processing only use reported values, and up to the superobbing temperature is the thermal variable used. The stability check uses potential temperature (because of the equation of state used) and so does the background check of reported level values and all the model level checks.

XBT depth correction and start-up transients

XBT (expendable bathythermograph) depth is calculated from the time since release of the probe. In general XBTs fall slightly faster than suggested by the manufacturers' equations. Based on comparisons with CTD data Hanawa et al (1995) suggest the linear correction Zcor = 1.0336 Zrep for T4, T6 and T7 XBTs. They note that even after correction 17.5% of reports still fall outside the required accuracy of 5 m or 2% of depth - mainly due to batch to batch variation of XBTs, but with a possible regional contribution (see below). For interpolation to standard levels Conkright et al (2002, WOD01 documentation, p 46) use the Hanawa et al correction - but not for T5 XBTs: 'When instrument model information was not provided with XBT data, profiles with all measurements shallower than 840 meters were assumed to need this correction (profiles with measurements deeper than this were assumed to be taken with a T5 instrument, which does not have a systematic error in its drop rate equation).' We use a similar exclusion (note also that

MetDB BATHY reports with integer callsigns are assumed to come from fixed moorings and are not corrected).

The other complication is that colder water has higher viscosity - reducing the probe fall rate. Most of the comparisons used by Hanawa et al (1995) were at relatively low latitudes. Thadathil et al (2002) examined CTD-XBT pairs in the Southern Ocean and concluded that XBTs there fall marginally *slower* than given by the manufacturers' equations - and therefore the Hanawa et al correction was inappropriate there. The North-East Atlantic sample of Hanawa et al (with a mean temperature of 10 degrees C) also supports the suggestion that the correction needed at high latitudes is less than that at low latitudes. To attempt to model this in a simple way but with some physical basis we calculate the mean temperature of the report over the top 300 m (Tmean). If Tmean is above 15 degrees then a correction factor of 1.0336 is applied, if Tmean is below 5 degrees the factor is 1.0, with a linear transition in between. This transition range is broadly consistent with the results quoted above, and is slightly cooler than suggested by the modelling of Thadathil et al. The correction factor is stored in the ENACT output file so it could be re-evaluated and changed if necessary at a later stage.

XBTs take a finite time to adjust to the water temperature after entering the sea. For this reason UNESCO (1997) recommended that temperatures in the top 3.7 m should not be used. Kisu and Hanawa (2002a) report that transient responses vary between 2 and 10 m based on instrument type. *No action at present. Based on the UNESCO recommendation should we flag any XBT temperatures at less than 4 m depth?*

Another occasional problem of XBT data is a drift; an error typically increasing by 0.1 degree per 100 m (sometimes known as bowing), see Kisu and Hanawa (2002b) and references. This may start at the surface or lower down, and is more frequent for some instrument types than others. It is fairly subtle and only likely to be detected where there is a deep mixed layer. There is no specific check for these errors at present.

Spike and other univariate checks

Constant value check: If 70% or more of the temperature (salinity) levels, covering at least 50 metres, are set to the same value then all of the temperature (salinity) values are rejected. There is also a specific rejection for temperatures less than 1 C, above 1000 m depth, in the Tropics. This is mainly to detect zero values which sometimes indicate missing data in TAO profiles.

Spike and step check: First the tolerances (TTol and STol) to be used at each depth are calculated (see Table 1). Then vertical differences of the temperature and salinity profiles are formed, where level I is at most 50 m deeper than level I-1 (100 m where level I is 350 m or more deep):

$DT(I) = T(I) - T(I-1)$ and $DS(I) = S(I) - S(I-1)$

If either DT(I) or DT(I-1) is larger in magnitude than TTol and $ABS(DT(I) + DT(I-1)) < 0.5*TTol$ then T(I-2) and T(I) are in good agreement with each other, but T(I-1) is rejected as a spike.

There is also a check for sharp, but smaller amplitude, temperature spikes:

if DT(I) or DT(I-1) is larger in magnitude than 0.5*TTol and at least one of the vertical gradients is larger in magnitude than 0.05 deg/m and $ABS(DT(I) + DT(I-1)) < 0.25*ABS(DT(I) - DT(I-1))$ then T(I-1) is rejected as a spike.

Remaining values of DT(I-1) with magnitude larger than TTol represent large steps in the profile, both T(I-2) and T(I-1) are flagged as suspect unless one of the following conditions applies:

a) T(I-1) agrees to within 0.5*TTol with a value interpolated from T(I-2) and T(I)

b) for depths of 250 m or less if $-3*TTol < DT(I-1)$

If the last DT value is larger in magnitude than TTol or the bottom temperature is zero then the bottom temperature is flagged as suspect.

If there are four or more temperature spikes or steps detected then the whole report (salinity as well) is rejected. In general spike values are rejected but other suspect values can be reprieved by the background check.

The salinity spike check is similar but simpler. If either DS(I) or DS(I-1) is larger in magnitude than STol and $ABS(DS(I) + DS(I-1)) < 0.5*STol$ then S(I-1) is rejected as a spike.

Remaining values of DS(I-1) with magnitude larger than STol result in both S(I-2) and S(I-1) being flagged as suspect unless a condition analagous to a) above holds.

If the last DS value is larger in magnitude than STol then the bottom salinity is flagged as suspect.

If a temperature spike is detected the corresponding salinity value is automatically rejected - the temperature check seems more sensitive and there is often a small "blip" in the salinity at the same level - possibly caused by the use of the temperature in deriving salinity from conductivity. The temperature spike check seems to be fairly well tuned, in that it picks up most subjectively determined spikes and nothing else.

Tolerance	Temperature	Salinity
Default	5.0 deg	1.0 PSU
Depth	TTolFactor	STolFactor
0	1.0	1.0
200 (*)	1.0	1.0
300 (*)	0.5	0.2
500	0.4	0.2
600	0.3	0.2

Table 1 Default tolerances and their variation in the vertical for the spike check. (*) these are extratropical transition depths, within 20 degrees of the equator the (linear in depth) transition is between 300 and 400 m. The 500 and 600 m changes in TTolFactor are implemented as step changes.

Track check

An index is set up which lists observations in callsign order, and within that by time, if the implied speed between two reports is excessive then one of them is wrong. The tricky bit is deciding which one, and having an extended sequence of reports helps. (There is an argument for having overlapping windows, but for practical reasons this is not done.) A series of tests are applied: looking for excessive speeds compared to neighbouring reports, looking for a distinct kink in the track and looking for a smooth distance/time relationship in approximately collinear tracks. More details are given below. There is no explicit check for reports at land (or ice) points, but reports at clearly inland positions will not have background values and will be rejected for that reason.

Exclusions. If there is more than one vessel with the same callsign then track checking is very much more difficult and the algorithm below would often reject all the reports. Given that only 1% or so of reports have track errors it is better in this case to not flag any of them. Certain default callsigns ('SHIP', '0' and ' ') are not checked, and if there is only one report with a particular callsign then it cannot be checked. Some sequences of ship reports contain between two and four different positions at each nominal date/time - presumably because a batch of reports has been compiled or transmitted together. If more than half the reports in the sequence have such repeat times then the sequence is not checked (for most oceanographic purposes time errors of a day or so are not significant). Even so repeat times can cause a significant number of reports from a few callsigns to be rejected. A few platforms have consistently high speeds - notably air-dropped XBTs - attempts to check these using a larger speed tolerance were not very satisfactory. Thus if the median implied speed is larger than MaxSpeed (see below) the reports are not checked.

MaxSpeed is set to 15 m/s for ships and to 2 m/s for buoys - both moored and profiling. Originally there was no intention to check moored buoys, but it was discovered that a few large position errors do exist (possibly due to the System Argos satellite positioning method). The latitude and longitude are converted to 3D cartesian coordinates - this facilitates calculation of distances and angles between successive observations (straight line rather than great circle distance is calculated - they will only be significantly different for obviously wrong tracks). To allow for rounding of both position and time the implied speed between observation K-1 and K is calculated as:

Speed(K) = (Dist(K)-0.5*DistRes)/MAX(DTime,TimeRes)

where DistRes = 20000 m (20 km) and TimeRes = 600 seconds. Where both displacements are larger than DistRes then Angle(K), the change in angle at K (going from K-1 to K to K+1) is calculated, this is approximately 0 degrees for a great circle track.

The largest speed in the whole track is found at position Kmax. If it is larger than MaxSpeed, or if it is larger than 0.8*MaxSpeed and either Angle(Kmax-1) or Angle(Kmax) is greater than 90 degrees then one of the positions Kmax-1 or Kmax is deemed to be wrong. The checks below determine which to reject, the offending report(s) is (are) omitted from the sequence, the distances and angles are recalculated and the process repeated until there are no excessive speeds in the track. If more than half of the reports in the track have been rejected then the whole track is rejected.

Speed(Kmax) SpeedTol: deciding between Kmax-1 and Kmax.

If Kmax=2 there is less checking that can be done. If the implied speed between 1 and 3 is less than SpeedTol and either Speed(3) SpeedTol or Angle(3) 45 then report 2 is rejected (ErrCat=2), otherwise report 1 is rejected (ErrCat=1). An analagous test is used if Kmax is the last report in the track.

ErrCat=4: if Speed(Kmax-1) SpeedTol then reject Kmax-1;

if Speed(Kmax+1) SpeedTol then reject Kmax.

ErrCat=5: if (implied speed between Kmax-1 and Kmax+1) SpeedTol then reject Kmax-1;

if (implied speed between Kmax-2 and Kmax) SpeedTol then reject Kmax.

ErrCat=6: if Angle(Kmax-1) 45 + Angle(Kmax) then reject Kmax-1;

if Angle(Kmax) 45 + Angle(Kmax-1) then reject Kmax.

ErrCat=7: if Angle(Kmax-2) 45 and Angle(Kmax-2) Angle(Kmax+1) then reject Kmax-1;

else if Angle(Kmax+1) 45 then reject Kmax.

ErrCat=8: if Speed(Kmax-1) < MIN(Speed(Kmax+1),0.5*MedianSpeed) then reject Kmax-1;

if Speed(Kmax+1) < MIN(Speed(Kmax-1),0.5*MedianSpeed) then reject Kmax.

Dist1 is calculated as the distance Kmax-2 to Kmax to Kmax+1;

Dist2 is calculated as the distance Kmax-2 to Kmax-1 to Kmax+1.

DistTol is the larger of DistRes and 0.1*(distance Kmax-2 to Kmax-1 to Kmax to Kmax+1).

ErrCat=9: if Dist1

if Dist2

PD1 = Dist(Kmax-1)/Dist2, this is the ratio of the distance from Kmax-2 to Kmax-1 over the distance from Kmax-2 to Kmax+1 omitting Kmax;

PD2 is analagous: the distance from Kmax-2 to Kmax over the distance from Kmax-2 to Kmax+1 omitting Kmax-1

PT1 (PT2) is the time from Kmax-2 to Kmax-1 (Kmax) over the time from Kmax-2 to Kmax+1; for smooth motion the ratios PD1 and PT1 (also PD2 and PT2) should be approximately equal.

ErrCat=10: if ABS(PD1-PT1) 0.1+ABS(PD2-PT2) then reject Kmax-1;

if ABS(PD2-PT2) 0.1+ABS(PD1-PT1) then reject Kmax.

ErrCat=100: if none of the above tests have been positive or if removing the rejected report, Krej, leaves an excessive speed between Krej-1 and Krej+1 then both Kmax-1 and Kmax are rejected. This is typically a fairly symmetric situation where Kmax-2 and Kmax-1 are consistent with each other and Kmax and Kmax+1 are also consistent with each other but not with the first two.

The ErrCat values can be printed out in diagnostic messages. Broadly speaking the tests above show diminishing returns - fewer positive results for the later tests.

Ingleby (1994) estimated that about 1.5% of surface ships had (small or large) position errors. The proportion seems to be about 1% for oceanographic reports from the MetDB and between 1 and 2.5% for reports from WOD01. What constitutes an error depends how time errors are regarded and at what distance spatial errors become insignificant.

Superobbing

The main motivation for superobbing is that in the TAO/TRITON moored buoy array in the tropical Pacific the TAO (Tropical Atmosphere Ocean) buoys report daily averages, whereas the TRITON reports (2000 onwards) are quasi-hourly. For data assimilation it is desirable to present the TRITON data as daily averages (in this context the superobbing does not modify the observation error estimates which should already be set appropriate for daily averages). Thus superobbing is only applied to moored buoys - they generally report on the same set of levels, simplifying the averaging. There are sometimes very dense sequences of XBTs for which superobbing might be desirable, but this would raise various practical issues.

The moored buoys are sorted by callsign and time (omitting any that failed the track check). They are grouped into windows (normally 24 hours) and for each callsign/window a superob is created. The input observations can be on different subsets of the buoy levels - hence there is an initial step to find the combined set of levels. For each combined (output) level a simple average is calculated (duplicates with the same level and time are not included, rejected values (spikes) are also not included). The standard deviation of the reported values at each level is also calculated and if this is excessive (see Table 2) then the averaged value is omitted. There can be quite a large SD within the thermocline partly due to the semi-diurnal tide (and very small SDs at other levels) so this test is set to be rather loose.

Depth (m)	Temperature	Salinity
0-40	0.7	0.5
41-200	2.5	0.5
201-250	2.0	0.5
251-500	0.8	0.5
500	0.4	0.5

Table 2. SD limits for superobs.

By default the superob overwrites the last of the input observations, but there is a namelist option OverallocateFactor (usually set to 1.1 if used) that allocats extra space in the main arrays in which to store the superobs without overwritting.

Conversions

Depth is converted to pressure for use in the calculation of potential temperature and density. This uses an iterative inverse to the algorithm of Saunders (1981), which includes the variation of gravity with latitude. Currently reports with pressure but not depth coordinates are not processed (for ENACT note that we will use the depth calculated and stored in WOD01 (p22) for such reports).

Temperature is converted to potential temperature (theta) using numerical integration (with a 1 decibar increment) from the reported level to the surface. This uses the Bryden (1973) expression for the adiabatic lapse rate as used in the Unified Model. (This gives consistency within our data assimilation system, and the errors in the conversion are probably smaller than those in the observations, but we may at some point want to change to the more accurate McDougall et al (2003) algorithm).

Stability check

The stability check makes sure that a profile does not have a density inversion (higher density above lower density). In some cases temperature increases with depth but this is offset by a decrease in salinity. Thus both temperature and salinty are needed for a rigorous stability check - reports with temperature only (XBTs and many moored buoys) are not checked. Currently only the original reports are checked, in future we may recheck the reports after averaging onto model levels.

The equation of state used to calculate density (rho) is that of McDougall et al (2003) - this uses theta (rather than T), salinity and pressure as input. When comparing levels k-1 and k, we calculate

Drho(k) = rho(theta(k),S(k),P(k)) - rho(theta(k-1),S(k-1),P(k))

the second term on the right hand side is the potential density of water at level k-1 referenced to level k. Negative values of Drho indicate inversions, but we allow small inversions and test for Drho(k) < -0.03 kg m**-3 (partly based on Conkright et al, 2002, WOD01 documentation, p 48-49).

There is a test for "spikes" in the density:

ABS(DRho(k-1)+DRho(k)) < 0.25*ABS(DRho(k-1)-DRho(k))

in this case theta and S at level k-1 are flagged, otherwise DRho at k and k+1 are tested in the same way (and level k flagged if the test is true). If neither of these tests are positive it is not clear whether level k-1 or k is erroneous and temperature and salinity at both these levels are treated as suspicious by increasing the PGE.

If there is a significant inversion at the bottom of the profile then the lowest values of temperature and salinity are flagged (errors in the level above are possible but less likely).

Occasionally there is a salinity profile that has been reported to only one decimal place - the 0.1 steps in salinity sometimes trigger the stability check. (The estimate of salinity observation error varies between 0.11 and 0.17. Should we reject reports with 0.1 resolution, or increase their observation error estimates and relax the stability check?)

An alternative approach is that of Jackett and McDougall (1995). They minimally modify density inversions so that buoyancy frequency is larger than a specified lower bound using a constrained least-squares problem. (Depends on local T-S relationship to give smooth profile.) The approach taken here is to reject dubious values rather than to try to correct them. However, examining the local dependence of density on T and S might allow the rejection of only one of them where an inversion is present.

Duplicate check

If switched on the duplicate check looks for pairs of reports within 0.2 degrees latitude/longitude and 1 hour, the report with the highest preference factor is retained (one is chosen at random if they have the same preference factor). The preference factor is set to the number of levels in the report, plus 10 if the report is not an XBT, plus 10 if the callsign is not 'SHIP'. (For ENACT preference will also be given to WOD01 data compared to WOCE and BMRC data.)

Background values and background check

The background is generally a short-range (typically 1 day) ocean forecast for FOAM, a climatological field for GloSea or a damped anomaly persistence forecast for ENACT (these will have different errors - see below). The background values on each model level are first interpolated horizontally to the latitude/longitude of each observation. In mid-ocean bilinear interpolation is used, but near the coast or bottom topography one or more of the four surrounding grid points may be missing - in this case the background value at the closest valid grid point is used. If all four surrounding grid points are missing (because of limited model resolution, or missing inland seas, eg Black Sea) then the background value is set as missing and the observed value will fail the background check. The profile on model levels is then vertically interpolated (linear in depth) to reported levels as necessary. Background error estimates are interpolated in the same way. For ENACT background values are on the 40 level GloSea grid, whereas background error estimates are on the 20 level FOAM grid.

The Bayesian background check is as described by Lorenc and Hammon (1988) or Ingleby (1998, OSDP2). It is applied to the reported level values - with increased background errors specified (below) to make the check less strict. After vertical averaging the background check is reapplied to the model level values and the output probabilities form the input to the buddy check.

For the Bayesian background check we specify

a) the initial PGE (probability of gross error) is set to 0.01 (ie assuming that 1% of reports are erroneous) except that it is set to 0.05 for bathythermograph temperatures

b) Values considered as suspect by the spike/step check or the stability check have their initial PGEs increased:

PGEnew = 0.5 + 0.5*PGE

c) the probability density of 'bad' observations is set to 0.1 for temperature and 0.25 for salinity (corresponding to error ranges of 10 degrees and 4 PSU)

Note that the background check is more sensitive to the specified rms (o-b) errors than it is to these parameters.

Observation and background errors - discussion

For high quality temperature observations accuracies (measurement error) of 0.01 degree or better are quoted, and they are reported to the nearest 0.01 degree in the TESAC code (lower quality temperatures are reported to the nearest 0.1 degree in BATHY code). Note that any errors in the depth assigned to a given temperature value will appear as a temperature error proportional to the vertical temperature gradient. Emery and Thomson (2001, p39) suggest that "modern CTDs are accurate to approximately 0.002 C in temperature, 0.005 psu in salinity and <0.5% of full-scale pressure in depth". For XBT temperature measurements they suggest an accuracy of 0.1 C or slightly more (p17) but error in depth estimation tends to dominate. (Bailey et al QC cookbook, p32 suggests 0.15 for XBTs) McPhaden et al (1998, Appendix B1) suggest errors less than 0.1 degree for subsurface temperatures from the TAO array. For XBT data (Appendix B4.2) they suggest temperature accuracy of 0.15 degrees or better and "After correcting for the systematic error the depth error is within the manufacturer's specifications (2% of depth or 5 m) in ~82% of XBT drops." However 15% of XBTs suffer instrument malfunctions before reaching 250 m. For sea surface salinity data (Appendix B4.3) an accuracy of 0.2 psu is quoted.

In data assimilation systems it is usual to define

observation error = measurement error + representivity error

Representivity error comes from features that may be sampled by observations, but are below the scale of the model background - thus it is a function of the model resolution (see Desroziers et al 2001 and references).

The estimates of observation error used in data assimilation appear somewhat ad hoc. BMRC (Wang et al 2002) use an rms observation error of 1 degree (correlated with spatial and temporal scales of 150 km and 5 days, respectively). Weaver et al (2002) use 0.5 degree for daily averaged TAO data and 1 degree for other observations (assumed uncorrelated). These estimates require representivity error to be much larger than measurement error (probably true up to a point), or a proportion of reports with much worse measurement error.

Martin et al (2002) estimated background and observation errors from O-B covariances as a function of distance, their figure 11 suggests observation error SDs between 0.5 and 0.7 degree for depths between 50 and 400 m, above and below this the estimates are significantly smaller - set to an SD of 0.1 (the estimated variances went negative for some levels). They used a moderately high resolution (one third of a degree) model of the North Atlantic ocean and represented the background error as the sum of synoptic and mesoscale errors, both represented by SOAR functions. They find the largest contributor to the O-B variance to be the mesoscale error, with a scale of 40-60 km. With a coarser resolution model some of the 'mesoscale error' would become representivity error. It is also possible that some of it is correlated observation error (especially as the closest pairs of reports will often come from the same ship or buoy). Both of these effects would help to explain why the other groups have larger observation error estimates.

Weaver et al (2002, section 4.5) show that forecast errors, evolved using 4D-Var, tend to be largest round the level of the thermocline - where vertical temperature gradients are a maximum. This seems reasonable, and probably applies to representivity errors as well as background errors. Regionally at least, observation errors (dominated by representivity error) may be approximately proportional to background error. For example Derber and Rosati (1989, p 1338) used larger observation errors in the northwestern Pacific and Atlantic since "In these regions there is a propensity to observe eddies not resolvable by the model".

The estimation of salinity errors has received little attention in the literature - especially background salinity errors. Bacon et al (2001) found that the four PALACE (ARGO) floats that they compared with CTD data had stable salinity sensors within or very close to the manufacturer's specification with a maximum drift of (0.0009 +/- 0.0004) per month. They suggest corrections based on such comparisons. Davis et al (2001) describe gradual drifts in salinity calibration of PALACE floats - of different signs depending on whether antifouling biocide was applied - as well as abrupt jumps of 0.1 or more (see their Figure 4). MacDonald et al (2001) suggest that some Japanese naval and fishery reports have poorly calibrated salinity values. As yet no special action has been taken to check or reject these values.

Information about climatological covariance scales in the tropical Pacific is given in Kessler et al (1996), these show pronounced East-West stretching near the equator. Atmospheric fields show a related but less marked East-West stretching in the tropics, and forecast error scales are generally shorter than climatological covariance scales (see Ingleby (2001) and references).

Observation and background errors - specification

Temperature observation error variances are taken to be half the global observation minus background (o-b) variances from Martin et al (2002). The forecast error variance was calculated by subtracting the observation error variance from the (o-b) variance in 10 degree lat/long squares, but was reset to equal observation error variance if smaller. The mesoscale and synoptic scale variances are then each set to be half the total forecast error variance. This method, whilst being fairly ad hoc, should at least give robust estimates of the variances and should capture spatial variations in the total error variances reasonably well. For salinity global (o-b) variances were taken from preliminary runs of the QC/processing described here, vertically smoothed by hand, and divided equally into observation and background error estimates. The background error estimates were then increased in areas where the temperature background errors are larger than average. In addition the background salinity rms error estimate is increased by 50% North of 75 N - the extent and timing of melting sea ice give considerable variability of near-surface salinity in the Arctic Ocean (and apparently in the Baltic Sea as well). There are some plots here and example tabulated values in Appendix 1. The default correlation scales are set to be 100km for the mesoscale and 400km for the synoptic scale for both temperature and salinity.

The statistics described above are used for FOAM, except that for the background check on reported level data the background error rms is multiplied by a factor of 1.5 (2.0 for a climatological background) and within 10 degrees of the equator, for both the original and averaged temperature data, there is an extra factor of 1.5 (3.0 for a climatological temperature background). (A climatological background is likely to have larger correlation scales as well but this is not modelled.)

ENACT uses a background that is a combination of climatology and persisted anomalies. This is certainly better than climatology - particularly in the tropical Pacific - but probably not much better in earlier years due to the scarity of reports. For this reason it is given 'climatological' background errors.

Vertical averaging

Currently Ops_OceanVertAverage assigns each unflagged reported level increment to the nearest model level and performs a simple averaging. For ENACT a later step maps back flags from the model level value to the reported level value(s) that contributed to it.

Some aspects of this simple averaging are not ideal, and in some conditions (a sharp, tropical thermocline, shallower in the background than in the report) the use of o-b increments can result in a small warm 'nose' at the bottom of the mixed layer. The modified Reiniger and Ross interpolation method described in Conkright et al (2002, p 46) is apparently designed to cope with sharp thermoclines but has not been investigated.

Other options For most ocean observation profiles (the obvious exception being moored buoys which are a series of point measurements rather than a complete profile) it would be more appropriate to use Ops_VertAverage which is used for radiosonde vertical averaging (documented in section 5 of OSDP5, see also Ops_SondeAverage). This basically joins the reported values using a series of straight line segments and calculates the average for each model layer that is at least half covered by the profile (for radiosonde temperature there are added complications in order to conserve layer thicknesses). A samll gap in the reported profile is filled in by interpolation, but with a 'large' gap no interpolation is done. This is especially appropriate if the reported values are turning points intended to represent the whole profile (some TESACs), and also assuming that the model values represent layer averages rather than point values or finite elements. Initial tests of this method gave slightly worse results (in rms o-b), this will be revisited when time allows; there is also the question of whether to average the reported values or the o-b increments.

For reports at fixed levels (including moored buoy data) it isn't clear to me which is better: a) throw increment to nearest model level, or b) interpolate and average reported profile. Note that some buoys have a bottom level value with a fairly large gap above; under b) as it stands only the upper levels would contribute to the averaged profile, with the detached value being ignored. (ECMWF have reported problems from interpolating across large gaps from TRITON buoys.) Various combinations of a) and b) would be possible, in principle the decision should depend on what the reported and model values represent.

Buddy checking

The averaged (model level) data is buddy checked using as input the PGEs from the background check of averaged values. Pairs of close profiles are found and the o-b increments at the same depth (model level) are compared and the PGEs updated. The comparison and updating are performed separately for temperature and salinity. At the end of the buddy check those values with PGEs = 0.5 are rejected. Note that the buddy check can either increase or decrease the PGE. The buddy check is essentially as developed by Lorenc and Hammon (1988) and described in more detail in Ingleby (1998, OSDP2) except that the correlation has been modified to have two different scales (following Martin et al (2002)) and to be anisotropic near the equator. Within 10 degrees of the equator the East-West scales are doubled but the North-South scales are unchanged, with the correlation patterns being eliptical.

The details of the covariance specification are as follows. Within 10 degrees of the equator the anisotropy factor c is 2.0 decreasing linearly to 1.0 (isotropy) at 15 degrees from the equator. The mesoscale and synoptic scales MesScale and SynScale (default values of 100 km and 400 km) are scaled similarly. The values of c from the two reports A and B are averaged. Then set Sin2 = sin**2 of angle of line joining A and B with EW. The scaled 'mesoscale' distance between A and B is

MesSDist = SQRT(1.0 + (c**2 - 1.0)*Sin2)*Dist/MesScale

where Dist is the actual distance.

The scaled 'synoptic' distance is defined similarly.

The time scale (CorScaleT) used is 5 days (432000 seconds) and we set

TimeD2 = ((TimeA-TimeB)/CorScaleT)**2

A SOAR function is used in the horizontal, and a Gaussian correlation function in time so that the background error covariance is:

Covar = SQRT(MesVarA * MesVarB) * (1.0+MesSDist) * EXP(-MesSDist -TimeD2)
+ SQRT(SynVarA * SynVarB) * (1.0+SynSDist) * EXP(-SynSDist -TimeD2)

where MesVarA is the mesoscale error variance at A etc. Currently the mesoscale and synoptic variances are taken to be equal.

Profiles within 400 km are allowed to buddy check each other (in rare cases this may result in inappropriate pairings - where there is land in between - or possibly between deep water and coastal profiles). Each profile is allowed to "choose" up to nine buddies (up to six within each five degree latitude band used for sorting) - because it can itself be chosen this means that each profile can have up to about 18 buddies. Buddies with the same callsign are not allowed - testing allowing such pairs showed some problems with biased observations (ie correlated errors). As for atmospheric data the 'damping factor' (see section 3 of OSDP2) is set to 0.5.

Future work: Error correlations are thought to be higher along density surfaces than at constant depth; this could be taken account of in the choice of which levels to compare from different profiles but there would be complications. Ideally the correlations would also reflect differences in bottom topography and any barrier (a peninsular or isthmus) between any two points (see Ridgway et al 2002). However modelling such effects (particularly in the analysis - where it would be more important than in the buddy check) is far from trivial and will not be attempted for now.

Final multi-level check

It was found that, because of the level by level nature of most of the checks, that a final check on vertical consistency of the QC decisions was required. If more than 50% of the temperature levels have been flagged then the whole report is rejected. If more than 50% of the salinity levels have been flagged then all the salinities are rejected. The percentage flagged is calculated from the sum of the percentages for reported levels and averaged levels - because flagged reported levels don't contribute to the averaged values.

There is then a step to reinstate flagged T or S values if they agree with values above/below - typically these will be values with slightly larger o-b differences, sometimes in the thermocline. A tolerance is defined: 0.5 degrees, or 0.1 PSU down to 200 m, changing linearly to 0.25 degrees, or 0.05 PSU below 300 m (the transition depths are 100 m deeper within 20 degrees of the equator). If a flagged model level value matches an adjascent unflagged value to within the tolerance, or if the values above and below are both unflagged and the value lies between them, then it is reinstated.

Vertical smoothing

After the buddy check, if an 'old style ACobs file' is required a filter VRTF (from AC code) performs vertical smoothing of the increments. With the current options VRTF copies the lowest increment to all levels below, and similarly for the highest increment, but these 'extrapolated' levels are flagged.

ENACT data sources and processing

The primary data source for ENACT is the World Ocean Database 2001 (WOD01) (time sorted data). This is reformatted into the ARGO NetCDF format. The WOD01 categories are mapped into WMO instrument types as follows:

```
MBT 1941- : mechanical bathythermograph data          60101   800
XBT 1967- : expendable bathythermograph data          60100   401
CTD 1967- : high resolution CTD data                  60400   830
MRB 1990- : Moored buoy data                          60300   820
PFL 1994- : Profiling float data                      60201   831
DRB 1998- : Drifting buoy data                        60500   995
APB 1997- : Autonomous Pinniped Bathythermograph      60700   997
UOR 1992- : Undulating Oceanographic Recorder         60600   996
OSD 1800- : low resolution (bottle) CTD data          60200   741
```

Additional data from WOCE, G Johnson (PMEL) and BMRC are also used towards the end of the ENACT period.

Some **vertical thinning** is done of reports with more than 150 levels: levels are selected every 100m below 1000m, every 50m between 500m and 1000m and all the remaining available levels are divided equally in top 500m (usually giving a spacing of 5m or less).

After quality control and processing the reported level data is written out for ENACT use. Again this uses the ARGO NetCDF format with some additions, notably the background values, the XBT fall-rate correction applied and extra QC information. For moored buoy data any superobs are written out, but the original reports which were combined into superobs are not. Similarly reports found to be surplus by the duplicate check are not written out.

Diagnostics

Various statistics are produced: rejection rates by category (Bathy, Tesac, Argo, Buoy) for both reported and model levels. Mean, rms, max and min statistics for both the averaged values and their differences from background by depth/model level and category (also a latitude band option).

Depending on the output level selected information on reject or other decisions are available from various checks. The following key words are useful to search for, or to grep: Track - track check; OOVC - vertical spike and step check; OOVS - vertical stability check; OOVF - final multi-level check; Whole - rejections of whole reports from various checks.

Listings of observations, including the flags are available. These can be inspected directly or displayed graphically using the PV-Wave program SeaView. This provides graphical display of obs locations, ship tracks, vertical profiles - with background - and QC decisions.

References

Bacon, S., L. R. Centurioni and W. J. Gould, 2001: The evaluation of salinity measurements from PALACE floats. Journal of Atmospheric and Oceanic Technology 18 1258-1266

Bryden, H.L. 1973, New polynomials for thermal expansion, adiabatic temperature gradient and potential temperature of sea water. DEEP-SEA RES., 20, 401-408

Conkright,M.E., J.I. Antonov, O. Baranova, T. P. Boyer, H.E. Garcia, R. Gelfeld, D. Johnson, R.A. Locarnini, P.P. Murphy, T.D. O'Brien, I. Smolyar, C. Stephens, 2002: World Ocean Database 2001, Volume 1:

Introduction. Ed: Sydney Levitus, NOAA Atlas NESDIS 42, U.S. Government Printing Office, Washington, D.C., 167 pp.

Derber,J. and Rosati,A. 1989: A global oceanic data assimilation system. JPO 19, 1333-1347

Desroziers,G., Brachemi,O and Hamadache, B. 2001: Estimation of the representativeness error caused by the incremental formulation of variational data assimilation. QJRMS 127, 1775-1794

Davis, R.E., J. T. Sherman and J. Dufour, 2001: Profiling ALACEs and other advances in autonomous subsurface floats. Journal of Atmospheric and Oceanic Technology 18 982-993

Emery,W.J. and Thomson,R.E. 2001: Data Analysis Methods in Physical Oceanography. Elsevier. xvi + 638 pp.

Hanawa, K., P. Rual, R. Bailey, A. Sy and M. Szabados, 1995: A new depth-time equation for Sippican or TSK T-7, T-6 and T-4 expendable bathythermographs (XBT). Deep-Sea Research I 42 (8), pp. 1423-1451.

Ingleby,N.B, 1994: Surface ship position errors and detection algorithms Forecasting Research Division Technical Report No. 101

Ingleby,N.B, 1998: Generic quality control. OPS Scientific Documentation Paper No. 2; Met Office.

Ingleby,N.B, 2001: The Statistical Structure of Forecast Errors and its Representation in The Met. Office Global 3D-Var. Quarterly Journal of Roy.Met.Soc., 127, 209-231, 2001.

Ingleby,N.B, and Lorenc,A.C. 1993: Bayesian quality control using multivariate normal distributions. Quarterly Journal of Roy.Met.Soc., 119, pp 1195-1225

Jackett DR, McDougall TJ, 1995: Minimal adjustment of hydrographic profiles to achieve static stability Journal of Atmospheric and Oceanic Technology 12 (2): 381-389

Kessler WS, Spillane MC, McPhaden MJ, and Harrison DE. 1996: Scales of variability in the equatorial Pacific inferred from the tropical atmosphere-ocean buoy array J CLIMATE 9 (12): 2999-3024 Part 1 DEC 1996

Kizu, S., and K. Hanawa, 2002a: Start-up transient of XBT measurement. Deep-Sea Res., 49(5), 935-940.

Kizu, S., and K. Hanawa, 2002b: Recorder-dependent temperature error of expendable bathythermograph. J. Oceanogr., 58(2), 469-476.

Lorenc,A.C. and Hammon,O. 1988: Objective quality control of observations using Bayesian methods. Theory, and a practical implementation. Quarterly Journal of Roy.Met.Soc., 114, pp 515-543

MacDonald, A.M., T. Suga and R. G. Curry, 2001: An isopycnally averaged North Pacific climatology. J. Atmos. Oceanic Tech. 18 394-420

McDougall TJ, Jackett DR, Wright DG, Feistel R, 2003: Accurate and computationally efficient algorithms for potential temperature and density of seawater. Journal of Atmospheric and Oceanic Technology 20 (5): 730-741

Martin,M.J., Bell,M.J. and Hines, A. 2002: Estimation of three-dimensional error covariance statistics for an ocean assimilation system. Ocean Applications Tech. Note 30. on-line version

Ridgway KR, Dunn JR, Wilkin JL 2002: Ocean interpolation by four-dimensional weighted least squares - application to the waters around Australasia Journal of Atmospheric and Oceanic Technology (9): 1357-1375

Saunders, P.M. 1981: Practical conversion of pressure to depth J. PHYS. OCEANORG. 11, 573-574

Stammer, D. 1997: Global characteristics of ocean variability estimated from regional TOPEX/POSEIDON altimeter measurements. Journal of Physical Oceanography, VOL. 27 NO. 8, 1997 1743-1769 Section 5 shows longer horizontal scales at low latitudes broadly in agreement with estimates of the Rossby radius of deformation for the first baroclinic mode (cf W Kraus et al 1990, JGR 95, 18049; P Y Le Traon various)

Thadathil P, Saran AK, Gopalakrishna VV, et al., 2002: XBT fall rate in waters of extreme temperature: A case study in the Antarctic Ocean. Journal of Atmospheric and Oceanic Technology 19 (3): 391-396

UNESCO, 1997. SOOPIP-I recommendations to OOPC, SMC and IGOSS/GOOS. First Session of the Joint IOC-WMO IGOSS Ship-of-Opportunity Programme Implementation Panel: Annex VI, Cape Town, South Africa, 16-18 April 1997.

Weaver,A.T., Vialard,J., Anderson,D.L.T. and Delecluse,P. 2002: Three- and four-dimensional variational assimilation with a general circulation model of the tropical Pacific Ocean. ECMWF Technical Memorandum 365

Wang,G., Kleeman,R., Smith,N. and Tseitkin,F. 2002: The BMRC coupled general circulation model ENSO forecast system. MWR 130, 975-991

Appendix 1: Error estimates

Level	Depth	ObErr Mean	BkErr Mean	BkErr Max	ObErr Mean	BkErr Mean	BkErr Max
1	5.0	0.777	0.997	1.998	0.141	0.214	0.400
2	15.0	0.811	1.042	2.091	0.148	0.231	0.395
3	25.0	0.881	1.133	2.026	0.155	0.241	0.385
4	35.0	0.947	1.242	2.172	0.158	0.253	0.395
5	45.0	0.986	1.232	2.086	0.163	0.239	0.378
6	55.0	0.994	1.275	2.241	0.165	0.252	0.402
7	65.0	0.991	1.172	2.068	0.167	0.227	0.386
8	75.0	0.981	1.286	2.172	0.168	0.260	0.403
9	85.0	0.969	1.173	2.063	0.169	0.232	0.386
10	95.0	0.958	1.239	2.072	0.170	0.251	0.394
11	105.0	0.937	1.212	2.021	0.170	0.256	0.400
12	115.0	0.916	1.113	1.910	0.170	0.230	0.400
13	125.0	0.893	1.058	1.845	0.170	0.230	0.397
14	135.6	0.869	1.078	1.813	0.170	0.237	0.411
15	148.5	0.842	1.091	1.787	0.170	0.262	0.422
16	165.9	0.806	0.993	1.763	0.170	0.232	0.437
17	190.2	0.754	1.010	1.772	0.170	0.265	0.451
18	223.6	0.704	0.881	1.787	0.170	0.234	0.461
19	268.1	0.662	0.913	1.842	0.168	0.263	0.480
20	325.8	0.611	0.880	1.863	0.167	0.264	0.504
21	398.6	0.558	0.837	1.804	0.165	0.265	0.491
22	488.2	0.498	0.771	1.676	0.162	0.258	0.500
23	596.4	0.439	0.692	1.420	0.156	0.249	0.479
24	724.5	0.387	0.588	1.220	0.145	0.238	0.438
25	873.8	0.361	0.569	1.192	0.127	0.245	0.402
26	1045.6	0.331	0.496	1.139	0.110	0.236	0.379
27	1240.7	0.303	0.465	0.878	0.110	0.228	0.368
28	1459.9	0.269	0.440	0.765	0.110	0.219	0.356
29	1703.7	0.262	0.351	0.766	0.110	0.180	0.350
30	1972.4	0.262	0.445	0.766	0.110	0.215	0.346
31	2265.8	0.262	0.319	0.766	0.110	0.153	0.342
32	2581.7	0.262	0.262	0.262	0.110	0.120	0.162
33	2914.9	0.262	0.262	0.262	0.110	0.116	0.139
34	3257.5	0.262	0.262	0.262	0.110	0.113	0.136
35	3602.5	0.262	0.262	0.262	0.110	0.112	0.137
36	3947.5	0.262	0.262	0.262	0.110	0.111	0.137
37	4292.5	0.262	0.262	0.262	0.110	0.111	0.135
38	4637.5	0.262	0.262	0.262	0.110	0.112	0.137
39	4982.5	0.262	0.262	0.262	0.110	0.110	0.111
40	5327.5	0.000	0.000	0.000	0.000	0.000	0.000

The values are those at the TESAC positions in January 1984,
the levels are the GloSea model levels - used for ENACT backgrounds.
They are the estimates of rms error in degrees and PSU.
The ObErr values do not vary (geographically or with observation type
- the latter because representivity error is assumed to dominate
measurement error).
BkErr values do vary, minimum values are the same as the ObErr column
(the values quoted here do not include the extra scaling in the tropics),
they are rms values for the combined mesoscale and synoptic errors.

Appendix 2: Technical issues

The ocean specific code has been written assuming the use of shared memory (although some of the routines that are common with the atmospheric processing work with distributed memory - to cope with the larger atmospheric data volumes). This makes the code (notably the track check, superobbing and duplicate check) simpler and easier to follow. The 'station list' code used by the atmospheric processing to specify observation errors, initial PGEs and any prior rejections is not used. Observation error estimates (constant globally) are supplied in the background error file, initial PGEs are set in the code and there are currently no prior rejections.

Temperature and potential temperature values are stored in Celsius - this contrasts with the atmospheric processing which uses Kelvin. The file containing background error estimates differs from the atmospheric error estimates by storing variances rather than standard deviations. It also contains separate geographically varying estimates for the synoptic and mesoscale components of the background error and their length scales. It is assumed that both temperature and salinity will always be extracted and processed (even if the latter only has missing values). There is currently no provision for reports which contain pressures but not depths - they are simply rejected - although this would not be difficult to change. Note that in WOD01 such reports have had depth calculated.