Norwegian Meteorological Institute

MET report

no. 26/2013
Oceanography

Assimilation of HF radar total current vectors in a realistic ROMS 4DVAR application

Ann Kristin Sperrevik, Kai H. Christensen & Johannes Röhrs

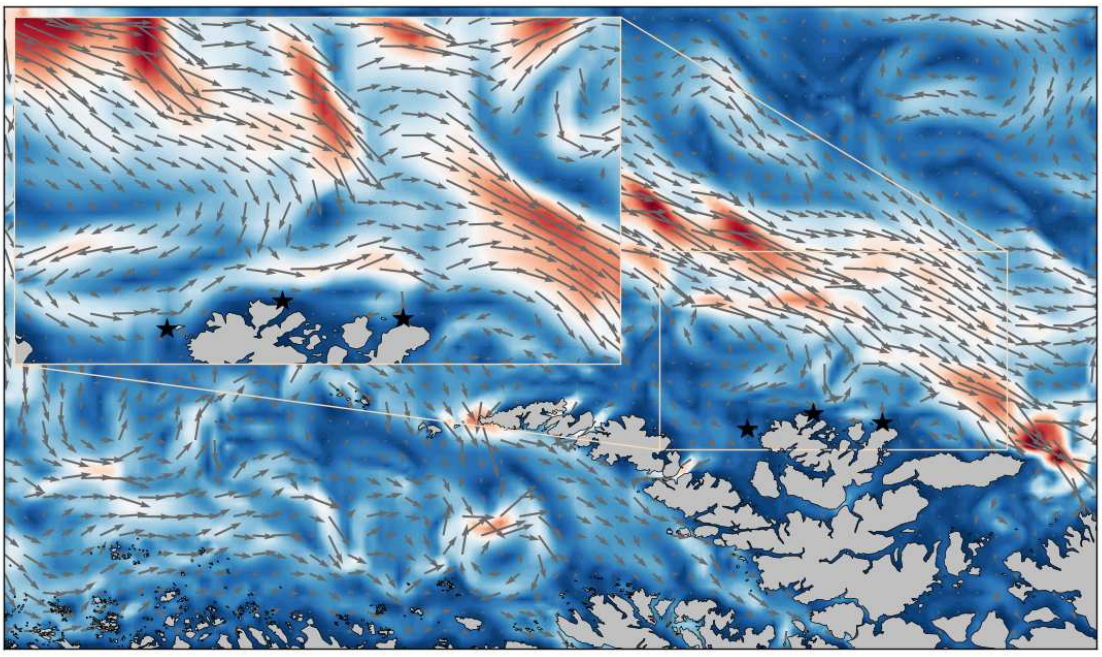

Norwegian Meteorological Institute

MET report

Title	Date
Assimilation of HF radar total current vectors in a realistic ROMS 4DVAR application	December 5, 2013
Section	**Report no.**
Oceanography	26
Author(s)	**Classification**
Ann Kristin Sperrvik, Kai H. Christensen & Johannes Röhrs	● Free ○ Restricted
Client(s)	**Client's reference**
Norwegian Clean Seas Association for Operating Companies (NOFO)	Petter Reed

Abstract

This report summarises the work done in acticity 1, task 2 and 3, of the ENI/NOFO HF radar project. The report describes a realistic model setup for Lofoten/Vesterålen, in which HF radar total vectors and CTD salt and temperature profiles are assimilated using the standard incremental 4DVAR scheme in ROMS. We first examine the impact of the different observation types on the forecast skill of the model. ROMS I4DVAR is run over fixed length (24 h) assimilation window with different set of observations as input, and the updated initial conditions are used to initialize a forecast simulation. We then proceed to examine how the model forecast skill is affected when data are assimilated over different number of days. Again, we initialize a forecast simulation from the updated initial conditions. Drifter trajectories, ADCP data and a high resolution satellite SST product are used to validate the model results. The main findings are:

- Assimilation of HF radar total vectors yields better results than assimilating CTD hydrography,
- assimilating both observation types yields the best results,
- assimilation of HF radar currents reduces SST bias through advection of water masses,
- longer periods of data assimilation is most beneficial, but unconstrained upstream conditions pose a limitation for model improvement.

Keywords

Oceanography, data assimilation, ROMS, 4DVAR, HF radar

Disciplinary signature

Responsible signature

Meteorologisk institutt
Meteorological Institute
Org.no 97127404
post@met.no

Oslo
P.O. Box 43 Blindern
0313 Oslo, Norway
T. +47 22 96 30 00

Bergen
Allégaten 70
5007 Bergen, Norway
T. +47 55 23 66 00

Tromsø
P.O. Box 6314
9293 Tromsø, Norway
T. +47 77 62 13 00

www.met.no
www.yr.no

Contents

1 Introduction

As part of a project on the use of HF radars for monitoring and data assimilation, an intensive field campaign was launched in the spring of 2013 in the Lofoten/Vesterålen area. Three mobile HF radar stations were deployed in Vesterålen, ovelooking the continental shelf. At the same time, the R/V Johan Hjort passed through the area on the annual cod stock assessment cruise, and hydrographic data were collected. The main aim of the work presented here has been to evaluate the impact of assimilating HF radar data in the operational ocean model used at the Norwegian Meteorological Institute, and to compare the results with the impact of assimilating more traditional hydrographic data from CTD. Data from an ADCP rig deployed before the cod stock assessment cruise, as well as from surface drifters deployed from R/V Johan Hjort, have been used as independent sources of velocity and trajectory information.

Section 2 of this report describes the general model setup, while Section 3 describes the initial experiments that were made for tuning the ocean model's data assimilation system. Section 4 contains a comparison of the impact of HF radar data versus CTD data, and also results from experiments where short versus long timeseries of data have been used. Finally, Section 5 contains a discussion and concluding remarks.

2 Model configuration

The basis for the realistic model simulations is the NorKyst800 (Albretsen et al. [2011]) version of the Regional Ocean Modeling System (ROMS, `www.myroms.org`), which is run operationally by the Norwegian Meteorological Institute (MET Norway). NorKyst800 covers the coast of mainland Norway and has a horizontal resolution of 800 meters. Ocean data assimilation at such high resolution is associated with extensive challenges, as outlined in Christensen [2013], hence we have therefore chosen to use a coarser horizontal resolution with 100x170 horizontal grid cells that are 2.4×2.4 km^2.

ROMS is a free-surface, hydrostatic, primitive equation model, and consists of a nonlinear (NL-ROMS), tangent linear (TLM) and an adjoint (ADM) model. There are seven dependent variables in ROMS: $(u, v, T, S, \zeta, \bar{u}, \bar{v})$, representing horizontal velocity in easterly direction, horizontal velocity in northerly direction, potential temperature, salinity, sea surface height, vertically averaged velocity in easterly

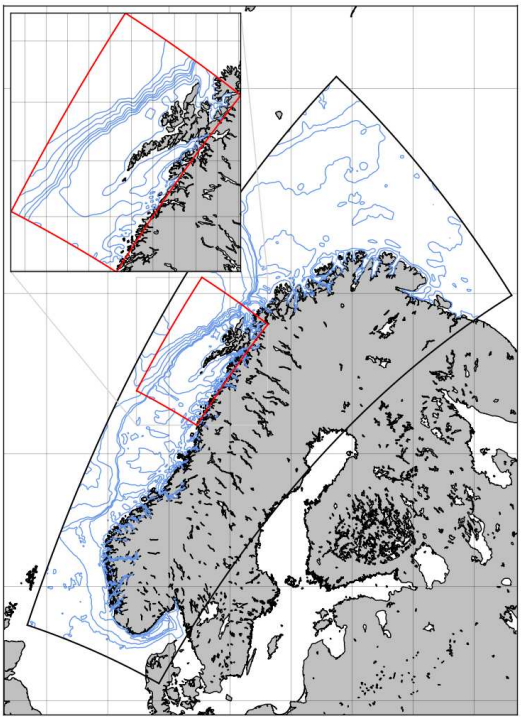

Figure 1: The model domain. The full NorKyst800 domain is shown with black border, while our subdomain Vesterålen-2.4 is shown with red border.

direction, vertically averaged velocity in northerly direction, respectively. In the ROMS code and ROMS output these variables are referred to as (u,v, temp, salt, zeta, ubar, vbar).

The simulations described in this study uses Chapman [1985] and Flather [1976] open boundary conditions for sea level elevation and barotropic currents. For tracers and baroclinc velocities, boundary conditions as described in Marchesiello et al. [2001] are used. During the assimilation clamped boundary conditions with a sponge layer is applied.

2.1 Model grid

The model domain is a subset of NorKyst800 centered around the Lofoten and Vesterålen archipelago. The bathymetry of the domain was taken from NorKyst800. Not all features present in the original bathymetry can be represented in a coarser resolution, hence a smoothing algorithm was applied to convert the data from 0.8 to 2.4 km resolution. The model has 35 terrain-following vertical layers. The model domain used in this study will hereafter be referred to as “Vester”å“len-2.4”.

2.2 Forcing and initialization of the model

The lateral boundary conditions in the Vesterålen-2.4 experiments are retrived from the operational runs of the NorKyst800 at MET Norway. Fields of sea surface elevation, temperature, salinity and currents have been collected at three hour intervals for the area of interest. Before the boundary condition fields are applied, they are subject to a smoothing algorithm, which averages the values within a box consiting of 3×3 grid points. This is done to remove fine scale features present in the NorKyst800 fields, which will be unresolved in the coarser Vesterålen-2.4. After the smoothing algorithm the fields are interpolated to the coarser grid.

River discharge data are taken from the NorKyst-800 system. The river runoff in this system is based on modelled estimates of discharge from NVE (Norges vassdrags- og energidirektorat). As there are no near real-time data available, the discharge data used in this experiment are daily climatological values based on data from the period 1962-2009.

The Vesterålen-2.4 applies the same atmospheric forcing as the operational runs of NorKyst800 at MET Norway. The forcing consists of air temperature and humidity 2 m above ground, 10 m winds as well as mean sea level air pressure, cloud cover and precipitation. The source of the data is MET Norway's operational weather forecast from UM4km.

Atmospheric forcing and boundary conditions were collected from the operational suites starting from February 15th and ending April 15th, 2013. A model simulation was initialized from the smoothed NorKyst800 fields and run throughout this period. All initial conditions for the different assimilation experiments stem from this spin-up run.

3 Configuration of the data assimilation system

ROMS contains a four dimensional variational data assimilation system that accepts observations in any of the dependent variables (ROMS-4DVAR). The 4DVAR system does not yet allow assimilation of radial currents from HF radars. In our experiments, data assimilation is only considered for interior model points and no adjustment of the surface/bottom fluxes or boundary conditions are made. We use incremental, strong-constraint 4DVAR. The data assimilation system and the associated parameters are described more in detail in Christensen [2013].

3.1 Background standard deviations

Circulation statistics (mean and standard deviations) are used to estimate model errors, and these should ideally be calculated from a free simulation of the application in question. Since we did not have any such simulations with Vesterålen-2.4, we have used data obtained in the BIOWAVE project as a substitute. In the BIOWAVE project a model application covering the same geographical domain, but with the original 0.8 km resolution, was run for a period starting in October 2010 and throughout March 2011. Daily fields from this run have been used to estimate the circulation statistics of Vesterålen-2.4. The same smoothing algorithm as described in Section 2.2 was used to remove fine scale features that cannot be resolved by the coarser grid. In addition, statistics from the free simulation initialized from NorKyst800 was calculated and used as a supplement to the BIOWAVE statistics.

3.2 Observations

The observations used for data assimilation in this project consists of CTD profiles of temperature and salinity, taken during the cruise with R/V Johan Hjort during March 2013 and total current vectors from three SeaSonde HF radars from CODAR stationed in Vesterålen [Christensen et al., 2013]. Before the observations were assimilated, the data were screened using tools from the ROMS developers to ensure optimal use in our ROMS-4DVAR application. Due to discontinuities in the time record of HF current data, no filtering of the tidal signal has been carried out. The ideal method for assimilating such observations would be to first remove the tidal signal of the observations, and then add the tidal signal of the model application in question (See Zhang et al. [2010] for a more in-depth description).

3.3 Tuning of ROMS-4DVAR

The findings in Christensen [2013] were used as a starting point for tuning the 4DVAR parameters. As in Christensen [2013] we examine the assimilation window length, number of inner/outer loops and the horizontal background error correlation scales.

The procedure is as follows: Each experiment results in an analysis, that is, a solution of ROMS-4DVAR. This analysis is used to initalize a 5-day simulation with

Exp.no	Corr. scale	Inner/outer	Assim window
1	10 km	8/2	72 h
2	10 km	8/2	48 h
3	10 km	8/2	24 h
4	10 km	5/4	24 h
5	10 km	6/3	24 h
6	10 km	20/1	24 h
7	10 km	12/2	24 h
8	**10 km**	**10/2**	**24 h**
9	5 km	10/2	24 h
10	15 km	10/2	24 h

Table 1: The numerical experiments used for tuning of the data assimilation system in the realistic model. From left to right the columns denote horizontal error correlation scale, number of inner and outer loops, and the length of data assimilation window. Bold face indicate the experiment with best skill.

NL-ROMS. During this last simulation, numerical floats are released in the positions occupied by the surface drifters deployed during the field campaign [Christensen et al., 2013]. Numerical floats are released every three hours, at the position held by the real drifters at that time. The depth of the floats is fixed, and is set to 65 cm below the sea surface for the iSLDMB drifters, and 10 cm for the iSphere drifters.

To investigate the horizontal dispersion in the model, we have also considered an ensemble of numerical floats. Each time a float is released, we release additional ten floats. The initial positions of these floats are perturbed in a random way by up to 1.5 grid points (i.e., about 3.5 km). The ensemble of modelled trajectories provides an idea of the variability in the current fields.

The numerical trajectories are compared with the observed trajectories following the methods described in Appendix A. The overall agreement is used to evaluate the performance of the ROMS-4DVAR configuration. We focus on the iSLDMB drifters in this evaluation, as the iSphere drifters are affected by the wave field, which is not accounted for in ROMS [Röhrs et al., 2012]. The results are shown in Table 1. Based on this outcome, the rest of the experiments use a horizontal error correlation scale of 10 km, 10 inner and 2 outer loops, and an assimilation window length of 24 hours.

4 Results

The experiments discussed in this section follow the same procedure as described in Section 3.3.

4.1 Impact of different observation data sets

The first series of experiments are focused on assessing what impact different observation sets have on the analysis and subsequent forecast. The ROMS-4DVAR system is run for one cycle with three different sets of observations. The first experiment assimilates only HF radar currents, the second only CTD hydrography, while the third assimilates all data (denoted HF, CTD and ALL in the following). The results from these three experiments are compared with results from a free model simulation in which no assimilation is done (denoted CTRL). As the simulations are started from the beginning of the assimilation window, the first 24 hours of these runs should be considered as analyses. The results from the analysis period are representative for the performance of the system when used in "reanalysis mode", i.e. when statistical measures are sought based on historical data.

We validate the results from the analysis part and the forecast part of the simulations separately. When considering the results, we must keep in mind that we have only assimilated data from a time period of 24 hours and that CTD observations are sparse compared with the number of HF radar current observations. Also, as these experiments are carried out for a specific time period, the results are influenced by the weather conditions and does not necessarily represent the variations in predictability of different flow regimes.

4.1.1 Drifter velocities

Using the methods of validation described in Appendix A, we evaluate the predicted velocities of the numerical floats. The results are given as a value between 1 and -1, where a value of 1 means perfect correlation in both speed and direction, a value of 0 means no correlation, while a value of -1 means that the data are anti-correlated, i.e. having opposite direction but the same speeds. The results of the vector correlation comparison are shown in Figure 2 and summarised in Table 2. Only the iSLDMB surface drifter data and equivalent numerical trajectories have been included.

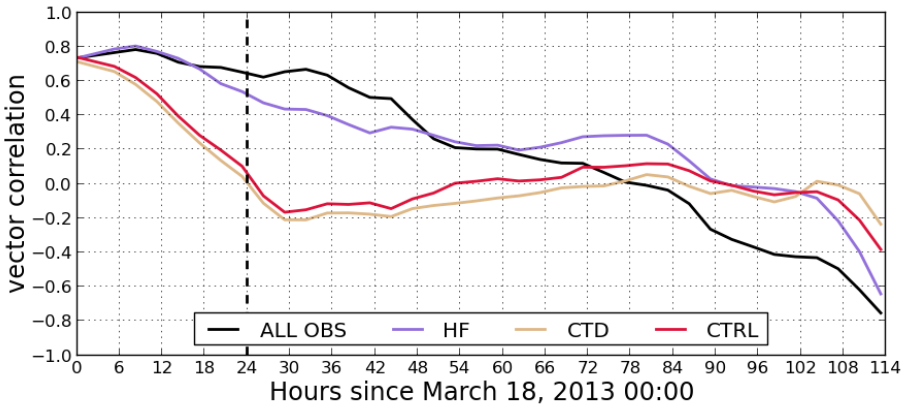

Figure 2: Vector correlation as a function of time. The vertical line indicates the shift from analysis to forecast. A value of 1 means perfect correlation in both speed and direction, a value of 0 means no correlation, while a value of -1 means that the velocities are anticorrelated.

The vector correlation of the CTRL decreases rapidly, and as the simulation enters into the forecast, modelled and observed drifter velocities are uncorrelated. The CTD experiment follows the control run closely. In fact, the correlation coefficient has a lower value than the CTRL throughout almost the entire simulation. This indicates that the number of CTD observations in the experiments are too sparse to constrain the circulation, and therefore have only minor impact on the subsequent model predictions. Assimilation of HF radar currents on the other hand significantly improves the drifter velocities during both the analysis and forecast part of the model predictions. The velocities remain correlated for two days into the forecast. An interesting fact is that the addition of CTD hydrography to the assimilated data set further improves the model predictions during the first part of the forecast, which is in contrast to the deterimental effect of only using CTD data. A likely explanation is that the addition of CTD observations acts as an complementary constraint of the circulation that is set by the HF radar observations. This shows how important it can be to constrain all state variables in order to achieve the best possible model predictions.

The significance of the improvement has been tested comparing ALL and CTRL

Observation set	Analysis		Forecast	
	Average	Median	Average	Median
HF	0.70	0.73	0.31	0.29
CTD	0.40	0.42	-0.12	-0.12
ALL	0.72	0.72	0.37	0.32
CTRL	0.45	0.46	-0.05	-0.06

Table 2: Average and median vector correlation between observed and predicted iSLDMB surface drifter trajectories during analysis window and first 48 h of forecast.

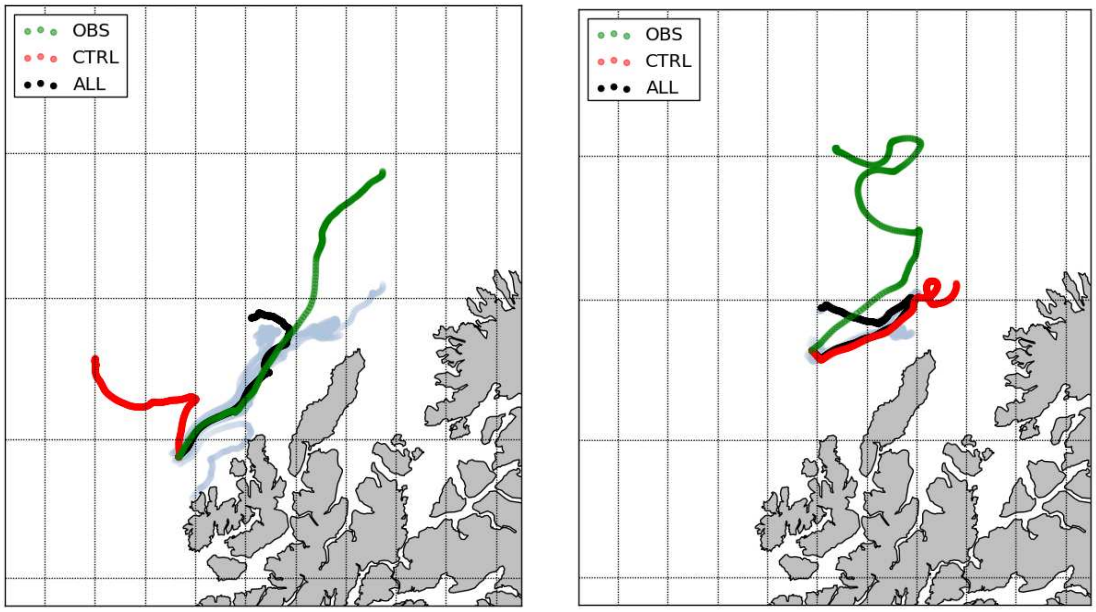

(a) Drifter No. 4 (b) Drifter No. 5

Figure 3: The panels show the trajectories of two different drifters as observed (green), predicted by CTRL (red) and predicted by ALL (black). In addition, the grey tracks show the pathways of the floats with perturbed initial positions. The numerical floats were released at the start of the analysis.

using a Wilcoxon rank-sum test. The improvement in current direction is statistically significant, while not so for the speeds. Our interpretation is that the major benefit of assimilating HF radar currents in these experiments is the correction of the current direction, e.g. adjustments in the positions of eddies and the coastal current. It should be pointed out that this test does not imply that no improvement in speed is obtained (Figure 2 indicates that the results are better), but that the impact cannot be statistically verified with the limited observational data available. As discussed more in detail in Sec. 4.2.2, we also know that the model is not as energetic as the observations, which means that we should also seek to reduce biases and errors by improving the physics in the model and the boundary and surface forcing.

4.1.2 Drifter trajectories

The impact of data assimilation is also evaluated by comparing observed and predicted trajectories. For this comparison we use the normalised cumulative Lagrangian separation (see Appendix A). Basically we compare not only the end points of the observed and modelled trajectories, but also the entire history of the drifter trajectories.

Observation set	Analysis		Forecast	
	ISLDMB	ISPHERE	ISLDMB	ISPHERE
HF	0.44	0.21	0.30	0.26
CTD	0.40	0.15	0.16	0.19
ALL	0.43	0.21	0.33	0.22
CTRL	0.42	0.16	0.18	0.21

Table 3: Skill score during analysis and forecast when comparing with the two different surface drifter types.

The results shown here are obtained by calculating the skill score for predicted trajectories of floats released at the start of the assimilation window for a period extending throughout the length of the window, which is 24 hours. The skill for the forecast part is found by evaluating the floats released at the very start of the forecast and their trajectories during the next 48 hours. Here, values close to one indicate good skill, while values close to zero indicate no skill. Results are shown in Table 3 for both types of surface drifters.

Assimilation does improve predictions of drifter trajectories, although the impact is more limited compared to the drifter velocities. The results show that the skill improves when we consider periods longer than a day. In these cases, data assimilation seems to constrain the ocean circuation in such a way that the predicted trajectories does not stray as far away from the observed paths as they do in a free simulation. This is likely a consequence of the CTRL not being energetic enough, as further discussed in Sec. 4.2.2. Two examples of modelled versus observed trajectories are shown in Figure 3.

4.1.3 Sea Surface Temperature

The sea surface temperature (SST) is well observed by satellite, which gives us an independent variable to validate against. Here, we have use an OSI-SAF SST product with 1.5 km resolution[1]. In addition, the iSphere buoys released during the field campaign were equipped with thermometers measuring the sea surface temperature along the drifter paths. As the model results in this section are based on the assimilation of observations within a 24 hour period, we do not excpect model improvements in regions far away from the observations. We therefore compare satellite SST with model results only in a subarea of the model domain, centered around the HF radar coverage area, as comparing with the entire model domain would possibly obscure the impact of the assimilation (the subdomain is

[1] `http://osisaf.met.no`

Observation set	Analysis				Forecast			
	ISPHERE		OSI-SAF		ISPHERE		OSI-SAF	
	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS
HF	1.10	-0.93	1.72	1.57	1.73	-1.62	1.30	0.70
CTD	1.05	-0.90	1.81	1.66	1.51	-1.43	1.35	0.74
ALL	1.10	-0.93	1.76	1.60	1.69	-1.58	1.35	0.74
CTRL	1.21	-1.12	1.77	1.54	1.56	-1.46	1.38	0.57

Table 4: SST error statistics. Root mean square error (RMSE) and bias when comparing model results with ISPHERE buoy SST and OSI-SAF SST.

outlined in Fig. 8). The iSphere buoys on the other hand, were released in the area covered by the HF radars. Therefore we do not need to limit the area from which to pick data for comparison.

The error statistics are summarised in Table 4. We find that assimilation in all cases seem to reduce the root mean square error (RMSE). When comparing iSphere buoy SST with the model results, we also find bias reduction, whereas the bias seems to increase when comparing with satellite SST. An interesting fact is that the bias is negative when comparing with iSphere, while it is positive when comparing with satellite data. The most likely reason for this discrepancy is the distribution of observations: the model has a warm bias in open waters, while it is too cold close to shore. It should also be noted that due to cloudy conditions during the validiation period, satellite observations are sparse within the HF radar coverage area.

4.1.4 ADCP measurements

During the field campaign there was also installed an ADCP rig within the HF radar coverage area, for a detailed description see Christensen et al. [2013]. Two separate instruments measured the upper ocean currents: one with high vertical resolution (0.5 m) in the upper 8 meters, and another with coarser resolution (1 m) in the upper 41 meters. These two instruments will be referred to as the “upper ADCP” and “lower ADCP”, respectively. Only data in the range 8-41 meters depth will be used from the lower ADCP.

We compare the results from the different model simulations with the speed and direction measured by the ADCPs. We average the ADCP currents over the depth column in their specific range (i.e., 0-8 m and 8-41 m depth), and convert to speed and direction. Modelled speed and direction are calculated in the same manner.

Figure 4 shows that the ADCPs generally measure a current with a NNE heading.

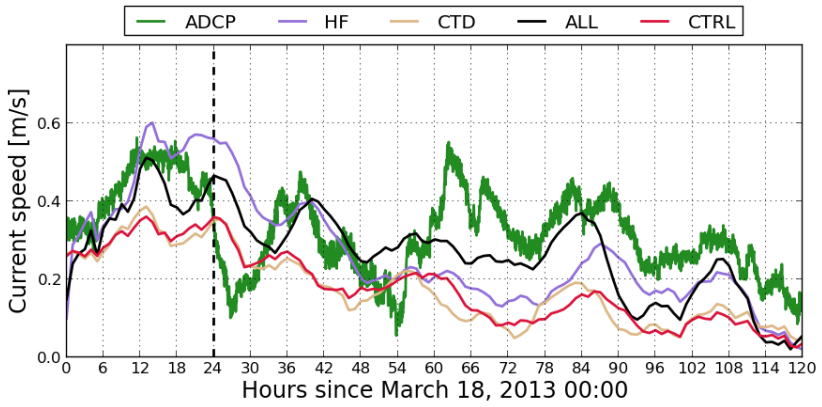

(a) Speed

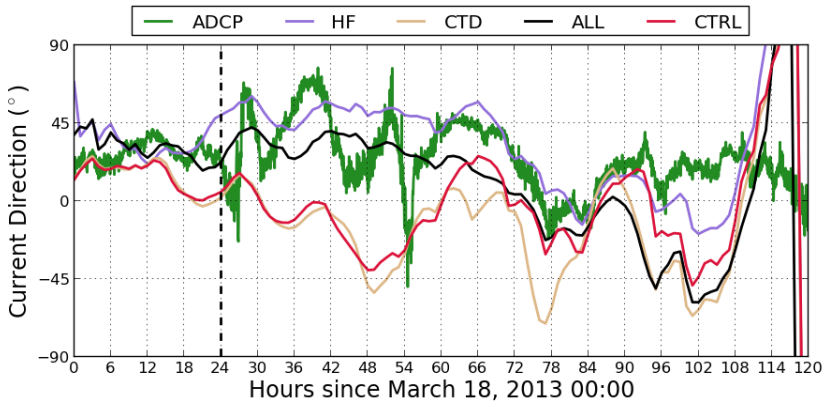

(b) Direction

Figure 4: The figure shows time series of vertical averaged current over the upper 8 meters of the water column (4(a)) speed and (4(b)) direction during the analysis and forecast periods.

The CTRL run on the other hand predicts a current towards NNW during the period in question. In addition to the discrepancy in direction, the current speed predicted by the control run is too weak during the entire simulation. Assimilation of CTD hydrography does not help the predictions, it seems to even worsen the predictions of current direction. Assimilation of HF radar current on the other hand, has a remarkable effect on both speed and direction. The effect of the assimilation is maintained for several days into the forecast. The combination of HF radar currents and CTD further improves the model during the analysis period, but the effect wears off earlier in the forecast compared to the case when only HF was assimilated.

To further investigate the effect of assimilation on the simulated currents, we also look at the distribution of current strength with regard to the current direction. Figure 5 shows current roses for the vertically averaged current in the upper 8 meters of the water column for the upper ADCP, CTRL, and ALL. During both analysis and forecast, the CTRL run has the main current direction basically correct, but

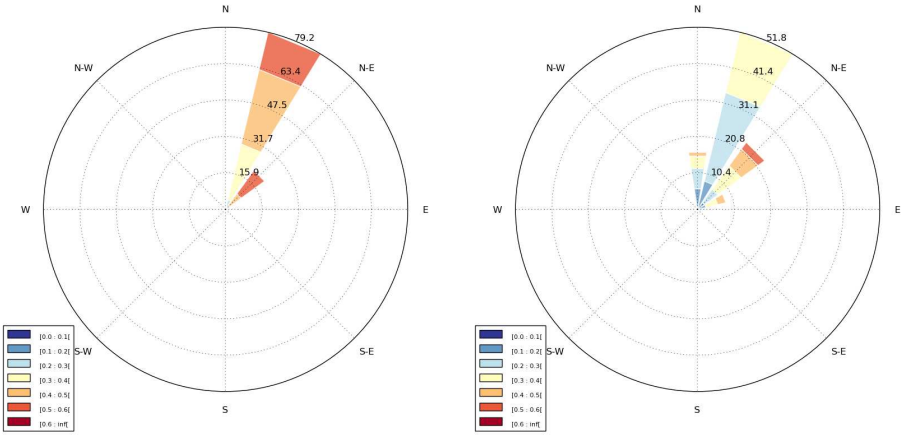

(a) ADCP during analysis

(b) ADCP during forecast

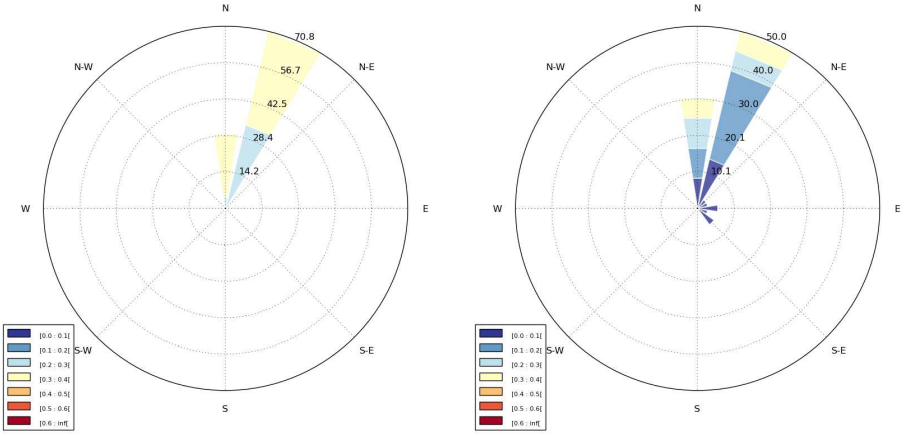

(c) CTRL during analysis

(d) CTRL during forecast

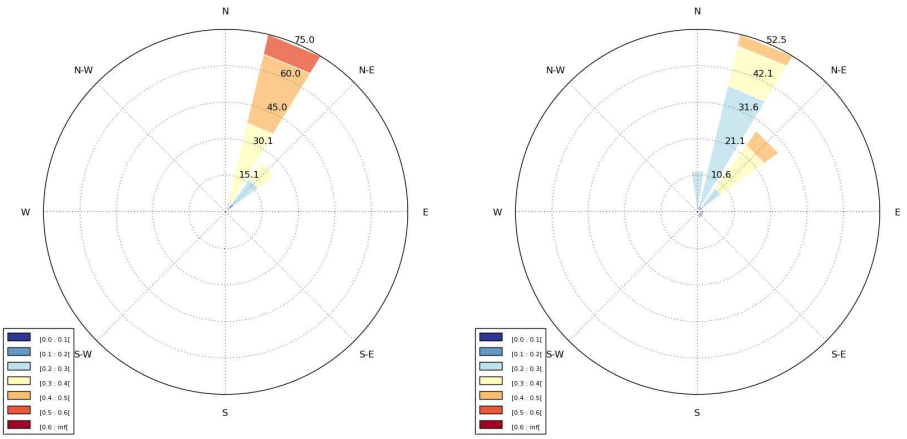

(e) ALL during analysis

(f) ALL during forecast

Figure 5: The left column shows speed and direction distributions during the analysis from ADCP 5(a), CTRL 5(c) and ALL 5(e). The right column shows the corresponding distributions during the forecast period. Blue colors indicate current speed in the range $0 - 0.3$ m/s, yellow and orange colors $0.3 - 0.5$ m/s and red colors current speeds stronger than 0.5 m/s.

there is a too strong northward signal. We again see that the predicted current speed is too weak, and there is little correlation in the distribution of speed with

regard to direction. After assimilation, the model performs better, both in terms of direction of the current as well as the distribution of current speeds. Error statistics for speed and direction are summarised in Table 5 and 6. With the exception of the CTD run there RMSE and bias are reduced in both in speed and direction throughout upper 8 meters. ROMS-4DVAR thus appear to efficiently spread the information from the surface observations downwards in the water column.

Observation set	Analysis				Forecast			
	Upper		Lower		Upper		Lower	
	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS
HF	0.08	0.01	0.10	0.02	0.15	-0.06	0.15	-0.08
CTD	0.14	-0.13	0.10	-0.10	0.19	-0.15	0.19	-0.16
ALL	0.08	-0.06	0.07	-0.05	0.12	-0.04	0.14	-0.08
CTRL	0.14	-0.13	0.10	-0.09	0.18	-0.14	0.18	-0.15

Table 5: ADCP current speed error statistics. Root mean square error (RMSE) and bias when comparing model results with ADCP current speed. All numbers are given in m/s.

Observation set	Analysis				Forecast			
	Upper		Lower		Upper		Lower	
	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS
HF	17.2	7.9	10.9	-3.0	44.4	13.1	22.2	5.8
CTD	17.2	-12.6	21.4	-19.6	57.7	-32.7	48.0	-38.3
ALL	12.8	3.9	9.0	-4.1	45.3	-16.2	32.5	-19.1
CTRL	16.5	-12.7	21.2	-19.8	48.3	-26.3	40.1	-29.8

Table 6: ADCP current direction error statistics. Root mean square error (RMSE) and bias when comparing model results with ADCP current direction. All numbers are in degrees.

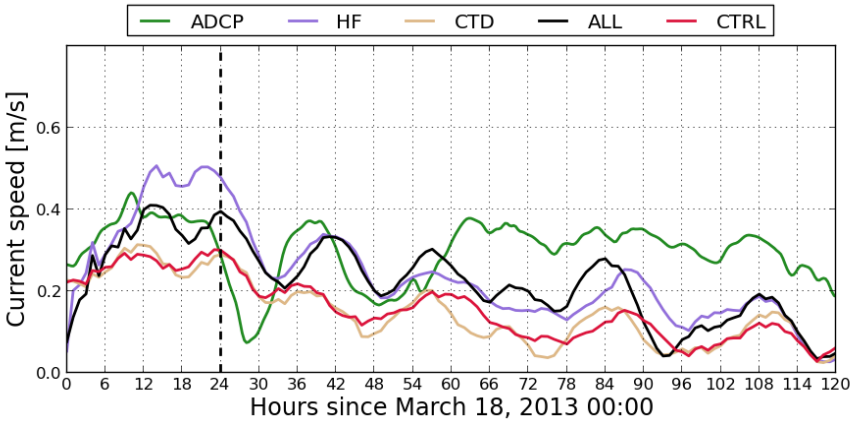

(a) Speed

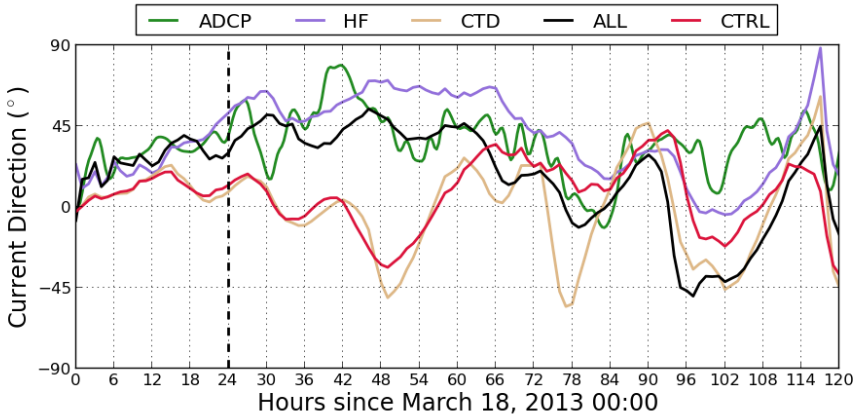

(b) Direction

Figure 6: The panels shows time series of vertically averaged current in the water column between 8 and 41 meters depth. The upper panel shows the speed and the lower panel shows the direction during the analysis and forecast periods.

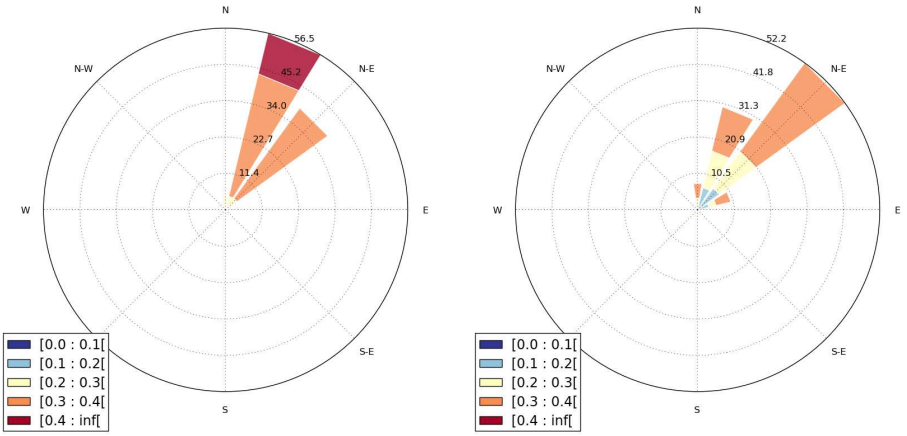

(a) ADCP during analysis

(b) ADCP during forecast

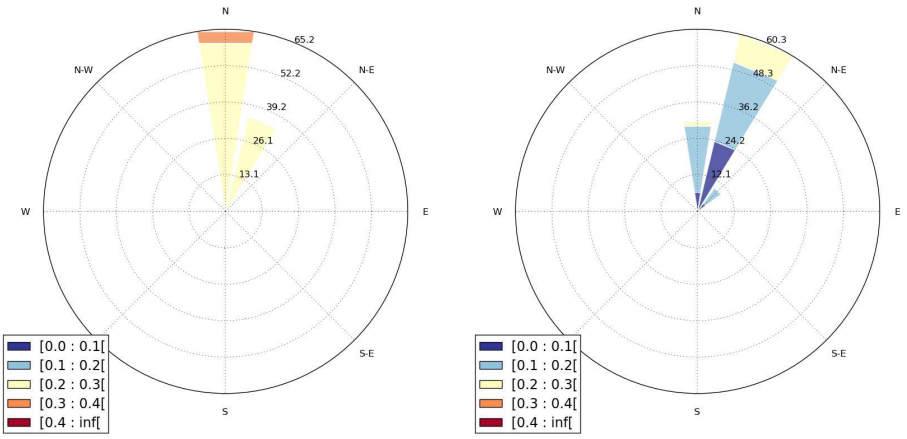

(c) CTRL during analysis

(d) CTRL during forecast

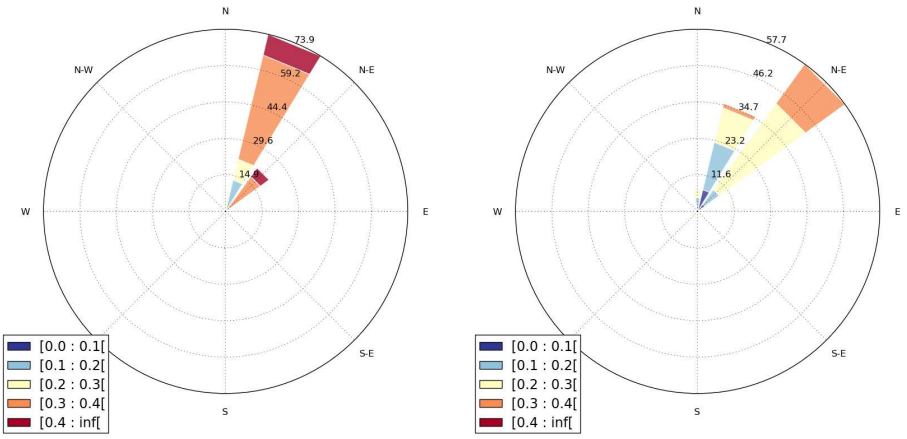

(e) ALL during analysis

(f) ALL during forecast

Figure 7: The left column shows speed and direction distributions during the analysis from the lower ADCP 7(a), CTRL 7(c) and ALL 7(e). The right column shows the corresponding distributions during the forecast period. Blue colors indicate current speed in the range $0-0.2$ m/s, yellow and orange colors $0.2-0.4$ m/s and red colors current speeds stronger than 0.4 m/s.

4.2 Impact of assimilation over a longer time period

After the initial experiments, the next step is to investigate what the impact is when we do not only run our 4DVAR system once, but a multiple of days in a row. The procedure is as follows: As before we run ROMS-4DVAR, and integrate NL-ROMS to the end of the assimilation window. This gives us a best estimate of the state at the end of the window. This estimate is then used as the initial condition for the next assimilation run. This procedure is then repeated a given number of times. Lastly, we run NL-ROMS for 5 days to obtain a forecast.

As we want to compare the different experiments, all NL-ROMS simulations must cover the same time period. We therefore initiate ROMS-4DVAR at four different dates. The first experiment is initiated 20 days ahead of the NL-ROMS simulation, the second 15 days, the third 10 days and the fourth 5 days before. These experiments will in the following be referred to as 20D, 15D, 10D and 5D, respectively. Additionally, NL-ROMS is run without any assimilation, and this simulation is our control run (CTRL). Both CTD hydrography and HF radar total vectors have been assimilated in all these experiments.

As the three HF radar antennaes started to submit data on March 9, our longest 4DVAR experiment is initiated this day. Twenty days later, when our experiments end, there are no longer any surface drifters left in the model domain. This leaves us with ADCP measurements and satellite SST to validate the results against.

4.2.1 Sea Surface Temperature

As in Section 4.1.3, we compare our results with the high resolution OSI-SAF SST product. Figure 8(a) shows that we have a warm bias in our model. Around the Lofoten/Vesterålen archipelago the model performs better, but close to the shore the model is too cold, especially around AndÃ¸ya. Figure 8(b) shows the temperature differences during the same period of time, but the model results are now taken from the 20D experiment. The warm bias is still clearly visible, but in the northwestern part of the domain, i.e. downstream from the HF radars, there is a noticeable improvement. Also in the southern part of the domain we see improvements. The cold bias near AndÃ¸ya has almost disappeared in this simulation.

The differences are quantified in Table 7. Both RMSE and bias are significantly reduced. An interesting fact is that the bias seems to be most reduced in the 5D and 10D experiments. However, examining the results more in detail reveals that

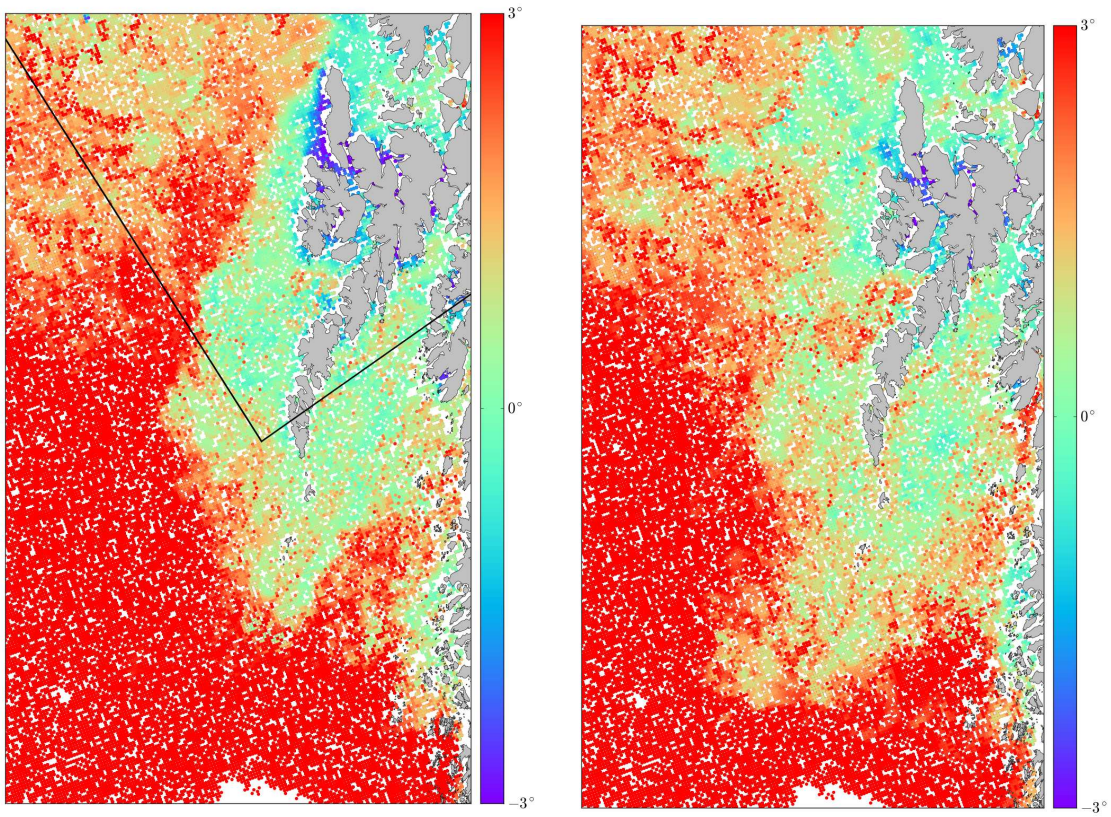

(a) CTRL (b) 20D

Figure 8: The panels show SST difference between model and satellite observations for 8(a) the control run and 8(b) the 20 days long simulation. Red color indicates warm bias in the model, and blue color indicates cold bias. Areas with green colors match the observations well. The subdomain used in the analysis is the area north of the black line. The color scale is $\pm$ 3 degrees.

the area of cold bias still persists in these runs, and thus influences the mean bias. Examining the mean warm bias and the mean cold bias separately for all runs, shows that both cold and warm biases are most reduced in the longest data assimilation run (20D), and gradually decreases with the number of days of data assimilation.

The probabilty density functions (PDFs) of the model temperature errors are shown in Figure 9 for both the full domain and the subdomain centered around Vesterålen. The warm bias is evident from the right skewed curves. However, we see that the PDFs based on the assimilation runs are slightly less skewed, and also narrower than the PDF of the control run. As temperature observations used for assimilation are limited in number, model improvement is likely to be due to improved circulation. Adding satellite SST to the assimilated observations might thus have great potential to further improve the model predictions.

Observation set	Full domain		Subdomain	
	RMSE	BIAS	RMSE	BIAS
CTRL	3.09	2.29	1.51	0.54
5D	2.69	1.91	1.32	0.24
10D	2.68	1.94	1.38	0.36
15D	2.69	2.03	1.32	0.59
20D	2.67	2.01	1.31	0.58

Table 7: SST error statistics. Root mean square error (RMSE) and bias when comparing model results with OSI-SAF SST. Statistics for the full model domain are shown in column 2 and 3, while statistics for the subdomain are shown in column 4 and 5.

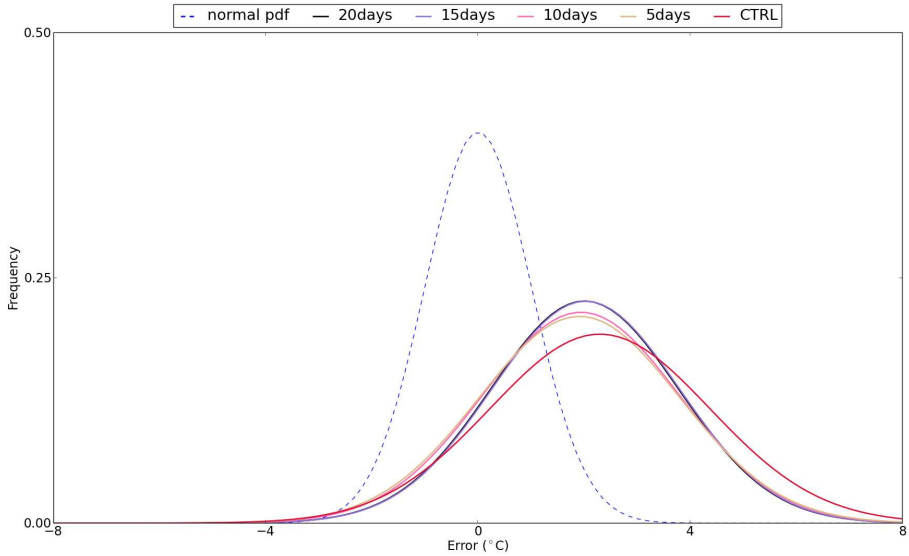

(a) Full domain

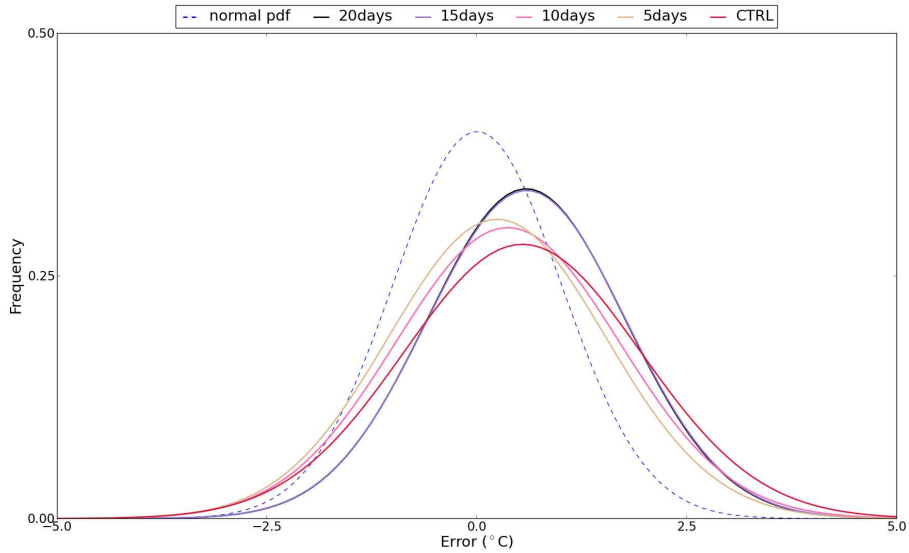

(b) Subdomain

Figure 9: Probability density functions (PDFs) of the model temperature errors for the entire model domain (top panel) and the subdomain centered on Vesterålen (bottom panel). In both figures the normal distribution PDF is indicated with a dashed line.

4.2.2 ADCP measurements

The model current directions in the upper 8 meters does not correspond well with the directions measured by the upper ADCP, in general it has an eastward offset, and has greater spread than the observations. However, we must stress that the ADCP measures the current at a specific location, while the model can only represent the average current over a region the size of a grid box with an area of 2.4×2.4 km^2. The HF radar surface currents also represent a spatial and temporal average, and hence do not measure the same as the ADCP. To illustrate this, we include current roses from the near-surface ADCP and the HF radar total vectors (Figure 10) at the location of the ADCP. Only simultaneous observation have been included in this analysis. As in the model simulations, the main direction measured by the HF radars is more to the northeast than the ADCP measurements. It is likely that the local bathymetry has some influence on the currents in the location of the ADCP, which is not captured by neither the model nor the ADCP.

The model simulations show better agreement with the measurements from the lower ADCP. For these measurements both speed and direction improve with data assimilation. As for the upper ADCP, speed predictions seems to benefit from data assimilation over a longer time period. However, there is not much difference when assimilating data for 10 or 20 days. In general the model simulations are not energetic enough when compared with the currents measured by the ADCP.

Direction of the predicted currents nevertheless benefit from data assimilation over a longer period of time, both in the near-surface and the deeper waters. The current roses in Figures 11 and 12 emphasize these results. It is clear that the data assimilation system intensifies the current speeds, and in the period investigated here, the directions of the lower ADCP correspond better with observations. The error statistics are summarised in Tables 8 and 9.

Observation set	upper		lower	
	RMSE	BIAS	RMSE	BIAS
CTRL	0.21	-0.17	0.16	-0.14
5 days	0.19	-0.14	0.16	-0.12
10 days	0.14	-0.06	0.12	-0.03
15 days	0.14	-0.07	0.13	-0.06
20 days	0.16	-0.10	0.13	-0.07

Table 8: ADCP current speed error statistics. Root mean square error and bias when comparing model results with ADCP current speed. All numbers are given in m/s.

Observation set	Upper		Lower	
	RMSE	BIAS	RMSE	BIAS
CTRL	26.6	-6.5	28.4	-19.9
5 days	39.7	8.2	31.1	-5.3
10 days	34.2	10.6	28.6	-0.5
15 days	35.1	18.5	29.5	2.6
20 days	31.3	9.4	22.4	-5.3

Table 9: ADCP current direction error statistics. Root mean square error and bias when comparing model results with ADCP current direction. All numbers are in degrees.

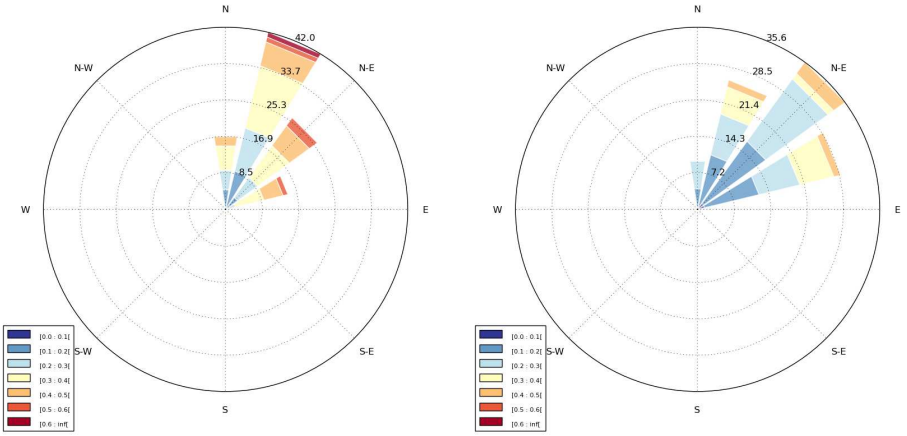

(a) ADCP (b) HF

Figure 10: Upper ADCP (left panel) and HF radar (right panel) during the observation campaign. Only simultaneous observations have been used.

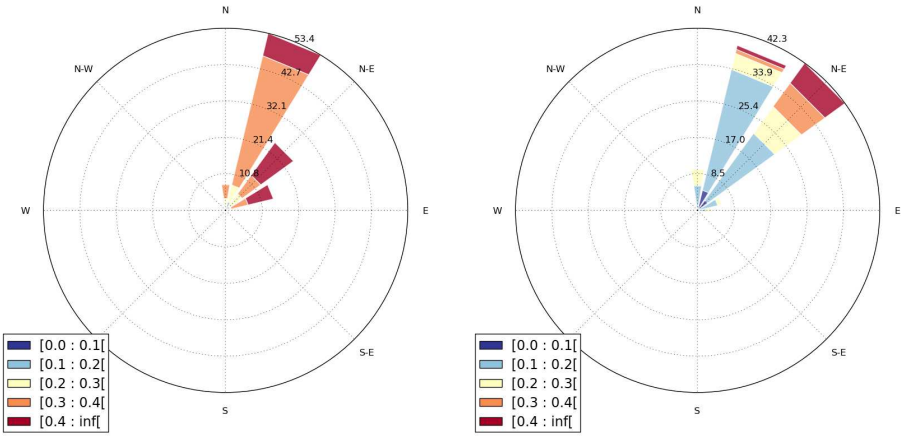

(a) Upper ADCP

(b) CTRL

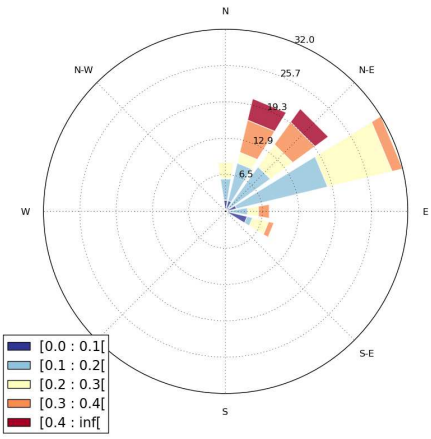

(c) 5D

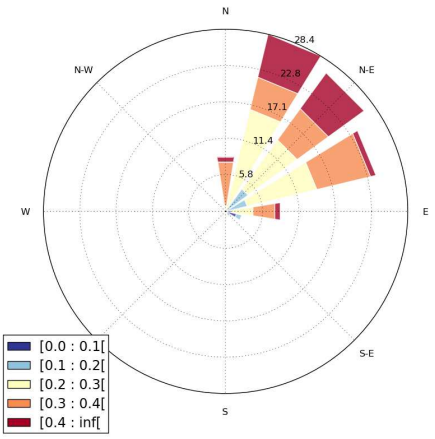

(d) 10D

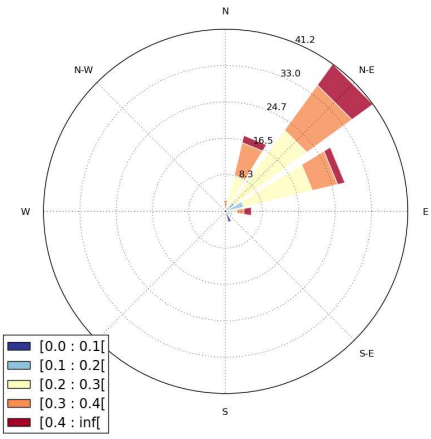

(e) 15D

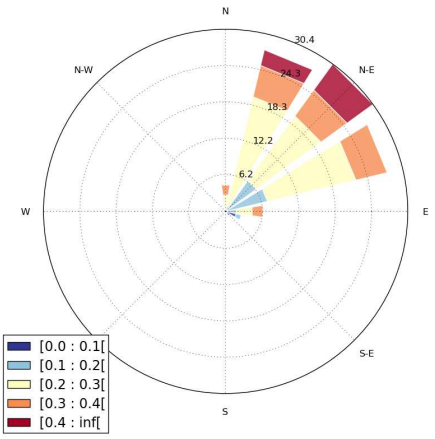

(f) 20D

Figure 11: The left column shows speed and direction distributions during the analysis from 7(a) upper ADCP, 7(c) CTRL and 7(e) ALL. The right column shows the corresponding distributions during the forecast period. Blue colors indicate current speed in the range $0 - 0.2$ m/s, yellow and orange colors $0.2 - 0.4$ m/s and red colors current speeds stronger than 0.4 m/s.

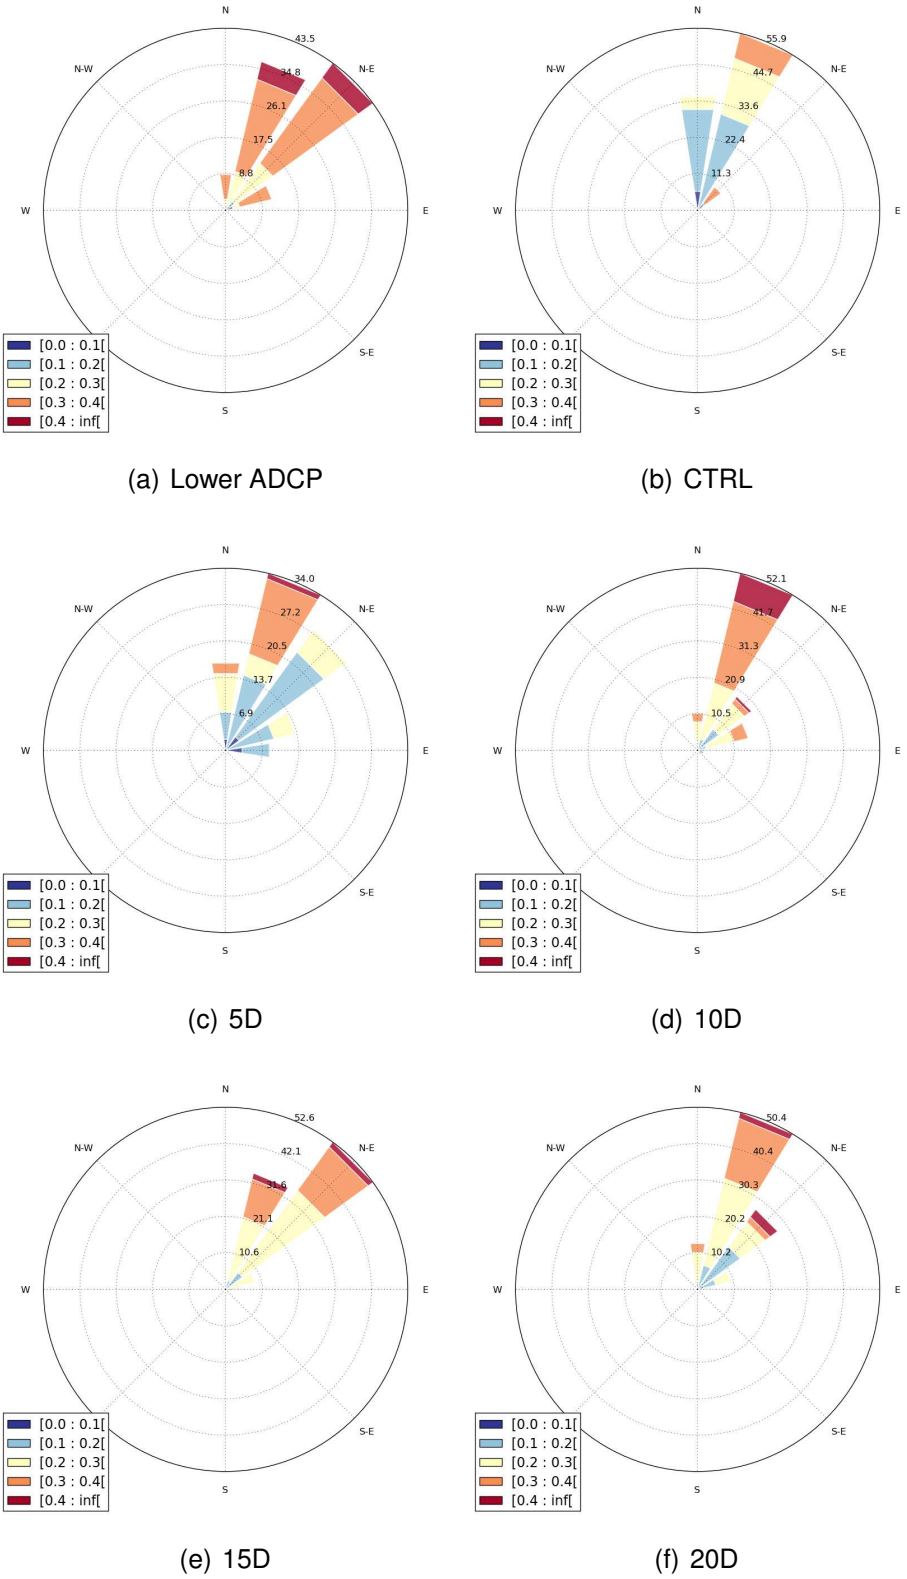

(a) Lower ADCP (b) CTRL

(c) 5D (d) 10D

(e) 15D (f) 20D

Figure 12: The left column shows speed and direction distributions during the analysis from 7(a) Lower ADCP, 7(c) CTRL and 7(e) ALL. The right column shows the corresponding distributions during the forecast period. Blue colors indicate current speed in the range $0 - 0.2$ m/s, yellow and orange colors $0.2 - 0.4$ m/s and red colors current speeds stronger than 0.4 m/s.

5 Concluding remarks

HF radar total vectors and CTD hydrographic profiles have successfully been assimilated into a high resolution version of the operational ocean model ROMS used at MET Norway. The data assimilation is based on strong constraint 4DVAR. Comparing with independent data obtained from surface drifters, upward looking ADCPs, and satellite SST, we find that the bias and root mean square errors in modeled velocities and temperatures are reduced.

A set of tuning experiments were made, and it was found that the best results were obtained using 10 km horizontal correlation scale, 24 hour assimilation window, and 10 inner and 2 outer loops. The ideal error correlation scale is twice as large as in the idealised experiments, which likely reflects that the realistic experiments contain wind and tidally driven dynamics on a larger scale than the case considered in the idealised case.

The impact of assimilating HF radar currents is larger than the impact of assimilating CTD hydrography. Excluding HF data and only assimilating CTD data is detrimental for the forecast, which means that the CTD observations collected during the 24 hour assimilation window are too few to constrain the solution. An interesting finding is that the impact of CTD data is positive when combined with HF radar data. Experiments using longer timeseries of observations shows that the results generally improve with the length of the time series, typically up to ten days.

Acknowledgements

This work was funded by the Norwegian Clean Seas Association For Operating Companies (NOFO) and ENI Norge A/S.

A Evaluation of drifter trajectories

Following Röhrs et al. [2012] we apply two different metrics to evaluate drifter velocities and trajectories. The first metric follow an analysis similar to that of Davis [1985], where drifter trajectories are split in pieces of 3 hours. Drifter velocities can then be calculated for each 3 hour trajectory. Correlation coefficients of simultaneous pairs of drifter velocities may then be calculated. We use the same definition for the vector correlation as in RÃ¶hrs et al.:

$$r = 1 - \frac{\langle (\boldsymbol{v}_i - \boldsymbol{v}_j)^2 \rangle}{\langle \boldsymbol{v}_i^2 \rangle + \langle \boldsymbol{v}_j^2 \rangle}. \tag{1}$$

The correlation coefficient takes the value 1 for equal velocities, e.g. the start and ending point are the same for both the real and modelled drifters, and -1 for opposite velocities.

The second metric used in this report to evaluate drifter trajectories is the normalised cumulative Lagrangian seperation presented in Liu and Weisberg [2011], defined as

$$s = \sum_{i=1}^{N} d_i / \sum_{i=1}^{N} l_{o,i}, \tag{2}$$

where d_i is the separation distance between observed and modeled trajectory endpoints at time i after initialization, $l_{o,i}$ is the length of the observed trajectory and N is the total number of time steps evaluated. A skill score S is then defined as

$$S = \begin{cases} 1 - s & \text{if } s \leq 1 \ , \\ 0 & \text{if } s > 1 \ . \end{cases} \tag{3}$$

High skill score means that observed and modelled trajectory have good agreement throughout the period under evaluation.

Bibliography

J. Albretsen, A. K. Sperrevik, A. StaalstrÃ¸m, A. D. Sandvik, F. VikebÃ¸, and L. Asplin. Norkyst-800 report no. 1: User manual and technical descriptions. Technical Report 2, Institute of Marine Researh, Bergen, Norway, February 2011. http://www.imr.no/filarkiv/2011/07/fh_2-2011_til_web.pdf/nb-no.

D. C. Chapman. Numerical treatment of cross-shelf open boundaries in a barotropic coastal ocean model. *J. Phys. Oceanogr.*, 15(8):1060-1075,1985.

K. H. Christensen. Assimilation of HF radar total current vectors in an idealized version of ROMS-4DVAR. Technical report, The Norwegian Meteorological Institute, Oslo, Norway, February 2013.

K. H. Christensen, A. K. Sperrevik, and J. RÃ¶hrs. The spring 2013 field experiment of the ENI/NOFO HF radar project. Technical report, The Norwegian Meteorological Institute, Oslo, Norway, October 2013.

R. E. Davis. Drifter observations of coastal surface currents during CODE: The method and descriptive view. *J. Geophys. Res.*, 90(C3):4741–4755, 1985.

R. A. Flather. A tidal model of the north-west European continental shelf. *Memoires de la Society Royal des Sciences de Liege*, 10(6):141-164, 1976.

Y. Liu and R. H. Weisberg. Evaluation of trajectory modeling in different dynamic regions using normalized cumulative Lagrangian separation. *J. Geophys. Res.*, 116(C9):, 2011.

P. Marchesiello, J. C. McWilliams, and A. Shchepetkin. Open boundary conditions for long-term integration of regional oceanic models. *Ocean Mod.*, 3(1-2):1 – 20, 2001.

J. Röhrs, K. H. Christensen, L. R. Hole, G. Broström, M. Drivdal, and S. Sundby. Observation-based evaluation of surface wave effects on currents and trajectory forecasts. *Ocean Dyn.*, 62(10-12):1519–1533, 2012.

C. J. Willmott. On the validation of models. *Phys. Geogr.*, 2(2):184–194, 1981.

W. G. Zhang, J. L. Wilkin, and H. G. Arango. Towards an integrated observation and modeling system in the New York Bight using variational methods. part I: 4DVAR data assimilation. *Ocean Modelling*, 35(3):119 – 133, 2010.