TECHNICAL MEMORANDUM

232

Modifications to the ECMWF WAM code

J. Bidlot, B. Hansen and P. Janssen

Research Department

January 1997

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen

MODIFICATIONS TO THE ECMWF WAM CODE

Jean Bidlot, Björn Hansen, Peter Janssen

1. INTRODUCTION

In the course of 1996, ECMWF migrated all its operational numerical codes from its Cray C90 to its new Fujitsu supercomputer. The wave model (WAM) was first modified to run on one single processor of the new computer. That way, we insured that the production of the operational forecasts was not interrupted.

To take advantage of the full computing capabilities of the new Fujitsu, the code was extended further to run on more than one processor. Due to the new computer architecture (parallel, distributed memory), it was necessary to write a set of new routines. These routines basically set the task on each processor and perform the exchange of information between the processors when needed.

To improve services to the Member States, we increased the resolution of the global wave model by introducing a 0.5° reduced lat-lon grid which effectively brings the spacial resolution of the wave model to around 55 km. A new topographic data set was used to define the 0.5° bathymetry, resulting in a realistic description of the coastlines.

Finally we corrected the code for reported errors and inconsistencies and made a few adjustments to improve the efficiency. More details on the implementation of all those changes are given in the next chapters.

2. OPTION TO RUN THE MODEL ON REDUCED LAT-LONG GRIDS

The numerics of WAM are limited by the well-known CFL criterion which imposes a restriction on the size of the integration time step based on the reciprocal of the smallest spatial step. In regular spherical coordinates, both step increments are constant in longitude and latitude, for that reason, the zonal spatial step can become rather small (in actual distance) as one approaches the poles. To alleviate this restriction, a new type of lat-long grid was developed. It maintains a constant latitudinal increment but adjusts the size of the longitudinal increment in such a way that the actual distance between grid points is almost constant. Such a grid is commonly known as reduced or irregular lat-long grid.

This grid option was implemented in WAM and required minor modifications to the code. It involves the computation of the number of grid points used along each latitude circle. This number is obtained by multiplying the number of points at the equator (determined as before by the longitudinal step) with the cosine of the latitude and rounding it off in such a way that the same parity is kept as the equatorial one (to insure that the grid points are aligned along the 0° and 180° meridians). This number of points per latitude is used to determine the actual grid spacing and to limit the extend of the zonal index of the 2-D arrays. The option to run the model on a reduced grid is set by running the grid initialisation program (PREPROC) with the new input grid option flag (IRGG) set to 1. The necessary grid informations are then carried through to the next model run steps.

The propagation scheme (in PROPAGS) was modified to account for the new irregular grid. In the previous regular grid, WAM uses a first order upwind scheme for which information from the 4 neighbouring grid points (one in each cardinal direction from any given grid point) is potentially needed. The same scheme is used for the irregular grid by choosing instead the closest grid point directly north and south of any active grid point and by introducing explicitly the grid spacing in the zonal direction.

The output field format was also altered. For integrated fields saved as binary files, the full 2-D array is still output, however for any given latitude, meaningful information is only contained in the first N columns of that array, where N is the number of grid points for that latitude. For GRIB outputs, the GRIB header was modified accordingly and the data were encoded following the definition of table 3 in the GRIB manual ("Encoding and decoding GRIB data", ECMWF M 1.9/3) (see subroutine GRIBPAC). Note that the MARS software was updated to allow for the retrieval of such fields on their original grid (i.e.: REPRES = LL) or on any specified regular lat-long grid by specifying explicitly the grid size (i.e.: REPRES = LL, GRID = 0.5/0.5).

The assimilation scheme was corrected by making sure that the correlation region centred around any sea point was limited to a given maximum distance rather than a set number of grid points in each direction (subroutine OIFIELD).

3. ADAPTATION TO THE NEW FUJITSU

3.1 Migration from the Cray to a non parallel version on Fujitsu.

We made all the necessary modifications to the code which contained a few Cray specific routines and statements in order to run it on one processor of the new Fujitsu. A list of the different modifications can be found in appendix A. One key difference between runs on Cray and Fujitsu lies in the default representation of real and integer numbers. Fujitsu uses a default 32-bit representation whereas 64-bit was used by Cray. For this reason, 10 digit long integer variables (for date representation) had to be converted into character variables and few default small numbers set to appropriate small values. Note that the numerics in WAM do not require a 64-bit representation as confirmed by comparing a few runs at that precision with the 32-bit counterparts. Finally, routines used to read and write Cray blocked binary files were replaced by less specific unblocked binary read and write routines.

This non-parallel version ran operationally from the end of September 1996 to December 5th 1996.

3.2 Message passing - run WAM in parallel.

In order to make full use of the computing power of the Fujitsu, programs should be run in parallel, i.e on more than one processor, also known as processing element (PE). The single numerical code which will be submitted to the parallel computer must be modified to contain instructions that will tell each processor what to do. Moreover, since the Fujitsu is a distributed memory computer, information contained on one PE is not shared with the others and must be exchanged with the others if there is a need for it. This information exchange is actually done in the message passing context, whereby one PE can send a message which is received by one or more other processors. In its very basic implementation, the message is nothing more than a one dimensional array of a given type containing values that are needed by the other processor(s) plus all the necessary informations about the sender and the receiver. Note that both send and receive have to be successful for a successful message exchange.

ECMWF has set up a very basic message passing library (MPELIB) which contains all the necessary routines to perform the exchange of information between different PE's as well as to set up the system for message passing. The WAM code was upgraded to allow for message passing, it was also decided to keep the option of running the code on a non distributed memory machine (i.e.: no message passing). An input flag determines what option is selected. Note that the message passing routines are written in such a way that they will also run with only one processor (even though no actual message passing takes place). The code was made as flexible as possible to accommodate any number of PE's.

In the current setting, the number of processors is determined at run time (in the batch script command header). Once the message passing program starts simultaneously on all assigned processors, it needs to initialise the parallel environment on each PE and the message passing protocol. It should also determine the total number of processors used as well as the logical PE number on which the code is run. In the case of WAM, this initialisation procedure is done in CHIEF (or any other start up main program which calls the wave code).

Another key issue when running in parallel is how to distribute as evenly as possible the task to be executed on each PE. In the WAM case, it comes down to how to split the global computation grid into even size regions, keeping in mind that information is only known locally on each PE and can only be exchanged with the other PE's via message passing, which is a slower process than computing.

Once the total number of PE's is known, MPDECOMP is called to set up an even decomposition of the total grid into one sub domain per PE. Since the global grid is mapped onto a one-dimensional sea point array following increasing latitude lines, those sub domains were chosen to be consecutive segments of the full sea point array (figure 1). The length of each segment is determined by distributing as evenly as possible the total number of sea points. Subsequently, each PE will only perform the integration of one sub domain, with the first PE responsible of the most southern region and so on. However, the upwind scheme, which solves the advection term, uses neighbouring grid points in the 2-D grid (in the 4 cardinal directions) that might belong to another subdomain.

For a given segment on PE N, information that might potentially be required from the previous sub domain on PE N-1 is actually entirely contained in a small sub segment spanning all active sea points of the previous region N-1 from the point directly south of the first point of the given region N (point A on figure 1) to the last sea point of the previous region N-1 (point B), as long as this first southern point is a sea point, otherwise the next one will be selected. A similar sub segment definition is used for information from the next region on PE N+1. The length and structure of those small sub segments are used to construct the messages that are exchanged between processors, before calling the advection routine. Note that it was also possible to reduce the amount of memory used on each PE since one only needs to allocate enough memory to cover the sub domain and the length of both messages. The fortran 90 function ALLOCATE was used. Note, however, that those newly allocated files cannot be part of a common block and were therefore passed as calling arguments from one subroutine to the next.

Using the basic send and receive command (see MPELIB library), it is possible to proceed to a global exchange of information between PE's (see MPEXCHNG). Basically, each PE builds a message following the structure defined in MPDECOMP and aside from the last PE, send it to the next PE. The same is done to send a message to the previous PE (except for the first PE). As all messages are sent, each PE ensures it receives one message from the previous PE and one from the next PE. Only at that point is the actual exchange of information between contiguous sub domains complete, and the advection routine can be called to advance the integration by one time step. If no data assimilation is required, then it is all the data exchange which is needed (all sources terms are local in nature) until the code reaches an output time.

At output time, integrated and spectrum fields are produced by transforming the 1-D array of sea point values into a 2-D array. This 2-D array must contain all sea point values (plus a default value for land) that must be collected (gathered) from all PE's onto one PE. As more than one integrated field can be output, it is feasible to assign the field collection task to different PE's (MPCRTBL creates the necessary table which assigns the gathering task to the PE's). Schematically, the gathering routines build a long message with all relevant information contained on one PE and send it to the designated receiving PE (see MPGATHERSCFLD and MPGATHERFL). Once a 2-D field is on one PE, it can be output with the proper format. Note that if the output option requires all 2-D fields to be saved in the same file, then all 2-D fields are collected onto one PE before output (see MPGATHERGRFLD). This is not the case if the GRIB output is directed to the field data base (FDB) since it accepts parallel output.

If data assimilation is performed, informations from the 2-D arrays for wave height, friction velocity and drag coefficient are needed on all PE's (see OIFIELD). A similar gathering procedure as for all outputs can be used to broadcast (send to all PE's) the 3 fields before the call to the optimum interpolation routine (see MPGATHEROIFLD).

At this time, all input to the wave model is done in parallel by accessing the same input files from all PE's. The standard input is also replaced by one input file which can be accessed from all PE's. The standard output is redirected to one logfile per PE which can be inspected separately. Finally at the end of program, the parallel environment is closed after closing all units still opened.

4. NEW TOPOGRAPHIC DATA SET - 0.5°X0.5° REDUCED GRID IMPLEMENTATION

We obtained a 5'x5' global topographic data ("Data Announcement 88-MGG-02, Digital relief of the Surface of the Earth. NOAA, National Geophysical Data Center,Boulder, Colorado, 1988."). This data set is in principle a terrain following data set in which negative numbers indicate ground below mean sea level or most of the time water depth except for land depressions. However, large lakes with mean surface level well above mean sea level are not readily trackable since their bottom elevation is still positive. For some of the largest lakes (US Great Lakes) it is possible to locally change the reference level (as found in atlas) and recover an approximate lake bathymetry. Unfortunately, the topographic data for most of the other large lakes and enclosed seas is limited to constant values which seems to correspond to an approximate mean level.

We intended to extract a 0.5°x0.5° bathymetry data from this raw data set with as much water as possible and a good description of the coastlines. In order to keep the details of the original data set, the 5" data were only averaged in 15" windows and the 0.5° data were determined by taking the closest 15" data point. The 15" data was also locally adjusted by specifying small enclosed areas for which a new reference level and possibly a new bathymetry are used instead (we are only interested in negative values that will be transformed into bathymetric values). This way, we were able to reject many land areas that are lying below mean sea level, to obtain a nice description of most coastlines, as well as to reject small islands that would otherwise interfere with the numerical propagation of waves (see below) and to impose a constant depth over large lakes and enclosed seas which were otherwise not described in the original data set (see appendix B). Nevertheless, the resulting 0.5°x0.5° data set still contains small undesirable features (water over land and vice versa) which were adjusted using the corrective input list to PREPROC.

It was found that the propagation scheme in its current discretisation (25 frequencies and 12 directions) could fail when propagating a wave system around small islands and resulting in a strong unrealistic decrease in wave height in the lee of the islands. This shadowing effect was particularly visible around the Hawaii and Solomon Islands. An example is a wave system of November 10, 1996 in the North Pacific which was propagating due south (see figure 2a). Within two days, pronounced, unrealistic shadowing effect show up south of Hawaii (figure 2b). More work needs to be done to investigate the possible corrections to this undesirable numerical feature, but in the meantime, the land points corresponding to those islands as well as a few others were set to sea points (appendix B). This correction removes the shadowing as seen in figure 2c. The grid may still contain a few erroneous features which can always be corrected at later time by modifying the corrective input list to PREPROC.

Due to the good coverage of all the seas around Europe as well other coastal regions, it was necessary to switch on the shallow water option. Note that in this latter case, the model is more sensitive to the shadowing effect since the smoothing of small islands may have resulted in shallow water sea points for which the wave propagation is mostly altered.

For the North Atlantic area, the resulting grid point distribution for a 0.5 reduced grid is shown in figure 3. The whole globe is covered by 116,489 active grid points.

5. BUG FIX

During this major model overhaul, small errors and inconsistencies in the WAM code were corrected. The following list will give a rapid overview of the corrections implemented in the new code.

IMPLSCH:

use of a new limiter on the spectrum increment at each time step. It was found that the new limiter was less constraining than the previous one and resulted in better wave growth for young wind sea cases for localised regions with short fetch (Hersbach and Janssen, personal communication). The new section 2.4 can be found in appendix C.

SEPWISH:

bug in the computation of the mean direction of swell and windsea. It was assumed by error that the frequency increment (DF) was a constant in the integration of the spectrum for the mean direction (Cavaleri, personal communication)(see appendix C).

AIRSEA:

It is physically more consistent to perform the bi-linear interpolation on the square root of the stress to obtain the friction velocity and use the drag coefficient to determine the roughness length. The modifications are given in appendix C.

STRESS:

Removal the viscous contribution to the roughness length and ensure that there is zero stress for zero wind (see appendix C).

OIFIELD:

It was found that the computation of the distance over which the satellite observations are taken into account at any given point was incorrectly determined and was set to correspond to a given number of grid spacing and instead to a physical distance projected on the grid. In a regular lat-long grid, this feature may result in the shortening of the zonal correlation distance of each observation when approaching the poles. One should instead compute the number of grid points in each direction which corresponds to the given maximum correlation distance and deduce the size of the influence region of each observation based on those numbers.

6. CODE IMPROVEMENTS

In an attempt to improve the performance of WAM on the new Fujitsu, we modified in PROPAGS the computation of the group speed as it does not depend on the direction but only on frequency and depth (in shallow water). The subroutine OIFIELD was optimised by removing a few unnecessary if statements and replacing them by appropriate do loops.

The main bottle neck when running on the new Fijitsu lies with the slow and inefficient output to permanent storage. Every effort should be made to limit the amount of output. In the current operational settings, the wave model should output to disk the full restart files (fortran units 12 and 15) at the end of the analysis and the 2-D spectra every 6 hour until day 5. Two new input lines were added to the WAMINFO file to limit the disk output of the full restart files to a prescribed date and the spectra until a given date (at intervals prescribed in the model input). Furthermore, those restart files were output as standard fortran units (fort.nn), to be copied later with an appropriate name to their storage location via a call to GSFILE. We can by-pass this copy by writing the files directly to their storage location with their appropriate names. A call to a new routine (GRSTNAME) yields the filename with the necessary storage path as obtained from the user defined environment variable YMPPATH. This latter filename is used to save the file directly to disk and if the path points to a disk mounted as a VFL file system (for large files) then the blocksize can be increased to write large chunks of data at a time and the disk space can be preallocated (reserved as a continuous physical disk area), which greatly improves the output performance. In the operational runs, the disk space of all restart files is kept preallocated at all time, however it can also be preallocated at run time from within the wave model by calling PREALLOC_FILE.

All outputs of integrated parameters are directed to the field data base (FDB) which can be accessed by any PE. For this reason, the integrated fields are scattered as evenly as possible on all PE's and their output to FDB should occur simultaneously, provided that the FDB server is defined by setting the environment variable FDB_SERVER_HOST before the model run and provided that the FDB server is initialised on all PE (see call to IINITFDB in CHIEF). The FDB software was also updated by Dragan Jokic and a faster version of FLD2FDB was written.

At this time, the binary spectrum files are still postprocessed to be saved in GRIB format in the FDB. Due to the size of each file (~140 MB), it is a lengthy process which has been shortened by only archiving in the FDB the spectra for the northern hemisphere. The postprocessing is also triggered from within the main model run after the last file is transferred to storage disk (after day 5 forecast).

7. PERFORMANCE

The effort spent on improving the output performance proved to be essential for the global performance of the model runs. Currently, the global model (0.5° reduced grid, shallow water) is run on 4 PE's and takes about 50 minutes to complete, where before, the non parallel version (1.5° regular lat-long grid, deep water) took about 42 minutes on one PE. The effect of all performance enhancements is even more apparent for the Mediterranean model. The same software as the new global model is used, except that the model is only run on one PE with the regular 0.25° lat-long grid. The new Mediterranean model needs about 15 minutes to complete compared to 23 minutes before (on the exact same grid).

Preliminary results show that the wave model at the new resolution can capture more of the wind variability. As a consequence, wave systems tends to have slightly sharper gradient in agreement with the wind distribution. The test job for the new operational model was unfortunately run only for a relative short period of over a month. During that period, a few adjustments took place. Nonetheless, scores for that period indicate that the new model is performing rather well. For example, figure 4 shows the evolution of the root mean square error as a function of forecast days for the new (EW) and the old version for wave heights, wind speeds and peak periods. Those statistics are obtained by comparing the model outputs with observations at a selected subset of ocean buoys. A basic quality control on the buoy monthly time series was used to remove spurious observations. The resulting buoy data were averaged over 6 hour time windows centred around the model output times to match the model representative scales. Statistics are presented for all 33 selected buoys, as well as for a subset of buoys per geographical areas (see figure legend). The wave statistics are rather similar in both cases. There is a slight increase in the wind speed RMSE when it is extracted on the 0.5° grid. This increased variability should translate into a higher variability of the wave fields. Note however, that the test run was restarted at mid period with the new bathymetry but with an initial wave field which was reset to the local wind (cold start). The effect of this cold start are not entirely removed as attested by the scores at Hawaii and elsewhere for the peak period. The new bathymetry was also beneficial for the North Sea area. The effect of which can be seen in the analysis and one day forecast (later forecasts are more affected by the use of the old bathymetry).

The evaluation of the quality of the wave analysis and forecasts is still under way. The higher resolution also means that coastal regions are better represented. Model outputs can now be realistically compared to more buoy observations. More attention will be given to the propagation problem around islands as it seems that even with the corrected bathymetry the observed and model wave heights around Hawaii (figure 2) might still differ by more than a meter.

8. ACKNOWLEDGEMENTS

We would like to thank Geir Austad, Georgios Konstandinidis, John Chambers, Dragan Jokic and John Hennessy who dedicated their time to implement the wave model in the operation department. Their contributions were also essential in improving the global performance of the model run on the new Fujitsu. Advice on parallel programming and use of message passing was kindly provided by Lars Isaksen and Sami Saarinen. We would also like to acknowledge the useful discussions with Neil Storer, Peter Towers and David Dent.

APPENDIX A

List of the modifications to run WAM on one processor on Fujitsu.

All integers variables used to represent 10 digit long dates were converted to character*10 variable. Please note that 0 is now '0000000000'.

Calls to BUFFERIN and BUFFEROUT are now handled by the calls to routines from the PBIO package.

All input and output files are linked to real Fortran unit (fort.nn) instead of using an Hollerith constant linking a character file name to a fortran unit.

There is no half precision exponential and logarithmic function, the standard EXP and ALOG were used instead of EXPHF and ALOGHF.

Calls to ISHELL were replaced by calls to SYSTEM.

There is no call the ECFILE routine, replace it by calls to SYSTEM('ecfile....)

Calls to ORDERS were replaced by calls to MYINDEX, a sorting routine written for that purpose.

All small numbers are now defined by variable EPSMIN, define as 0.1E-32.

All arrays used for coding and decoding grib data were updated to the necessary size.

All the scripts were modified to include Fujitsu compiling and linking commands. Fujitsu compiler requires that all subroutines for which there is a call statement be defined even if they will never be triggered at run time. The code still contains multitask commands (that could be used if the programs were to be run again on Cray) for which there is no definition. In order to compile, dummy definition for those few routines were provided in which the routine is defined by its name followed by a return-end statement. Similarly, all calls to the MAGICS library routines were removed since there is no plotting on Fujitsu.

APPENDIX B

Lakes and enclosed seas set to a constant depth:

NAME	NEW DEPTH (m)
Lake Victoria	50
Lake Turkana	50
Lake Tanganika	500
Lake Malawi	300
Lake Onezhskoye	25
Caspian Sea	200
Aral Sea	30
Lake Balquash	200
Lake Baikal	500
Lake Winnipeg	50
Great Slave lake	50
Lake Keathabasca	50

Islands artificially removed and considered as sea:

Hawaii Islands

Solomon Islands

Galapagos Islands

Chatham Islands

Sakalin Islands

Natuna Besar Island

Mascarene Islands

Kerguelen

South Georgia

APPENDIX C - Corrections of bugs

IMPLSCH: The new section 2.4 should read as:

```
C*   2.4 COMPUTATION OF NEW SPECTRA.
C      --------------------------
C
C     INCREASE OF SPECTRUM IN A TIME STEP IS LIMITED TO A FINITE
C     FRACTION OF A TYPICAL F**(-4) EQUILIBRIUM SPECTRUM.
C
 2400 CONTINUE
      DELT = IDELT
      XIMP = 1.0
      DELT5 = XIMP*DELT
      DO 2401 M=1,NFRE
        DELFL(M) = 5.E-07*G*FR(M)**(-4.)*DELT
        DO 2402 K=1,NANG
          DO 2403 IJ=IJS,IJL
            GTEMP1 = MAX((1.-DELT5*FL(IJ,K,M)),1.)
            GTEMP2 = DELT*SL(IJ,K,M)/GTEMP1
            FLHAB = ABS(GTEMP2)
            FAC   = USNEW(IJ)*FMEAN(IJ)*DELFL(M)
            FLHAB = MIN(FLHAB,FAC)
            FL3(IJ,K,M) = FL3(IJ,K,M) + SIGN(FLHAB,GTEMP2)
            FL3(IJ,K,M) = MAX(FL3(IJ,K,M),FMIN)
 2403       CONTINUE
 2402     CONTINUE
 2401   CONTINUE
C
```

SEPWISH: the following do loops should read as:

```
DO 2004 M=1,NFRE
        DF = DFIM(M)
        FD = DF/FR(M)
DO 2005 IJ=IJS,IJL
          F3SW  = FL3(IJ,K,M)*FL1(IJ,K,M)
          F3WIS = FL3(IJ,K,M)-F3SW
          SF3SW(IJ) = SF3SW(IJ)+F3SW*DF
          SF3WIS(IJ) = SF3WIS(IJ)+F3WIS*DF
          ESWELL(IJ) = ESWELL(IJ)+DF*F3SW
          FSWELL(IJ) = FSWELL(IJ)+FD*F3SW
          ESEA(IJ)   = ESEA(IJ)+DF*F3WIS
          FSEA(IJ)   = FSEA(IJ)+FD*F3WIS
2005      CONTINUE
2004    CONTINUE
```

AIRSEA: The modifications are given below.

```
C*    1. DETERMINE TOTAL STRESS FROM TABLE.
C        ---------------------------------
C
 1000 CONTINUE
      DO 1001 IJ=IJS,IJL
         XI     = TAUW(IJ)/DELTAUW
         I      = MIN ( ITAUMAX-1, INT(XI) )
         DELI1  = XI - FLOAT(I)
         DELI2  = 1. - DELI1
         XJ     = U10(IJ)/DELU
         J      = MIN ( JUMAX-1, INT(XJ) )
         DELJ1  = XJ - FLOAT(J)
         DELJ2  = 1. - DELJ1
         US(IJ) = (SQRT(TAUT(I,J  ))*DELI2
     1           + SQRT(TAUT(I+1,J  ))*DELI1)*DELJ2
     2           +(SQRT(TAUT(I,J+1))*DELI2
     3           + SQRT(TAUT(I+1,J+1))*DELI1)*DELJ1
 1001 CONTINUE
C

C*    2. DETERMINE ROUGHNESS LENGTH.
C        --------------------------
C
 2000 CONTINUE
      DO 2001 IJ=IJS,IJL
         SQRTCD = MIN(U10(IJ)/US(IJ),100.0)
         Z0(IJ) = XNLEV*EXP(-XKAPPA*SQRTCD)
 2001 CONTINUE
```

STRESS: removal the viscous contribution to the roughness length and ensure that there is zero stress for zero wind:

```
DO 1200 ITER=1,NITER
          X     = ZTAUW/TAUOLD
          UST    = SQRT(TAUOLD)
          Z0     = ALPHA*UST**2/(G)/(1.-X)**XM
F     = UST-XKAPPA*UTOP/(ALOG(XNLEV/Z0))
          DELF   = 1.-XKAPPA*UTOP/(ALOG(XNLEV/Z0))**2*2./UST*
   *            (1.-(XM+1)*X)/(1.-X)
          UST    = UST-F/DELF
          TAUOLD =  MAX(UST**2., ZTAUW+EPS1)
 1200      CONTINUE
        TAUT(I,J)  = TAUOLD
 1100   CONTINUE
C
C*   END DO LOOP OVER INDICES OF TAU-TABLE
C
 1000 CONTINUE
C
C*   FORCE ZERO WIND TO HAVE ZERO STRESS
C
    DO 2000 I=0,ITAUMAX
      TAUT(I,0)=0.0

2000   CONTINUE
```

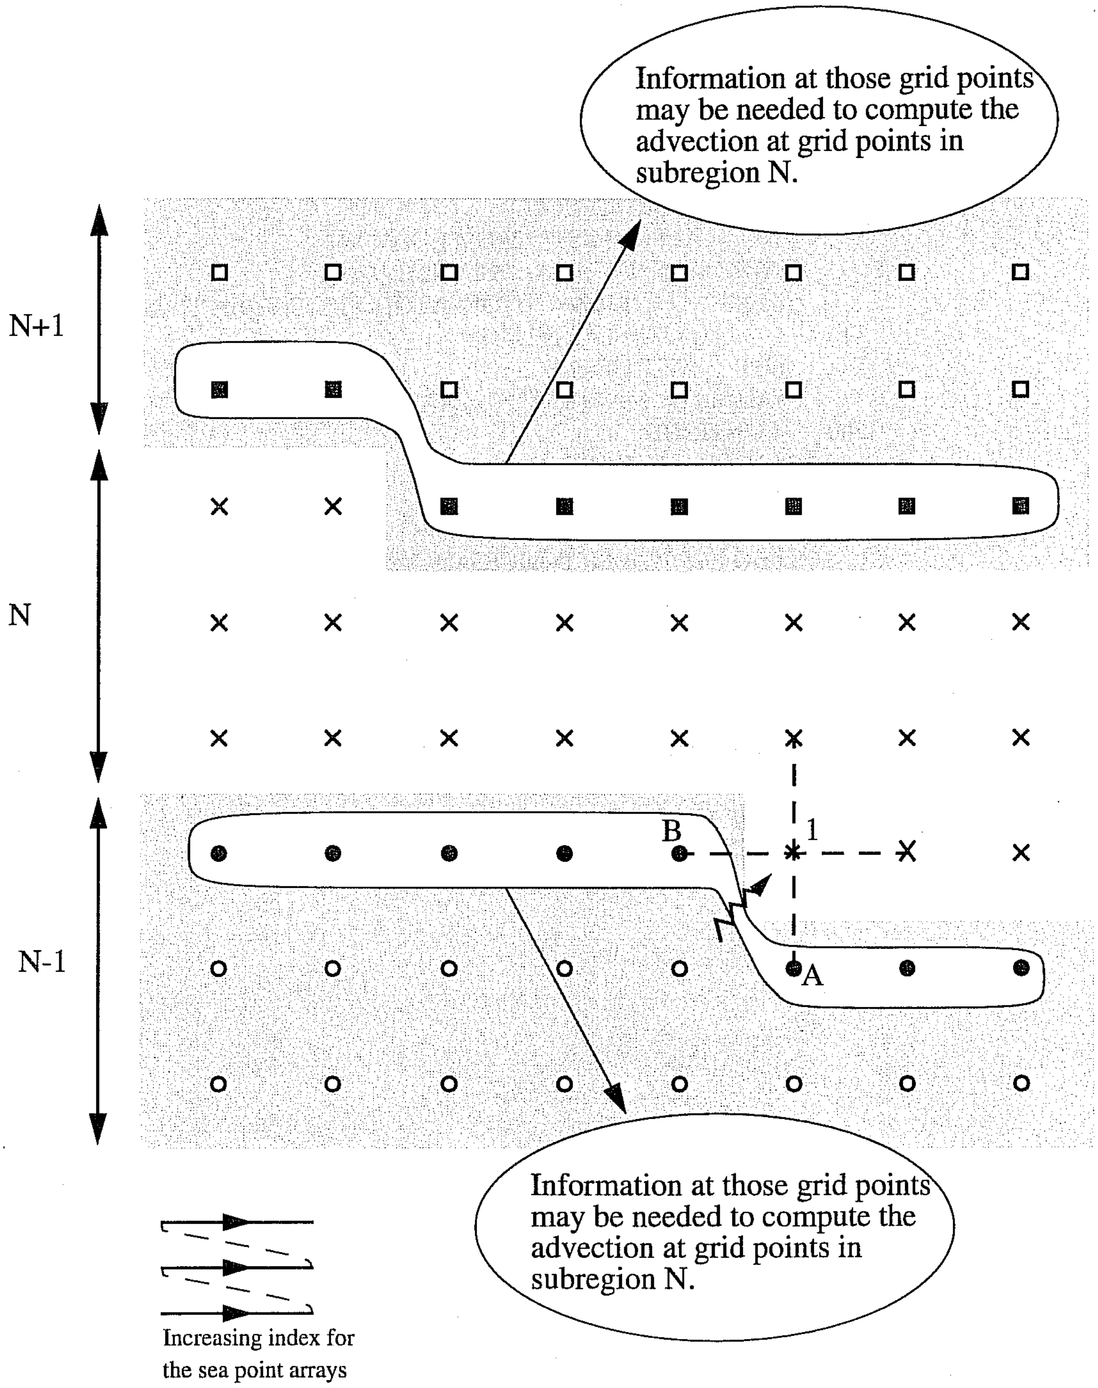

Figure 1: Schematic representation of how the global grid is divided into sub-regions (represented by different symbol shapes). The computation of each sub-region is assigned to a different processor (N-1, N, N+1). Information about grid point values from neighbouring sub-region must be obtained via message passing. For example, if wave energy is propagating towards grid point 1 as indicated by the wiggly arrow, then grid point values from points A and B are required to evaluate the contribution to the first order advection scheme. Thus the message that will be sent to N from N-1 should include all possible A's and B's or all encircled points (darkened circles). Similarly, for the message from N+1 (all encircled darkened squares).

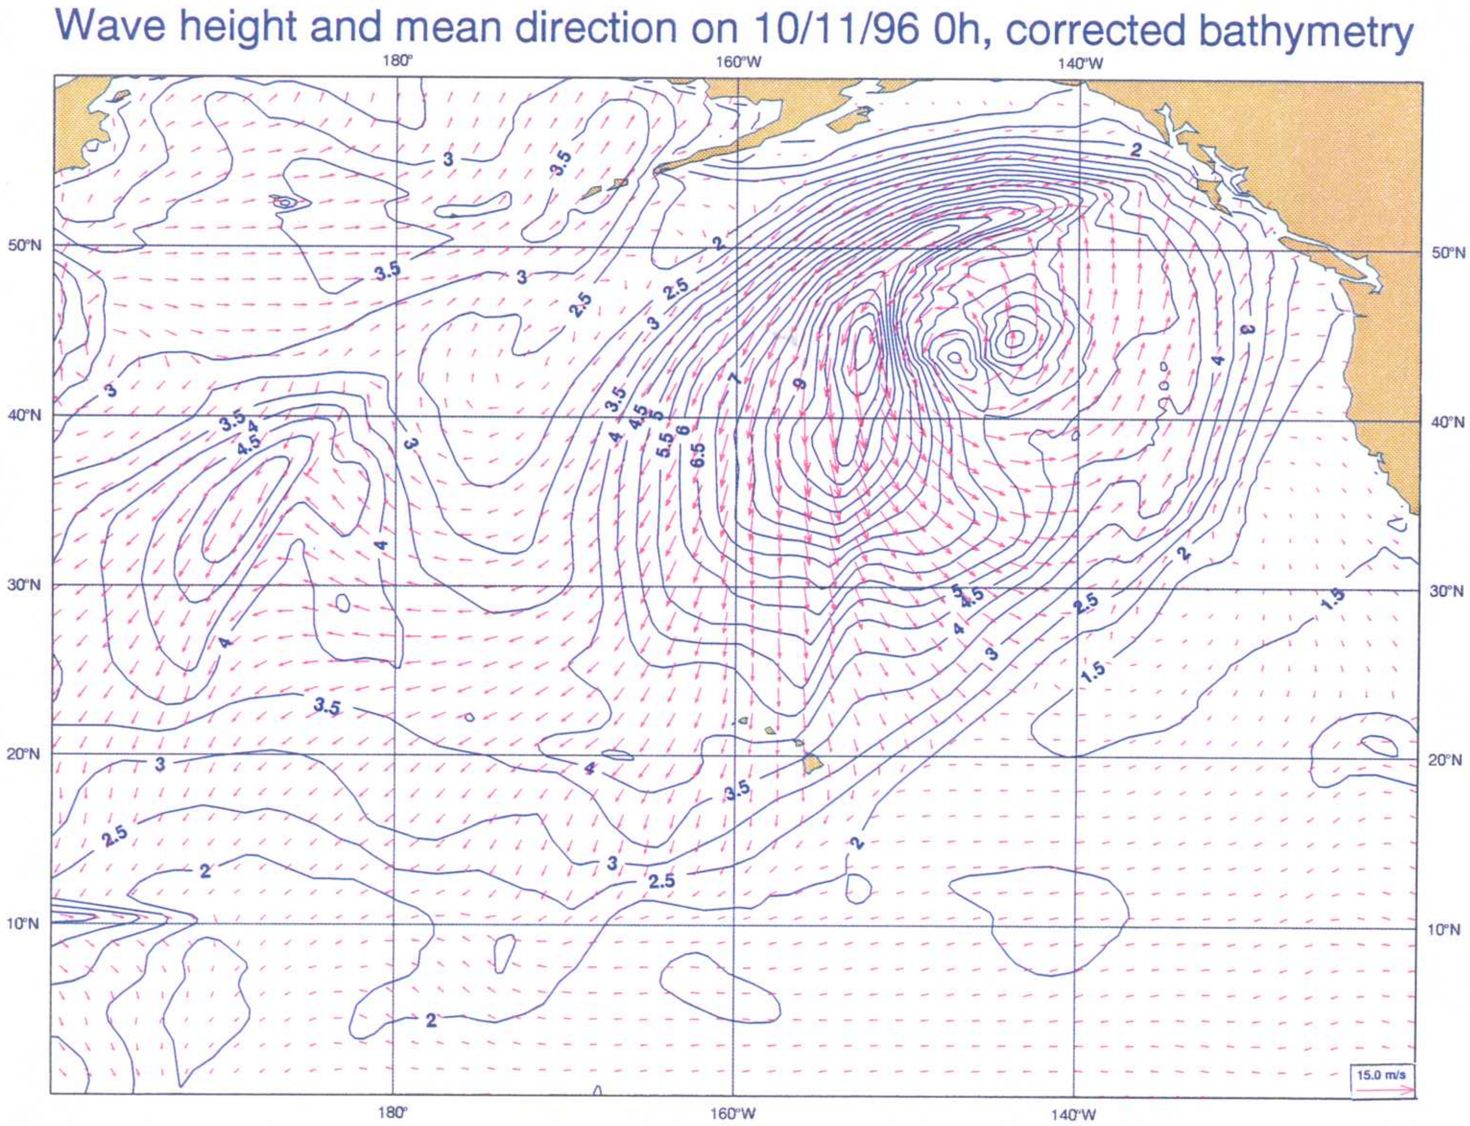

Figure 2: Analysed wave height (contour lines) and mean wave direction (arrows) over the North Pacific Ocean.

(a) On November 10, 1996, 0Z, a strong depression has produced an intense wave system propagating southward. This system reached the Hawaii Islands in the following 2 days.

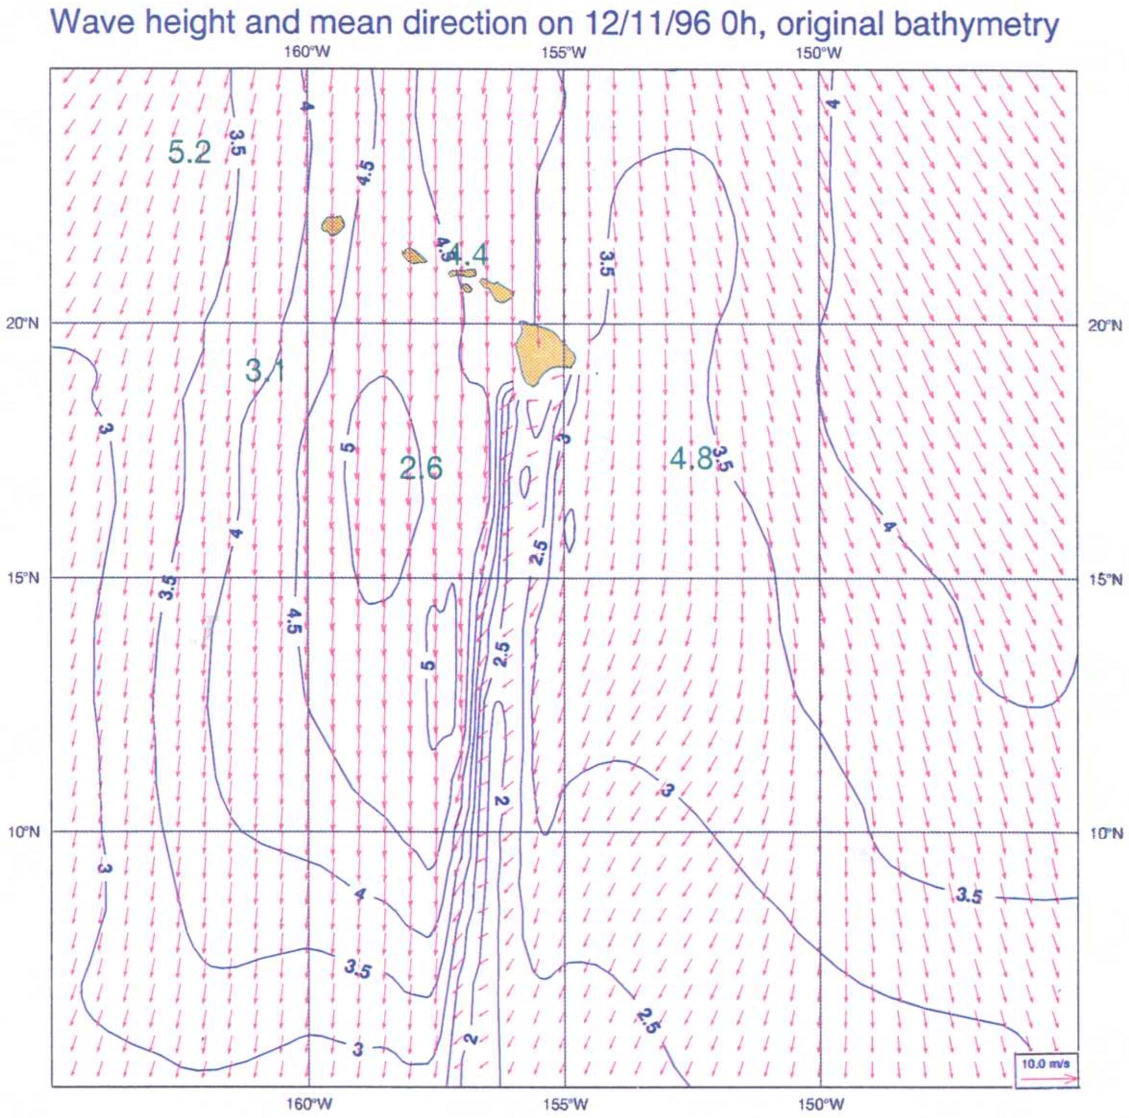

Wave height and mean direction on 12/11/96 0h, original bathymetry

Figure 2 continued

(b) Wave field around Hawaii on November 12, 1996, 0Z, with the original bathymetry which retained Hawaii.

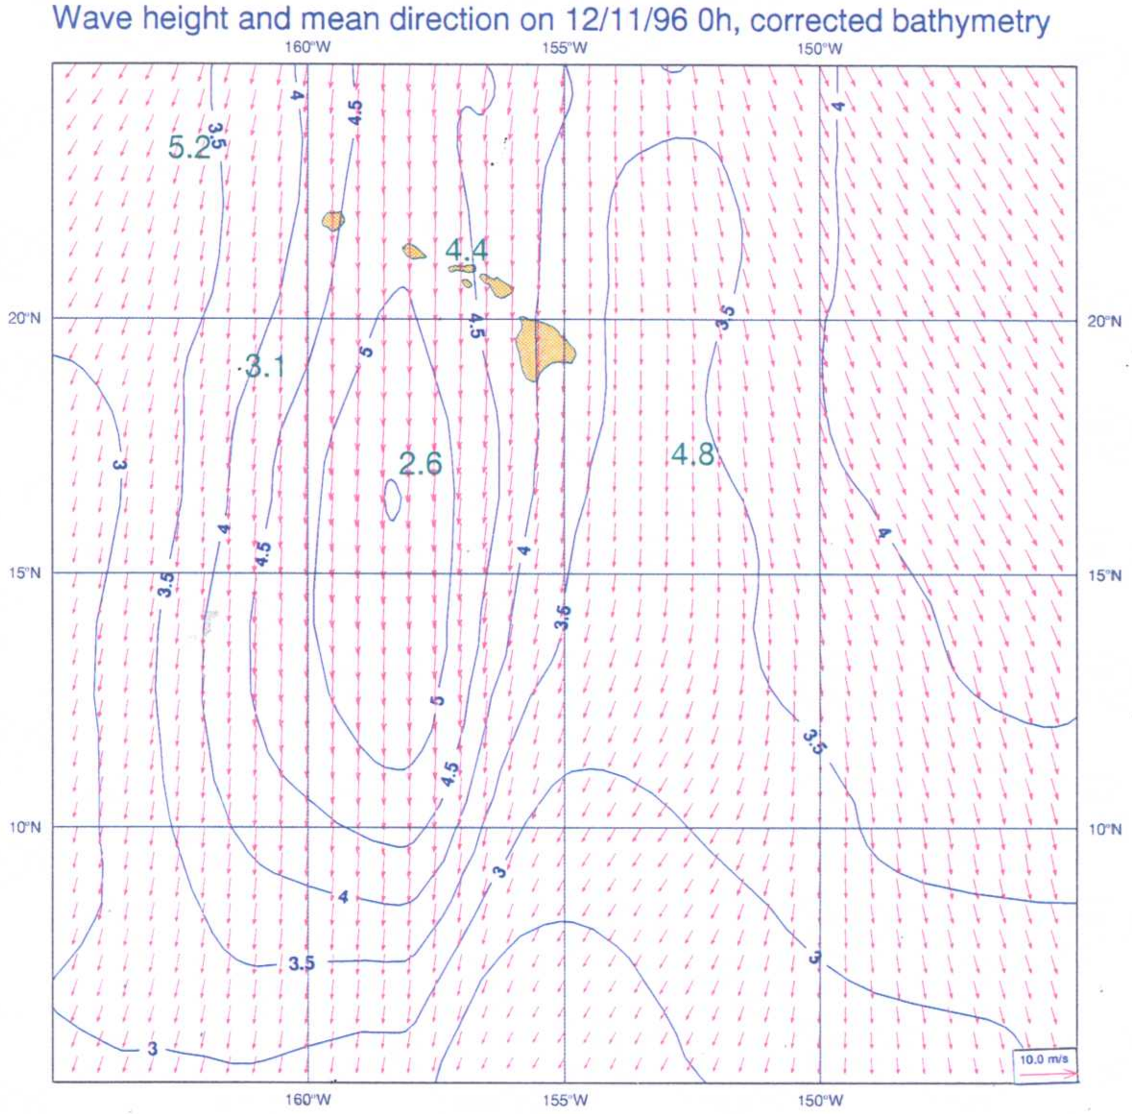

Wave height and mean direction on 12/11/96 0h, corrected bathymetry

Figure 2 continued

(c) Wave field around Hawaii on November 12, 1996, 0Z, with the corrected bathymetry in which Hawaii has been removed. The four large green numbers indicate the NDBC moored buoy observations.

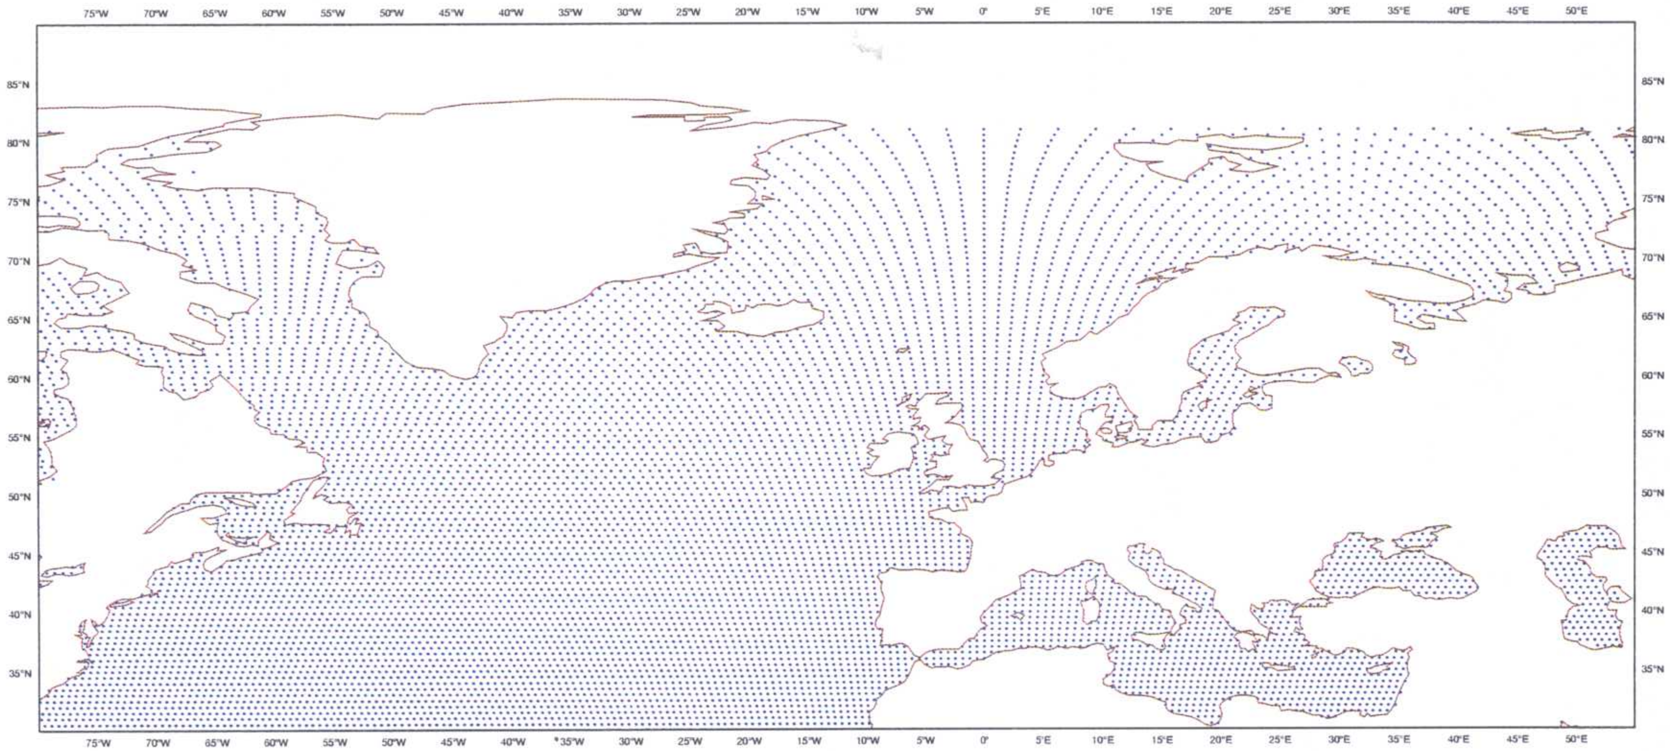

Figure 3: 0.5° reduced lat-lon grid used operationally at ECMWF for the North Atlantic.

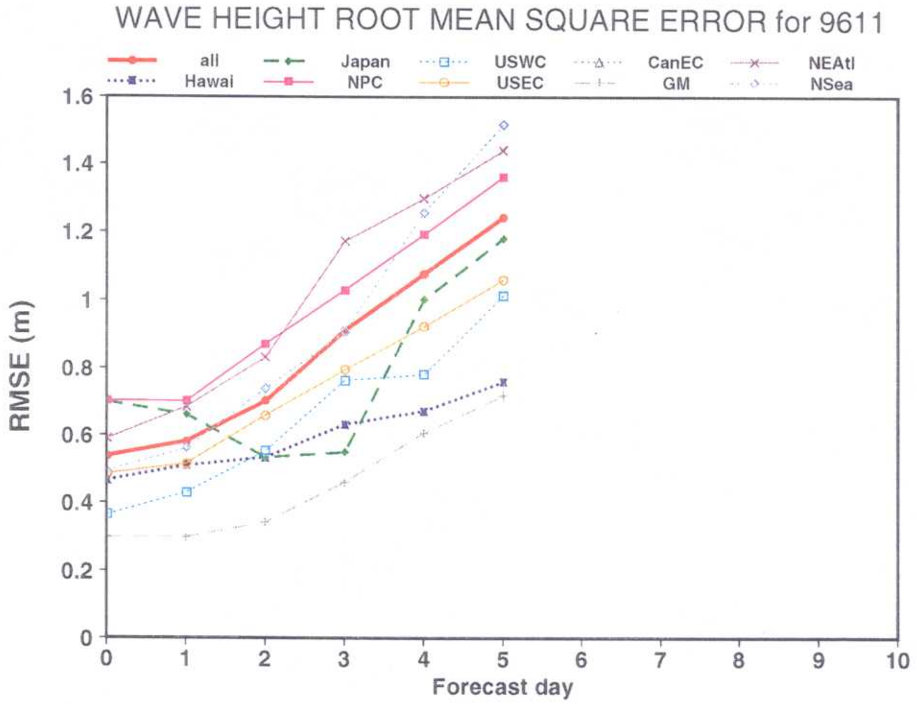

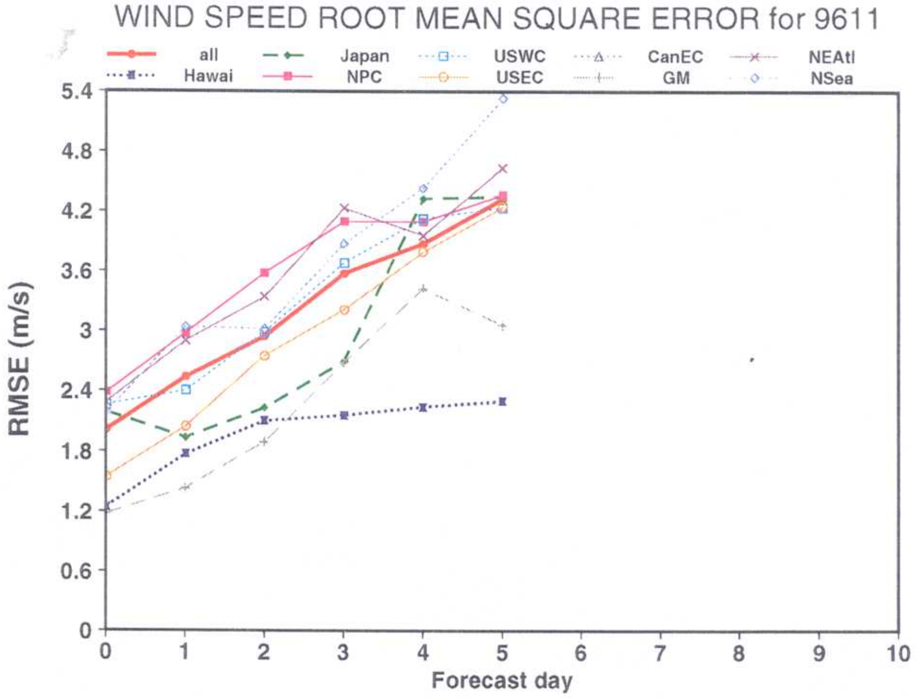

Figure 4: Root mean square errors as function of forecast day for November 1996. The statistics were obtained from the comparison of model results with 6-hourly averaged moored buoy data. The RMSE are given for all 33 buoys and as well as for Hawaii (51001, 51002, 51003, 51004), Japan (21004, 22001), the North PacifiC area (46001, 46003, 46035, 46184), off-shore from the US West Coast (46002, 46005, 46006, 46036, 46059), off-shore from the US East Coast (41001, 41002, 44004, 44008, 44011), the Canadian East Coast (no data), the Gulf of Mexico (42001), the North East Atlantic (62029, 62081, 62105, 62106, 62108, 62163, 64045) and the North SEA (62026, 62109, 62112, 62165, 63111). (a) Results for the 1.5°x1.5° operational grid

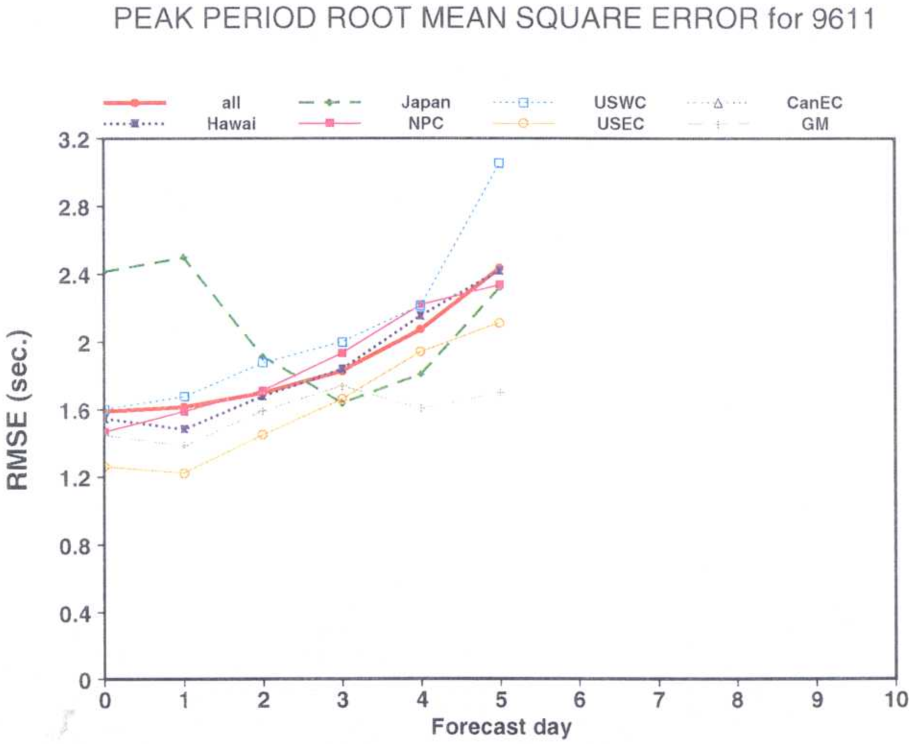

Figure 4a continued

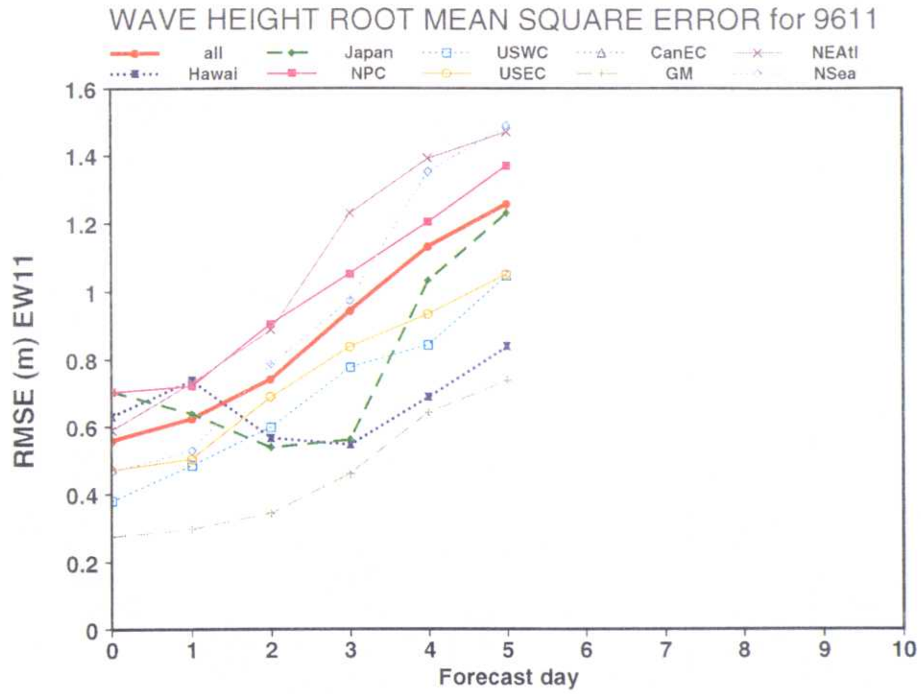
WAVE HEIGHT ROOT MEAN SQUARE ERROR for 9611
all
Japan
USWC
CanEC
NEAtl
Hawai
NPC
USEC
GM
NSea
RMSE (m) EW11
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
0
1
2
3
4
5
6
7
8
9
10
Forecast day

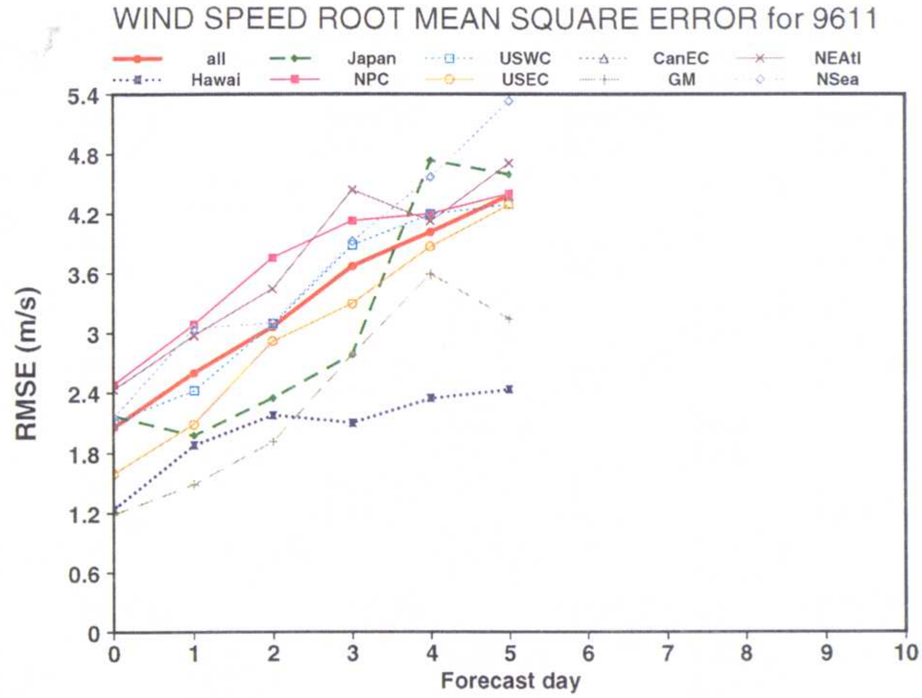
WIND SPEED ROOT MEAN SQUARE ERROR for 9611
all
Japan
USWC
CanEC
NEAtl
Hawai
NPC
USEC
GM
NSea
RMSE (m/s)
5.4
4.8
4.2
3.6
3
2.4
1.8
1.2
0.6
0
0
1
2
3
4
5
6
7
8
9
10
Forecast day

Figure 4b

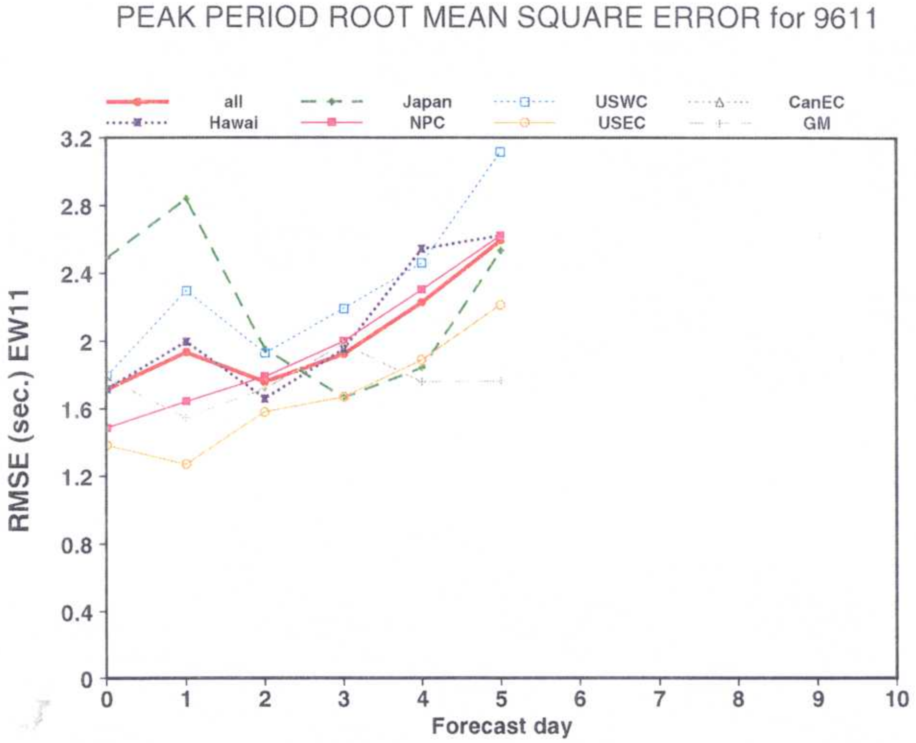

Figure 4b continued