Multi-grid parallelisation on SMC grids in WW3 wave model

Jian-Guo Li
Met Office, FitzRoy Road, Exeter, UK
Email*: *Jian-Guo.Li@metoffice.gov.uk

Abstract: Spherical Multiple-Cell (SMC) grid is an unstructured grid, supporting flexible domain shapes and multi-resolutions. It retains the quadrilateral cells as in the latitude-longitude (lat-lon) grid so that simple finite-difference schemes can be used. Sub-timesteps are applied on refined cells and grid cells are merged at high latitudes to relax the CFL restriction. A fixed reference direction is used in polar regions to solve vector polar problems. The SMC grid was implemented in the WAVEWATCH III (WW3) wave model in 2012 as an alternative for the lat-lon grid and updated in the latest WW3 Version 7. The WW3 model is parallelised by wave spectral component decomposition (CD) in MPI mode, which has a limit on number of MPI ranks. Hybrid MPI-OpenMP parallelisation may extend the node usage, but the OpenMP scalability flattens out beyond a few threads. The multi-grid option in WW3 provided another parallelisation option that combines CD with domain decomposition (DD), which is used in atmospheric and ocean models. This combined CD-DD parallelisation is initially applied only to the lat-lon grid. Although the curvilinear and the unstructured triangle cell grids are later added into the multi-grid option, they are limited to run in series mode with a parent lat-lon grid model, which may be split into sub-grids and run in the multi-grid parallel mode. The SMC grid module in WW3 has recently been extended to use this multi-grid option to expand node usage in hybrid parallel mode. The flexible domain shape of the SMC grid allows optimised domain splitting and minimised boundary exchanges. The SMC sub-grid setup can be pre-calculated and tuned to save model run time. A new subroutine is added into WW3 to set up the one-to-one spectral boundary exchanges among SMC sub-grids without any interpolation as the unstructured SMC grid allows boundary cells for one sub-grid to be duplicated in its donor sub-grids. In addition, mixed resolution SMC sub-grids can be used, and they ease the computing load balance among sub-grids. The combined CD-DD parallelisation is tested on mixed resolution SMC sub-grids with various hybrid node-thread combinations. Results indicate that switching from MPI to hybrid MPI-OpenMP mode can halve the runtime of our present SMC 3-6-12-25 km global wave forecasting model. Using the hybrid CD-DD method on 3 SMC sub-grids may reduce the elapsed time further by 30%. The elapsed time for one model day is reduced from about 3 min on 12 nodes in single grid MPI mode to less than 1 min on 180 nodes in hybrid multi-grid mode. Besides, the hybrid multi-grid method reduces memory demand on one computing node and allows future model updates for higher resolutions. The UK Met Office is planning to update its global wave forecasting model with 3 SMC sub-grids at 1.25-2.5-5-10-20 km resolutions and preliminary results indicate that the runtime reduction with multi-grid option is more effective when the spatial resolution is increased. For the 1.25-20 km 5 level global grid, the fastest single grid runtime is 205 s for one model day on 270 nodes with 18 OpenMP threads, while the 3 sub-grids run takes only 99 s on 240 nodes with 6 threads, over 50% reduction of model runtime.

1. Introduction

The WAVEWATCH III model (WW3DG 2019) is a widely used 3rd generation ocean surface wave spectral model, which was initially developed by Tolman (1991) and later refined by a group of international wave modellers and computer scientists. Today, the WW3 has become an international community model with comprehensive physics and numerical features. Nevertheless, the WW3 model development is still lagging the computer hardware advance and hindering full use of available computing resources. For example, our present Cray XC40 supercomputer in the UK Met Office has over 10 thousand of computing nodes with nearly half million of processors (cores), but our operational WW3 wave model (Li and Saulter 2014) can use only 10 nodes (360 cores) efficiently, not to mention the advent of high-performance computers based on GPUs.

The WW3 model parallelisation is constructed by Tolman (2002) with a unique wave spectral component decomposition (or CD for short hereafter) and data transport via Message Passing Interface (MPI). The CD method solves the whole model domain propagation of each wave spectral component on a single processor, allowing parallelisation on processors as many as the wave spectral components. As there are about a thousand of components in a normal wave energy spectrum, this method was good enough for the available computers by the turn of this century. This limit is, however, quite restrictive for modern supercomputers which have processors on the order of a million. In addition, the CD method is only effective when two or three spectral components are placed on one MPI rank because of the different computing loads among different frequency components. Low frequency swell components may need more than tripled computing time than high frequency wind sea components because the swell propagation speed is much faster than that of the wind sea. So, mixing a swell component with a wind sea component on one MPI rank mitigates the computing load imbalance. For instance, our UK Met Office global wave model is configured with 1080 spectral components (30 frequencies by 36 directional bins) and its best performance is achieved with 360 MPI ranks on 10 computing nodes, which is only a tiny fraction of our present available supercomputer resources. Hence different parallelisation methods to expand the WW3 model node usage on supercomputers are needed.

Hybrid parallelisation with combined MPI and OpenMP is a straightforward way to stretch the CD parallelisation. Theoretically, it can expand one MPI rank from a single processor to as many processors as on one computing node with shared memory. In practice, the OpenMP method does not show improved scalability beyond a few OpenMP threads. Tolman introduced the OpenMP directives into WW3 model for regular lat-lon grids but did not progress much on this line because of the poor scalability with OpenMP threads. This might be partially due to a bug in the original hybrid code where OpenMP threads were not properly initialised. The hybrid scalability was improved after fixing the bug, but it was still not ideal for large numbers of OpenMP threads. For instance, the Cray XC40 supercomputer used in the UK Met Office has 36 processors on one computing node, but our wave model runtime flattens out beyond 9 OpenMP threads, resulting in an effective limit of 90 nodes for our wave model. What caused this poor OpenMP scalability in WW3 is still not clear.

Domain decomposition (DD) method is a commonly used tool for parallelisation in atmospheric and oceanic models (Tint et al 2017). Although domain decomposition has a history of over a hundred years, it did not become a major tool in numerical models until the last few decades when parallel supercomputers become available. This DD method also found its way in ocean surface wave models, such as SWAN (Booij et al 1999) and WW3. Most of DD applications so far are on rectangular sub-grids, including the latitude-longitude (lat-lon) grid. The advantage of using a rectangular sub-grid is that the decomposition is simple, particularly for homogeneous model domains. One problem related to DD on rectangular grids is computing load balance and it is especially severe in oceanic models where the ocean boundaries are irregular. This is also the case for ocean surface wave models as revealed by Tolman (2008) in one of his WW3 model parallelisation work. This sub-domain imbalance problem is mitigated in unstructured finite-element grids where the sub-domain splitting is more flexible than in rectangular grids. One such an attempt in WW3 model is reported by Abdalali et al (2020). However, finite-element method on unstructured grids is more expensive in computation than finite-difference one on lat-lon grids and is not favoured for large operational models, as the timing results in Abdalali et al (2022) revealed as well.

Tolman (2013) introduced another way to extend the WW3 parallelisation by combining the CD approach with the DD method. This is facilitated by introduction of a so-call multi-grid mode in WW3 version 3.14. In this multi-grid mode, the whole model domain is split into sub-domains and linked up by boundary exchanges like the normal DD method. The difference of the combined CD-DD from the DD method is that each sub-domain is treated like a regional model and is further parallelised by the CD method. The CD-DD method has one advantage over the pure DD method because of the reduced number of sub-domains. The multi-grid method is initially used to replace the NOAA/NCEP nested forecasting system (Chawla et al 2013) and later tested for parallelisation on regular lat-lon sub-grids (Tolman, 2013). Global grids at five spatial resolutions of 2°, 1°, 30′,15′, and 5′, respectively, were decomposed into as many as 24 sub-grids with a special decomposing tool for lat-lon grids. The results revealed that the multi-grid method scalability was almost linear when the number of sub-grids was small, and it flattened out when the number of sub-grids passed an optimal number. In addition, model initialisation time became noticeable when sub-grid number increased due to the optimisation of the computing loads among sub-grids. Increased halo points also contribute to the computing and communicating loads. Another problem with a large model is the memory limit on one computer node. As the MPI method allocates distributed memories to all MPI ranks, data shared among all ranks are duplicated for each rank on one computing node even if the node has shared memory. For a large model, the allocated portion for one rank may not be enough and one way to by-pass the problem is to reduce the number of ranks on one computer node. For instance, the 5′ global grid had to be run with 3 ranks per node, leaving the rest processors on each computing node idle. This is clearly not an efficient usage of supercomputers. If the hybrid MPI-OpenMP method was used, the rest processors might be brought back for some work.

Spherical multiple-cell (SMC) grid (Li 2011, 2012) was initially developed to handle the polar problem in the regular lat-lon grid. The SMC grid relaxes the Courant-Friedrichs-Lewy (CFL) restriction of the Eulerian advection time-step by merging longitudinal cells towards the Poles as in the reduced grid (Gates and Riegel 1962). Round polar cells are introduced to remove the polar singularity of the spherical coordinate system so the whole globe could be covered. Polar errors of vector components, defined by the local east reference direction and assumed to be scalars, are removed by replacing the local east with a fixed reference direction at high latitudes (Li 2016). The SMC grid has combined the unstructured technique with rectangular cells to gain the flexibility of domain shapes and to retain the simplicity of finite-difference schemes. Multi-resolution is enabled in SMC grid by a similar way as the adaptive mesh refinement (Berger and Oliger 1984) except that the refinement is static. Sub-timesteps are used for refined cells for efficient integration and flux-form finite-volume formulations are used to handle multi-resolution cells. Boundary condition at coastlines and the periodic condition for global models are merged into inner cell loops for the convenience of vectorization and parallelisation. The SMC grid was introduced into the WW3 V4.18 model in 2012 as an alternative for the regular lat-lon grid. It is now used for most of UK Met Office wave forecasting models, including the 4-level (3-6-12-25 km) resolution global forecasting wave model (Valiente et al 2023), the North Atlantic ensemble wave model (Bunney and Saulter 2015), and the 2-level (1.5-3 km) UK regional wave forecasting model on rotated SMC grid (Bruciaferri et al 2021). SMC grids have also been used for wave forecasting in other weather centres (Hou et al 2022, Zieger and Greenslade 2021) and Arctic wave climate studies (Casas-Prat et al 2018, Li et al 2019). A 2-D global model of shallow water equations is also tested on the SMC grid (Li 2018) and boundary conditions are updated when the model is converted into a regional model (Li 2021). The SMC grid has potential applications in atmospheric and oceanic dispersion models, and it may reduce computing time by replacing transport schemes on regular lat-lon grids in chemistry models (Li 2019).

The SMC grid module in WW3 has been extended for hybrid parallelisation in the last WW3 V6.07 public release (WW3DG 2019) and multi-grid option is added recently to further expand its computing node usage (Li 2022). The unstructured SMC grid allows more flexible domain shapes than the rectangular lat-lon sub-grids and it uses pre-split sub-grids rather than on-line domain decomposition. This article presents the parallelisation work on SMC sub-grids with the combined CD-DD method in hybrid MPI-OpenMP mode, plus a simplified boundary exchange scheme to reduce the communication load among sub-grids. A brief introduction of the SMC grid is provided in section 2 for readers who are not familiar with it. The hybrid multi-grid updates for the SMC grid in WW3 is described in section 3. Multi-grid configurations of 2, 3 and 4 SMC sub-grids are presented in section 4, including sub-grids for our present 4-level (3-6-12-25 km) global model and our planned 5-level (1.25-2.5-5-10-20 km) future global model. Test results are compared with single grid ones and spectral buoy observations. Timing results of the hybrid multi-grid runs are presented in section 5, and it is followed by a summary and conclusions.

2. A brief description of the SMC grid

SMC grid is an unstructured mesh with rectangular cells. The grid consists of a list of cells at multiple sizes of given unit increments of latitude ($\Delta\varphi$) and longitude ($\Delta\lambda$) as illustrated in Fig.1. The solid red lines mark actual cells in the grid cell list and the red dots indicate their cell centres. A unique 5-element integer array is assigned to each cell to hold its south-west corner *x*-, *y*-indices (i, j), cell side *x-, y*-increment (Δi, Δj) and the cell floor bathymetric height (b), which is rounded to an integer in unit of metre to keep the whole cell array to be integers. The *x*- and *y*-indices are measured in unit or size-1 cell increments ($\Delta\lambda$ and $\Delta\varphi$) so cell centre latitude and longitude are given by

$$\phi_j = \phi_0 + (j + 0.5 * \Delta j)\Delta\phi; \qquad \lambda_i = \lambda_0 + (i + 0.5 * \Delta i)\Delta\lambda \tag{1}$$

where λ_0 and φ_0 are the longitude and latitude of the *x-y* coordinate origin. For a global grid, the origin of the grid can be set at zero-meridian on the Equator so both λ_0 and φ_0 are zero. Cell sizes may be doubled gradually away from highest resolution areas, resulting in multiple resolutions at size 1, 2, 4, 8, …, hence the name of spherical multiple-cell grid. Cells are merged longitudinally at high latitudes to relax the CFL restriction when their physical mesh lengths are less than half of the equatorial mesh length. There is no restriction on the domain shape as long as required domain area is covered by tightly paved cells. Neighbouring cells are not required to be side-by-side in memory as unstructured techniques are used. In fact, all cells are sorted by their sizes in a final cell list so that sub-timesteps could be used for refined cells in nested loops. For ocean surface wave models only sea point cells are used for wave propagation.

Spatial relationships among the unstructured cells are precalculated and stored in face arrays. For each cell face, a 7-element face array is used to store its location and size indices and the upstream, centre, and downstream cell identity numbers. The neighbouring cell information is enough for use of the 2nd and 3rd order upstream non-oscillatory advection schemes (Li 2008) on the SMC grid. The face arrays are also sorted by their face sizes for use of sub-timesteps. These cell and face arrays are used for calculation of advection, diffusion, spatial average, and directional gradient on the SMC grid. Finite-volume fluxes are calculated first on cell faces in nested sub-loops and fluxes into each cell are then accumulated in a temporary variable for a net-flux, which is used for cell update in corresponding cell sub-loops. For cells with multi-resolution interfaces, the fluxes through different sized faces are automatically averaged over sub-timesteps for large cells.

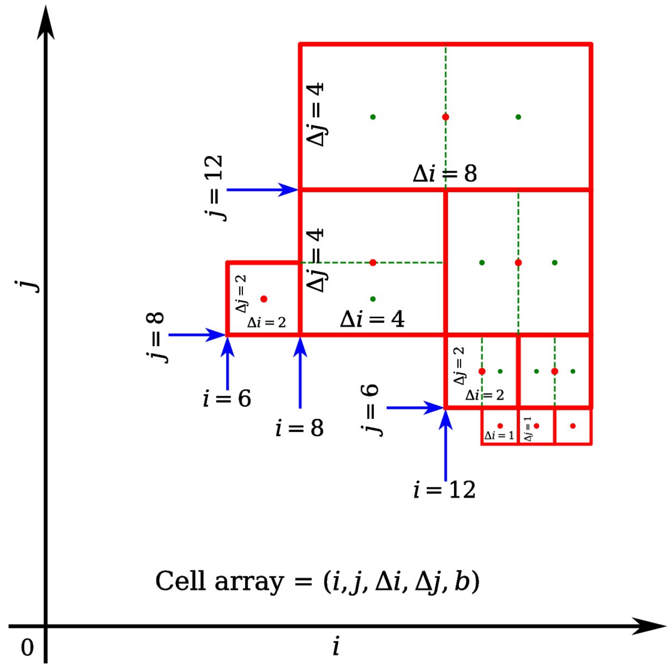

Fig.1. Cell arrays for a SMC grid and virtual cells over size changing interfaces.

For refinement interface, such as a cell face between a coarse and two refined cells, SMC grid assumes a first-order approximation that the coarse cell is uniform within its cell area and it could be divided into two halved cells of the same properties as the coarse cell, such as velocity or wave energy spectrum. This approximation avoids interpolations between different sized cells, whose centres are unaligned due to doubled size increasement. The dashed green lines in Fig.1 indicate how the coarse cells marked by the solid red lines are divided into halved cells. The green dots indicate the centres of those halved virtual cells, which are aligned with the refined cells. Finite-difference across such a refinement interface is then approximated by averaged differences between the halved cells and the refined cells. More details on the refinement interface approximation are available in Li (2021). The sub-timestep loops and the halved virtual cell approximations for multi-resolution SMC grid allow a smooth transfer between different resolution zones with optimised efficiency.

Zero boundary condition at coastline is used for ocean surface wave models and it is conveniently applied on the SMC grid by bounding an empty cell beyond a coast cell in its face array. In fact, any cell face not bounded by other cell is treated as 'coastline' and the zero-boundary condition is applied. For a global model, a periodic condition is usually applied in the longitude direction and this is achieved on the SMC grid by simply linking up the last longitude cell with the first one in their face array. These two boundary conditions convert all boundary faces into equivalent inner faces, so no extra boundary loops are required for coastlines and period boundaries. For regional model with open boundaries, zero-boundary condition is also applied on its edge faces. However, SMC grid updates a list of boundary cells with boundary conditions provided by a large or parent model in an extra boundary cell loop for a regional model. This is like regular lat-lon grid models except for that SMC grid boundaries may be irregular shaped like coastlines. This irregular boundary feature makes the SMC grid more flexible than the regular lat-lon grid with the DD method.

3. Hybrid multi-grid updates for the SMC grid in WW3

The SMC grid hybrid MPI-OpenMP option was included in the WW3 V6.07 public release (WW3DG 2019). It was later noticed that a bug in hybrid MPI-OpenMP initialization had some influence on the hybrid scalability. The OpenMP threads were not properly initialised by the `MPI_INIT` function, which is designed for the pure MPI mode. After replacing it with the `MPI_INIT_THREAD` function for hybrid initialization, the scalability was improved. Some hybrid tests on both the regular lat-lon and SMC grids were done after the bug was fixed. Results indicated that the elapsed time was reduced by about 15% on average by this line change. Further tuning of the SMC module by exclusion of flux updates for zero boundary cells has noticeable reduction in model elapsed time. The net improvement for the SMC grid hybrid mode is about 30% in comparison with the last public release code. This will be discussed further in the timing result section.

The WW3 model CD parallelisation (Tolman 2002) is only applied on the propagation terms, such as the advection and diffusion terms. Other point-based terms, such as wind generating source terms and non-linear wave interactions, are parallelised by localised sea points. The CD scheme removes land points from wave spectral storage and distributes the sea-point wave spectra evenly among MPI ranks. Sea-point based terms are calculated by looping over the local sea points on each MPI rank. For hybrid MPI-OpenMP mode, further parallelisation of the local loop among OpenMP threads on each MPI rank is done with additional OpenMP directives. The sea-point based OpenMP parallelisation is also applicable to the SMC grid because SMC grid uses sea-point based schemes for all terms.

The SMC grid was introduced into the WW3 model in 2012 as an alternative for the regular lat-lon grid. As SMC grid is framed on sea-points it does not need to expand the sea-point wave spectra to full global grid for propagation and to convert them back afterwards at each advection timestep as for the regular lat-lon grid. One extra benefit is that it could extend the wave model domain to cover the whole Arctic ocean in response to the Arctic sea-ice retreat in recent summers. The multi-resolution feature also allows the SMC grid to merge global and regional models into a single domain model (Li and Saulter 2014). However, the computing node limit of the CD method becomes an obstacle for our Met Office wave forecasting model even if the hybrid mode is activated. There are thousands of computing nodes available now, but our wave model can only use less than a hundred nodes efficiently in hybrid mode. The CD-DD combined method is one we recently explored to extend the node usage for our SMC grid wave forecasting model.

Bringing the SMC grid into the WW3 multi-grid framework is not straightforward because the original WW3 multi-grid option is designed for regular lat-lon grids. From version 4.18, there is an option to use curvilinear and unstructured triangle-cell grids in the multi-grid mode, but they are restricted for lower ranked grids, that is, these grids can only be used in series mode rather than in parallel mode as equal ranked grids. Besides, the core component for communication among sub-grids is an interpolation in space to accommodate sub-grids at different spatial resolutions. Although this interpolation scheme follows an energy conservative method (Jones 1999), it requires large overlapping zones more than two boundary rows as the propagation scheme required. This increases the communication loads especially for high resolution models.

Three different features are introduced for the SMC multi-grid option in this update. One is that the DD procedure is pre-calculated instead of online within each model run. This is possible because the SMC grid uses pre-calculated sea-point cells instead of online removal of land points as for the regular lat-lon grid. Besides, SMC sub-grid boundary points are also specified in advance by a boundary cell list so there is no need to look for boundary points online as in the regular lat-lon grid. Each SMC sub-grid can be tested independently just like a regional grid with given boundary conditions.

The second feature of this SMC multi-grid update is that the boundary exchange among SMC sub-grids is simplified as whole wave spectral swapping. No spatial interpolation is required, thanks to its multi-resolution feature, which allows the overlapping boundary zone to be exact duplication of neighbouring sub-grid cells. For each sub-grid, consecutive cells beyond its boundary line and within the distance of three base resolution cells are added to the sub-grid as boundary cells. Around small islands or at step joints of oblique boundary lines, extra boundary cells are added to ensure the effective width to be three-cell wide. The overlapping zone between two neighbouring SMC sub-grids is then approximately 6 base resolution cells wide. Wave spectra for those boundary cells are updated by the wave spectra from the corresponding cells in its neighbouring sub-grid at main timesteps, which is set to be less than the time required by the fastest wave to travel across one base resolution cell. For instance, the long wave component at frequency 0.05 Hz has a group speed about 15 m s^{-1} in deep waters, which needs about 28 min to cross one base resolution cell of 25 km wide. The boundary conditions are updated at the main timestep of 15 min in this study. However, inherent numerical diffusion in the advection scheme and the explicit diffusion scheme used for smoothing purpose could pass the zero-boundary effect across the overlapping zone if the number of advection sub-timesteps within one main timestep is larger than the number of boundary cells across the overlapping zone. This boundary diffusion effect could be reduced by increase of the boundary line width. In this study, 2- and 3-base cell wide boundary lines are tested, and the 3-cell width is much better than the 2-cell case. To minimise WW3 model code changes, the communicating arrays for the lat-lon sub-grids are reclaimed for SMC sub-grids and a new subroutine is added to set up communication arrays for equal ranked SMC sub-grids without interpolations, leaving the lat-lon grid interpolation scheme unchanged.

The third change is to use sea-point forcing winds for SMC sub-grids. The default WW3 wind forcing is assumed to be on lat-lon grids, and this is kept for the SMC grid as well. During model runs, the regular grid wind is interpolated to sea points for either a regular or the SMC grid because only sea points are used for source terms. As SMC grid cells are already sea-point only, the wind forcing can be interpolated on to SMC cell points before model runs. This change not only contributes to the model run time reduction by avoiding the online spatial interpolation, but also makes it possible to merge different resolution winds for multi-resolution SMC grids. For instance, in this multi-grid test, using sea-point wind alone may reduce the model run time by about 5% in comparison with runs forced by lat-lon grid winds. The timing plots in the following section are all based on sea-point wind runs.

It is also worth to emphasize the differences between the SMC multi-grid option and the existing unstructured triangle cell DD methods in WW3. The triangle cell DD method (Abdolali et al 2021) replaces the original CD parallelisation, instead of combining with it. The spectral-spatial domain is split into several sub-domains in the triangle cell DD parallelisation. Spectral component propagations in all these sub-domains are simulated synchronously and the information exchange between adjacent sub-domains are performed within the MPI environment. In such an environment, the decomposed N sub-domains can be solved by N processors to achieve the parallel computations. Besides, an implicit propagation scheme is added to the triangle cell grid module and it accelerates the model integration. The triangle cell DD procedure is also done online in each model run as the multi-grid mode for regular lat-lon grids. The SMC multi-grid option follows the CD-DD combination method but simplifies the boundary exchange and sets up the sub-grids prior to model runs.

The latest WW3 model has over a hundred of modules and more than 20 thousand lines of source code. This SMC multi-grid update has modified over 30 modules and added 3 new subroutines. All these changes have merged into the WW3 trunk in July 2021 and are freely available to the public on the open-source web site: https://github.com/NOAA-EMC/WW3/. Interested readers may refer to this web site for details of the update.

4. SMC sub-grids parallelisation tests

The SMC36125 global grid (Li 2016) is used as a reference grid for this study, which has a base resolution of 25 km over most open ocean surfaces and refined to 12, 6 and 3 km resolution cells near coastlines and in the European waters. The global grid is split into 2, 3 or 4 sub-grids, respectively for multi-grid tests. Fig.2a shows the first sub-grid of the 3 sub-grids case, which covers the North Atlantic and the Arctic and will be referred to as the Altn36125 grid. The European area is refined to 12, 6 and 3 km in the global SMC36125 model so that it could merge our European regional model. Its latitude and longitude increments at the finest resolution or for a size-1 cell are $\Delta\varphi = 0.029296875°$ and $\Delta\lambda = 0.0439453125°$, respectively, which are roughly 3 km at mid-latitudes. The longitudinal increment is chosen so that merged cells at high latitudes are exact multiples of size-8 cells. The cell floor bathymetric depth at each cell is indicated by its edge colour. There are 138013 cells in the Atln36125 sub-grid, including 17248, 25891, 34436 and 60438 cells in size 1, 2, 4 and 8, respectively. It also contains 1321 size-8 cells in the Arctic part within the red circle shown in Fig.2a.

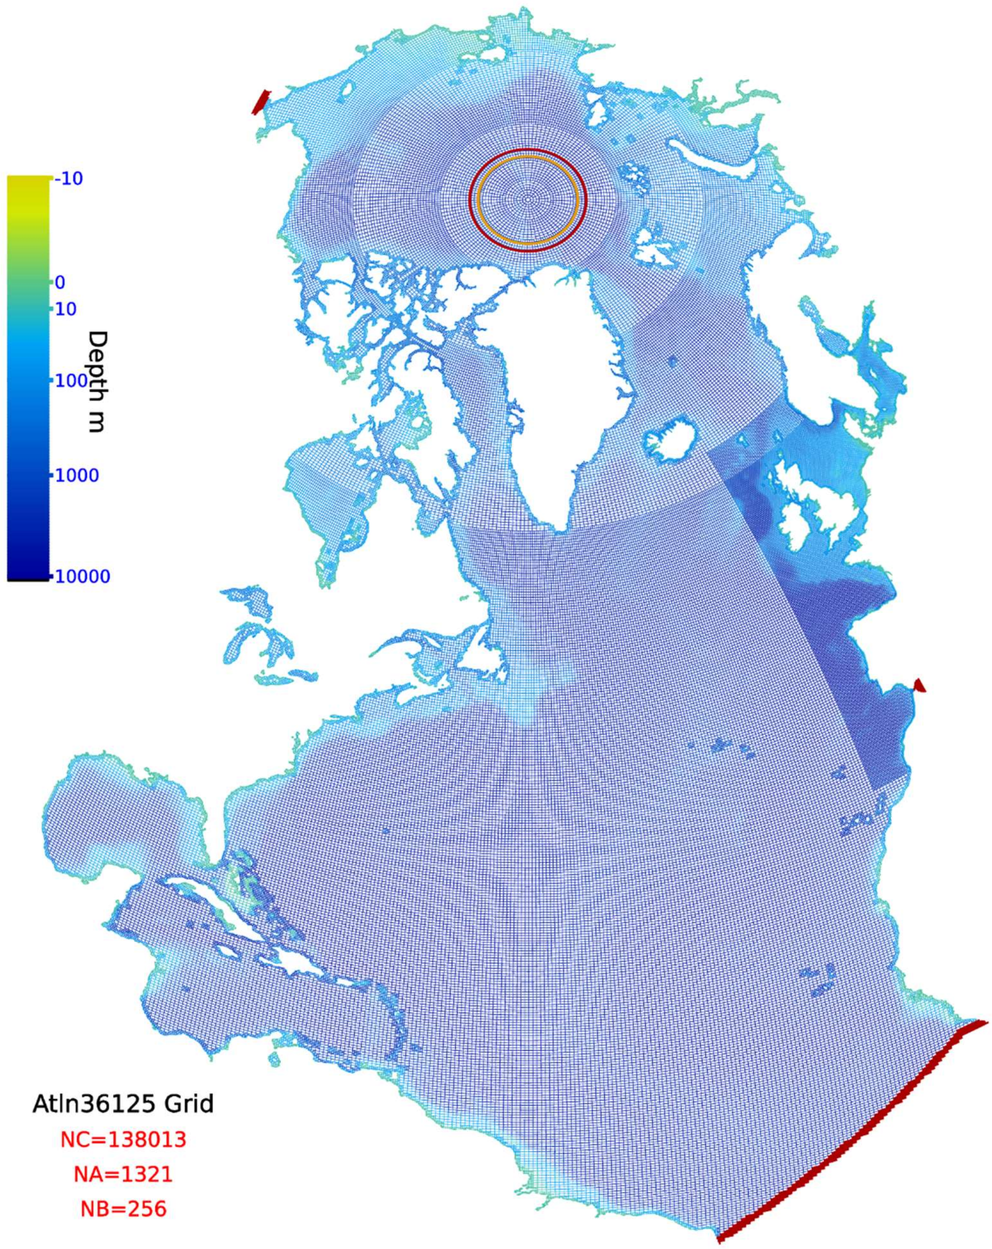

Fig.2a. The Atln36125 sub-grid for SMC multi-grid test, including the whole Arctic.

The orange circle marks the overlapping zone for linking the Arctic part with the rest. The filled red cells in the Bering Strait (on top of Fig.2a), in the Gibraltar Strait (to the right edge in Fig.2a), and along the oblique line connecting the South America and Africa continents (at bottom of Fig.2a) are boundary cells to link the Atln36125 grid with neighbouring sub-grids. The sub-grid generally follows natural coastlines and makes open water boundary as short as possible to minimise sub-grid communication. As a result, this sub-grid has only 581 boundary cells. These boundary cells are duplicated in neighbouring sub-grids.

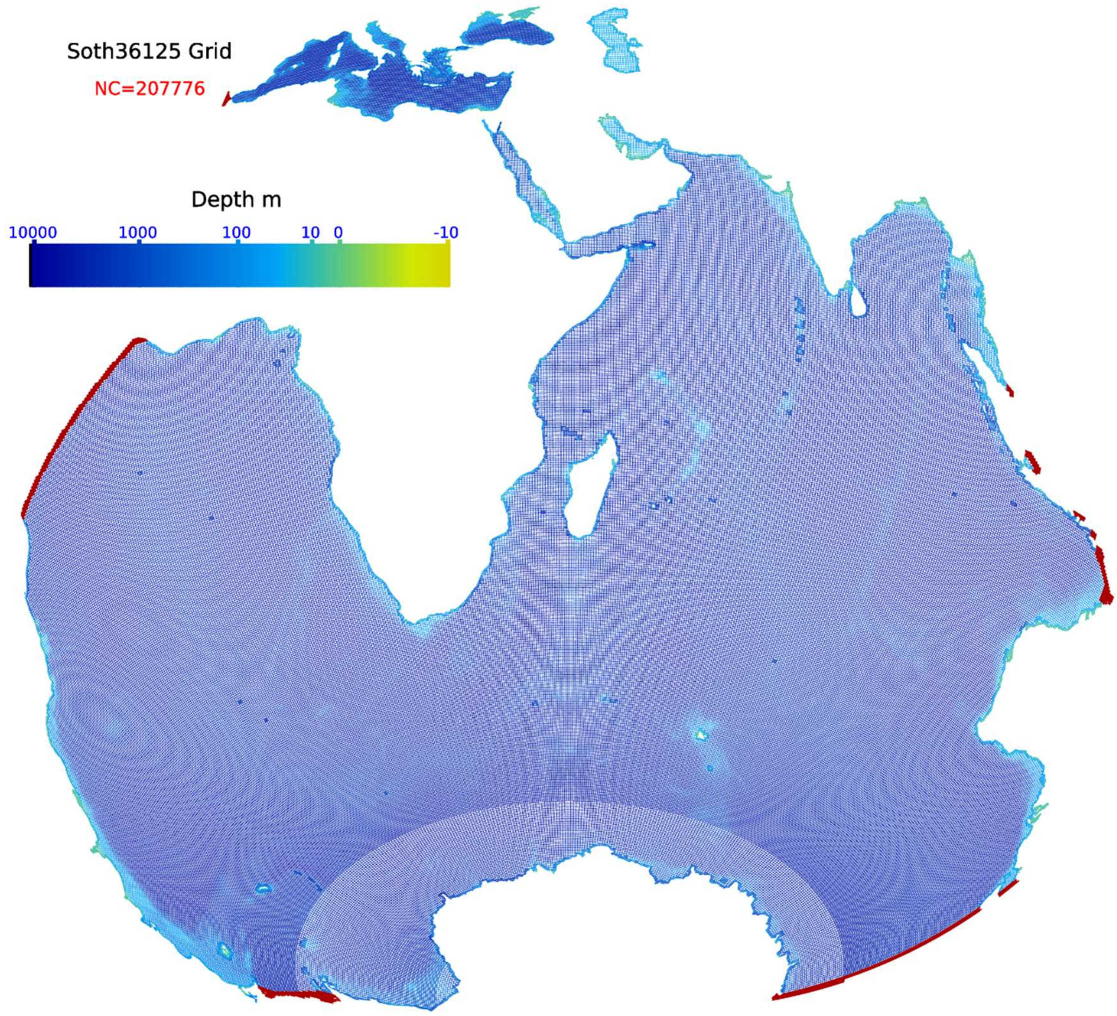

Fig.2b. The Soth36125 sub-grid for SMC multi-grid test.

The sub-grid in Fig.2b covers the South Atlantic, the Indian Oceans and the Mediterranean Sea. It will be referred to as the Soth36125 grid. Its water boundaries are optimised by crossing the Drake Passage in the west and the minimum distance from Antarctic to Australia in the east. Above Australia the boundary line crosses the Timor Sea and then follows the Indonesia islands and Malaysia peninsula. There are 207776 cells in total (19114, 20174, 24922, 143566 for size 1, 2, 4, 8, respectively) and 1888 boundary cells in the Soth36125 grid. Fig.2c shows the sub-grid that covers the Pacific Ocean, hence the name of Pacf36125 grid. It has 251981 cells in total (566, 23200, 26212, 202003 for size 1, 2, 4, 8, respectively) and 1238 boundary cells.

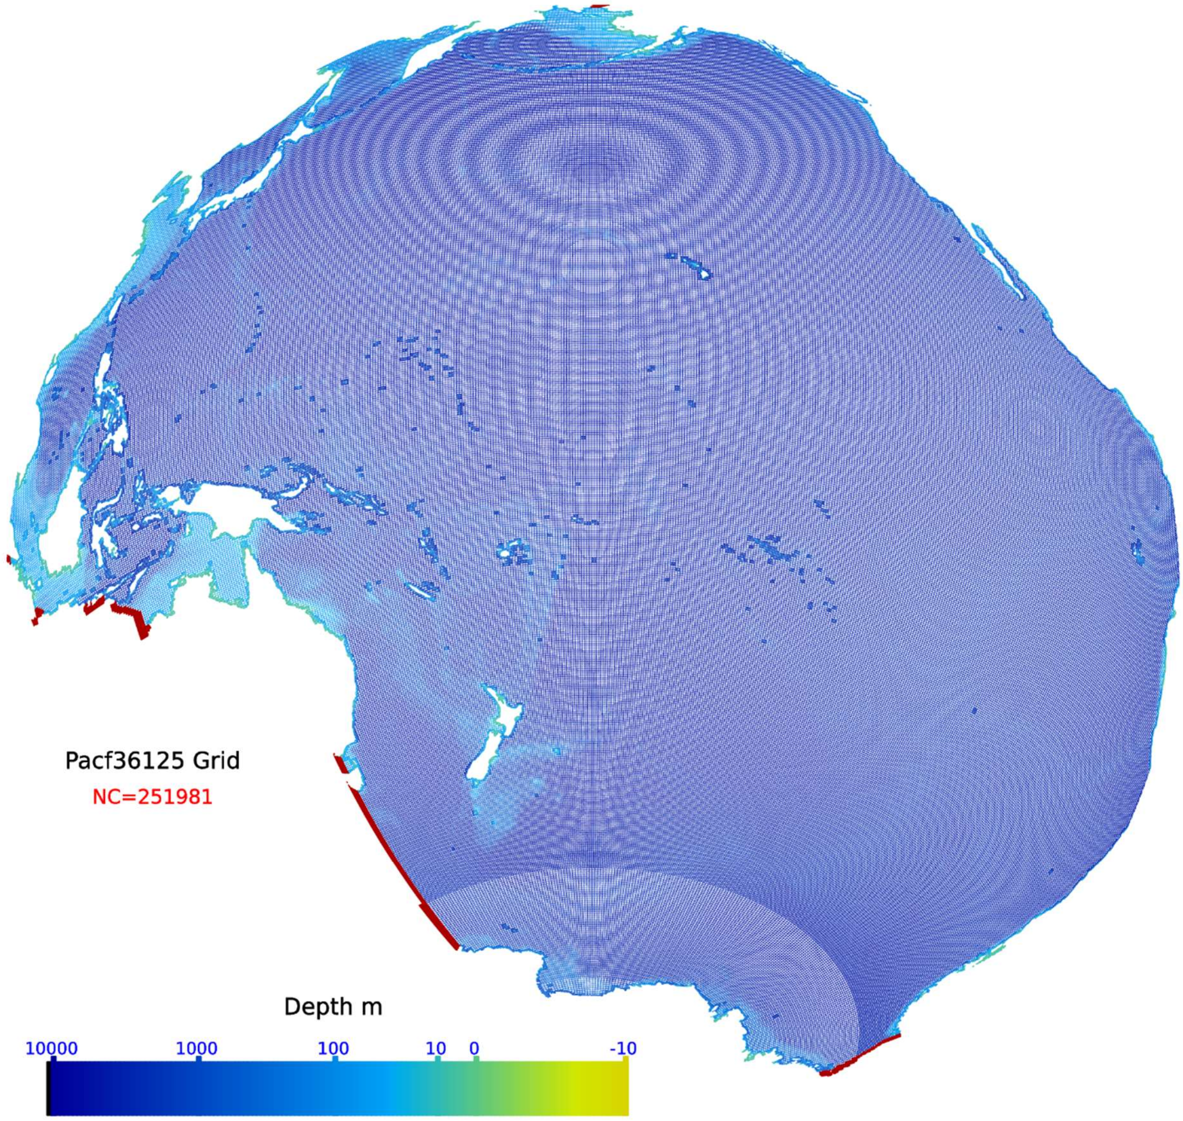

Fig.2c. The Pacf36125 sub-grid for SMC multi-grid test.

SMC grid uses sub-timesteps for propagation over refined cells, that is, for each advection timestep of the base resolution cells (size-8 cells in SMC36125 grid), refined size-4 cells will have two sub-timesteps, size-2 cells have 4 and size-1 cells have 8 sub-timesteps. As a result, the actual propagation computing load will be proportional to an equivalent cell number estimated by $N_E = 8\times N_1 + 4\times N_4 + 2\times N_2 + N_8$, where N_m represents the size-m cell number. The equivalent cell number for the Atln36125 grid is 370853. For the Soth36125 and Pacf36125 grids, the equivalent cell numbers are 427018 and 351755, respectively. Note that the Pacf36125 sub-grid has the smallest equivalent cell number despite that its total cell number is the largest among the 3 sub-grids. This is because it has fewer size-1 cells than the other two sub-grids. Also note that computing loads of sea-point based source terms are proportional to the actual sea-point number or the total cell number. Hence the balance of the net computing loads will be a compromise between the equivalent and the total cell numbers. The net computing loads of these SMC sub-grids can be estimated by the elapsed times of running each sub-grid independently as a regional grid on same number of computing nodes. The elapsed times for the Atln, Soth and Pacf36125 sub-grids are 44.98, 50.77 and 55.00 s, respectively on 20 computing nodes of our Cray XC40 machine. These timing results are used to determine their resource ratios for hybrid multi-grid runs. As the 3 sub-grids have nearly balance computing loads, a simple 1/3 ratio is assigned to each sub-grid. These cell numbers and timing information for the 3 sub-grids are listed in Table 1 for clarity. The single SMC36125 grid parameters are also listed there for comparison.

Table 1. Information of the single SMC36125 grid and the 3 sub-grids split from it. NC is the total number of cells, N1-N8 are the numbers of size-1 to 8 cells in each grid. NE is the equivalent cell number and NB the number of boundary cells for the grid. MGR is the multi-grid distribution ratio and ET the elapsed time of one model day integration for each model on 20 nodes (1440 cores).

Grid\Param	NC	N1	N2	N4	N8	NE	NB	MGR	ET(s)
Soth36125	207776	19114	20174	24922	143566	427018	1888	0.333	50.77
Pacf36125	251981	566	23200	26212	202003	351755	1238	0.333	55.00
Atln36125	138013	17248	25891	34436	60438	370853	581	0.333	44.98
SMC36125	594063	36617	68305	84652	404489	1139949	0	1.00	105.4

The flexibility of the SMC grid allows not only optimised grid splitting but also convenient sub-grid merging for visualization. Fig.3a illustrates the 3 sub-grids merged into a single global grid, which is identical to the global SMC36125 grid except for the overlapping boundary cells (filled red in Fig.3a).

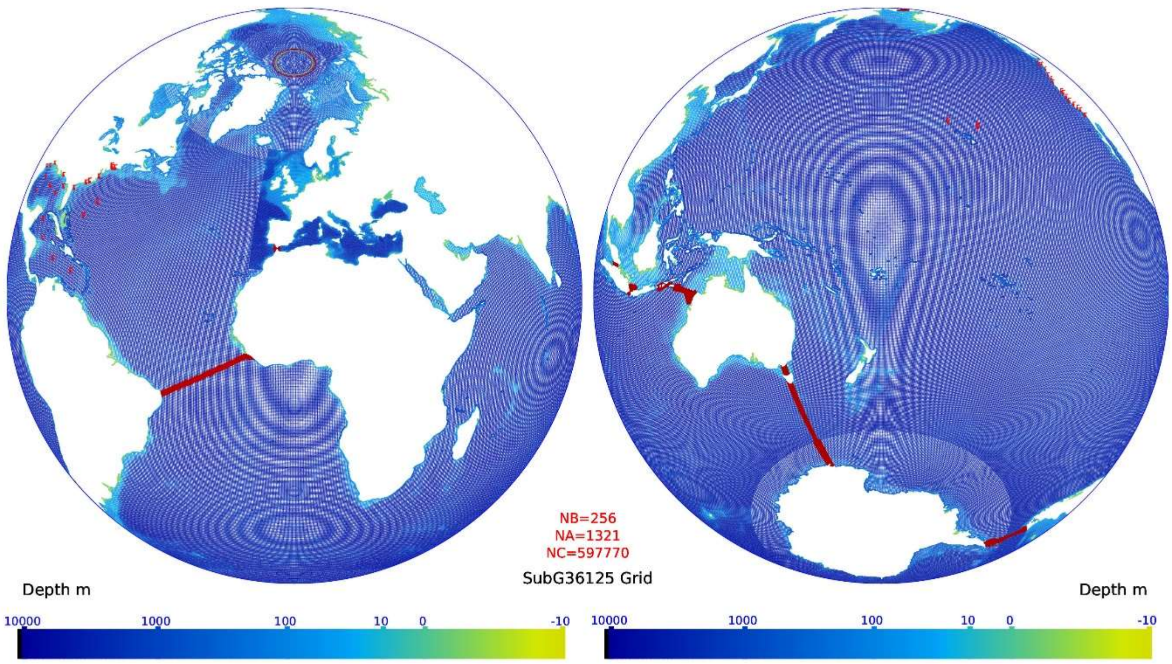

Fig.3a. The merged 3 sub-grids, equivalent to the global SMC36125 grid except for boundary cells.

The hybrid MPI-OpenMP updates are first tested with the SMC36125 single grid, and its results are taken as references for later hybrid multi-grid tests. The hybrid multi-grid option is tested with 2 SMC sub-grids at first. The 2 sub-grids roughly split the world oceans into two parts by a parallel line across the Southern Atlantic Ocean at about 22°S and a boundary line across the Bering Strait as in the 3 sub-grids case shown in Fig.2. The 3 sub-grids configuration shown in Fig.2 is found to be better than the 2 sub-grids one by the measure of model run time, which will be discussed in the next section. Outputs from sub-grids can be used as independent regional model outputs or merged into a single global one. This provides more flexible options for postprocessing of model results than a single global grid output.

Fig.3b shows a typical significant wave height (SWH) plot, merging results from the 3 sub-grids. SWH is an integrated scalar field of the wave spectrum over all directions and the full frequency range. It is a practical quantity to be measured and a convenient variable for assessment of wave model performance. The merged SWH field looks fine and there is no visible interruption at the boundary zones in comparison with the global SMC36125 grid result.

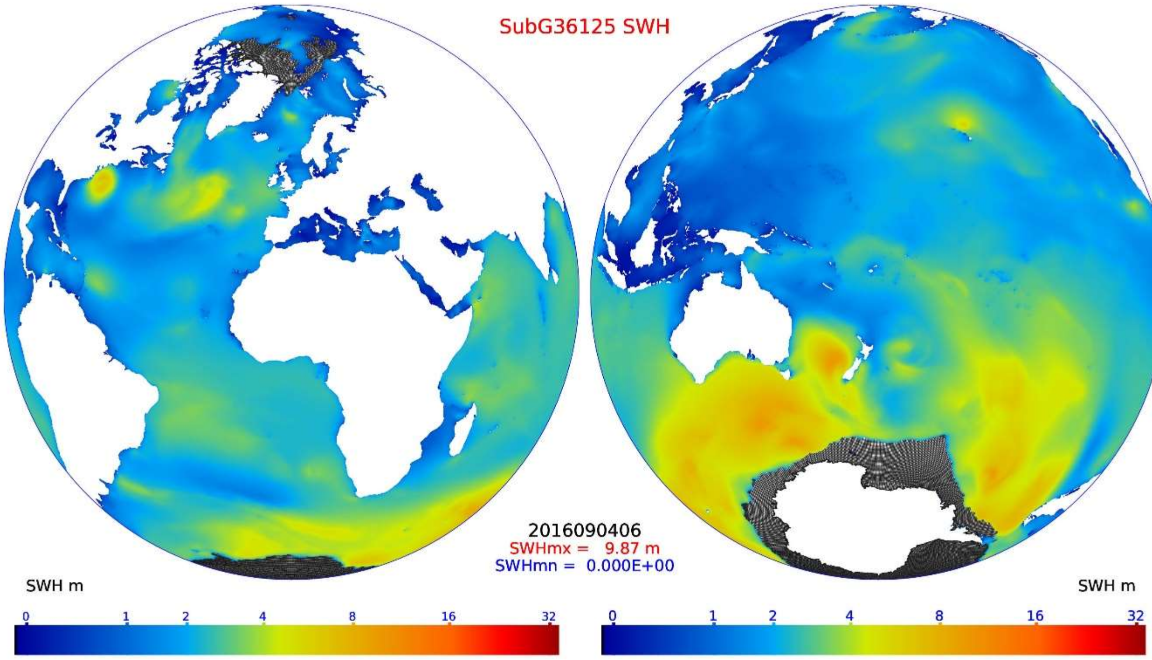

Fig.3b. A merged SWH field from the 3 sub-grids at 0600 on 20160904.

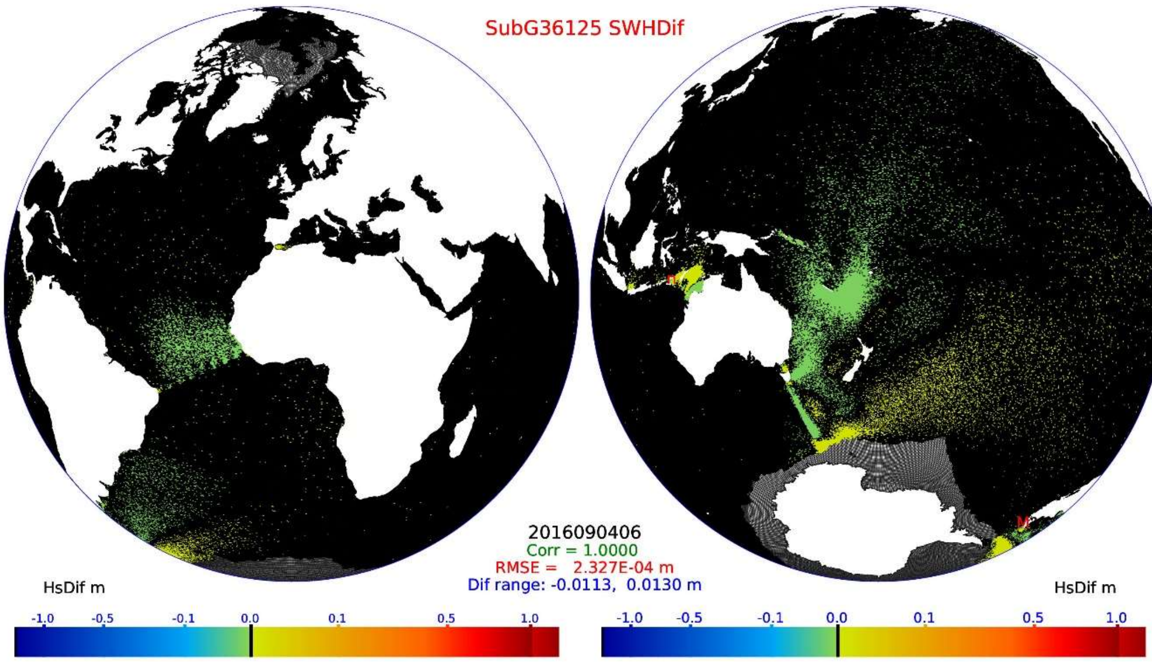

Fig.4a. SWH differences between 3 sub-grids and SMC36125 grid models at 0600 hr on 20160904.

A close comparison of the SWH fields between the 3 sub-grids and the single global SMC36125 grid reveals some subtle differences as shown in Fig.4a. It shows the cell-by-cell SWH difference between the 3 sub-grids and the single SMC36125 grid models at 0600 hr on 4 September 2016, the same time as the SWH field shown in Fig.3b. The range of the SWH difference is between -0.011 and

0.013 m though most of them are close to zero. The non-zero differences appear in the downstream boundary area, indicating that they are caused by the boundary conditions. The largest difference (0.013 m) is near the boundary line across the Drake Passage, as shown by the 'M' letter in Fig.4a. A possible reason for this difference is that the effective boundary width is narrowed by a small island within the boundary zone. The minimum SWH difference -0.11 m is in the boundary zone across the Timor Sea as marked by the 'n' letter in Fig.4a. The errors are also caused by small islands within the boundary zone. Widening the boundary zones around small islands may reduce the SWH difference and the boundary tuning has been done for the present result. Otherwise, the maximum SWH difference could exceed 0.88 m. There are also some broken lower SWH patches above the south Atlantic boundary line (refer to Fig.2a) and this broken pattern coincides with the oblique boundary line steps as its effective wides is reduced at these step joints. These SWH differences indicate that consistent boundary width should be maintained for all sections, particularly for oblique ones and around small islands. Apart from those minor differences, the sub-grid SWH field is almost identical to that of the single global grid. The overall root mean square error (RMSE) is quite small (2.3E-4 m for this time) and the correlation coefficient is perfectly 1.0.

Fig.4b shows the SWH difference range and overall RMSE in September 2016, sampled at 6 hr intervals. The maximum SWH difference range is between -0.08 and 0.04 m and the maximum RMSE is about 3.65E-4 m. These results confirm that the 3 base cell wide boundary lines plus widened boundary zones around small islands are good enough if the main timestep is set to be shorter than the long wave travel time over one base resolution cell. Boundary lines of 2 base resolution cell wide are also tested and they lead to increased SWH differences in comparison with the 3 cell wide ones. It is envisaged that wider boundary line may further reduce the boundary effect, but the communicating load will be increased due to increased number of boundary cells.

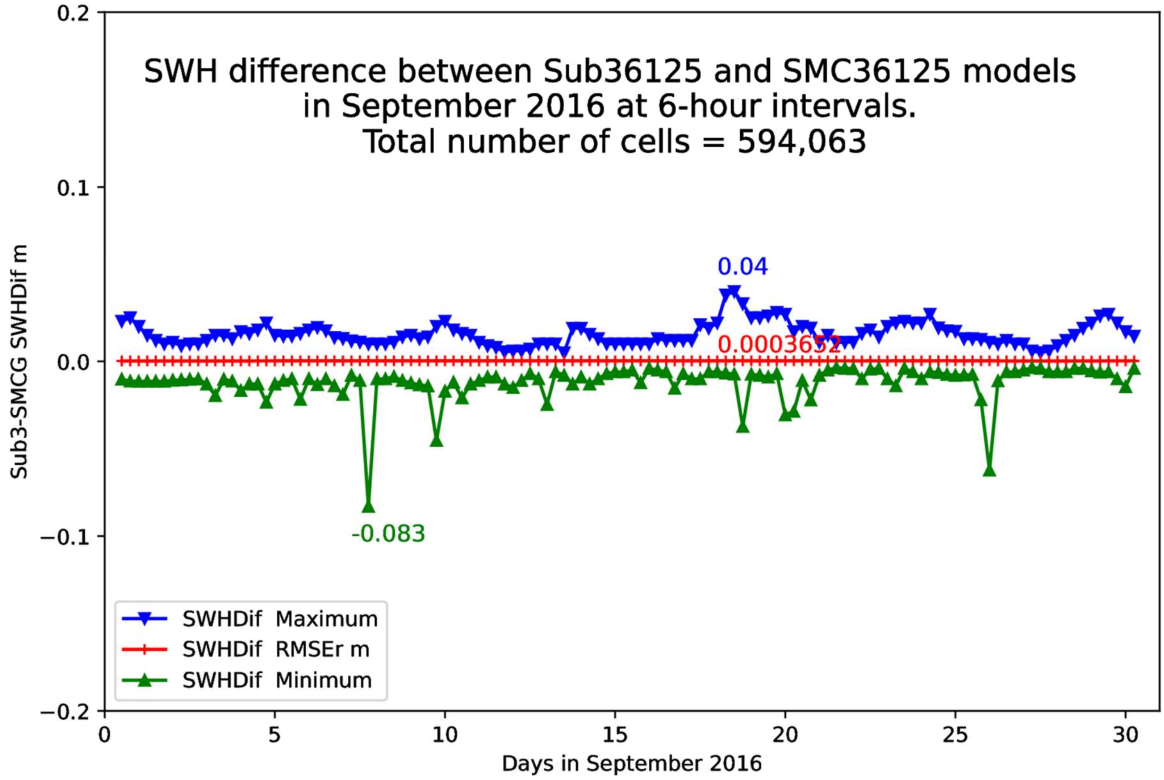

Fig.4b. Range of SWH differences between 3 sub-grids and SMC36125 grid and their global averaged RMSE at 6-hr intervals in September 2016.

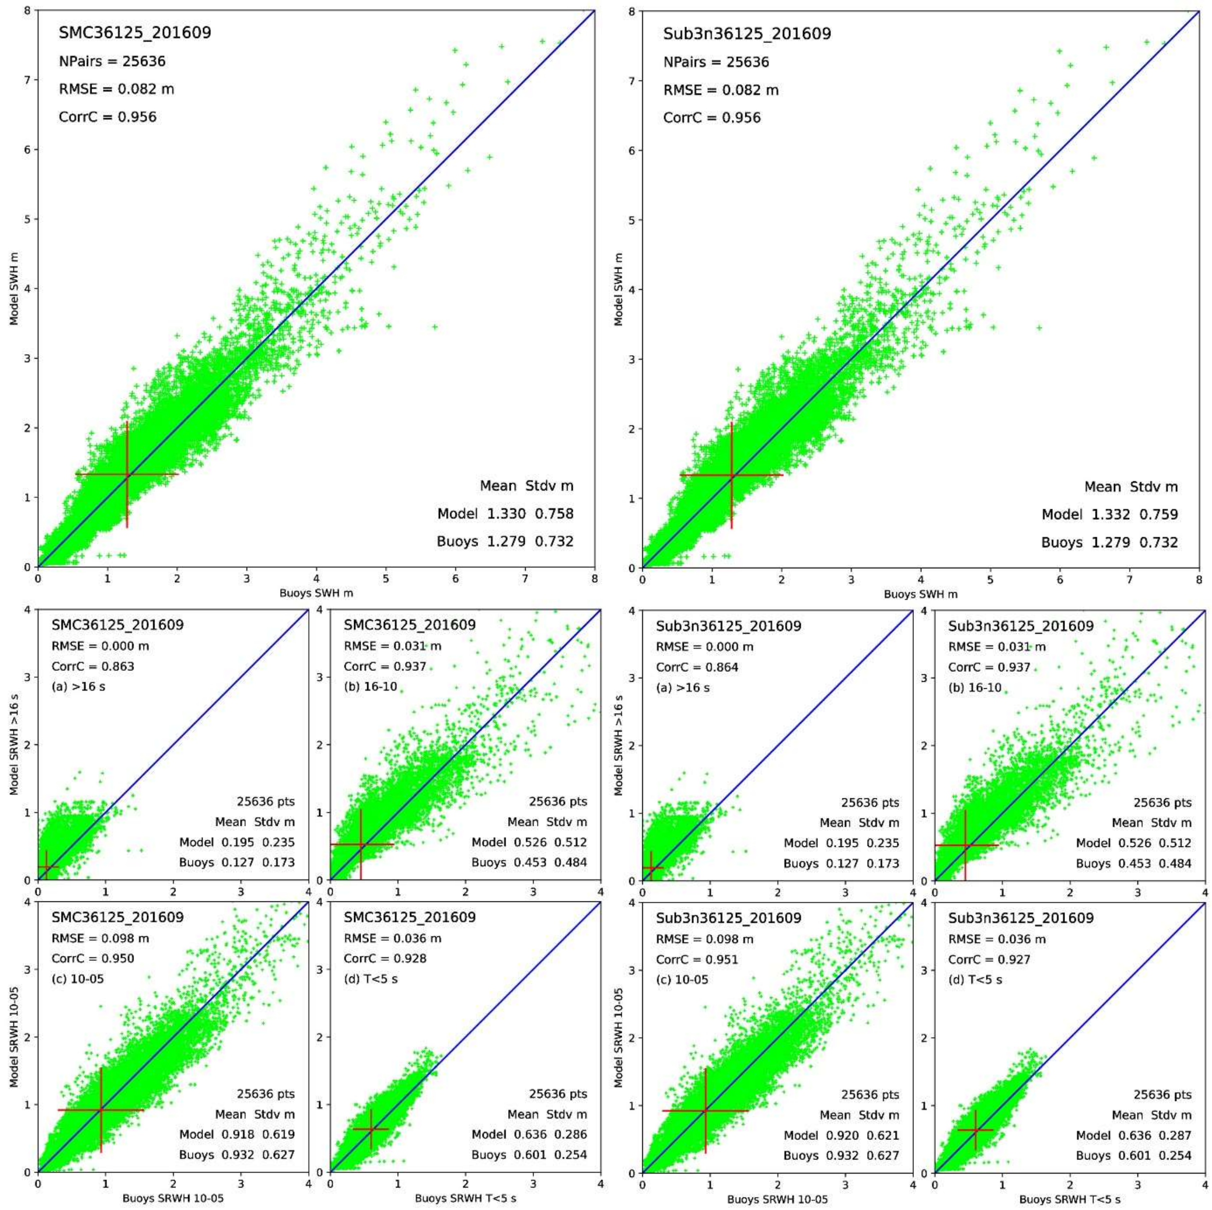

Fig.5. SWH (top row) and SRWH (bottom quadruple panels) comparisons of the SMC36125 (left) and the 3 sub-grids (right) via spectral buoy observations. SRWH bins are indicated by wave period ranges.

Spectral buoy observations are also used for validation of the 3 sub-grids and comparison with the single global grid. The spectral buoy data are from the NDBC web site (www.ndbc.noaa.gov) and buoy locations are indicated by the red 'r' markers in Fig.3a. Most of the spectral buoys are distributed along the American coastlines, with a few around the Hawaii Islands. Fig.5 shows the comparisons of the single SMC36125 grid and the 3 sub-grids results with spectral buoy observations. About 36 spectral buoys were active in the one-month study period (201609) and there are total of 25636 buoy spectra passed quality checks.

The reason to choose spectral buoy observations for this validation is that they can reveal the spectral performance of the model, particularly the wind sea and swell characteristics. The so called 4-bin sub-range wave heights (SRWH) are used for the spectral assessment. Definition of the SRWH is the same as the SWH except that the frequency integration is chopped into 4 sub-ranges (Li and Saulter 2012). Generally, the first SRWH bin (wav period T > 16 s) can be treated as pure swell energy and the last bin (T < 5 s) is all wind sea. These two bins can be used for model tuning of the wave dissipation and wind generating terms, respectively. The two intervening bins (10 s < T < 16 s, and 5 s < T < 10 s) may be treated as mixture of swell and wind sea. As most of the spectral buoys are located far away from the sub-grid boundary lines, the boundary effect will only affect the long-distance swell field at the buoy sites. Hence, the SRWH swell bin will be a good indicator whether the sub-grid boundary conditions are set up properly.

Panels on the left half of Fig.5 are the SWH and SRWH scatter plots for the single SMC36125 grid and panels on the right half are for the 3 sub-grids. The statistical parameters of the total SWH are identical for the two models (RMSE = 0.082 m and correlation coefficient CorC = 0.956). The statistic numbers of the SRWH swell bin (T > 16 s) are nearly identical as well except for the small difference in the correlation coefficients (0.863 via 0.864). Note that the swell energy is quite small as the scatter points are almost clustered near the origin, the correlation coefficient is not quite reliable to confirm that the 3 sub-grids are better than the single global grid. The next bin for wave period between 10 and 16 s also contains some swell energy and the two models also show identical statistics in this SRWH bin. The other two SRWH bins are dominated by wind sea and their statistical numbers are almost identical as expected because the sub-grid boundary effect would not affect the wind sea near these buoys. It could be concluded that this buoy spectral comparison does not reveal any fault in the sub-grid boundary condition, and the 3 sub-grids model yields almost identical statistical numbers as the single grid model.

The UK Met Office is planning to update its wave forecasting model with a 5-level (1.25-2.5-5-10-20 km) and 3 sub-grids model and a test model has been developed recently, using the latest GEBCO_2022 bathymetry data (GEBCO Compilation Group 2022) and refined SMC grid generating tools (Li 2023). The merged 3 sub-grids are shown in Fig.6a. It consists of a 5-level Atlantic and Arctic grid (Atn125120) and two 4-level sub-grids (Sth251010 and Pcf251020). Only the European coastlines are refined to the highest resolution of 1.25 km while other coastlines are refined up to 2.5 km. The splitting boundary lines are simplified for better computing load balance among the 3 sub-grids. As most of our customers are in the European area, the Atn125120 grid is designed to cover the essential area as shown in Fig.6b. The Sth251020 sub-grid covers the whole Southern Ocean for a better representation of the circular flow. The rest are put into the Pcf251020 sub-grid, including the Great Lakes and the Caspian Sea, thanks to the flexible domain feature of the SMC grid.

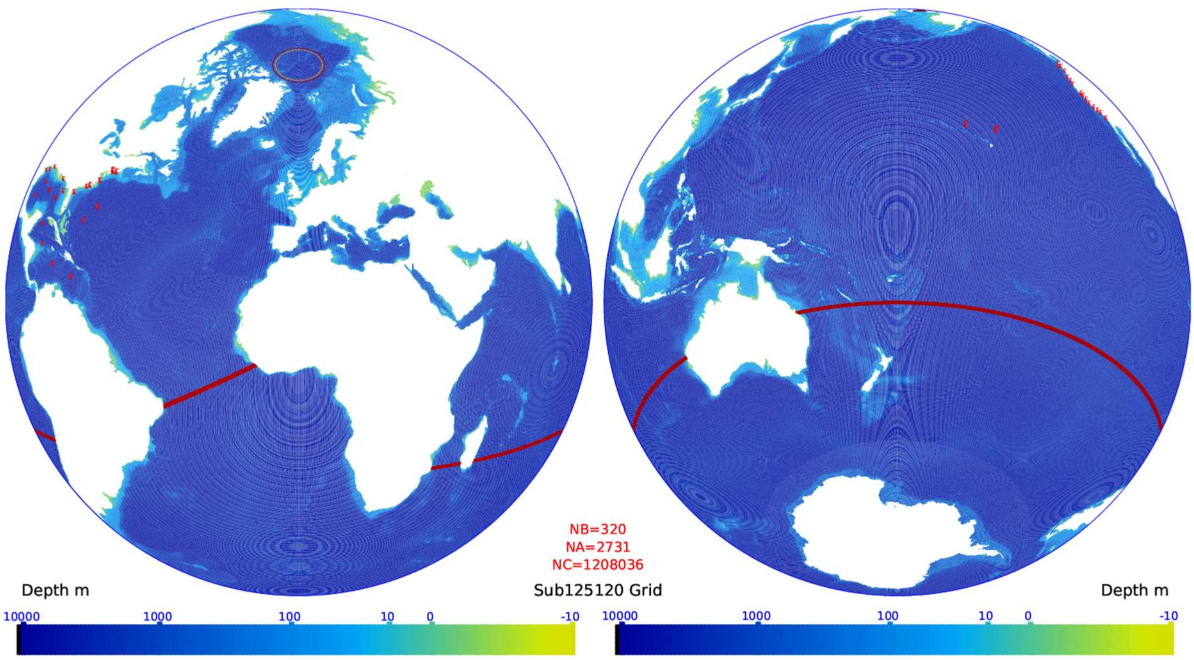

Fig.6a. The planned UK Met Office multi-grid global wave model grid at 1.25 to 20 km resolutions.

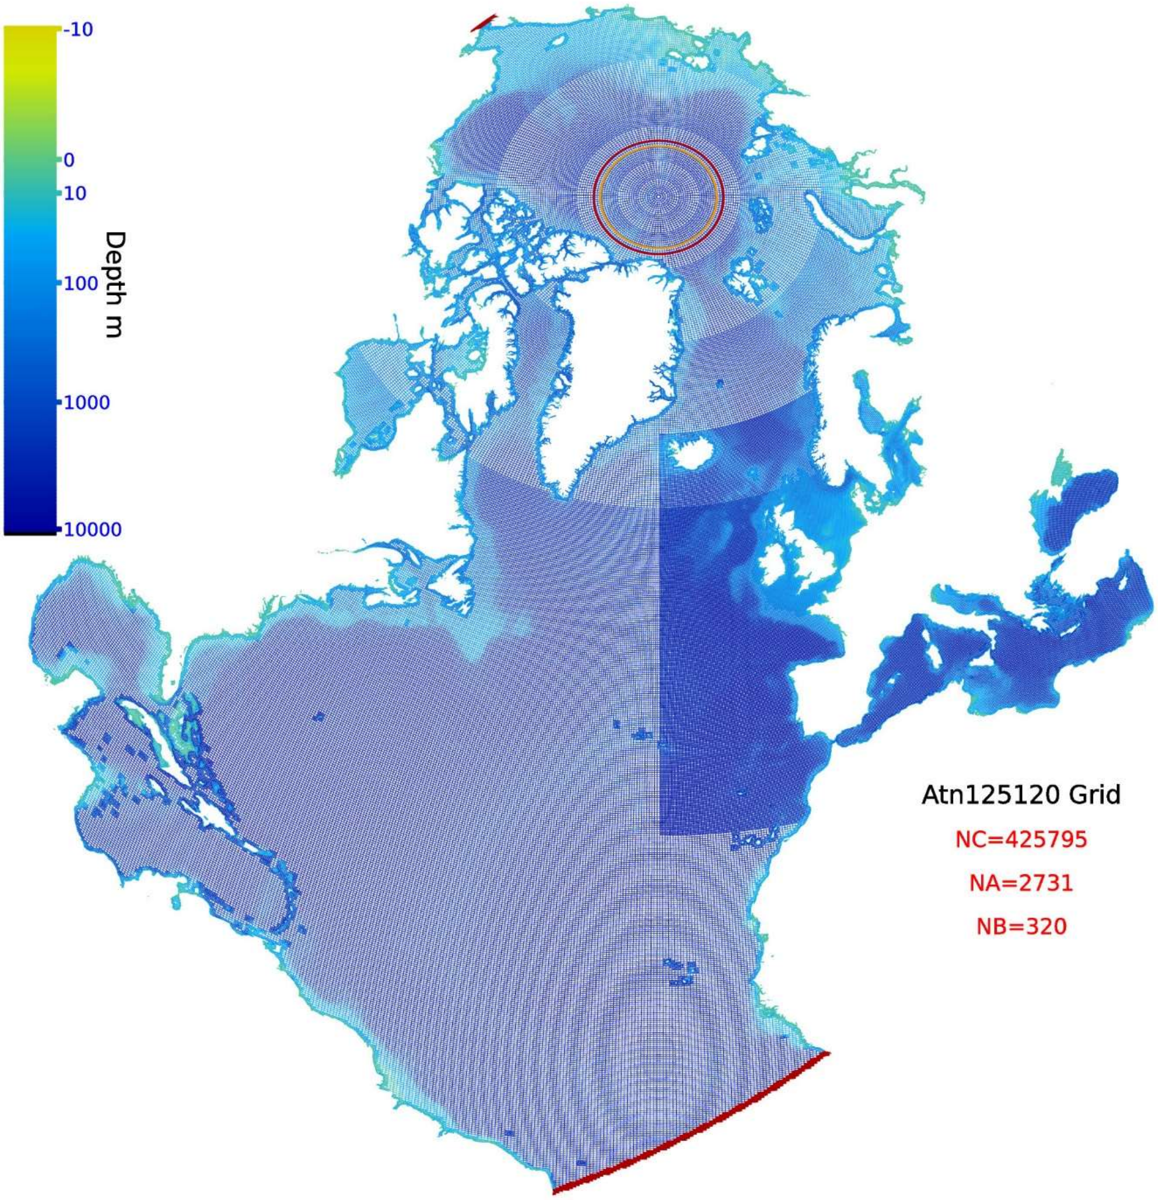

Fig.6b. The Atlantic and Arctic sub-grid at 1.25 to 20 km resolutions.

The planned new model has doubled the sea point number of our present global forecasting model (from about 0.6 to over 1.2 million) with improved resolutions both in our key European area and over the whole global oceans. Comparison with spectral buoy data reveals slightly better agreement than the present global model and computing performance shows better hybrid scalability, though model runtime is increased in comparison with the present global model. This will be further discussed in the next section.

5. Timing of hybrid multi-grid runs

One main target of this multi-grid study is to find out whether the multi-grid option could expand the node usage of our wave model efficiently. For this purpose, different node configurations are tested, and their timing results are summarised in 3 scalability diagrams. Fig.7a shows the elapsed times for the single SMC36125 grid MPI and hybrid tests against number of MPI ranks. For pure MPI (or one OpenMP thread) runs as shown by the blue curve in Fig.7a, the best time for one model day integration is about 170 s on 12 nodes with 430 MPI ranks. The elapsed time is an average of over 30 daily runs to remove machine fluctuations. For the hybrid runs (`OMP_NUM_THREADS` = 3 or 9 in red and green), the scalability curve flattens out around 360 MPI ranks, or 1/3 of the total wave spectral components (30×36 = 1080). These results indicate that the CD method, in either pure MPI mode or hybrid MPI-OpenMP mode, is most efficient when about 3 spectral components are allocated on one MPI rank for the propagation computing. Most hybrid tests are then configured with fixed 360 MPI ranks to study the OpenMP effect, but their elapsed times are clustered into one vertical line in Fig.7a.

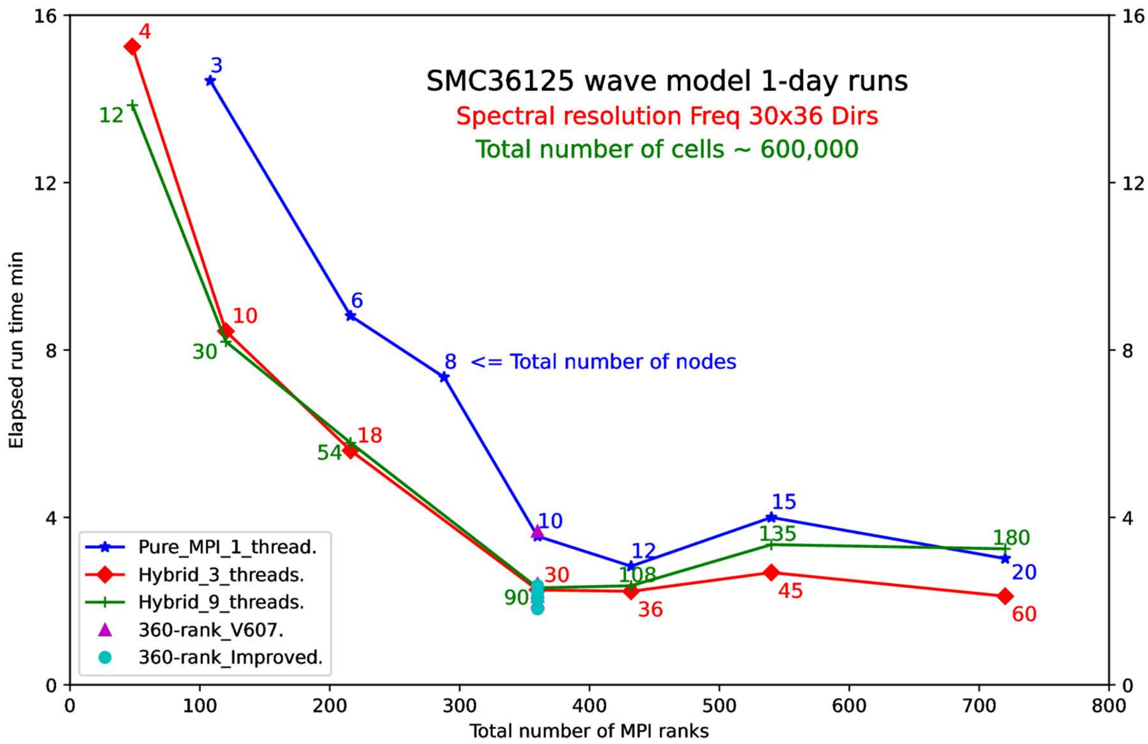

Fig.7a. WW3 SMC36125 model run times against number of MPI ranks by MPI and MPI-OpenMP hybrid parallelisation methods. Marked numbers show total number of nodes.

Fig.7b redraws the single grid elapsed times against number of computing nodes for those with 360 MPI ranks to illustrate the OpenMP effect. The top curve in magenta is for the SMC36125 grid model with the WW3 V6.07 public release code prior to the recent updates. The elapsed time first decreases with increased OpenMP threads and reaches the minimum elapsed time of about 2 min with 6 threads. Then it reverses to increase when the thread number is increased to 9 and 18. The blue curve shows the elapsed times of the single grid model with all the recent updates and the scalability is clearly improved. Nevertheless, the runtime ceases to decrease after the best time of about 83 s with 9 threads. This result implies that the hybrid single grid model can only use as many as 90 computing nodes efficiently. Comparing with the best time of 170 s on 12 nodes in pure MPI mode, the hybrid mode has halved the elapsed time (50%), including about 30% reduction by the recent updates.

Why the OpenMP scalability drops beyond a few of threads in WW3 is still not clear. One possible cause might be the machine bandwidth as one anonymous reviewer suggested. An optimisation work by other developers also supports this hypothesis as they found that removing some unnecessary cache memory loading in the SMC module may speed up their forecasting model. Their optimized code has been merged into the NCEP WW3 development branch in 2022.

The hybrid multi-grid option is first tested with a configuration of 2 SMC sub-grids. The MPI ranks are set to be 360 for each sub-grid so the hybrid run with 2 OpenMP threads requires 40 computing nodes. Its elapsed times for one model day run are shown in Fig.7b by the green line. The

2-thread run on 40 nodes takes almost the same time as the single SMC36125 grid hybrid run with 3 threads on 30 nodes. When the node usage of the 2 sub-grids run is expanded to 60 nodes with 3 OpenMP threads, its elapsed time almost matches the single grid case on the same number of computing nodes with 6 OpenMP threads. The 2 sub-grids model finally beats the single grid model when it is expanded on 120 computing nodes with 6 threads. The elapsed time reduction is not significant, about 15% from the single grid best time of 83 s to about 71 s for the 2 sub-grids. Further expansion of the node usage to 180 nodes with 9 OpenMP threads for the 2 sub-grids model does not show any more reduction but a slightly increase of model run time (~ 73 s). One reason is probably due to the imbalanced computing load for the 2 sub-grids, particularly for the sea-point based source terms (192547 via 401194 cells), though the equivalent cell numbers of the 2 sub-grids are close (591716 via 542626).

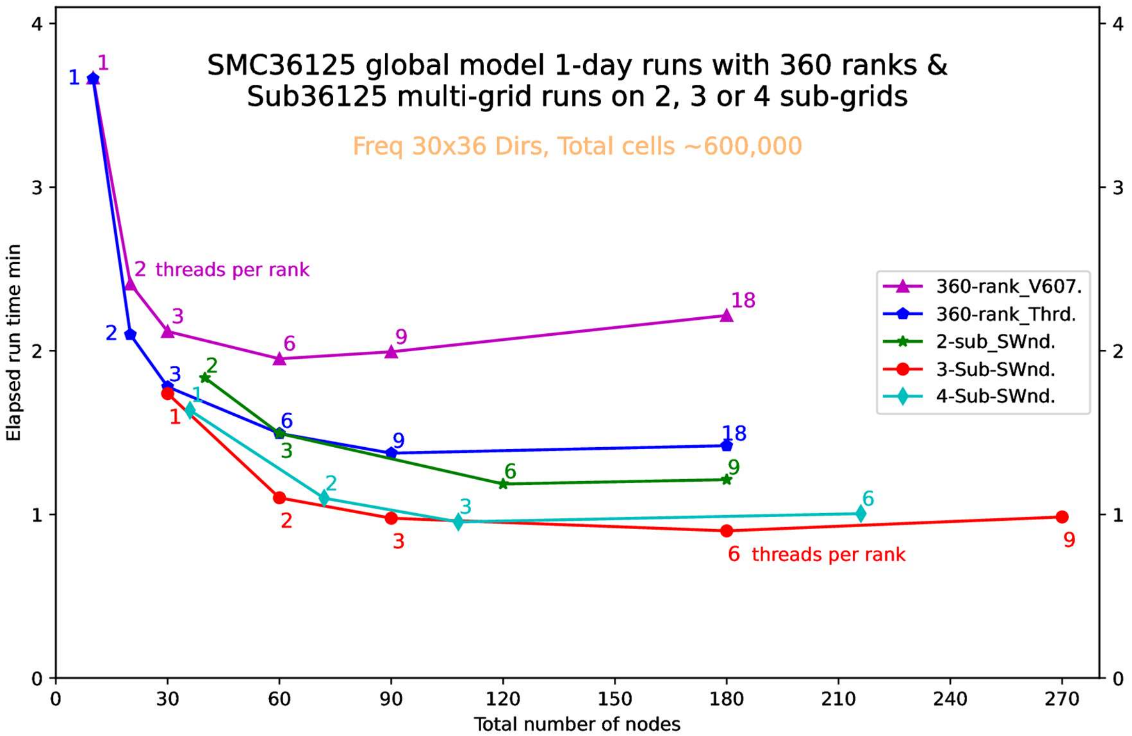

Fig.7b. Comparison of WW3 SMC36125 and sub-grids run timings with fixed 360 MPI ranks against total number of computer nodes. One node consists of 36 dual cores. Marked numbers show number of OpenMP threads per MPI rank.

The computing load balance is improved in the 3 sub-grids configuration (shown in Fig.2) by careful splitting of the sub-grids. As total cell number in each sub-grid is reduced in comparison with the 2 sub-grids, communication load for the CD method within each sub-grid is also reduced. The combined effects make the 3 sub-grids configuration better than the 2 sub-grids one. The elapsed times for the 3 sub-grids runs are shown by the red curve in Fig.7b. For the pure MPI case with 1 OpenMP thread on 30 nodes, the 3 sub-grids elapsed time is almost the same as the single SMC36125 grid on the same number of computing nodes. After the OpenMP is activated in hybrid mode, the 3 sub-grids model outperforms the single grid and reaches a best time of about 54 s with 6 threads on 180 nodes. The reduction rate is about 35% in comparison with the single grid best time of 83 s on 90 nodes. The elapsed time of the 3 sub-grids model with 3 threads on 90 nodes is about 59 s, which is about 31% less than the best time of the single grid model on the same number of 90 computing nodes with 9 OpenMP threads. Further expansion of the 3 sub-grids node usage with 9 OpenMP threads on 270 nodes leads to increase of the elapsed time (~ 59 s), close to that on 90 nodes with 3 threads.

A 4 sub-grids case is also tested. It uses almost the same sub-grids as in the 3 sub-grids case except that the Mediterranean Sea is separated from the Soth36126 grid to form the 4-th sub-grid. As the Pacf36125 sub-grid is still the most expensive one as in the 3 sub-grids case, there is no further reduction of the elapsed time as shown by the cyan curve in Fig.7b, even though more computing nodes are used than the 3 sub-grids. The resource ratios for the 4 sub-grids are tuned by their computing loads. The 4-th sub-grid with the lightest computing load is allocated 216 MPI ranks while the other 3 sub-grids use the optimised 360 ranks each. The 4 sub-grids model is then run on 36, 72, 108 or 216 nodes with 1, 2, 3 or 6 threads, respectively.

Two new features are also tested with the 4 sub-grids case. The first one is replacing the 4th sub-grid with a regular lat-lon Mediterranean Sea grid. It is set to be a lower ranked sub-grid and its boundary condition is provided by the 3 SMC sub-grids, using the wave spectra from the closest SMC grid cells rather than the existing interpolation scheme for regular lat-lon grids. This option makes it possible to merge a regional lat-lon grid model and a global SMC multi-grid model into a multi-grid frame, removing the need to save boundary condition files for the regional model for a separate run.

Another tested feature is mixed resolution sub-grids. The Pacf36125 3-level sub-grid in the 4 sub-grid model is replaced with a 4-level (6-12-25-50 km) sub-grid. Because the base resolution cell is enlarged from 25 km to 50 km, it effectively reduces the computing load of the Pacf sub-grid, and the overall elapsed time is reduced to about 46 s on 200 nodes. As this configuration uses a reduced solution sub-grid in the Pacific region, it is then not fully comparable with the previous sub-grid configurations, which are split from the same global SMC36125 grid. Nevertheless, this mixed resolution option will be useful for merging high resolution regional grids into global multi-grid systems.

Performance of the planned future UK Met Office global wave forecasting model is also tested on our Cray XC40 supercomputer. Its parallel runtimes are compared with the present global model in Fig.8. The blue curve in Fig.8 is for the present SMC36125 single grid model and the red one is for its multi-grid configuration with 3 sub-grids, same as shown in Fig.7b. The top green curve is for the new 5-level (1.25-2.5-5-10-20 km) model in a single grid (SMC125120), which is identical to the merged 3 sub-grids shown in Fig.6a. The total number of cells for the single global grid is 1203277, or 4755 cells less than the merged sub-grids (1208032). Its quickest runtime for one model day is about 205 s on 270 nodes with 540 MPI ranks and 18 OpenMP threads per rank. On 45 nodes with 3 threads per rank, it takes about 380 s elapsed time for one model day. Note this model is quickest with 540 MPI ranks or two spectral components on one MPI rank, and its OpenMP scalability is also improved in comparison with the present SMC36125 model. Nevertheless, the 5-level model is slower than the present model due to the doubled total cell number (1.2 via 0.6 million sea points).

The new 5-level grid is tested with two multi-grid configurations, one with 2 sub-grids and another with 3 sub-grids, which is shown in Fig.6. The 2 sub-grids configuration minimises boundary exchanges by merging the Pacific and Southern Ocean sub-grids into a single 4-level grid at 2.5-5-10-20 km resolutions. The Atn125120 sub-grid is almost the same as the in the 3 sub-grids case shown in Fig.6b except for that the Great Lakes and Caspian Sea are included as well. As the 2 sub-grids are not balanced in computing load, the multi-grid distribution ratio is set to be 0.4 for the Atn125120 and 0.6 for the merged Pacific and Southern Ocean sub-grid. It was tested on 75, 150, and 225 nodes with 3, 6, and 9 OpenMP threads, respectively. The Atn125120 sub-grid takes 360 MPI ranks while the merged sub-grid uses 540 ranks. The one model day runtime of the 2 sub-grid model is reduced in comparison with the single global grid and its quickest runtime on 225 nodes is about 142 s. Comparing with the quickest runtime of the single grid model 205 s, the 2 sub-grid model runtime is about 30% shorter even though it uses less computer nodes (225 against 270 nodes).

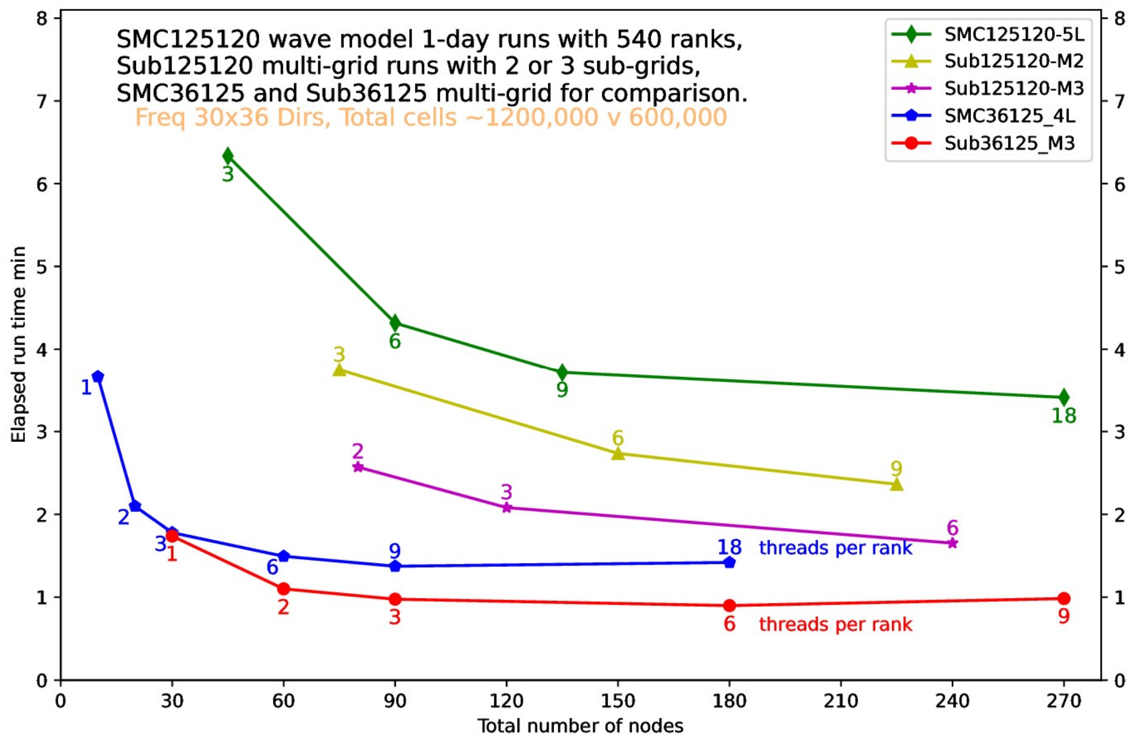

Fig.8. Comparison of the 5-level SMC125120 single and multi-grid global model and the present SMC36125 model performance against total number of nodes. One node consists of 36 dual cores. Marked numbers show number of OpenMP threads per MPI rank.

The magenta curve in Fig.8 is for the 3 sub-grids case of the new 5-level grid. The Sth251020 sub-grid in this configuration is the smallest among the 3 sub-grids and is assigned a multi-grid distribution ratio of 0.25. The other two sub-grids share the rest resources with an equal ratio of 0.375. The multi-grid model is run on 80, 120, and 240 nodes with 2, 3, and 6 OpenMP threads. The Sth251020 sub-grid uses 360 MPI ranks while the other two sub-grids take 540 ranks each. Its quickest runtime on 240 nodes is 99 s for one model day, about 52% shorter than the single grid on 270 nodes. It is also quicker than the 2 sub-grids case (142 s on 225 nodes). It is obvious that the multi-grid option for this new 5-level grid is more efficient than for the present SMC36125 grid, which only has about 30% reduction when run in 3 sub-grids mode. The most likely reason is the doubled sea points as the OpenMP scalability is improved when spatial points increase.

The multi-grid option allows further increase of sub-grid number and there is room to accommodate higher resolution models with the hybrid CD-DD method. The present study does not attempt more than 4 sub-grids because the job queue on our present supercomputer is long for large jobs requiring over 200 computing nodes. Besides, the model series computing part (for initialisation and output) has become significant when the parallel time is reduced to less than 1 min. It may be worth to test more sub-grids with the hybrid CD-DD method for higher resolution models in the future.

6. Summary and conclusions

The SMC grid is an unstructured grid and retains finite-difference schemes on conventional latitude-longitude grid. It uses merged cells at high latitudes to relax the CFL restriction on Eulerian transport schemes and redefines vector components by a fixed reference direction near the poles to avoid vector polar problems. Flux form finite-volume formulations are included for multi-resolution interfaces and sub-timesteps are used for refined cells for computing efficiency. A brief description of the unstructured SMC grid is presented and its recent updates for hybrid and multi-grid options in WW3 model are explained. The SMC multi-grid option is explored for parallelisation by combining the existing spectral component decomposition (CD) with domain decomposition (DD) in hybrid (MPI-OpenMP) mode. The hybrid method can reduce the model elapsed time by as much as 50%, including 30% by recent updates. The combined hybrid CD-DD method can gain a further reduction of more than 30% for our Met Office global SMC36125 wave model. It extends the model node usage from the best 12 nodes (single grid in pure MPI mode) to 180 nodes (3 sub-grids with 6 OpenMP threads) and reduces the elapsed time of one model day integration from about 3 min to less than one min. Timing tests on our planned global model with doubled sea-points than the present global model reveal improved OpenMP scalability, and its hybrid multi-grid mode may halve the model runtime of a single global grid model. The multi-grid option can merge regional models into one global system without saving boundary files and makes post-processing more flexible than nested systems. The hybrid method also reduces the memory demand on one computing node, creating room for higher spatial and/or spectral resolutions in wave forecasting models.

Acknowledgements

The author is grateful to Drs Andrew Saulter, Breogan Gomez and Paul Selwood (UK Met Office) for their encouragement and support of this work. The author's gratitude also extends to Dr Hendrik Tolman (NOAA/NCEP), who provided information of the early parallelisation in WAVEWATCH III, and those involved in merging the updates into the WAVEWATCH III public website.

References

Abdolali, A., A. Roland, A. van der Westhuysen, J. Meixner, A. Chawla, T.J. Hesser, J. M. Smith, M.D. Sikiric, 2020: Large-scale hurricane modeling using domain decomposition parallelization and implicit scheme implemented in WAVEWATCH III wave model, *Coastal Engineering* **157**, 103656.

Berger, M., J. Oliger, 1984: Adaptive mesh refinement for hyperbolic partial differential equations, *J. Comput. Phys*. **53**, 484-512.

Booij, N., R.C. Ris, L.H. Holthuijsen, 1999: A third-generation wave model for coastal regions, part I, model description and validation. *J. Geophys. Res.* **104**(C4), 7649–7666.

Bruciaferri, D., M. Tonani, H.W. Lewis, J.R. Siddorn, A. Saulter, J.M. Castillo Sanchez, N.G. Valiente, D. Conley, P. Sykes, I. Ascione, N. McConnell, 2021: The impact of ocean-wave coupling on the upper ocean circulation during storm events. *J. Geophys. Res. Oceans* **126**, e2021JC017343.

Bunney, C., A. Saulter, 2015: An ensemble forecast system for prediction of Atlantic-UK wind waves. *Ocean Modelling* **96**, 103-116.

Casas-Prat, M., X.L. Wang, N. Swart, 2018: CMIP5-based global wave climate projections including the entire Arctic. *Ocean Modelling* **123**, 66-85.

Chawla, A., H.L. Tolman, V. Gerald, D. Spindler, T. Spindler, J.H.G.M. Alves, D. Cao, J.L. Hanson, E.M. Devaliere, 2013: A Multigrid Wave Forecasting Model: A New Paradigm in Operational Wave Forecasting. *Weather and Forecasting* **28**, 1057-1078.

Gates, W.L., C.A. Riegel, 1962: A study of numerical errors in the integration of barotropic flow on a spherical grid. *J. Geophys. Res*. **67**, 773-784.

GEBCO Compilation Group, 2022: GEBCO 2022 Grid. doi:10.5285/e0f0bb80-ab44-2739-e053-6c86abc0289c.

Hou Fang, Gao Zhiyi, Li Jianguo, Yu Fujiang. 2022: An efficient algorithm for generating a spherical multiple-cell grid. *Acta Oceanologica Sinica*, **41(5)**, 41-50, doi: 10.1007/s13131-021-1947-3.

Jones, P.W., 1999: First- and second-order conservative remapping schemes for grids in spherical coordinates, *Mon. Wea. Rev*. **127**, 2204–2210.

Li, J.G., 2008: Upstream nonoscillatory advection schemes. *Mon. Wea. Rev*. **136**, 4709-4729.

Li, J.G., 2011: Global transport on a spherical multiple-cell grid. *Mon. Wea. Rev*. **139**, 1536-1555.

Li, J.G., 2012: Propagation of ocean surface waves on a spherical multiple-cell grid. *J. Comput. Phys*. **231**, 8262-8277.

Li, J.G., 2016: Ocean surface waves in an ice-free Arctic Ocean. *Ocean Dynamics* **66**, 989-1004.

Li, J.G., 2018: Shallow-water equations on a spherical multiple-cell grid. *Q.J.R. Meteorol. Soc*. **144**, 1-12. doi: 10.1002/qj.3139.

Li, J.G., 2019: An efficient multi-resolution grid for global models and coupled systems. *Adv. Sci. Res*. **16**, 137-142. doi: 10.5194/asr-16-137-2019.

Li, J.G., 2021: Filling oceans on a spherical multiple-cell grid. *Ocean Modelling*, **157**, 101729. doi: 10.1016/j.ocemod.2020.101729.

Li, J.G. 2022: Hybrid multi-grid parallelisation of WAVEWATCH III model on spherical multiple-cell grids. *J. Parallel Distrib. Comput*., 167C, 187-198. doi: 10.1016/j.jpdc.2022.05.002

Li, J.G. 2023: Spherical Multiple-Cell (SMC) grid utility tools. Freely available at open-source site: https://github.com/ww3-opentools/SMCGTools/

Li, J.G., A. Saulter, 2012: Assessment of the updated Envisat ASAR ocean surface wave spectra with buoy and altimeter data. Remote Sensing Environ., 126, 72-83. doi: 10.1016/j.rse.2012.08.018.

Li, J.G., A. Saulter, 2014: Unified global and regional wave model on a multi-resolution grid. Ocean *Dynamics* **64**, 1657-1670. doi: 10.1007/s10236-014-0774-x.

Li, J., Y. Ma, Q. Liu, W. Zhang, C. Guan, 2019: Growth of wave height with retreating ice cover in the Arctic. *Cold Regions Sci. Techn.* **164**, 102790. doi: 10.1016/j.coldregions.2019.102790.

Valiente, N.G, A. Saulter, B. Gomez, C. Bunney, J.G. Li, T. Palmer, C. Pequignet, 2023: The Met Office operational wave forecasting system: the evolution of the regional and global models. *Geosci. Model Dev*., **16**, 2515-2538. doi:10.5194/gmd-16-2515-2023.

Tint, O., M. Acosta, M. Castrillo, A. Cortes, A. Sanchez, K. Serradell, F.J. Doblas-Reyes, 2017: Optimizing domain decomposition in an ocean model: the case of NEMO. *Procedia Computer Science* **108**C, 776-785.

Tolman, H.L., 1991: A third-generation model for wind waves on slowly varying, unsteady and inhomogeneous depths and currents, *J. Phys. Oceanogr*. **21**, 782–797.

Tolman, H.L., 2002: Distributed memory concepts in the wave model WAVEWATCH III, *Parallel Computing*, **28**, 35–52.

Tolman, H.L., 2008: A mosaic approach to wind wave modeling, *Ocean Modelling*, **25**, 35–47.

Tolman, H.L., 2013: Scaling of WAVEWATCH III on massively parallel computer architectures. Part I: Hybrid parallelization. NOAA/NWS/NCEP/MMAB, *Tech. Note* **310**, 41 pp.

WAVEWATCH III Development Group (WW3DG), 2019: User manual and system documentation of WAVEWATCH III version 6.07. NOAA/NWS/NCEP/MMAB *Tech. Note* **333**, 465 pp. https://github.com/NOAA-EMC/WW3/releases/tag/6.07.

Zieger, S., D.J.M. Greenslade, 2021: A multiple-resolution global wave model -AUSWAVE-G3. Bureau Research Report No 51. 74pp. Available online at web site: http://www.bom.gov.au/research/publications/researchreports/BRR-051.pdf.