
OrcaWave – working with meshes

Introduction

OrcaWave is a frequency domain diffraction analysis program which allows the user to calculate loading and response for one or more wetted bodies due to surface water waves via potential flow theory. The form of each body is communicated to OrcaWave via a panel mesh. The accuracy of the diffraction calculation is heavily dependent on how well the body mesh represents the true form of the body. As a result, mesh files often contain a large number of fine panels and this can be computationally expensive. Inevitably users will look to optimise their body mesh to improve results quality whilst minimising the computational effort. This document summarises some of the common working practices when dealing with panel meshes in OrcaWave.

Building a panel mesh

Body mesh

OrcaWave is not capable of discretising a solid body geometry into a panel mesh. For simple bodies with regular form, a user could choose to write a simple script which arranges panel geometry data in the expected format. However in many cases, marine vehicles and equipment have irregular form and dedicated meshing programs are required. OrcaWave supports a large range of mesh file formats, you can find a list of supported mesh file formats in the OrcaWave help ([Data | Mesh file formats](#)).

Control surface, damping lid, free surface panelled zone

In addition to the body mesh, OrcaWave users may choose to supply several other mesh files for the purpose of accessing alternative calculation methods, this can also increase the range of results reported. These additional mesh files should be supplied independently of the body mesh. Since v11.4 OrcaWave has been capable of building control surface and free surface panelled zone meshes using a small number of user specified parameters.

Mesh quality

Mesh sensitivity study

Whilst OrcaWave has a series of validation checks which aim to highlight key problems with a body mesh, OrcaWave cannot give a definitive answer on whether a body mesh is good enough. Therefore it is good practice to undertake a mesh sensitivity study. Typically this will involve comparison of a number of key results under different wave conditions whilst systematically refining the body mesh. At some point, each OrcaWave result will converge to a consistent value. At this point the quality of the result will not improve, even with further mesh refinement.

Mesh refinement typically involves a reduction in the panel size across the body. In doing so, the meshed body becomes a better representation of the underlying geometry. Refining the mesh also results in an improvement in the estimate of loads induced by waves of short wavelength, since the potential is calculated more regularly across the body and the loads induced by short wavelengths are better resolved. Furthermore refinement to the body mesh has influence on the hydrostatic calculation. Simply put, a coarse mesh will likely result in a poor estimate of displacement volume and centre of buoyancy. The hydrostatic properties reported in the OrcaWave results influence the hydrostatic stiffness matrix, displacement RAOs, mean drift loads and QTFs. Therefore a converged set of the hydrostatic properties will also contribute towards convergence of the OrcaWave hydrodynamic results.

OrcaWave – working with meshes

Useful meshing practices

Introduction

The total number of mesh panels is one of the main factors influencing run time. Furthermore the memory required to undertake the calculation is proportional to the number of body panels squared. Consequently some users undertaking a mesh refinement study may see their run time increase significantly before they observe results convergence. To help tackle this increase in run time, we will describe a range of useful meshing practices.

Mesh symmetry

It is possible to significantly accelerate a diffraction analysis by taking advantage of any global symmetry that exists in the model. A typical example of useful mesh symmetry is the analysis of a ship shaped body. Ship shaped vessels often have port-starboard symmetry and a user can simplify the diffraction analysis by providing mesh panels for one half of the body only. Provided the user specifies the symmetry plane on the Bodies tab, OrcaWave will mirror the mesh and make the relevant simplifications when undertaking calculations on those mirror image panels. These simplifications result in a reduction in run time and the memory required to run the calculation.

When you specify a body mesh in OrcaWave, you have the opportunity to translate and/or rotate the mesh. In doing so, a specified global symmetry plane may become invalid. In this scenario the mirror image panels will still be generated and visible in the mesh view, but the improvement in run time and memory usage may be lost. You can look for confirmation that OrcaWave has recognised the global symmetry plane on the Bodies tab, see Figure 2. Finally for users working with global symmetry planes in versions of OrcaWave prior to v11.4, there was an underlying assumption that the body would be defined using body panels positioned in the positive quadrant/half of the OrcaWave global axis system. Failure to abide by this assumption may have compromised the validation checks.




Figure 1 – Global symmetry plane, base panels from mesh file in red, mirrored panels highlighted in white.

OrcaWave – working with meshes

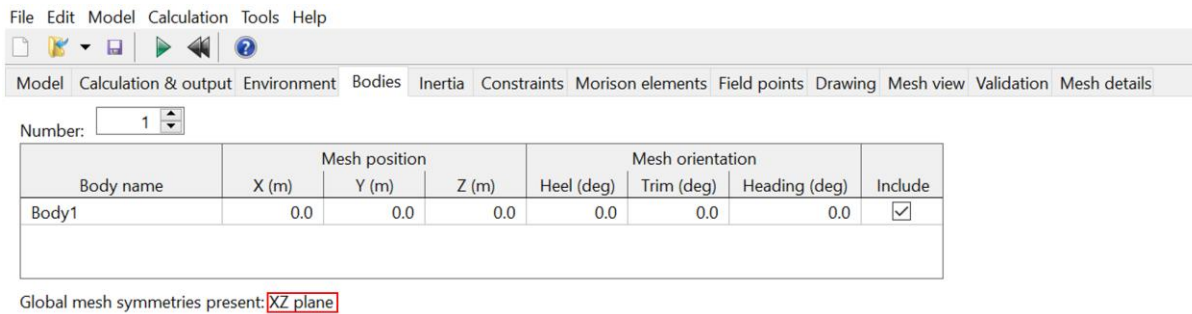


Figure 2 – Confirmation that global mesh symmetry has been recognised.

Targeted mesh refinement

The calculation of second-order loads includes an integral of pressure at the waterline (no such integral is needed for first-order loads). OrcaWave estimates the pressure at the waterline using the pressure at the centroid of each body panel adjacent to the water surface. One of the reasons quadratic load results improve with mesh refinement is these panel centroids move closer to the water surface, improving the estimate of pressure at the waterline. Bearing this in mind, we can save time by prioritising the reduction in depth of the body panels adjacent to the waterline, rather than refining the mesh uniformly across the whole body.

Discretising the body using a greater number of panels around the circumference of the waterline also stands to improve the calculation of the quadratic load. However this is usually less significant and does not usually motivate targeted mesh refinement.




Figure 3 – Example of targeted panel refinement close to the waterline.

OrcaWave – working with meshes

Control surface mesh

Quadratic loads can be calculated using the pressure integration method, the control surface method, and/or the momentum conservation method. The control surface method requires an additional mesh which when combined with the body mesh encloses a finite volume of water around the body. The control surface method infers the quadratic loads acting on the body by evaluating the fluid potential on the control surface.

In general, a control surface mesh will consist of panels of similar size to those used in the body mesh but variation on this relationship could form part of a mesh sensitivity study. The separation between the outer boundary of the control surface and the body is not significant, collision between the control surface and the body panels should be avoided and so a spacing of 5-10 panels can be used. Often control surface meshes are defined in the shape of a cuboid however this is for meshing convenience and is not compulsory.

If you are only concerned with the horizontal components of the quadratic load (surge, sway, yaw), the control surface panels on the water surface can be omitted. This can speed up the calculation, however please note the vertical components of the quadratic load will still feature in the results despite being miscalculated. If you are interested in quadratic loads in all six DOFs and the body mesh is surface piercing, control surface panels on the water surface should butt up to the body mesh, i.e. they should not occupy space inside the body. If the control surface method is used in a multibody analysis, each body must be assigned an independent control surface mesh. Care should be taken to make sure each control surface mesh does not collide with the body panels of another body.

Whilst introducing another mesh increases the calculation time for quadratic loads, the control surface method can achieve convergence using a body mesh with coarser discretisation than would be necessary to achieve convergence with the pressure integration method. Furthermore the control surface method has the same peak memory requirement as the pressure integration method. Therefore a user can run an analysis using both quadratic methods without an increase in the total memory requirement.




Figure 4 – Control surface mesh (yellow panels).

OrcaWave – working with meshes

Case study

Method

To demonstrate the effectiveness of these methods, we have chosen to undertake a mesh sensitivity study based on a floating body representative of a very large crude carrier (VLCC). The body is over 300m long and so we can expect a large number of panels will be required to appropriately resolve the wave loads, particularly those at small wave periods/wave lengths.

This case study involved several mesh sensitivity studies run in parallel. The first was undertaken utilising uniform mesh refinement. The second adopted a method of targeted mesh refinement close to the waterline. The third and fourth followed the same process of targeting mesh refinement at the waterline, however the quadratic load calculation method was changed from pressure integration to control surface and then momentum conservation respectively. Each study made use of the port-starboard global symmetry condition.

The target size for body panels was varied from 10m to 1m. When targeting refinement at the waterline, panels close to the waterline were given a target size of 40% of that assigned to the rest of the body. These two regions were then blended together over a consistent distance. When modelling with a control surface, a cuboid shaped mesh with panel sizes similar to that of the body mesh were used. Separation between the body and control surface was fixed for consistency. A consistent set of user defined interior surface panels were used in all cases. This avoided an unnecessary increase in the number of interior surface panels as the body mesh was refined.

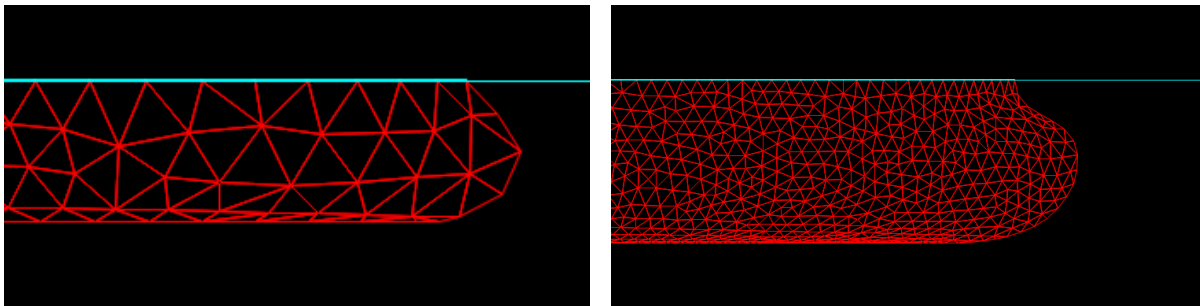


Figure 5 – Uniform mesh refinement




Figure 6 – Targeted refinement at the waterline.

OrcaWave – working with meshes

Load RAO convergence

Figure 7 shows the variation in normalised load RAO due to refinement of the body mesh for all six degrees of freedom when subject to a 5s, 45deg wave. The figure shows that the load RAOs for this particular vessel remained fairly consistent for panel sizes less than 8m and converged to a reasonably consistent solution for panel sizes less than 3m. Targeting mesh refinement at the waterline did have an impact on the results, however the effect on the RAOs was not significant. Any differences in the results were due to better representation of the underlying body geometry at the waterline and improved resolving of the wave loads.

Mean drift load convergence

If you were to consider the results from Figure 7 in isolation, you might conclude that a mesh consisting of panels with a target panel size of several metres is sufficient to model the vessel accurately. However Figure 8, which is a comparison of normalised mean drift load for a 5s, 45deg wave, demonstrates that a larger number of finer panels was required before reaching convergence. We also see significant difference in the rate of convergence when comparing the different meshing and calculation methods. In most DOFs, the results show the control surface calculation method arrives at a converged solution using a larger panel size compared to the pressure integration method. Furthermore the results demonstrate that targeted mesh refinement at the waterline can also be used to accelerate convergence.

Similarly Figure 8 also shows that the momentum conservation calculation method is effective at reaching a converged solution for larger panel sizes. However the method is limited in that it is only capable of reporting results in the horizontal DOFs (surge, sway and yaw).

Figure 8 also highlights how a comparison of the two **or more** quadratic calculation methods can act as a good check of convergence. Conversely a user running a mesh sensitivity study using only one method could misinterpret the slowly converging mean drift load for a converged and steady result. We chose to terminate the study at panel sizes of 1m, even though many of the cases were yet to fully converge. This was because run times were very long at this stage and the typical trends had been demonstrated.

To illustrate the kind of time savings that are sometimes available, we have included data in Table 1 comparing the calculation of sway mean drift load. The CPU time relates to the time taken to solve the diffraction analysis for a single wave direction and single wave period. The results demonstrate that in this instance, the momentum conservation and control surface calculation methods were capable of returning a result with less error, using less memory all within a shorter run time. Remember this analysis involves a 300m vessel. Panel number, memory requirement and run time are larger than would be the case for a smaller body.

OrcaWave - working with meshes

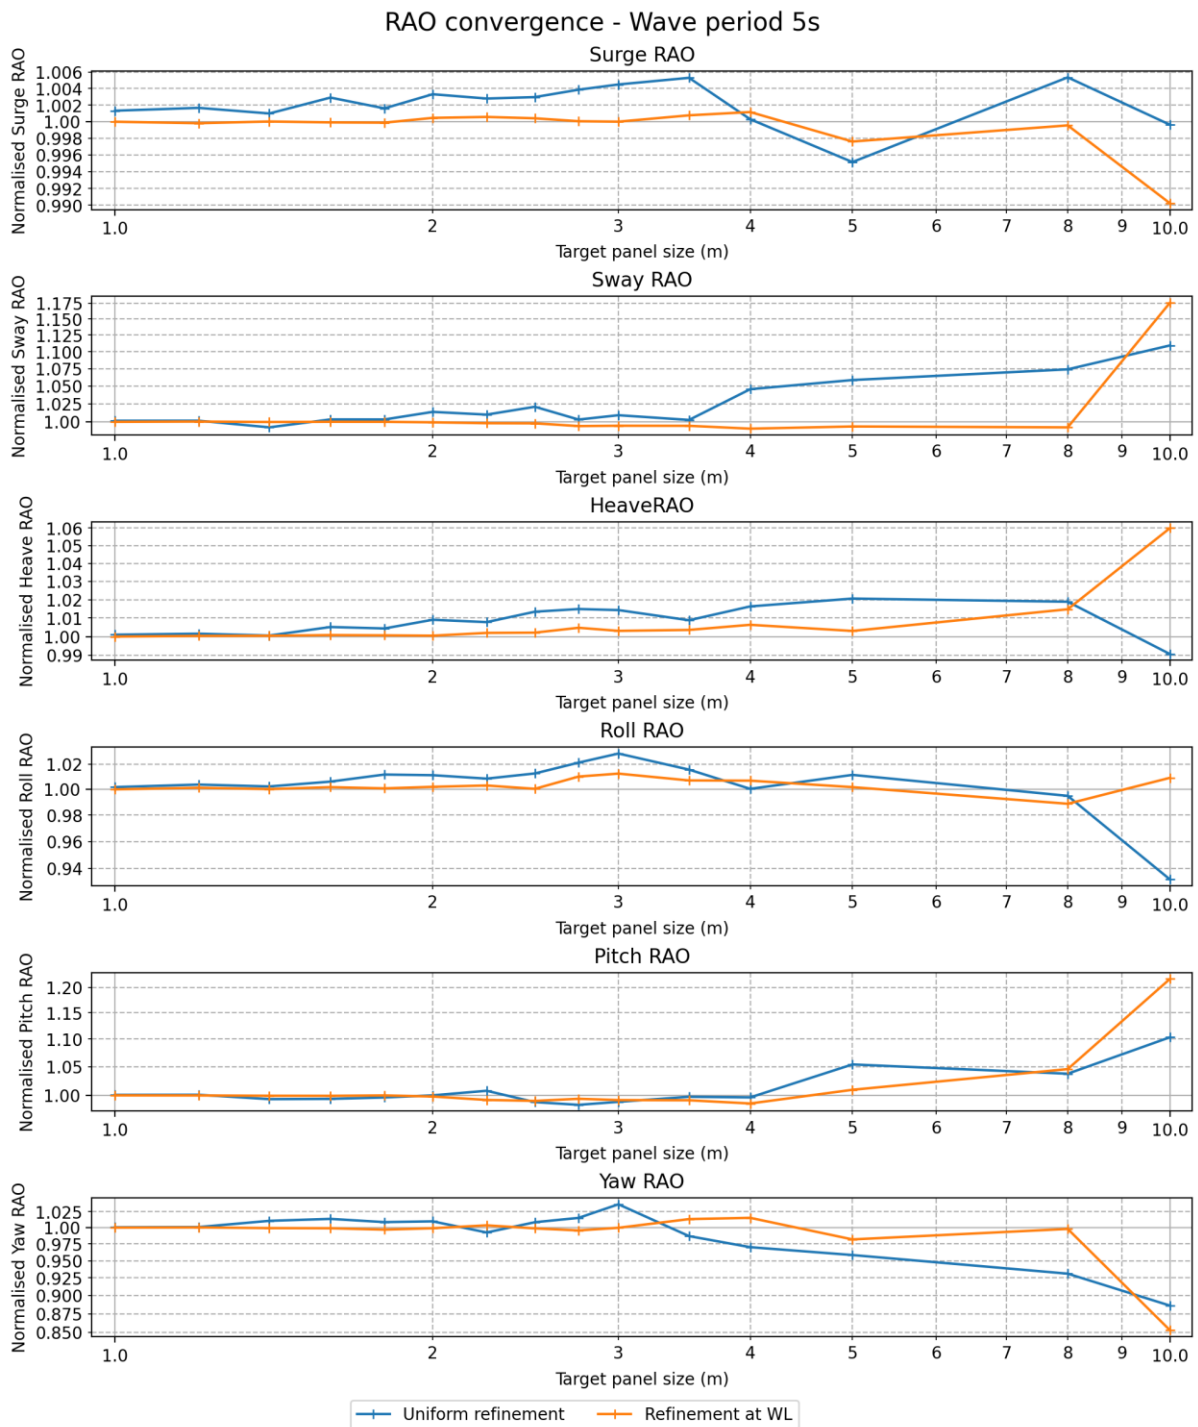


Figure 7 – Load RAO comparison

OrcaWave - working with meshes

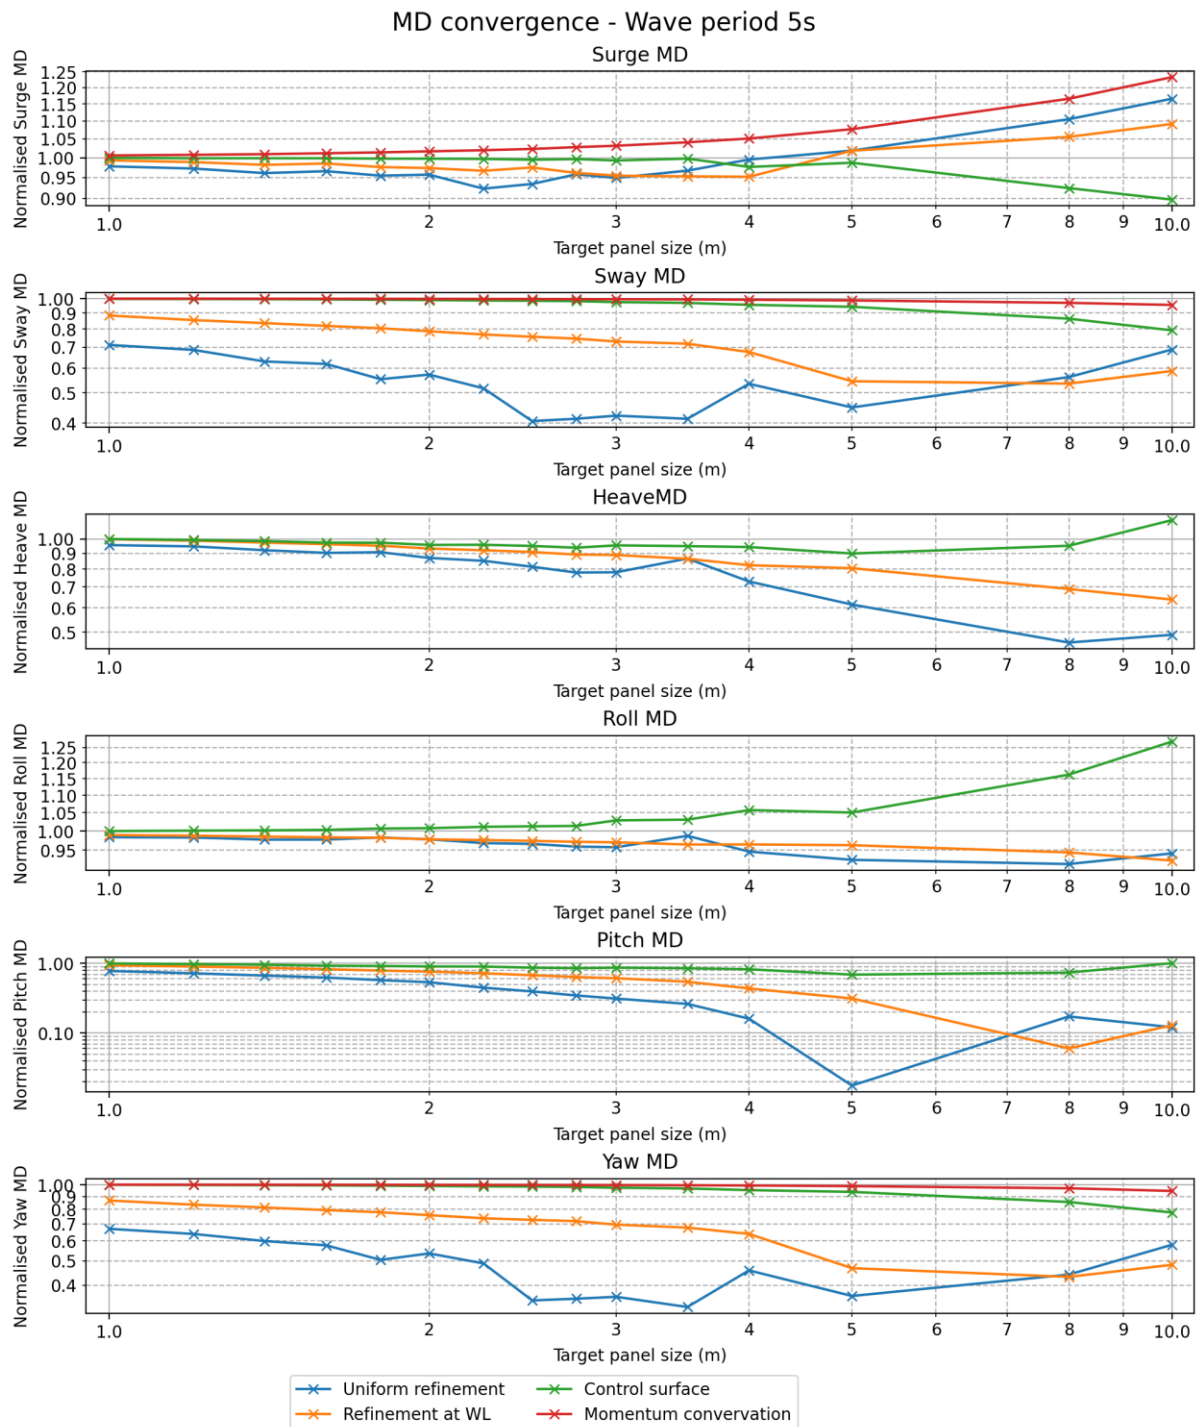


Figure 8 – Mean drift load comparison

OrcaWave – working with meshes

<u>Mesh & calculation method:</u>	<u>Uniform refinement</u>	<u>Refinement at WL</u>	<u>Control surface & refinement at WL</u>	<u>Momentum conservation & refinement at WL</u>
Quadratic method:	Pressure integration	Pressure integration	Control surface	Momentum conservation
Symmetry plane:	XZ	XZ	XZ	XZ
Target panel size:	1m	1m	2m	2m
Number of body panels:	67418	84346	21458	21458
Number of interior surface panels	1138	1138	1138	1138
Number of control surface panels	0	0	28810	0
Total number of panels	68556	85484	51406	21458
CPU time (hh:mm:ss):	03:26:47	07:40:58	00:17:57	00:09:31
Memory requirement:	190 GB	297GB	20.2 GB	20.2 GB
Error:	30%	12%	< 1%	< 1%

Table 1 – Run time & memory requirements, sway DOF mean drift load.

Further considerations

This case study only goes as far as calculating mean drift loads, however a sensitivity study including full QTF data would inevitably take more time. With that in mind, it is worth considering if the calculation workflow can be modified to improve run time. For example, the calculation of quadratic load results typically requires less memory in comparison to the calculation of first order results. Currently OrcaWave forces the user to set a maximum thread count that applies to the entire analysis, consequently you may find that there are periods during the diffraction analysis where the memory capacity of your machine is under utilised. If you were to split your analysis into two parts (first order and quadratic loads), you could increase the thread count for the second part of the calculation once the memory demand per thread has decreased, thus improving the overall run time. This is precisely the sort of calculation that OrcaWave restart analyses are capable of. Furthermore this approach could be scripted using OrcFxAPI to remove any down time between calculation of the first order loads (parent model) and quadratic loads (child model). For more information on using OrcaWave restart analyses see the OrcaWave help.

Similarly there are opportunities to reduce the memory requirement and save run time through careful selection of the diffraction solve type. There are three diffraction solve types. As you move from Potential formulation only through to Full QTF, a more comprehensive set of results becomes available as well as a wider selection of calculation methods. However, shifting to a more complex diffraction solve type normally incurs a penalty in terms of memory requirement. Some results and calculation methods are available via multiple diffraction solve types. Therefore users may find they can reduce the memory requirement and run time of their analysis by including diffraction solve type as a variable in their sensitivity studies where applicable. To demonstrate, table 2 has been included as a comparison of run time and memory requirement when calculating

OrcaWave – working with meshes

sway mean drift load. The analyses were made using a consistent calculation method, but varying diffraction solve type.

Quadratic method:	Momentum conservation	Momentum conservation
Diffraction solve type	Potential formulation only	Potential and source formulation
CPU time (hh:mm:ss):	00:06:30	00:09:31
Memory requirement:	9.3 GB	20.2 GB
Error:	<1%	< 1%

Table 2 – Run time & memory requirements, sway DOF mean drift load. Varying diffraction solve type.

Warning

We have chosen to present results relating to load RAOs and mean drift loads for one particular wave condition. The results demonstrate how different calculation methods and meshing practices can be used to accelerate convergence however please recognise these results are only representative of this particular body. The effectiveness of the methods discussed will vary with wave period, body geometry and result type. Furthermore a thorough mesh convergence study would consider a range of results types and degrees of freedom under a range of different wave periods/directions.

Additional information

Interior surface panels

OrcaWave relies on the simplification that we can estimate wave induced loads acting on a body by discretising the body into panels and calculating the potential at each panel. For surface piercing bodies, this simplification can lead to the introduction of irregular frequency effects. Irregular frequency effects are a numerical error and can have unwanted effects on the numerical results.

Interior surface panels are used to suppress irregular frequency effects. Interior surface panels are panels placed in the plane of the water surface inside the body mesh. If they are included with the body mesh, OrcaWave will automatically detect them. If they are omitted, OrcaWave can be instructed to add interior surface panels automatically on the Bodies tab, provided the body waterline is closed. This is particularly convenient when modelling a variety of load cases where the mesh position is manipulated in OrcaWave between runs. Interior surface panels are not required for fully submerged bodies where there is no intersection with the water surface. Interior surface panels are sufficient to remove irregular frequency effects from the entire analysis, meaning you do not have to take any additional steps when undertaking a full QTF calculation.

Dipole panels

Up to and including OrcaWave v11.2, there was an assumption that all mesh panels were wet on one side and dry on the other. In v11.3 dipole panels were introduced to represent a panel that is wet on both sides. Dipole panels have been introduced to help users avoid the numerical issues associated with modelling thin structures. They are intended to be used to approximate thin structures such as heave plates and strakes as zero-thickness elements. Dipole panels should be included with the body mesh file, they can then be identified by listing the relevant panel IDs on the Bodies tab of the OrcaWave user interface. It can be useful to view the mesh panel IDs in the mesh view when trying to identify dipole panels. Please note OrcaWave deals with two types of panel IDs, the first is the mesh file panel indices and these are the identifiers used to reference

OrcaWave – working with meshes

dipole panels. The second set of panel indices relate to the results where symmetry planes, interior surface panels and several other OrcaWave features can lead to a greater number of panels in the analysis than were included in the original body mesh.

Dipole panels do not strictly have to share vertices with ordinary body panels. For example when modelling strakes, you may wish to discretise the strakes using a finer panel size than is necessary elsewhere on the body mesh. However care should be taken to ensure that no unrealistic gaps between the body panels and the dipole panels are introduced.

Damping lid mesh

A damping lid is a panelled zone positioned on the free surface used to suppress fluid resonance. Typically this might occur in a moonpool or in a narrow gap between two vessels. If present, fluid resonance may be underdamped due to a lack of friction in the fluid (inviscid potential flow) and will manifest as narrow spikes in the hydrodynamic results graphs. In the case of a moonpool, a damping lid will typically cover the free surface inside the moonpool. The damping lid panels will butt up to the body panels but they will not extend through into the hull.

Underdamped body resonance can also appear as spikes in the OrcaWave hydrodynamic results, however unlike fluid resonance, body resonance develops in the calculation of displacement RAOs. Therefore spikes in load RAOs, added mass and damping will usually indicate underdamped fluid resonance (assuming irregular frequency effects are not present, see *Interior surface panels*).

Free surface panel mesh

A free surface panel mesh is required in order to determine the free surface forcing part of the quadratic transfer functions (QTFs). Typically the free surface panelled zone will have a circular outer boundary and will cover an area of the free surface. The panelled zone mesh can extend to cover areas of the free surface occupied by a moonpool but should not extend inside the body mesh. The free surface panelled zone is surrounded by an annular quadrature zone and an annular asymptotic zone. Advice on the procedure for sizing the free surface panelled zone can be found in the OrcaWave help ([Data | QTFs](#)).

Non-planar panels

OrcaWave is only capable of conducting a diffraction analysis using meshes consisting of planar panels. If OrcaWave detects a non-planar panel, a warning will be flagged on the Validation tab and OrcaWave will manipulate the panel vertices to force it to be planar. Often this process can introduce gaps into the mesh and this can lead to other issues e.g. failure to identify the waterline and failure to draw interior surface panels.

OrcaWave is capable of using quadrilateral and triangular mesh panels. Triangles are inherently planar and so this issue only arises when non-planar quadrilateral panels are present in the mesh. One solution is to split each non-planar quad into two triangular panels. In fact, OrcaWave is capable of doing this for you if you tick the Divide non planar panels tick box on the Calculation & Output tab. The downside of this approach is the number of panels will increase and this can have a significant effect on run time for meshes with large numbers of non-planar panels present. Ideally you would choose a meshing method that produces planar meshes from the outset. That way you can maintain control of the size and number of panels in your analysis.