L06 Full QTF diffraction analysis

Introduction

In this example, we model a tension leg platform (TLP) following the general specification of the ISSC TLP prototype (Taylor, R. E. & Jefferys, E. R., 1986). TLPs consist of a floating platform connected to the seabed via vertical mooring tendons (also referred to as tension legs or tethers). Used in offshore oil and gas production since the early 1980s, TLPs have many attractive features: the high axial stiffness in the tension legs is designed to eliminate virtually all vertical motion (heave, roll and pitch) whilst allowing some horizontal movement. They are suited to a great range of water depths (30 – 1500m) which has made them increasingly attractive in the offshore wind sector, allowing floating wind turbines to be installed further from the shoreline.

A phenomenon often associated with TLPs is “springing” – in which high frequency resonant motion of the structure is excited by the environmental conditions. The shortest natural periods discussed in the literature are 2–4 seconds whereas incident water waves typically have periods of 7–15 seconds. Thus, a full quadratic transfer function (QTF) calculation is invariably required in order to model the effect of high-frequency second-order wave loading. In this example, we demonstrate how to build an OrcaWave model and set the necessary QTF input data to compute sum frequency QTFs.

The full QTFs are given by the sum of the OrcaWave results for quadratic load, which is a second-order load associated with the first-order wave potential, and potential load, which is the load associated with the second-order potential. Both components of the full QTF are equally important, but we focus particularly on the potential load in this report. Quadratic loads are covered in less detail in this report because they are discussed in other OrcaWave Examples (L01, L02 and L03) and the technical note “OrcaWave – working with meshes”.

Platform details

Our TLP structure is based on the ISSC TLP (Taylor, R. E. & Jefferys, E. R., 1986). It has two planes of symmetry and is held at a draught of 35m. The main dimensions are shown in Figure 1 and further details are given in Table 1.

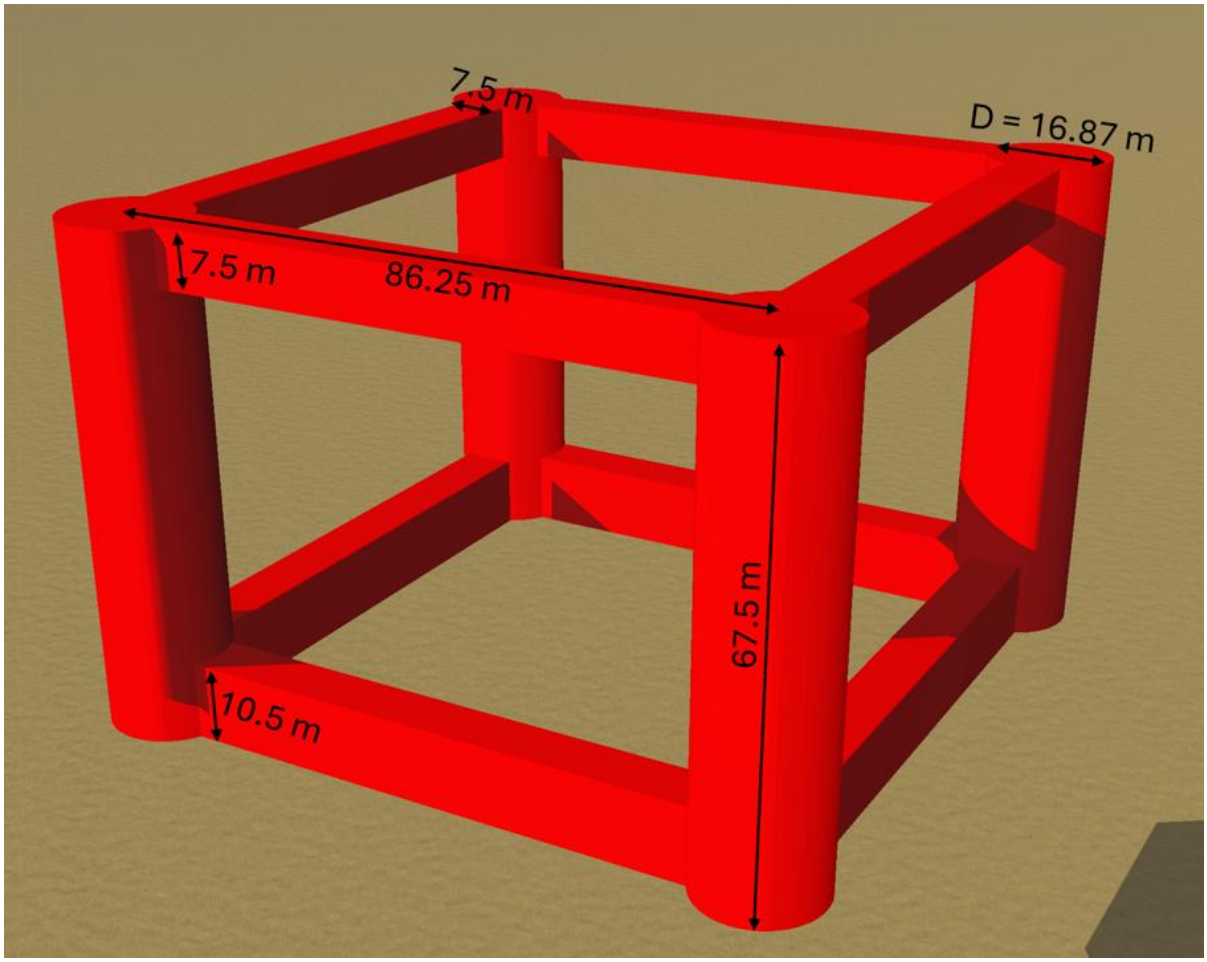

Figure 1: The main dimensions of the ISSC TLP

Parameter	Value
Mass	40500 te
Draught	35 m
Vertical position of centre of mass	38 m
Roll moment of inertia	82.37 x 10^6 te m^2
Pitch moment of inertia	82.37 x 10^6 te m^2
Yaw moment of inertia	98.07 x 10^6 te m^2

Table 1: Specifications of the ISSC TLP (Zhou B. Z. & Wu G. X., 2015).

To begin our OrcaWave diffraction analysis we need a panel mesh to represent the submerged surface of the TLP. OrcaWave accepts mesh files in a variety of formats (including WAMIT .gdf, AQWA .dat and Hydrostar .hst). For a full list, please see the Mesh file formats page of the OrcaWave help. In this example we use the gdf format. In order to conduct a convergence study, body meshes with a range of panel sizes will be used to assess the level of discretisation error in our results and to verify convergence (as described in the "OrcaWave – working with meshes"

technical note). For this reason, we have written a mesh-generating Python script that accepts, as user-input, a target panel size.

The panel size is uniform throughout the mesh except near the waterline. At the waterline it is beneficial to perform vertical refinement to create panels with smaller vertical height. This helps to reduce discretisation error in second-order load calculations, both quadratic loads and potential loads, arising from line integrals around the body waterline. Our Python script therefore performs this vertical refinement near the waterline by giving the panels a vertical height that is 10% of the panel height used throughout the rest of the body mesh.

An example mesh (*L06 ISSC TLP 1.25m panels.gdf*) is provided in the zip folder and is illustrated in Figure 2. Note that Figure 2 also demonstrates how we exploit the symmetry of the system – we need only provide OrcaWave with one quarter of the overall structure and the complete floating platform is then inferred from the symmetry information that we supply (as shown in Figure 3). Taking advantage of global symmetry in this way reduces the memory requirement and accelerates the calculations in OrcaWave.

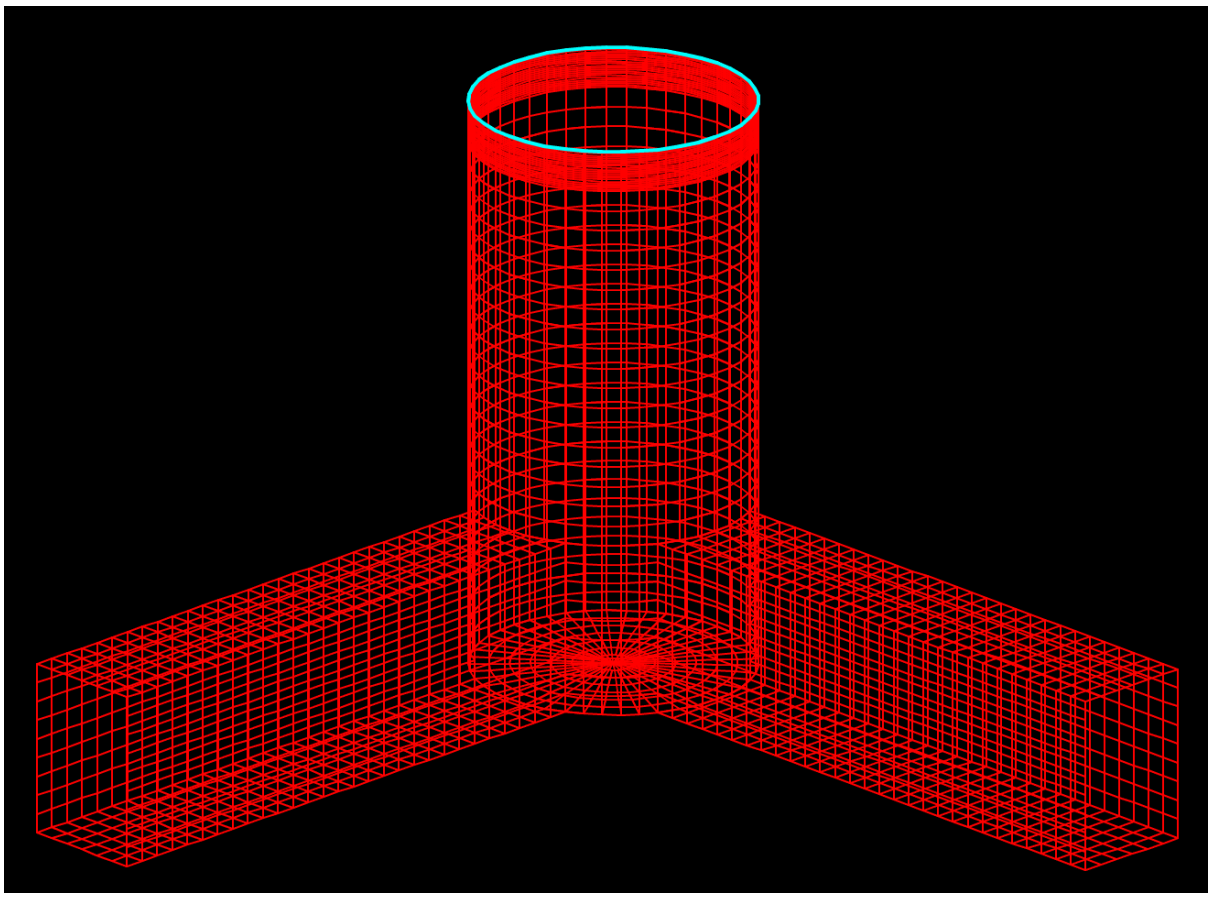

Figure 2: Example of a TLP mesh with target panel size of 1.25 m. Note the vertical refinement near the waterline (shown in blue).

Body mesh options

Format	Length units	Symmetry	Dipole panels	Body number	Import dry panels
Wamit gdf	m	xz and yz planes			

Figure 3: Specifying the mesh file's symmetry via the Bodies page in OrcaWave.

In order to obtain realistic QTF results, the displacement RAOs calculated in OrcaWave must also be realistic. To achieve this the mooring tendons must be represented within the OrcaWave calculation. This could be done by:

- Using spring/dampers within OrcaWave (available from version 11.5 onwards) to model the tendons.
- Building a model in OrcaFlex with line objects representing the tendons. The OrcaFlex mooring stiffness report tool can then generate an external stiffness matrix for use in OrcaWave.
- Using fixed degrees of freedom in OrcaWave to suppress vertical motion (by fixing the *heave*, *roll* and *pitch* degrees of freedom on the Constraints page).

For simplicity and since the focus of this example is the QTF input data rather than the TLP model itself, we use the final option in this example, as shown in Figure 4.

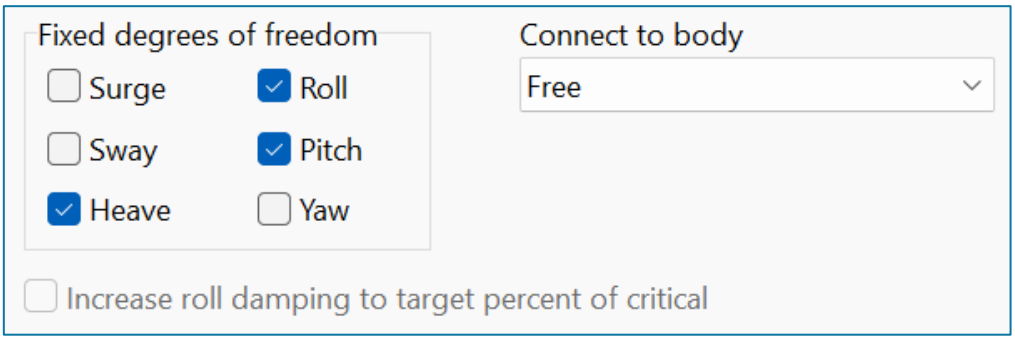

Figure 4: Fixing the Heave, Roll and Pitch degrees of freedom on the Constraints page to represent the effects of the mooring tendons.

Environment data

As mentioned in the introduction, TLPs have been installed in water depths ranging from 30m up to 1500m. In this example we have a water depth of 450m, which is similar to the water depths for TLP installations in the Gulf of Mexico.

Previous studies have involved wave periods in the range 7–15 seconds. Including second order effects, one study included periods of 4.5–14.1 seconds. Here, we consider incident wave periods in the range 8–24 seconds, giving rise to second order sum frequency effects with periods as short as 4 seconds. In this example we focus on high-frequency springing, and we neglect difference frequency effects.

We study unidirectional waves with headings of 0°, 30°, 60° and 90°, exploiting further the symmetry of the platform which means calculations for wave headings beyond 90° are unnecessary. The restriction to unidirectional waves is achieved by setting the range of *QTF crossing angles* to zero degrees on the *Environment* data page in OrcaWave, as shown in Figure 5.

With first-order incident wave periods in the range 8–24 seconds, the full QTF calculation will involve second-order wave periods in the range 4–12 seconds. However, whilst fine-tuning the model's input data we need not consider each possible period, or pair of periods. We will need to perform very many calculations whilst searching for input data which gives robust converged results, so it is of key importance to minimise the run time to make the convergence study

manageable. Therefore, many of our calculations will be limited to a representative range of wave periods and directions, as discussed in later sections.

Note that we include *Intermediate results* in the *Output options* in OrcaWave. This enables us to use restart analyses as another means to reduce model run time, by avoiding repetition of first-order calculations each time we perform a full QTF calculation.

First order convergence study

In this section, we discuss the convergence studies for Load RAOs and Mean drift loads. Whilst each study could be performed in isolation, it is more efficient to make OrcaWave models that address both simultaneously. We have provided the file *L06 First order convergence study.owd* as an example, and in the sections below, we describe how the data pertinent to each study is set.

Load RAOs

Prior to calculating the full QTFs, we need to perform a convergence study to ensure that the body mesh is adequately refined to accurately calculate the first-order results (added mass and damping, load RAOs and displacement RAOs). As mentioned in the 'Environment data' section above, we don't need to include the entire range of incident wave headings and periods (0°, 30°, 60°, 90° and 8–24 seconds). Instead, to reduce the calculation time during the convergence study, we restrict the environmental data to a representative subset. We choose to include incident waves with headings of 0° and 30°, and with periods of 4, 8, 16 and 24 seconds, as shown in Figure 5.

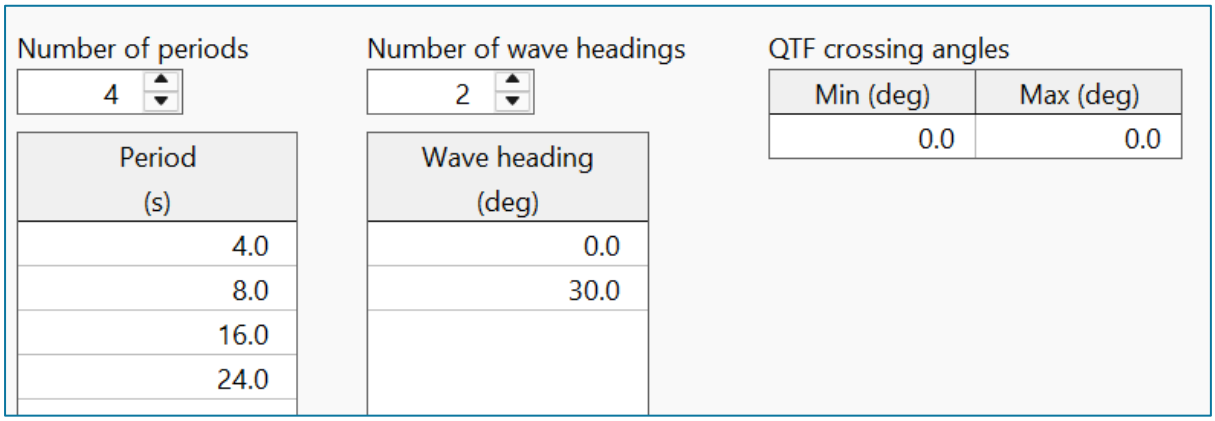

Figure 5: Setting the Environment data. Note the restriction to unidirectional waves via the QTF crossing angles.

Note that the shortest incident wave period we consider is 8 seconds, and this will therefore be the shortest period for the first-order velocity potential in our model. Nevertheless, we include a 4 second wave in our first-order convergence study because we must confirm that the mesh is sufficiently refined to calculate a second-order velocity potential at this period, which corresponds to the shortest sum-frequency wave in our final model.

We set up our OrcaWave models to solve a first-order problem using **both** the Haskind and diffraction calculation methods to calculate the load RAOs. Comparing the relative error in the results from these two calculation methods can be a simple and instructive way to assess the level of discretisation error in first-order diffraction results. Specifically, this comparison measures the discretisation error in the potential formulation.

Figure 6 and Figure 7 show the results from such a comparison. The relative error plotted is calculated by taking the difference between the results for Haskind load RAOs and diffraction load RAOs (using complex arithmetic). This difference is divided by our best available estimate of the vector magnitude of the force (for surge, sway and heave) or moment (for roll, pitch and yaw). The two calculation methods for load RAOs are different, but in general are considered to be equally accurate. Therefore, for a given body mesh, our best available estimate of the true load RAO is taken to be the arithmetic mean of the Haskind and diffraction load RAO results.

As shown in the plots, the percentage relative error decreases as the body mesh panel size is reduced. For the purposes of this example, we will assume that a maximum error of 2% in these results is acceptable, and therefore we select the body mesh with a target panel size of 1.25m.

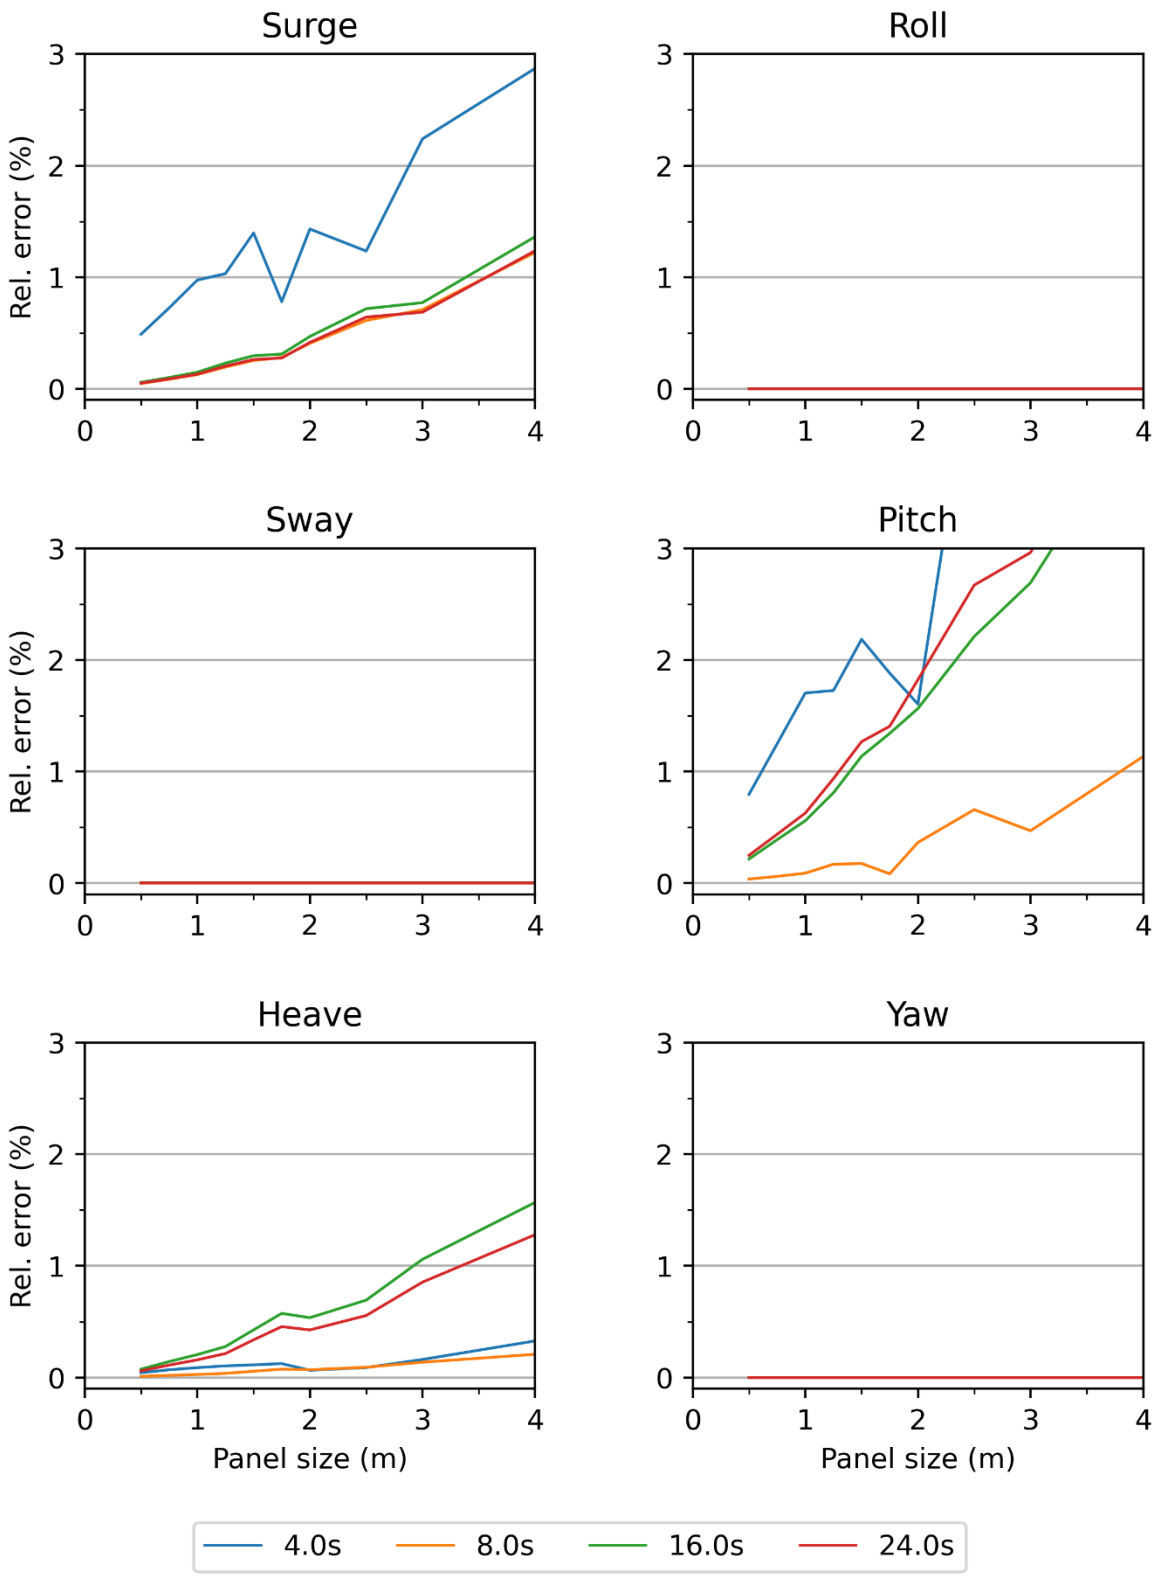

Figure 6: Relative error between Haskind and diffraction load RAOs for all 6 DOFs for incident wave heading of 0°.

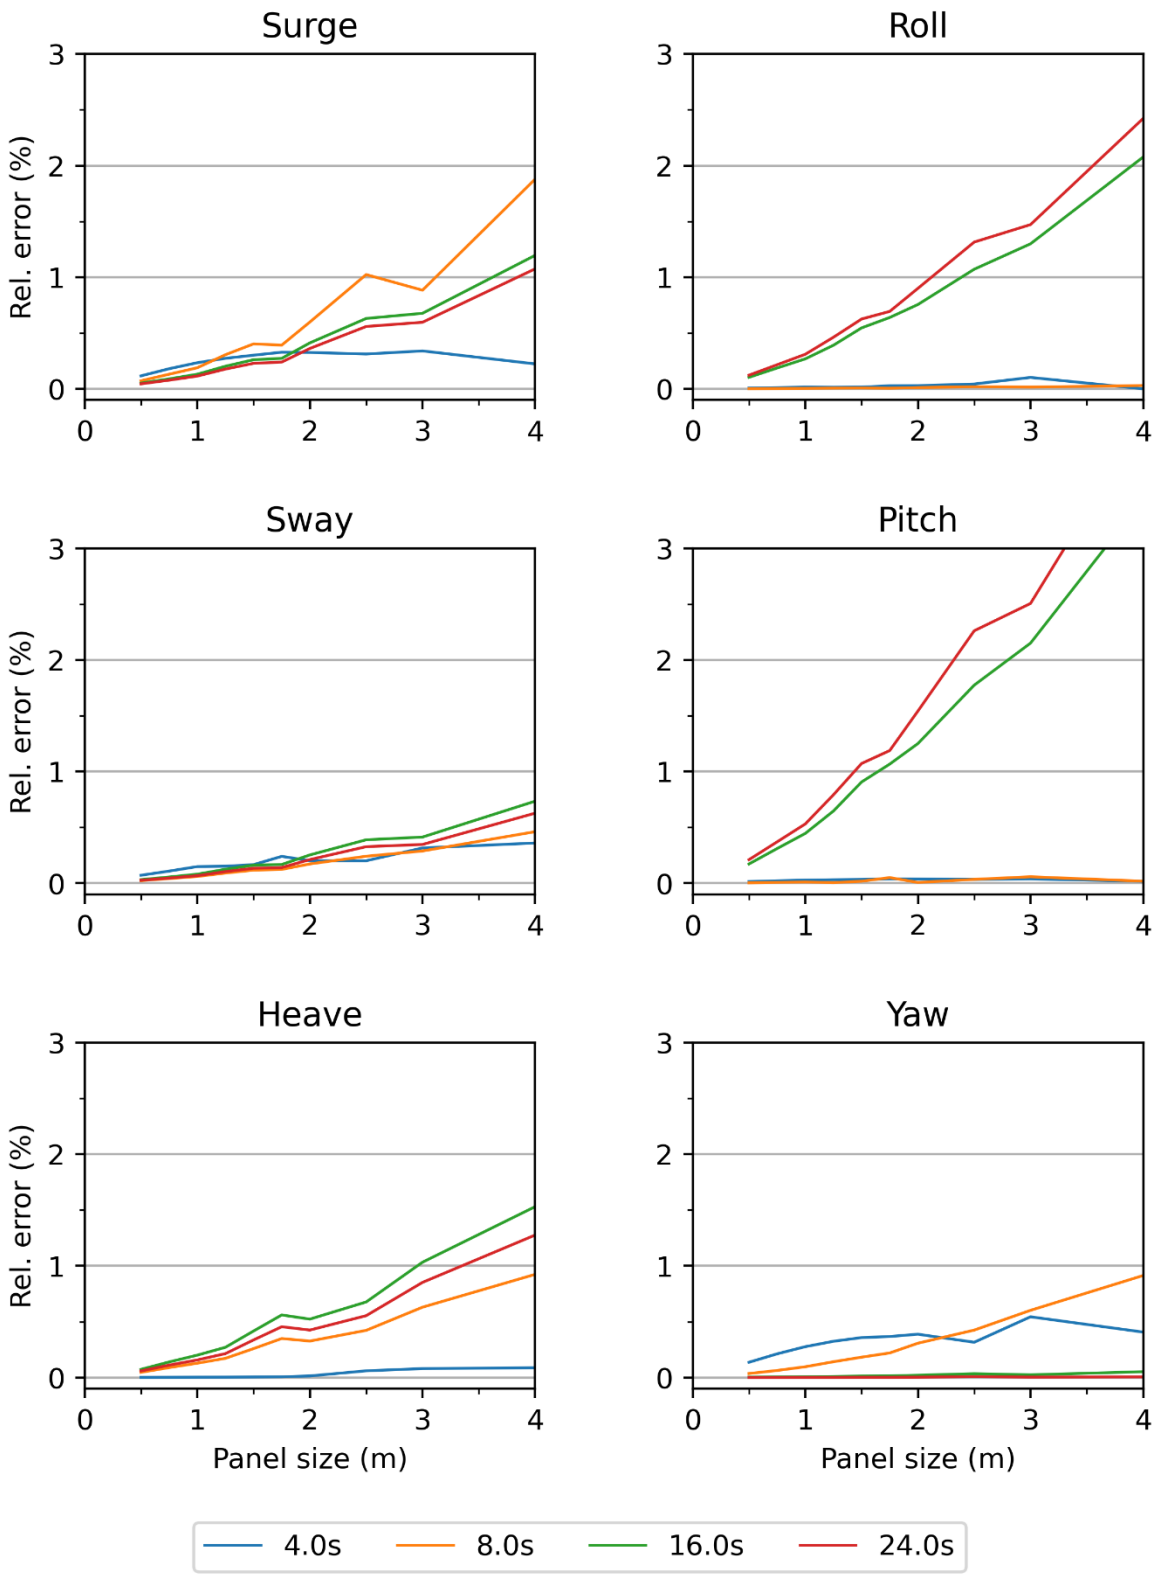

Figure 7: Relative error between Haskind and diffraction load RAOs for all 6 DOFs for incident wave heading of 30°.

Mean drift loads

We extend our convergence study by ensuring that the body mesh is adequately resolved to calculate accurate mean drift loads. Mean drift loads are a subset of the time-varying quadratic loads, being the quadratic loads for difference frequencies of zero. Aside from sometimes being included in the statics phase of an OrcaFlex calculation, it is also important to assess the accuracy of the mean drift loads because they can give an early indication of the likely accuracy of the time-varying quadratic loads, which form part of the full QTF calculation.

Again, we restrict the environmental data to decrease calculation time during the convergence study. We consider unidirectional incident waves with headings of 0º and 30º and with periods of 8, 16 and 24 seconds (a 4 second incident wave is not relevant here, as that period will only arise from sum frequency contributions in our full model).

Mean drift loads are available from first-order OrcaWave models. Three different calculation methods are available:

- Pressure integration (PI). Note that this method requires the solve type to include the source formulation.
- Control surface integration (CS).
- Momentum conservation (MC).

Our models use all three calculation methods. The necessary mesh for the CS method is defined on the *Bodies* page of the OrcaWave model. We elect it to be automatically generated using panels of equal size to those in the body mesh, with the separation from the body equal to five panel lengths.

It is especially important to include mean drift loads in the first order convergence study because obtaining accurate results is known to be challenging in some cases. Furthermore, since both the PI and CS methods involve a line integral around the body waterline (similar to those which form part of a full QTF calculation), checking for convergence at this stage provides an early opportunity to assess the suitability of the body mesh for the full QTF calculations to follow.

Since our full QTF analysis will consider unidirectional waves, we focus here on unidirectional heading pairs: (0º, 0º) and (30º, 30º). In such situations, the mean drift load results are purely real numbers. Furthermore, since mean drift loads become less significant for longer incident waves, for brevity we focus on the shortest incident wave period (8 seconds for this part of the study).

Figure 8 and Figure 9 show comparisons of the mean drift load results computed with the three available methods. Note that:

- Only surge, sway and yaw results are available from the MC method, but results for all six components are available from the CS and PI methods.
- It is usually found that the CS and MC methods are quicker to converge, as body mesh panel size is reduced, than the PI method.

The above points mean that we cannot compute a simple relative error metric in the same way as we did for load RAOs. Instead, we have plotted the dimensional load results for each method. As an alternative strategy to display the results in the context of an appropriate scale, we have ensured that the vertical axis has the same limits for all three force components (surge, sway and heave) and all three moment components (roll, pitch and yaw). This avoids giving undue prominence to relatively small components of the load (e.g. the sway force in Figure 8).

Figure 8 and Figure 9 show good general agreement between the three methods. They are consistent with the usual expectation that the CS and MC methods converge faster than PI, and they show that agreement between all three methods improves as the body mesh becomes finer.

We note that the maximum relative error between the methods, e.g. as seen in the heave force components, is larger than the relative errors we saw between the load RAO methods in Figure 6 and Figure 7. This is typical behaviour, and is caused by the more complicated calculations performed to compute these results:

- Both the PI and CS methods involve a line integral around the body waterline which can be a significant source of discretisation error. As discussed above, we use vertical refinement in the body mesh to mitigate this error.
- The PI method uses results from the source formulation (specifically, tangential fluid velocity on the body surface) which are prone to greater discretisation error than the results of the potential formulation.

For the purposes of this example, we will assume that the accuracy of the CS results in Figure 8 and Figure 9 is acceptable for the body mesh with a panel size of 1.25m. The relative error compared to the PI method is clearly greater than the 2% metric we used when comparing load RAO methods. If a smaller relative error were desired, then Figure 8 and Figure 9 demonstrate that this could be achieved by using a finer body mesh.

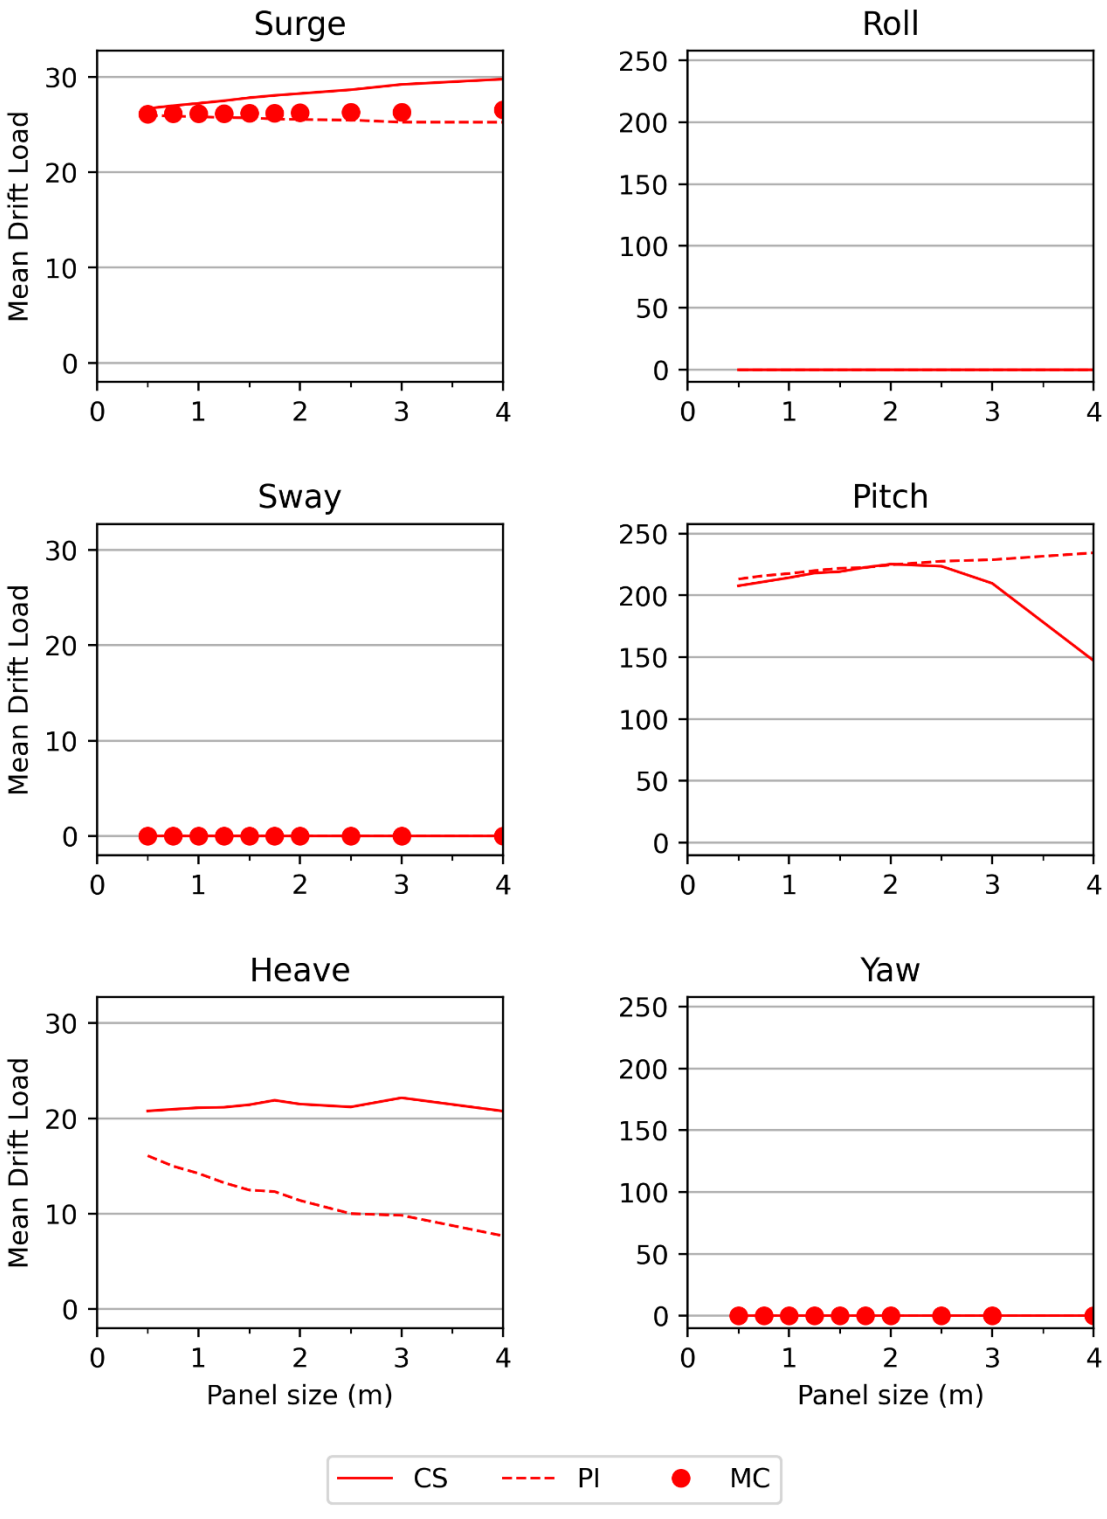

Figure 8: Mean drift load values calculated using the control surface integration method (CS), the pressure integration method (PI) and the momentum conservation method (MC) for the (0°, 0°) heading pair and the 8 second wave period. The units for the left-hand plots are kN/m² whilst the right-hand plots have units of kN.m/m².

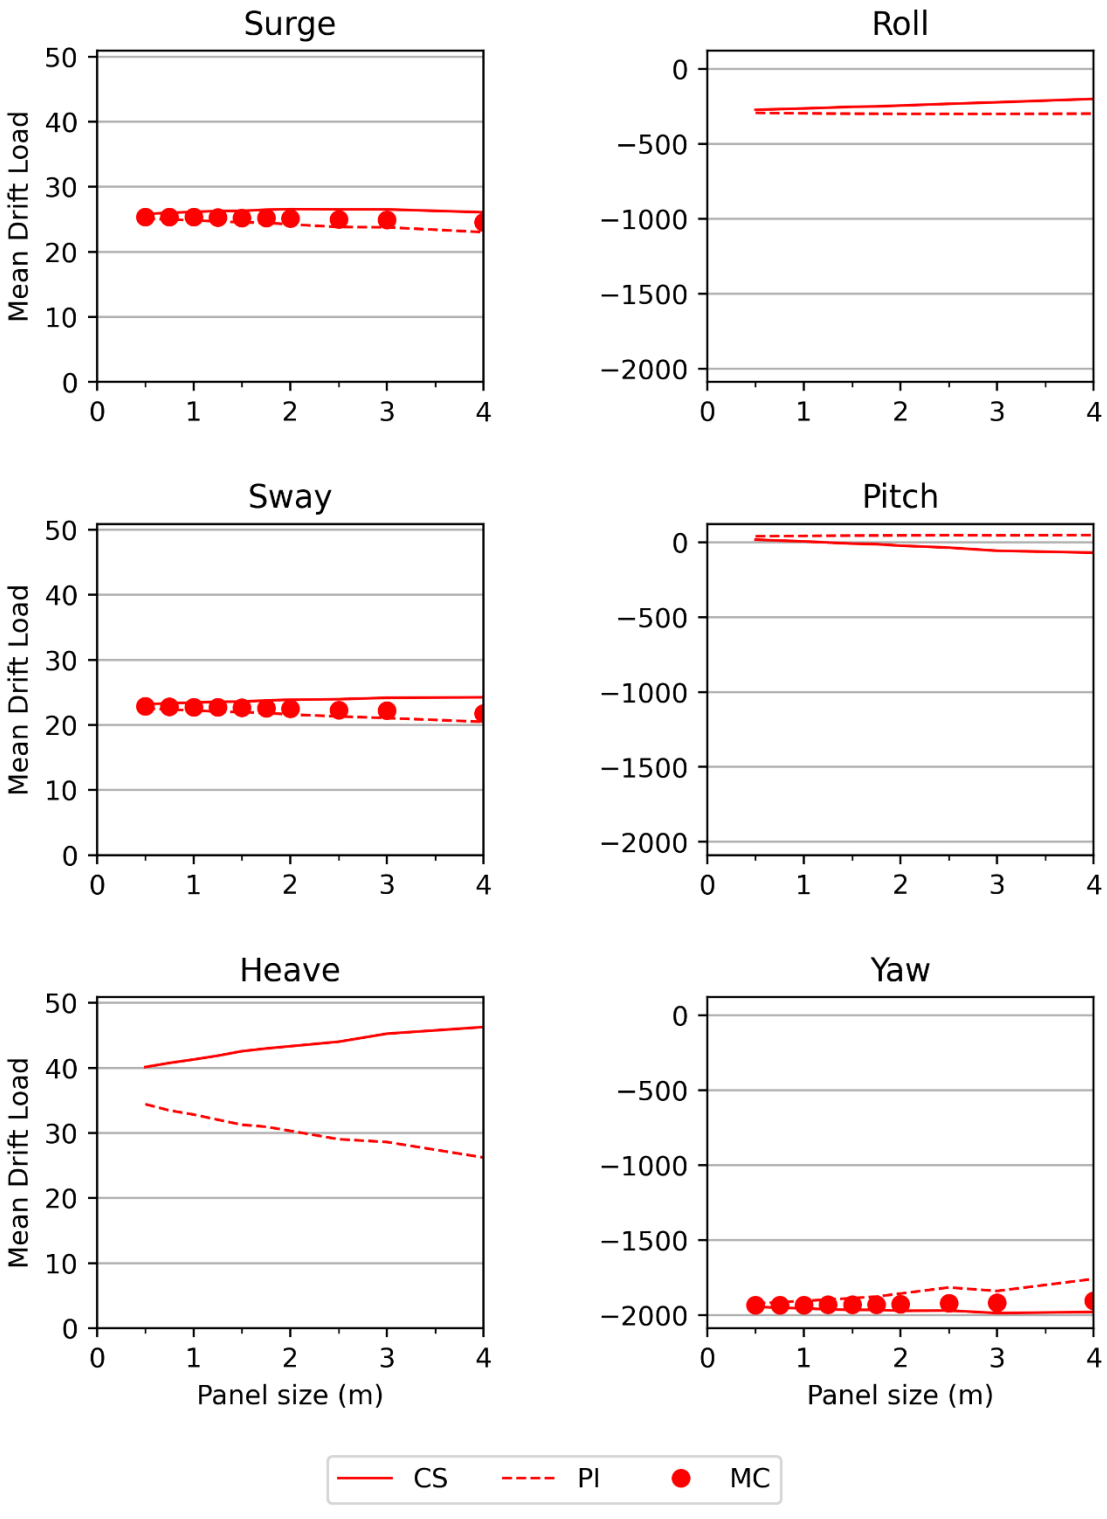

Figure 9: Mean drift load values calculated using the control surface integration method (CS), the pressure integration method (PI) and the momentum conservation method (MC) for the (30º, 30º) heading pair and the 8 second wave period. The units for the left-hand plots are kN/m² whilst the right-hand plots have units of kN.m/m².

Time-varying quadratic loads

With the first-order convergence study complete, we now turn our attention to the time-varying quadratic loads. As mentioned in the introduction, time-varying quadratic load results are added to the potential load results to obtain the full QTFs. Thus, to generate these results, we need to upgrade the diffraction *solve type* to a full QTF calculation.

As detailed on the help page, there are many parameters associated with a full QTF calculation. These parameters define the free surface used in the calculation of the potential loads which we discuss in detail in the next section. Since the time-varying quadratic loads are associated with the first-order wave potential, they are independent of the free-surface data and we can assess their convergence separately, before moving on to study the potential loads.

The quickest way to generate time-varying quadratic load results is to minimise the time spent calculating the potential load. Therefore, as can be seen in the file *L06 Time varying quadratic loads.yml*, we neglect free surface forcing by providing no mesh file for the free surface panelled zone (as shown in Figure 10) and include only the direct QTF calculation method.

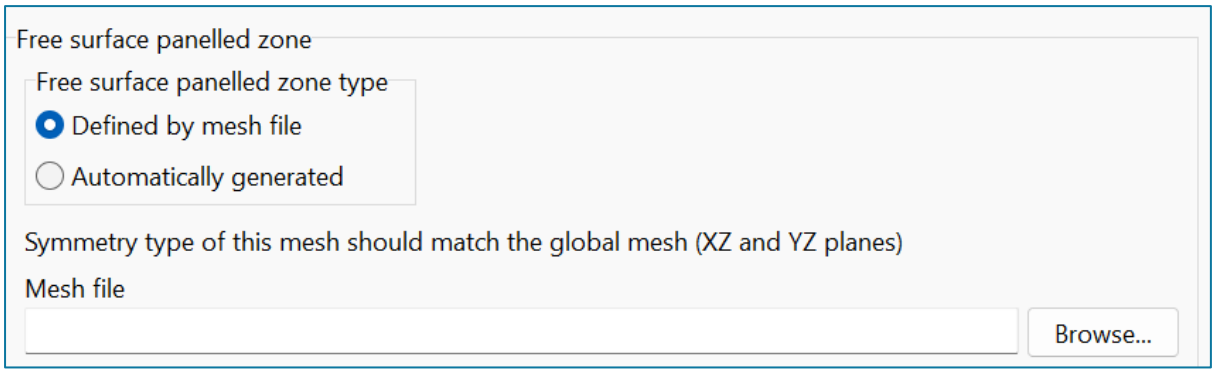

Figure 10: Leaving the mesh file blank neglects the free-surface forcing contribution to potential load calculations.

Time-varying quadratic load results are then available from the CS and the PI methods. For the CS method, we use the same mesh as was used when generating mean drift loads above, i.e. automatically generated by OrcaWave using panels of equal size to those in the body mesh, with separation from the body equal to five panel lengths.

Figure 11 compares results from both methods, for both the real and imaginary parts. We again ensure that the vertical axis has the same limits for all three force components (surge, sway and heave) and all three moment components (roll, pitch and yaw). For brevity, we present the results for the (30°, 30°) heading pair and for the second order wave period of 9.6 seconds (the combination of the first-order 16 second and 24 second waves). This plot is representative of the general results and demonstrates good agreement between the two methods. The CS method converges more quickly than the PI method, and in general, agreement between the two methods improves as the body mesh becomes finer.

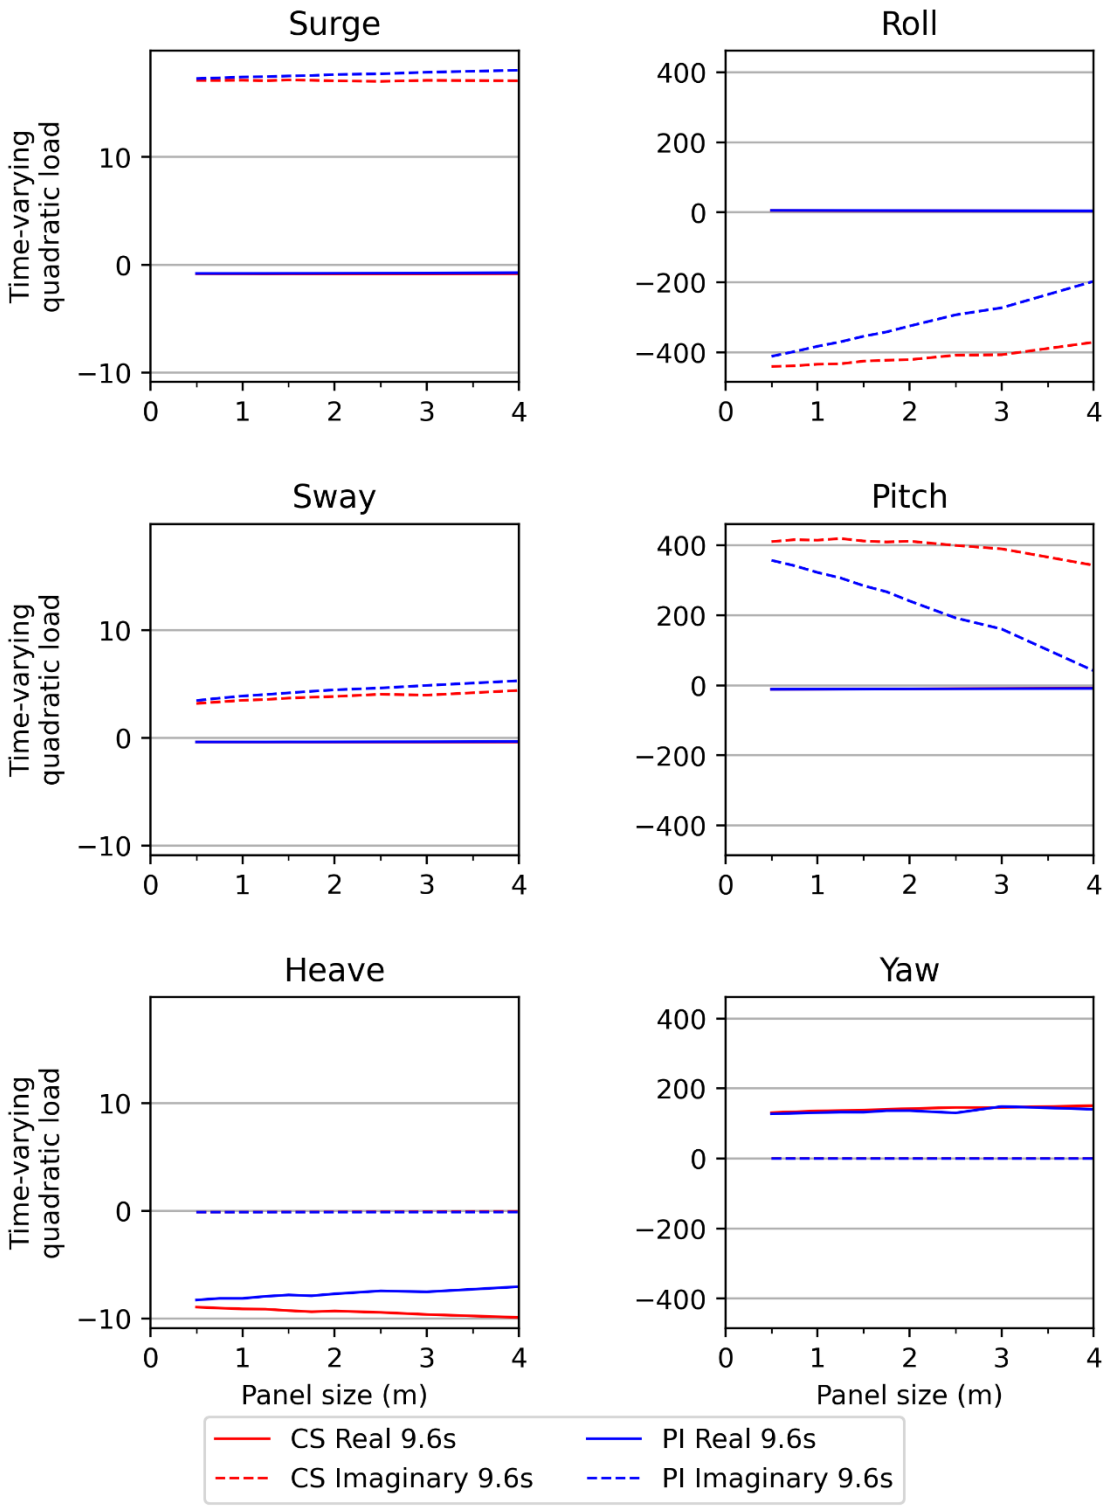

Figure 11: Real and imaginary parts of the time-varying quadratic load results calculated using the control surface integration method (CS) and the pressure integration method (PI) for the (30°, 30°) heading pair and the second-order wave period of 9.6 seconds. The units for the left-hand plots are kN/m² whilst the right-hand plots have units of kN.m/m².

QTF potential loads

In this section, we explain in detail how we select appropriate parameters to discretise the free surface during the calculation of potential loads. Figure 12 gives a pictorial representation of the key terms. The online help (includes links to the relevant theory sections of the documentation and also provides some useful guidance for specifying the QTF data (in the form of a "recipe") which we follow here.

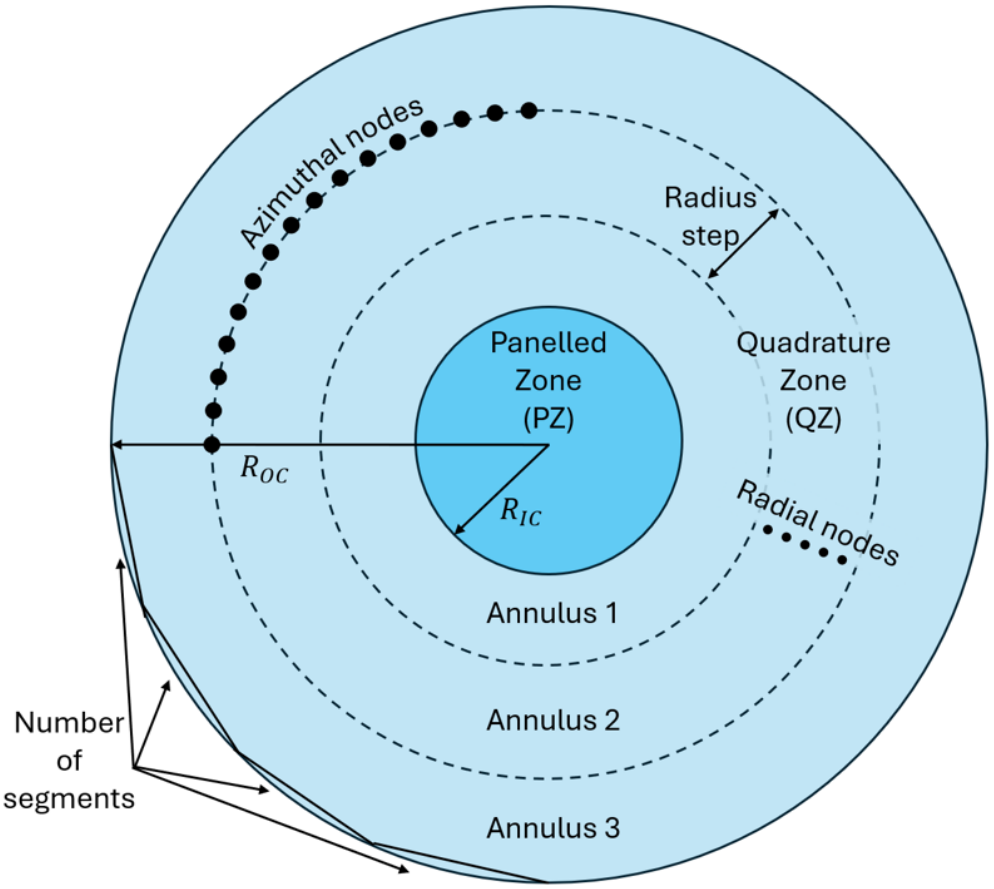

Figure 12: Diagram to show the main QTF data, based on a system with two planes of symmetry and with three annuli in the quadrature zone (QZ). The Number of outer circle segments is four in this diagram.

The process of testing and refining the QTF parameters is described below. It involves comparisons of many different experimental models, thus it is desirable to ensure that these models are optimised for speed. This is done by:

- Restricting once again the wave frequencies and headings (we consider unidirectional incident waves with headings of 0° and 30°, and with periods of 8, 16 and 24 seconds with the maximum QTF period set to 12 seconds).
- Minimising calculations that are not relevant to the potential loads that we are investigating in this section. For example, using only the Haskind method for the load RAOs and only the PI method for the quadratic loads.

- Using restart analyses to ensure that the first order results are not needlessly re-calculated. We restart our experimental models from a first order model (*L06 Parent model for full QTF experiments.owd*) set up as described above.

The indirect QTF calculation method takes longer to run than the direct method, however we nevertheless choose to include both direct and indirect QTF calculation methods as we shall use both methods to assess the accuracy of our results.

As recommended in the help page (Data | QTFs), we begin by setting up an initial experimental model (*L06 Potential loads (first experimental model).yml*), with parameters defined according to the given recipe. We then systematically refine each parameter in turn, to assess the sensitivity of the potential load results.

Based on the preceding sections we use the same body mesh, with panels of 1.25m length and vertical refinement at the waterline. Table 2 shows the QTF data used in the first experimental model, calculated according to the recipe. We will use this first model as a basis for our investigations into the suitability of the QTF data.

QTF data item	Calculation	OrcaWave input
Panelled zone panel size	= matching the body mesh	1.25m
Inner radius (R_{IC})	Body radius + 10 × panel size $= 72.9 + 10 \times 1.25 = 85.4$ m	85m
Outer radius (R_{OC})	Body radius + min(depth, λ_{max}) $= 72.9 + \min(450, 896)$	523m
Number of annuli	$\frac{R_{OC} - R_{IC}}{2\lambda_{min}} = \frac{523 - 85}{2 \times 25}$	9
Radius step	$\frac{R_{OC} - R_{IC}}{\text{number of annuli}} = \frac{523 - 85}{9}$	48.6m
Radial nodes	$= 12$	12
Azimuthal nodes	$\frac{12\pi R_{IC}}{2^2 \lambda_{min}} = 32.044$	32
Number of outer circle segments	$\frac{12\pi R_{OC}}{2^2 \lambda_{min}} = 197.166$	197
Expansion order	$= 24$	24

Table 2: The QTF data inputs for the first experimental model, based on the recipe given in the help documentation.

Some of the input data in Table 2 may need to be changed to more appropriate values – this will be assessed by a series of sensitivity tests. The sensitivity tests are based on the general premise that, if the model is well-resolved, results will be unaffected by refining a given data item.

Sensitivity tests of the experimental model

Following the guidance in the OrcaWave Help page, we perform sensitivity tests for the data items in Table 2.

Before discussing the test results, we note that in our experience there are two distinct patterns in the variation of potential load results as different QTF data items are varied:

1. Some data items cause the results to vary in a smooth and continuous way. As the data is refined (i.e. changed to a value which is expected to give more accurate results, for example increasing the outer radius) the potential load varies continuously. Ideally, the results are seen to approach a stable limiting value as refinement continues.
2. Other data items cause the result to vary in a discrete way. As the data is refined the potential load results change, but once a threshold value is reached the results suddenly stop changing. Further refinement (for example increasing the number of radial nodes) has no significant effect on the potential load results whatsoever.

In the discussions below, we measure the sensitivity of the potential load results to changes in the QTF data by calculating a "*largest percentage change*" as follows. Using complex arithmetic, we calculate the difference between the new and original results (from *L06 Potential loads (first experimental model).yml*), for each of the direct and indirect calculation methods for the (30°, 30°) heading pair. The vector magnitude of each difference is divided by the corresponding vector magnitude of force (for surge, sway and heave) or moment (for roll, pitch and yaw) from the original results. We then retain whichever percentage change is largest (for both direct and indirect methods, for both force and moment results, and over all wave periods). For brevity, we give only the largest percentage change in this report.

Sensitivity test 1: Radial nodes

To assess the sensitivity of the potential load results to changes in the number of *radial nodes*, we increase this data item from 12 to 16. Despite this significant increase by a factor of 1/3, the potential load results are practically unaffected (largest percentage change of 7.8e–5%, calculated as described above). We can therefore conclude that the free surface quadrature zone (QZ) in this model is sufficiently well-resolved in the radial direction with 12 nodes.

Sensitivity test 2: Azimuthal nodes

To assess the sensitivity of the potential load results to changes in the number of *azimuthal nodes*, we increase this data item from 32 to 42 (again, an increase of approximately 1/3). Once more, we observe no discernible change in the potential load results (largest percentage change of 5.8e–6%). Thus 32 *azimuthal nodes* are sufficient for this model.

Sensitivity test 3: Expansion order

To assess the sensitivity of the potential load results to changes in the *expansion order*, we increase this data item by 1/3, from 24 to 32. This time, there is a small change in the potential load results

for the second order wave period of 4 seconds (the largest percentage change value is 0.37%). In the context of the 2% maximum error we accepted in the first order convergence study, this value is small. Thus, we conclude that an *expansion order* of 24 is sufficient to ensure the free-surface asymptotic zone is well resolved in this model.

Sensitivity test 4: Outer circle segments

To assess the sensitivity of the potential load results to changes in the *number of segments* in the outer circle, we increase this data item from 197 to 263. The results showed little sensitivity, with the largest percentage change of 0.03%. Thus, we conclude that 197 segments are sufficient to resolve the circumference of the outer circle.

Sensitivity test 5: Outer radius

To assess the sensitivity of the potential load results to changes in the *outer radius* (R_{OC}), we increase this data item. This could be done by either increasing the *number of annuli* in the free surface QZ, or else by increasing the value of the *radius step*. We choose the former, in order to maintain the node density of the original model within the QZ. We note here that, in contrast to the preceding sensitivity tests, R_{OC} cannot be refined in isolation. The *number of segments* in the outer circle must be changed accordingly (i.e. using the formula given in the recipe and reproduced in Table 2).

R_{OC} is increased from 523m to 717.6m by increasing the *number of annuli* from 9 to 13. Correspondingly, the *number of segments* in the circumference of the outer circle is increased from 197 to 271. The potential load results show little sensitivity to these changes, with a largest percentage change of 0.47%. Thus, we conclude that the outer circle is sufficiently large when R_{OC} is 523m.

Sensitivity test 6: Inner radius

To assess the sensitivity of potential load results to the extent of the panelled zone, we increase the *inner radius* (R_{IC}). When increasing R_{IC}, we could allow R_{OC} to increase correspondingly, or we could reduce the *radius step* to ensure that R_{OC} remains constant. Since we have confirmed that the results are insensitive to changes in R_{OC}, both methods are equally valid as a sensitivity test. For simplicity we choose the latter (to avoid having to recalculate the number of outer circle segments, which would be necessary if R_{OC} were to change). We note that we need not recalculate the number of *azimuthal nodes* when R_{IC} is increased, since we have already confirmed that this number was sufficient throughout the entire QZ, including this larger radius. Similarly, we need not recalculate the *number of annuli* since increasing R_{IC} (with R_{OC} fixed) will increase the density of the collocation points in the QZ, thus will not jeopardise the accuracy of the results.

We increase R_{IC} from 85m to 100m and recalculate the *radius step* value to be 47m. In response to these changes, the potential load results show little sensitivity, with a maximum percentage change of 0.43%. Therefore, we conclude that the panelled zone is sufficiently large when R_{IC} is 85m.

Discussion

Discretisation error in the full QTF results

With the sensitivity tests now complete, we can be confident that our input data for the free surface (see Table 2) is sufficient for calculating reliable potential load results – that is, none of those data items require further refinement. In combination with the quadratic load results, we can conclude that the model will give reliable results for full QTFs.

However, we do not yet have a measure of the discretisation error in the potential load results, or the overall accuracy of the full QTFs. We can assess this by comparing the potential load results for the direct and indirect QTF calculation methods. Much like the two load RAO calculation methods, and the three quadratic load calculation methods, the direct and indirect QTF calculation methods are theoretically equivalent to each other. The difference between them gives a measure of the discretisation error in the results. Figure 13 shows the maximum relative error (relative to the arithmetic mean of the results) as a function of body mesh panel size. This indicates that the relative error is of the order of 10% for the 1.25m body mesh. Figure 13 also clearly demonstrates that the relative error is a direct function of body panel size, and hence if more accurate results were desired then this could be achieved by using a finer body mesh.

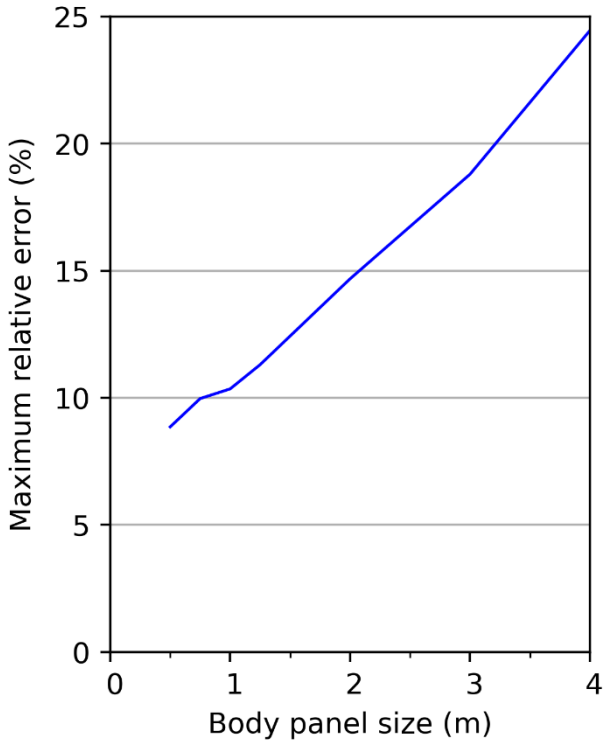

Figure 13: Maximum relative error between direct and indirect potential loads. The maximum is taken over all 6 DOFs and all periods for the (30°, 30°) heading pair. For each body panel size, the QTF data are computed following the recipe suggested in the OrcaWave help.

We note that the relative error for the 1.25m body mesh in Figure 13 is larger than the 2% relative error observed between the load RAO methods in Figure 3 and Figure 4. This typical behaviour is analogous to that discussed for mean drift loads in the convergence study above – it is caused by the more complicated calculations performed to compute these results. In particular, whilst the direct potential load is computed using the second-order potential formulation, the indirect potential load uses results from the second-order source formulation (specifically, tangential fluid velocity on the body surface) which are prone to greater discretisation error.

Based on the above, we believe that the potential load results are likely to be more accurate from the direct method than the indirect method, for any given body mesh. This is analogous to the statement that we believe the mean drift load results are likely to be more accurate for the CS method than the PI method, for any given mesh.

Taking a broader perspective, our studies of load RAOs, mean drift loads, time varying quadratic loads and potential loads all point to a consistent conclusion – the quality of the body mesh is the most important ingredient in determining the accuracy which it is possible to achieve. Further refinement of QTF data, or CS data, cannot improve results beyond a limit imposed by the body mesh.

Impact of QTF data on model run time

The recipe for QTF data given in the OrcaWave Help attempts to give a well-resolved free surface calculation, in order to provide a useful starting point when model building. For many models we anticipate that the recipe will provide conservative parameter values for the QTF data, as was the case in this example. Of course, the recommended sensitivity tests should always be conducted to confirm that each parameter is sufficiently refined. But if they are all sufficiently refined and model run time is an issue, it may be desirable to optimise the model run time by seeking next to *coarsen* the QTF data. Once satisfied with the final QTF data, we can reintroduce the full complexity of the production model (the full set of wave frequencies and headings, etc.).

As explained in the online help, there are various ways in which run time might be reduced, namely by:

- Reducing R_{OC}
- Reducing R_{IC} and/or increasing the *panel size* in the panelled zone
- Reducing the number of *radial nodes* and/or *azimuthal nodes* in the QZ, and
- Using only the *direct method* for potential load calculation.

We can seek to optimise model run time by varying each of the above. The approach is similar to the 'Sensitivity tests of the experimental model' section above, but now we know that the potential load results from the experimental model are reliable so we are attempting to find coarser free-surface data that does not significantly change these results.

In addition to adjusting the QTF data, as described above, another approach to improve model performance is to use the Computational strategy of OrcaWave's solver. This option may further optimise performance for large models.

To demonstrate the impact of QTF data on model run time, we consider our *L06 Potential loads (first experimental model).yml* (which uses the recipe data precisely) and adapt it in two very different ways. Firstly, we remove the QZ entirely and increase R_{IC} such that the panelled zone extends all the way to the outer circle radius of 523m. Secondly, we optimise the QTF data following the approach described above. Table 3 shows the wall clock time for each model (as reported on the *Diagnostics* sheet on the *Tables* page of OrcaWave). Each model gives equally accurate potential load results (the greatest percentage change between the results is 0.56%), but the run time differs significantly:

- The presence of the QZ (even with conservative QTF data) reduces calculation time by more than a factor of 10 in this example.
- Optimising the QTF inputs can reduce calculation time further still, by almost a factor of four in this example.

Model version	Wall clock time (s)
No QZ, $R_{IC} = 523$ m	25,169
L06 Potential loads (first experimental model).yml	1,585

Optimised experimental QTF model	420

Table 3. Comparison of run time data for different models to demonstrate the impact of QTF data.

Table 3 demonstrates that the values chosen for QTF data can dramatically affect model run time – even though these three models all give equivalent (equally accurate) results. Of course, it would be meaningless to optimise model run time unless we had first confirmed the reliability of our model and estimated the accuracy of its results – therefore those more important tasks have been the principal focus of this example.

References

Taylor, R. E. & Jefferys, E. R. (1986). Variability of hydrodynamic load predictions for a tension leg platform. Ocean Engineering, 449-490.

Zhou B. Z. & Wu G. X. (2015). Resonance of a Tension Leg Platform excited by Third-Harmonic Force in nonlinear Regular Waves. Philosophical Transactions A.