The OrcaFlex VIV Toolbox

User Guide to the Time Domain Models

1 Introduction

1.1 VIV models

The time domain version of the VIV Toolbox currently contains four VIV model options:

- Milan Wake Oscillator (MWO)
- Iwan and Blevins Wake Oscillator (IBWO)
- Vortex Tracking 1 (VT1)
- Vortex Tracking 2 (VT2)

This note discusses how these models might be best used.

Section 2 deals with general questions of segmentation, time step and log interval.

Section 3 then deals individually with each of the models, and the appropriate values which should be assigned to the governing parameters (where these are adjustable by the user).

Section 4 discusses the differences between the VIV models, and the circumstances under which one model might be preferred to another.

Section 5 deals with validation.

1.2 A caveat…!

Users should be aware that VIV analysis is not by any means a settled or mature technology. Some well-known VIV models have acquired a degree of acceptability by prolonged exposure and frequent use: others are less familiar. None can claim universal superiority. The VIV analyst must be prepared to consider carefully the characteristics of each particular case (e.g. three-dimensionality, directions of flow and body motion, susceptibility to wake washback, etc.) and maintain a high level of scepticism about the results obtained from any of the models. Wherever possible, we recommend repeating critical cases using more than one VIV model – for example a wake oscillator model, a vortex tracking model and, for a steady state case, a frequency domain model.

The Vortex Tracking models require special care, and should be viewed for the present as research tools which require further calibration against real world data. Both models have a number of governing parameters which can be adjusted by the user, and we would strongly recommend anyone considering using these models for project work to carry out sensitivity analysis on critical cases, and where possible, carry out a calibration exercise against similar cases where the VIV behaviour is known. The vortex tracking models are discussed further in Sections 3.2 and 3.3.

2 General issues affecting all models

2.1 Segmentation

Should be sufficiently fine to represent the highest expected mode shape. If the mode shapes are approximately sinusoidal, then the Nth mode has N half sine waves in the "free" length.

> *The "free" length means that part of the line which is free to vibrate. For a tensioned string it is the entire length. For a catenary it is the length from the hang off to touchdown.*
>
> *Note that OrcaFlex calculates modes for oscillations in all three degrees of freedom – X,Y,Z. For a vertical tensioned string in still water, the horizontal modes occur first and come in matched pairs. So OrcaFlex modes 1 and 2 correspond to Mode 1 as usually defined, OrcaFlex Modes 3 and 4 correspond to Mode 2, and so on. Things get even more confusing when axial modes appear, but this is usually too far up the list to matter for a tensioned string. Take care!*

Segment length should not exceed L/(n*N) where L is free length, N is mode number (as usually defined) and n is the number of nodes required per half wave. n = 10 gives a good representation; fewer may be acceptable but less than 5 is probably insufficient.

More segments = longer run time. The skill is in finding the fewest you can get away with without loss of accuracy. Try sensitivity analysis.

2.2 Outer time step, δtO

This is the update step for the vortex force. It should be short compared to the expected period of vortex shedding (the Strouhal period). The Strouhal period T_s is given by $T_s = D/(S*U)$ where D is the diameter of the line, U is the relative speed of the flow normal to the line, and S is the Strouhal number. Typically, S ~ 0.2 for a circular cylinder.

$\delta t_O < T_s/100$ is a good general rule, but on occasion a shorter outer time step is necessary for accuracy. $\delta t_O < T_s/200$ is probably a safer recommendation. Sensitivity analysis is recommended – use the longest outer time step which doesn't cause loss of accuracy. (Example B01 – Drilling Riser is a case in point. Analysis using $\delta t_O < T_s/200$ gives plausible and consistent results; results for $\delta t_O = T_s/100$ are unrealistic.)

Shorter outer time steps = more calculations = longer run time, particularly for the VT2 model. Other models are less sensitive – the vortex calculations for the wake oscillators are not particularly demanding, and the VT1 model has its own vortex time step.

2.3 Inner time step, δti

Should be set according to the usual rules. The OrcaFlex documentation recommends $\delta t_i < 1/10*T_{Nmin}$ where δt_i is inner time step and T_{Nmin} is the shortest natural period of the model. This is usually adequate but on occasion a shorter time step is needed. Try it and see. The criterion is stability of integration – if the simulation crashes, then the time step is probably too long; if not, it is probably OK (but keep a sharp lookout for anomalous behaviour – on rare occasions, a time step which is slightly too long can give stable but incorrect results).

Clearly, the inner time step cannot be longer than the outer time step, but it is quite acceptable for the inner and outer time steps to be equal.

2.4 Log interval

Must be short enough to capture the maxima of the vibration. Assume a sinusoidal vibration with period T_V (usually $T_V \sim T_s$) and m log points per cycle. Log interval $\delta t = T_V/m$. Peak displacement may be missed by up to $0.5*\delta t$ giving a maximum % error of $100*(1 - \cos[2\pi*0.5*\delta t/T_V]) = 100*(1 - \cos[\pi/m])$. m = 15 gives 2% maximum error.

The issue here is mainly SIM file size – short log intervals and long runs give big files.

2.5 Filter period

OrcaFlex uses a digital filter to separate the mean motion of the node from the VIV motion. Filtering is needed with both wake oscillator models in order to allow non-VIV motion of the node to contribute to VIV, without letting the VIV motion feed back into the velocity input into the wake oscillator model. For the vortex tracking models, the filter has no effect on the calculated forces and motions but is used to determine the in-line and transverse directions for reporting purposes. For more detail, see the documentation.

In steady flow with constant current, no waves and no imposed motions of the line, the filter period can be set to 'infinity' for Iwan + Blevins and both vortex tracking models. However, the Milan wake oscillator equations include a "spring" term based on the distance of the node from its "mean" position. The mean position of the node can vary (at a low frequency) as the simulation proceeds, and a suitable filter period allows the oscillator model to account for this. For the Milan model, a filter period of 10-20 times the expected period of VIV is recommended.

For unsteady flow, the filter period for all models should be set to a value which is:

a) at least 10 times greater than the expected VIV period, and

b) at least 10 times smaller than the period of oscillatory flow

This may not be possible. If this is the case, choose the best compromise – mid-way between the two.

2.6 Damping

VIV can be sensitive to the amount of structural damping present in the line. OrcaFlex defines target damping values for axial, bending and torsional degrees of freedom: for VIV, the relevant parameter is bending damping. The relationship between the target damping value in the OrcaFlex data and the % critical damping at a particular response frequency is discussed in the OrcaFlex Help topic titled Line Theory: Structural Damping.

A small change in structural damping can trigger a significant change in the nature of the VIV response. Sensitivity analysis is strongly recommended.

3 Issues specific to individual VIV models

3.1 Wake oscillator models

For both MWO and IBWO models, the default values in OrcaFlex are the “best fit” values as defined in the published references. To facilitate further research and possible re-calibration of the models, these values can be changed by the user. However, for normal project use, we recommend keeping to the default values.

3.2 Vortex tracking (1)

The VT1 model was extensively calibrated by the original researchers and “best fit” values were assigned to the governing parameters. Unfortunately, our implementation of the VT1 model does not give the same results as reported by the originators for reasons which remain unclear despite detailed investigation. See Section 5.1 below.

Four parameters are available for modification by the user:

> ***Smear factor***
>
> *This setting controls the degree to which vortices are smeared. The original Sarpkaya and Shoaff vortex tracking model used point vortices, i.e. the vorticity was concentrated at a single point. This is what arises in the formal mathematical solution of the inviscid Navier Stokes equations (i.e. those ignoring fluid viscosity) but it means that each vortex is a singularity, since the vorticity density at the point itself is actually infinity. In reality viscosity in the fluid spreads the vorticity to some extent, and we have found that the model is more stable if the vortices are smeared to some extent.*
>
> ***Vortex decay constant and thresholds (3 parameters)***
>
> *These data items set the rate of vortex strength decay in both vortex tracking models. Vortices are created at the separation points, with an initial vortex strength determined by the tangential velocity at the separation point. The strength of each vortex then decays at a rate that depends on how far the vortex is away from the centre of the disc, in the relative flow direction.*

Further details are available in the documentation. Our interim recommendation pending re-calibration is to use default values for all four parameters.

3.3 Vortex tracking (2)

The original Sarpkaya and Shoaff vortex tracking model contains a number of special features which work well when the line is subjected to steady flow leading to the formation of a classic vortex street, but less well under more confused conditions. The VT2 model was developed to deal with this situation. The Help file says:

The VT2 model shares many features of the VT1 model but differs in the following ways:

- *The VT2 model uses a constant time step, equal to the simulation outer time step. (The VT1 model uses a variable time step.)*
- *The VT2 model does not use sheet detachment to separate the vortices coming from one side of the disk into separate sheets.*
- *The VT2 model does not use entrainment, nor rediscretisation to keep the vortices at equally spaced arc lengths along the sheet.*
- *The VT2 model uses a coalescing algorithm to control the number of vortices being tracked.*

The VT2 model has the following user-adjustable coefficients:

> ***Smear factor***
>
> ***Vortex decay constant and thresholds (3 parameters)***
>
> *All as for VT1 model – see previous section.*
>
> ***Creation clearance***
>
> *Determines the distance from the cylinder surface at which new vortices are created.*
>
> ***Coalesce same; Coalesce opposite***
>
> *Determine the readiness with which vortices coalesce. Larger numbers mean greater readiness to coalesce, hence fewer active vortices and shorter run time.*

Further details are available in the documentation.

Experience suggests that small changes in smear factor and creation clearance have opposite effects: provided the two values are kept the same there seems to be little net change in results.

Recent work suggests that, in some cases, the coalescence thresholds can be increased considerably, with little effect on results and great benefit in reduced run time. The run time for an VT2 model is dominated by the VIV calculation, and the time for the VIV computation at each node increases as the square of the number of vortices present. Increasing the coalescence threshold values means that vortices are more readily amalgamated which reduces the total number present and speeds the computation. Recent experience with a 100 element model showed that setting the threshold as high as 0.15 had little effect on the VIV response and reduced run time by two orders of magnitude.

We recommend leaving the smear factor, decay parameters and creation clearance at their default values, and carrying out sensitivity analysis to determine the highest coalescence values which can be used without adversely affecting results. Separate threshold values can be used for coalescence of same sign and opposite sign vortices. Our experience suggests there is little to be gained by using different values.

4 Selection of an appropriate VIV model

When considering which VIV model to use, the following questions are relevant:

- Is the flow steady or time-varying? Steady flow usually means steady current or tow at constant speed. Time-varying may mean waves or long period oscillatory movement of the line.
- Is there any likelihood of reversing flow, where the line moves to and fro in its own wake? This can occur in waves, or where part of a line is force-oscillated in still water, either by an external excitation or as a result of VIV on a distant part of the same line.
- How similar are the flow conditions at different points on the line? If conditions are similar at all points, then single mode response is likely. If not then multi-mode response may occur.
- Is in-line VIV expected to be an important issue?

The answers form the basis for choosing a suitable model.

The **Wake Oscillator** models have the following characteristics:

- Quick to run – VIV analysis using the WO models is only slightly slower than analysis without VIV modelling.
- VIV can be slow to start up; a small initial disturbance sometimes helps to get the oscillation going. Recent changes have improved performance in this respect, but start-up is still slower than that seen in physical tests.
- Once oscillation has started, the models appear to respond realistically to changes in inflow, but there is nothing in the governing mathematics to arrange this, so if rates of change are important, then WO results should be viewed with reserve.
- Steady flow cases frequently (but not always) settle to a single mode response with constant amplitude, even in highly sheared flow conditions where multi-mode response may be expected.
- Blind trials have shown that the wake oscillator models are generally at least as accurate as other methods for prediction of cross-flow VIV.

The **Vortex Tracking** models have the following characteristics:

- Much slower to run – typical run time may be an order of magnitude longer than for a wake oscillator.
- VIV starts quite quickly, and responds quickly to changing inflow. Since the vortex wake is modelled explicitly, there is reason to expect that transient response might be fairly well represented.
- Even in steady flow conditions, a vortex tracking model rarely settles to a fully steady state with constant amplitudes; there is always some “beating” or other variability present. This is consistent with model tests and other real world observations.
- VT models show both in-line and cross-flow VIV.
- Trials suggest that VT models can overstate VIV amplitude by 50-100%.

The **VT1** model:

- is usually faster to run than the VT2 model.
- has special features which are appropriate in flow conditions which lead to the formation of a classic von Karman vortex street (e.g. a fixed cylinder in steady flow) but can cause trouble where the flow is more confused.
- requires further calibration (in progress).

The **VT2** model:

- is more general than the VT1 model, and has a larger number of user-adjustable parameters.
- requires further calibration (in progress).

Reversing flow causes special difficulty. Typical examples are cases where part of a line is forced to oscillate in still water, either by VIV on another part of the line, or by external excitation (e.g. touchdown region of a deep water SCR connected to a platform executing heave oscillations).

The wake oscillator models and the VT2 model can cope with this scenario, but with some reservations:

- In principle, the VT2 model should work, but since there is no mean flow to remove old vorticity we rely on decay and coalescence.
- The wake oscillators were calibrated for steady flow, so in principle they should not apply to reversing flow. However experience shows that they can produce plausible results. This may be fortuitous, or there may be a more fundamental reason, perhaps related to decay of old vorticity.
- With reversing flow the VT1 model often fails after a time. We believe this is because the special features of the algorithm (cutting the vortex sheets; defining a wake line and absorbing vorticity inside it) become ill-defined in reversing flow, and the response becomes increasingly chaotic.

The best modelling options for reversing flow are the wake oscillators and the VT2 model. See Section 5.2.1 below for a specific example.

5 Validation

5.1 Current status

The **Wake Oscillator** models both reproduce published models and in both cases we have confirmed that the OrcaFlex model replicates the published behaviour. The validation evidence presented in the original published work therefore applies to the OrcaFlex models.

It was our intention that the same should be true for the **VT1** model, but as noted above, this has not proved to be the case. Specifically, the original Sarpkaya and Shoaff paper gives lift and drag forces on a fixed disk which we have been unable to replicate. Our implementation of the VT1 model differs in a number of respects from the original. We have checked the effect of each such modification and found only minor changes in results, insufficient to explain the overall differences we observe.

The **VT2** model is based on the VT1 model but includes important changes as stated in the documentation.

5.2 Trials

We were involved in two "blind trial" exercises in which a variety of VIV codes were used to predict measured responses. A number of validation documents for the VIV Toolbox have been generated which report the results of these trials in detail. What follows is a preliminary assessment.

5.2.1 Model Tests at Delft (Chaplin)

Details of the model tests and predictions for all participating software packages were presented at a workshop in Trondheim on 25-26 October 2004. You can see a copy of this presentation at the following link: Chaplin-VIV-Trondheim-2004.pdf

A top tensioned riser was exposed to a uniform flow over the lower half of its length whilst the upper half was in still water. Tests were carried out for a range of flow speeds and curvature was measured at 32 locations. For each test, a chart shows in-line and cross-flow displacement and curvature distributions along the riser as measured and as predicted by 11 software packages including the MWO and VT2 models. Other packages are 4 CFD models and 5 frequency domain models.

The results show that the MWO model predicts cross-flow response with good accuracy throughout the speed range, and is at least as accurate as any other method.

The VT2 model over-predicts both cross-flow and in-line response typically by 50-100%, but gives a good qualitative picture of in-line response. The only other models giving any indication of in-line response are the full CFD models which take an order of magnitude longer to run.

It is worth noting here that the upper part of the riser was oscillating in still water, driven by energy input from the lower part. The upper part would therefore have experienced continuous "wash back" with the riser moving to and fro in its own wake. The success of both MWO and VT2 models in handling this difficult problem is particularly encouraging.

5.2.2 Schiehallion Drilling Riser (BP)

A top tensioned drilling riser was instrumented by BP at the Schiehallion field and results were used for a blind trial as part of the Norwegian Deepwater Programme (NDP). Accelerations in two directions normal to the riser were measured at five locations near the bottom end over an extended period. Current profiles were also measured. The measurements have not been published and remain confidential to BP/NDP.

Participating software packages included three CFD codes and the four time domain models from the OrcaFlex VIV Toolbox. Overall, the results were similar to the Chaplin comparisons: the two wake oscillator models behaved well and gave predictions which were as accurate as or more accurate than the CFD packages whilst the two vortex tracking models overestimated amplitudes by 50-100%.

R M Isherwood

1 December 2004

Revised 28 October 2005

Revised 22 August 2005

Revised 24 April 2019