

Modeling the Effects of Reynolds Number and Added Mass on the Vortex-Induced Vibration of Drilling Risers

by
Chadwyck T. Musser

S.B., Mechanical Engineering, Massachusetts Institute of Technology, 1997
S.B., Music, Massachusetts Institute of Technology, 1997

Submitted to the Department of Ocean Engineering
And the Department of Mechanical Engineering
In Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Ocean Engineering
and
Master of Science in Mechanical Engineering

At the MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2000

© 2000 Massachusetts Institute of Technology. All rights reserved.


Author **Signature Redacted**
Department of Ocean Engineering
May 24, 2000

Certified by **Signature Redacted**
J. Kim Vandiver
Professor of Ocean Engineering
Thesis Supervisor

Certified by **Signature Redacted**
Samir Nayfeh
Professor of Mechanical Engineering
Thesis Reader

Accepted by **Signature Redacted**
Nicholas Patrikalakis
Kawasaki Professor of Engineering
Chairman, Departmental Committee on Graduate Students
Department of Ocean Engineering

Accepted by **Signature Redacted**
Ain A. Sonin
Chairman, Committee on Graduate Students
Department of Mechanical Engineering



ARCHIVES

Modeling the Effects of Reynolds Number and Added Mass on the Vortex-Induced Vibration of Drilling Risers

by
Chadwyck T. Musser

Submitted to the Department of Ocean Engineering
And the Department of Mechanical Engineering
In Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Ocean Engineering
and
Master of Science in Mechanical Engineering

Abstract

This thesis addresses several important issues to be taken into account when modeling the vortex-induced vibrations on drilling risers. Modeling using modal analysis is compared to modeling using frequency-response functions to justify that modal analysis is an acceptable analysis method for these systems. The response predicted using both methods is shown to be very similar.

The importance of using a Strouhal number that is a function of local Reynolds number is demonstrated using simulator examples compared to measured data. An optimized Strouhal number curve as a function of Reynolds number is obtained by taking permutations of a typical expected curve and comparing prediction results with the permutations of the curve to the measured data. The optimal curve and riser response predictions using this optimal curve are shown.

The issue of power-in overlap is presented. Various methods for assigning conflicting power-in regions to the appropriate mode are discussed. A simplified model to predict single-mode dominance to be used as a pre-simulation tool is introduced and is compared with full simulation results.

The variability of the riser natural frequencies and the change in added mass that causes this are discussed. An iterative solution is presented to predict the variable natural frequencies from a recursive model for assigning added mass based on reduced velocity regions. Comparisons between the natural frequency predictions using this model and the actual measured natural frequencies are shown.

The possible methods in which lift coefficients may be assigned along the riser power-in locations are discussed. Comparisons are presented of response predictions to each other and to measured responses for various cases. Recommendations are made as to which lift coefficient assignment methods work best.

Thesis Supervisor: J. Kim Vandiver
Title: Professor of Ocean Engineering

Thesis Reader: Samir Nayfeh
Title: Professor of Mechanical Engineering

Acknowledgements

I would like to thank Professor J. Kim Vandiver, my thesis advisor, for his encouragement and support during my graduate program. His knowledge and enthusiasm for this subject has made working on this thesis a very enjoyable and educational experience.

I would also like to thank my parents, Ron and Netch Musser, my first teachers. Their guidance and constant support have made it possible for me to always follow my dreams. I was also inspired by my brothers, Eli and Nathan, the rest of my family and friends, and particularly my grandfather, Bob Gagich, the best engineer I ever knew.

This work was sponsored by an industry consortium and the Office of Naval Research. The industry sponsors included AGIP, Aker Engineering, ARCO, BP / Amoco, Chevron, Exxon Production Research, Marintek, Mobil, Norsk Hydro, Petrobras, Shell, Statoil, and Texaco.

Contents

1 Introduction	pg. 6
2 Comparison of Modal Analysis to Frequency Response Function Analysis	pg. 16
3 Reynolds-Number-Dependency of Strouhal Numbers	pg. 29
4 Resolving Conflicts Between Overlapping Power-In Regions	pg. 42
5 Convergence of Natural Frequencies and Added Mass Coefficients	pg. 49
6 Comparison of Response-Dependent Lift Coefficient Models	pg. 57
7 Conclusion	pg. 64
Bibliography	pg. 67
Appendices	pg. 69
Appendix A: Riser Properties.....	pg. 69
Appendix B: Lift Coefficient Reynolds Number Modification Factor.....	pg. 70
Appendix C: Permutations of Strouhal Number Curve.....	pg. 71

List of Figures

1-1	Helland-Hansen Riser Set-up with Accelerometer Positions.....	8
1-2	Directionality of Current Profile for Typical Case – May 22 2200h.....	10
1-3	Current Profile in Principal Direction for Typical Case – May 22 2200h.....	11
1-4	Current Profile in Normal Direction – May 22 2200h.....	11
1-5	Total Current Profile for Typical Case – May 22 2200h.....	12
1-6	Direction of Current Profile for Unusual Case – May 3 0000h.....	13
1-7	Current Profile in Principal Direction for Unusual Case – May 3 0000h.....	14
1-8	Current Profile in Normal Direction – May 3 0000h.....	14
1-9	Total Current Profile for Unusual Case – May 3 0000h.....	15
2-1	Example of Damping Allocation Along the Riser for Typical Case.....	22
2-2	Raw Lift Coefficient as a Function of Response Amplitude.....	24
2-3	1 st -Mode-Dominant Case, Agreement Between FRF and Modal Analysis.....	25
2-4	2 nd -Mode-Dominant Case, Agreement Between FRF and Modal Analysis.....	27
2-5	Case with “Worst-Case” Difference Between FRF and Modal Analysis.....	28
3-1	Smooth Cylinder Strouhal Number as a Function of Reynolds Number.....	31
3-2	Rough Cylinder Strouhal Number as a Function of Reynolds Number.....	32
3-3	Medium Roughness Cylinder Strouhal Number as a Function of Re. Number.....	33
3-4	Basis Strouhal Number Curve from Achenbach Data.....	34
3-5	Optimized Strouhal Number Curve as a Function of Reynolds Number.....	35
3-6	First Mode Response Comparison Between Different St #-Re # Curves.....	36
3-7	Second Mode Response Comparison Between Different St #-Re # Curves.....	37
3-8	May 5 0200 – Illustrates Need for Reynolds-#-Dependent Strouhal Number.....	39
3-9	May 5 0200 – Power-In Locations from Simulator for Constant St # = 0.25.....	40
3-10	May 5 0200 – Power-In Locations for Optimized St #-Re # Curve.....	41
4-1	Comparison of $U^3 \cdot D^2$ Response Prediction to Full Prediction – June 5-7.....	45
4-2	Comparison of $U^3 \cdot D^2$ Response Prediction to Full Prediction – May 3-7.....	46
5-1	Measured Natural Frequencies from May 3-7.....	50
5-2	Initial Mode 2 Power-In Locations for Default Added Mass – May 5 0200h.....	52
5-3	Second Step Mode 2 Power-In Locations for Iterative Added Mass.....	53
5-4	Converged Mode 2 Power-In Locations for Iterative Added Mass.....	55
5-5	Comparison of Measured and Predicted Natural Frequencies – May 3-7.....	56
6-1	Lift Coefficient as a Function of Max. Modal A/D for Different Methods.....	58
6-2	Lift Coefficient for Large A/D Response for Different C_L Methods.....	60
6-3	Response Predictions for Large A/D Case for Different C_L Methods.....	61
6-4	Lift Coefficient for Midrange A/D for Different C_L Methods.....	62
6-5	Response Predictions for Midrange A/D for Different C_L Methods.....	63

Chapter 1

Introduction

Long and slender marine structures such as drilling risers are often subjected to a phenomenon known as vortex-induced vibration, or VIV. This behavior occurs when the flow around a cylindrical object excites the structure near one or more of its natural frequencies. If a significant portion of the structure is excited at or near one of its resonance frequencies, “lock-in,” or wake synchronization with a resonant frequency, may result. [1]

In structures such as drilling risers vortex-induced vibration has the potential to damage the structure and eventually lead to failure of the structure when the amplitude of response is high enough to cause significant fatigue damage. It is therefore important to design these structures taking into account the likely current profiles to which they will be exposed during their operative life and the resulting vortex-induced vibration and corresponding fatigue to which they will be subjected due to the combination of riser properties and current profiles. An accurate method of predicting this vortex-induced vibration response is then a very useful tool in the development of these structures.

Currently there are several commercially available vortex-induced vibration response predictors, but there are several important modeling issues which they do not currently fully address. This thesis attempts to delve more deeply into exploring such programming issues as modal analysis versus frequency-response function analysis (Chapter 2), power-in overlap conflict resolution (Chapter 4), and lift coefficient assignment (Chapter 6). In addition, armed with the recently-acquired VIV data from the Helland-Hansen riser measurements, the effects of Reynolds number as a function of

Strouhal number (Chapter 3) and Added Mass and Power-In Location (Chapter 5) on the response of the riser are also studied.

Before drawing direct comparisons between the simulation techniques described in this thesis and the Helland-Hansen measurements, it is worth displaying and explaining the actual set-up and the limitations of the data collected from the Helland-Hansen riser. [2] The riser was connected to a drilling vessel just offshore of Norway in an area with a water depth of 685 meters. The riser had a buoyant covering so that there was an effective diameter of 1.13m over nearly the entire length of the riser. The top 100m of the riser and several small 15m sections toward the bottom of the riser are the only parts of the riser without this buoyant covering, so that the local diameter in these sections is 0.53m. A more complete description of the Helland-Hansen riser physical properties may be found in Appendix A.

Helland-Hansen VIV Monitoring

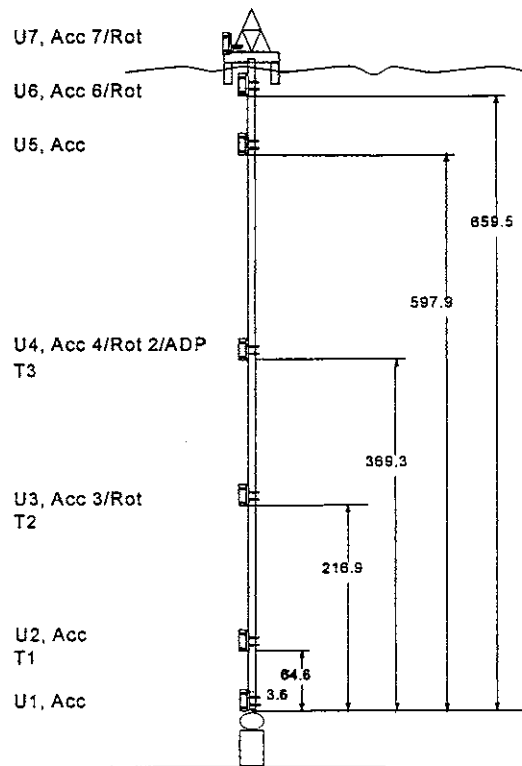


Figure 1-1: Helland-Hansen riser set-up with accelerometer positions

As shown in Fig. 1-1, there were six accelerometers located along the length of the riser. Three of the sensors, U3, U4, and U6, were capable of measuring rotation as well as transverse vibration. The ability to measure angular acceleration, and therefore angular rotation, was used in the processing of the measured data to try to remove the measured transverse acceleration caused by gravitational effects as the sensor rotated during VIV-induced vibration. [2] The measured riser response frequency and amplitude were extracted from these 6 transverse accelerometer and 3 angular accelerometer measurements. It is important to note that after the first month of testing, unit 1 failed, followed a few days later by unit 3. An accurate modal decomposition became difficult

for many of the higher-mode-response cases with only 4 sensors, and this had to be taken into account when examining measurements taken after May 21, 1998. The top tension measured for the drilling riser was also not constant, but was not measured after the initial set-up, so the measured value used in the simulation is most likely slightly different from the actual tension for some of the simulated cases.

The most important input parameter to a VIV-response simulator is a continuous and accurate current velocity profile. On the Helland-Hansen riser the current was measured by an acoustic doppler current profiler mounted in a vertical mooring near the drilling vessel and a rotor-type current meter to measure the current near the sea floor. [2] There was a gap in coverage above the bottom sensor. A linear interpolation was performed to obtain a continuous current velocity for use in the simulations. The measurement in the top 100 meters was inaccurate because of noise from the vessel thrusters, so the current profile in this region is taken to be a uniform flow profile at the current speed of the first accurate data point below 100 meters. Fortuitously, as will be shown in Chapter 4, this region has almost negligible impact in the computation of the VIV response regardless of velocity profile because of its relatively small compared to the rest of the riser. The current sensor at the ocean floor ceased functioning on June 10, 1998 before data acquisition had been completed. The last 23 days of data were collected without use of this sensor.

A final issue about the current profile is the directional variation of the current over the riser. Most VIV-response predictors assume that all current is flowing in the same direction, regardless of whether the current profile is a uniform flow or sheared flow. The research staff at Marintek evaluated the directional properties of the Helland-

Hansen data as follows. For each velocity measurement, a principal flow direction was established, based on the average flow direction over all of the velocity sensors. From this principal direction the current profile was decomposed into a principal direction and a normal direction. In many cases the variation from the established principal flow direction did not exceed more than 10 to 20 degrees for any given location along the riser. In such cases a uniform flow direction comes close to approximating the actual measured flow magnitude and direction.

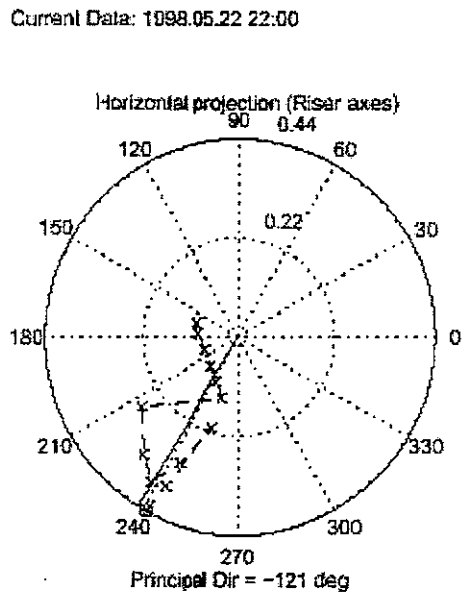


Figure 1-2: Directionality of Current Profile for Typical Case – May 22 2200h

Figure 1-2 displays the velocity profile magnitude and direction for a typical case in the measurement data set (May 22 2200h). [3] The polar plot in the figure simultaneously displays current magnitude and direction. This plot reveals that the

measurement points over the entire length of the riser all have a very similar directionality.

This is confirmed by observing the principal current profile in Figure 1-3 (velocity in the established principal direction indicated) and the normal current profile in Figure 1-4 (velocity component normal to the established principal direction). Over the length of the riser, the current magnitude normal to the established principal current direction is negligible compared to the current in the principal direction.





Figure 1-3: Current Profile in Principal Direction for Typical Case – May 22 2200h

Figure 1-4: Current Profile in Normal Direction - May 22 2200h

The total measured velocity magnitude for this case, shown in Figure 1-5, is very similar to the current profile shown as the current in the principal flow direction (Fig. 1-3). This similarity and the observation that most velocity profiles measured exhibited a directionality similar to this case provide a justification for using a single uniform flow direction over the length of the riser in the VIV-response simulator. To be conservative and encompass the highest flow velocity situation possible, the flow in this single principal direction was assigned the local magnitude of the total velocity profile.

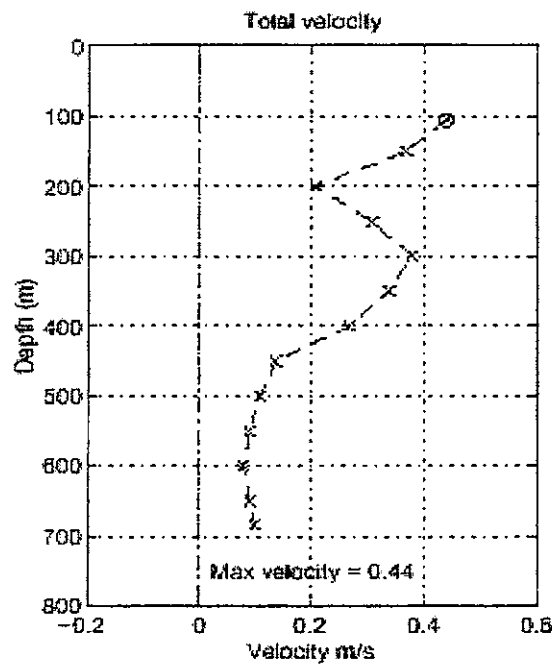


Figure 1-5: Total Current Profile for Typical Case – May 22 2200h

There were some cases in which there was more significant directional variation along the riser length. The uniform flow direction assumption is not completely valid for

such cases and negatively affects the ability of the simulation to predict the riser response. It is worth examining the same set of measurements for a case where the directionality is not as neatly uniform as in the majority of cases, similar to the one just presented in Figs. 1-2 – 1-5. Fig. 1-6 shows the polar plot for a case from the measured data (for date May 3, 0000 h) where the velocity profile for the lower portion of the riser is strongly aligned in the normal direction to the established principal current. Examining the velocity profiles in the established principal direction (Fig. 1-7) and normal direction (Fig. 1-8) shows that the current magnitude is accurately described near the top of the riser by the velocity in the principal direction but near the bottom is better depicted by the velocity in the established normal direction.




Figure 1-6: Directionality of Current Profile for Unusual Case – May 3 0000h




Figure 1-7: Current Profile in Principal Direction for Unusual Case – May 3 0000h




Figure 1-8: Current Profile in Normal Direction - May 3 0000h

Even in such a case as this, using the total velocity profile and assuming it acts in a uniform flow direction provides a reasonable and conservative answer, even though it may not come as close to modeling the physical reality of the system as it does in the more mono-directional cases such as those in Fig. 1-2 – 1-5. The power-in zone for a particular mode, which is the portion of current velocity that is exciting a particular mode of the riser, often acts over a short length of the riser, so that the change of direction of

the current profile within a certain power-in region is rarely very large, even in a case where the current variation over the entire riser is large. Such a simulation is conservative because a single uniform direction for the velocity always predicts a maximal VIV-response because the entire excitation acts on the in the same plane. The total velocity magnitude over the length of the riser for the multi-directional current profile is shown in Fig. 1-9.




Figure 1-9: Total Current Profile for Unusual Case – May 3 0000h

Chapter 2

Comparison of Modal Analysis to Frequency Response Function Analysis

In cases of vibrational analysis in which many modes may be excited but only a few critical modes are of interest, modal analysis is an effective way of decreasing the computation time with very little sacrifice in precision. This is in contrast to predicting response from frequency-response functions, which predict the response with no compromise in accuracy, but requires more computation time. Consequently, for simulations in which a drilling riser is modeled as having a very high number of degrees of freedom but has only a few critical modes of interest that may actually affect structural fatigue life, modal analysis presents itself as the more appealing solution as long as the assumptions key to an accurate modal analysis are valid. The basic assumptions and advantages for using normal modal superposition are explained in great detail in Meirovitch. [4]

One of the primary goals of this thesis is to compare VIV-response predictions using both modal analysis and frequency-response function (FRF) analysis. The objective was to determine if the modal analysis approach could yield nearly identical results to the frequency-response function approach. In the vast majority of cases analyzed, the two methods did in fact yield nearly identical results.

The frequency-response function component of the VIV-response simulator was based on calculating the steady-state harmonic response of the riser to an assumed steady and harmonic input at each location along the riser. The riser was discretized into approximately 100 evenly-spaced sections, each of which was assigned lumped-

parameter values for mass, damping, and stiffness, as will be discussed. This discretized riser is then described by the equation of motion, as discussed in Rao [5]:

$$[m] \{\ddot{x}(t)\} + [r] \{\dot{x}(t)\} + [k] \{x(t)\} = \{F(t)\} \quad (2-1)$$

where $[m]$ is the total mass matrix, $[r]$ is the total damping matrix, $[k]$ is the stiffness matrix, $\{x(t)\}$ and its derivatives are the displacement, velocity, and acceleration at the discretized points along the riser, and $\{F(t)\}$ is the dynamic excitation due to vortex shedding and applied at the discretized points.

For an assumed steady harmonic input and output,

$$\{F(t)\} = \{F\} e^{i\omega t} \quad (2-2)$$

$$\{x(t)\} = \{x\} e^{i\omega t} \quad (2-3)$$

Both the excitation force and the displacement of the response consist of a constant magnitude that oscillates harmonically with time. When this assumption may be made, then:

$$\{\dot{x}(t)\} = i\omega \{x\} e^{i\omega t} \quad (2-4)$$

$$\{\ddot{x}(t)\} = -\omega^2 \{x\} e^{i\omega t} \quad (2-5)$$

and so Eq. 2-1, using Eq. 2-2 to Eq. 2-5, may be expressed as:

$$(-\omega^2 [m] + i\omega[r] + [k]) * \{x\} = [Z] * \{x\} = \{F\} \quad (2-6)$$

so that the mass, damping, and stiffness matrices for any given excitation frequency ω may be represented as an impedance matrix $[Z]$. The total riser displacement at the discretized locations can then be found by inverting this impedance matrix $[Z]$ and multiplying by the force magnitude vector $\{F\}$:

$$\{x\} = [Z]^{-1} * \{F\} \quad (2-7)$$

This is a fairly straightforward procedure for systems without nonlinearities, but a drilling riser in water exhibits some characteristics that are nonlinear in nature and require special consideration before the above equations can be used to compute response for a given excitation. These considerations specific to drilling risers under hydrodynamic forces are discussed in detail within the SHEAR7 theory guide by Vandiver and Li. [6] Since the riser is in water, the diagonal mass matrix $[m]$ is equivalent to:

$$[m] = [m_{air} + am(\text{mod } e, V_r)] \quad (2-8)$$

where m_{air} is the mass of each discretized riser segment in air including contents and am is the external added mass. The added mass is an additional amount of mass added to each segment to take into account that the riser is in water and each segment has not only its own inertia to overcome when vibrating, but must also accelerate an amount of the

water in its immediate vicinity. [1] As described in Eq. 2-8, this added mass is not constant, but is rather a function of mode number and reduced velocity. One particular location along the riser may then exhibit different values of added mass depending on the mode being excited, which in turn changes the natural frequency of the mode. An iterative procedure to address this phenomenon is the focus of Chapter 5.

The damping matrix $[r]$ was modeled as having two components:

$$[r] = [r_{st} + r_{hy}(\text{mode}, V_r, A/D)] \quad (2-9)$$

r_{st} is a constant structural damping and r_{hy} is the hydrodynamic damping, which, like added mass, is a function of mode number and reduced velocity V_r , but is also a function of response amplitude. The structural damping was modeled as corresponding to a damping ratio of $\xi = 0.003$, where:

$$r_{st} = 2\xi m \omega \quad (2-10)$$

The structural damping was negligible compared to the hydrodynamic damping. Within a modal power-in region the hydrodynamic damping is accounted for within the lift coefficient and so is not separately included in the damping matrix. Within the modal power-in region the structural damping is the only damping component in the total damping matrix $[r]$.

The hydrodynamic damping is based on a model developed by Madan Venugopal [7] and confirmed by Vikestad, Larsen, and Vandiver [8] and used in the SHEAR7 VIV-response prediction programs. Hydrodynamic damping in the power-in region for a particular mode is accounted for in the assignment of the lift coefficient. Outside of the

modal power-in zone, the local hydrodynamic damping is specified by one of two formulas, depending on whether the local reduced velocity is greater than the reduced velocity in the power-in region or less than the reduced velocity in the power-in region.

The hydrodynamic damping for a low-reduced velocity region is:

$$r_h(z) = r_{sw} + C_{rl} \rho D V \quad (2-11)$$

where $r_h(z)$ is the local hydrodynamic damping, C_{rl} is an empirical constant experimentally evaluated to be 0.18, ρ is the fluid density (=1027 kg/m³ for salt water at 5 °C), D is the riser diameter, V is the local current velocity, and r_{sw} is a still water damping term defined as:

$$r_{sw} = \frac{\omega \pi \rho D^2}{2} \left[\frac{2\sqrt{2}}{\sqrt{\text{Re}_\omega}} + 0.2 \left(\frac{A}{D} \right)^2 \right] \quad (2-12)$$

where ω is the vibrational frequency in radians per second, A is the amplitude of vibration response, and Re_ω is a vibration Reynolds number defined as:

$$\text{Re}_\omega = \omega D^2 / \nu \quad (2-13)$$

Where ν is the kinetic viscosity, as mentioned earlier (equal to $1.56 \times 10^{-6} \text{ m}^2\text{s}^{-1}$ for salt water at 5 °C).

There are several key things to note about this low-reduced velocity hydrodynamic damping, which is directly proportional to the local current velocity and local riser diameter for most cases. The still water damping term (in Eq. 2-11 and 1-12) r_{sw} is response-dependent, necessitating an iterative method for calculating response and

hydrodynamic damping. For cases in which the local current velocity is sufficiently large, the second term of Eq. 2-11, which has no dependence on response amplitude A , will be the dominant term of the low-reduced velocity hydrodynamic damping. This leads to the advantageous result that the response-dependent portion of the hydrodynamic damping in the low-reduced velocity region is negligible for all cases where the velocity is not close to zero.

The high-reduced velocity hydrodynamic damping model is simpler and is described by:

$$r_h(z) = C_{rh} \rho V^2 / \omega \quad (2-14)$$

where C_{rh} is another empirical constant empirically evaluated as 0.2. The hydrodynamic damping in the high-reduced velocity region is then completely predictable without knowing the amplitude of riser response for a given mode. The damping in the high-reduced velocity region is proportional to velocity squared but is also divided by modal frequency (in Eq. 2-14) so that near the modal power in region the high-reduced velocity damping is fairly similar to the low-reduced velocity damping in magnitude.

The total damping in the system, as represented in the damping matrix, is then accounted for by a summation of these structural and hydrodynamic damping contributions. A typical case of total damping along the riser is represented by a constant structural damping with the addition of a response-dependent low-reduced velocity hydrodynamic damping near the bottom of the riser (where the current is relatively slow) and the addition of a high-reduced velocity hydrodynamic damping component above the modal power-in zone. Fig. 2-1 displays a MATLAB calculation of the local total

damping along the riser for a typical velocity profile measured (May 6 at 0000 h). In this example the response amplitude is small and has little effect on the low-reduced velocity damping:

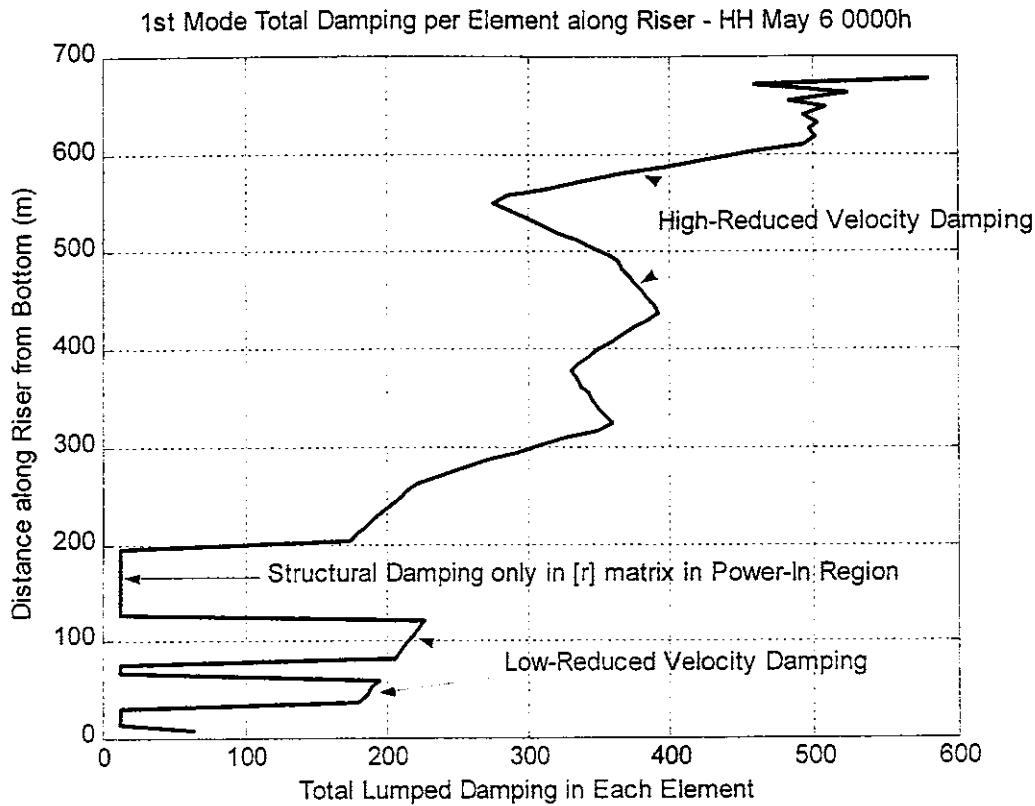


Figure 2-1: Example of 1st Mode Damping Allocation along Riser for a Typical Case

The stiffness matrix is the least complicated component of the impedance matrix since it does not exhibit dependence on mode number or response. For a tension-dominated drilling riser the bending stiffness is negligible compared to tension effects, and the lumped-parameter stiffness for each element may be modeled as the average tension in the section divided by the discretized section length. The tension linearly decreases from the top to the bottom of the riser, so the stiffness of each element also

decreases in the same manner. Since each element is connected only to the elements directly above and beneath it, the total stiffness matrix $[k]$ used in Eq. 2-1 and 2-6 is tridiagonal.

Both the magnitude and frequency of the excitation force vector are directly obtained as functions of the current velocity profile. The relationship between excitation frequency, current velocity, riser diameter, and Strouhal number (Eq. 2-15) yields the excitation frequency at a discrete location along the riser:

$$f = \frac{St * V}{D} \quad (2-15)$$

The excitation depends on a lift coefficient found by using a lookup table derived from Gopalkrishnan's extensive measured data on forced-motion cylinders. [9] This lift coefficient accounts for hydrodynamic damping in the modal power-in region. The total lift force at each discretized element on the riser is found from:

$$P(x, t) = \frac{1}{2} \rho D V^2(x) C_L(V_r(x), A / D) \sin(\omega_r t) \quad (2-16)$$

where P is the lift force, ρ is fluid density, D is riser diameter, V is current velocity, and C_L is the lift coefficient, determined by:

$$C_L(V_r(x), A / D) = \gamma_L C_{L0} \quad (2-17)$$

where γ_L is a modification factor based on local Reynolds number (see App. B) and C_{L0} is found directly from a lookup table for raw lift coefficient as a function of response amplitude and reduced velocity as derived from the experimental data of Gopalkrishnan.

[9] Fig. 2-2 shows the dependence of C_{L0} on the dimensionless response ratio A/D :




Figure 2-2: Raw Lift Coefficient as a Function of Response Amplitude

A current topic of research is whether the A/D used to assign lift coefficient should be based on the maximum modal A/D or a local A/D at the discretized points along the riser. For small response amplitudes the average lift coefficient obtained is

similar for both cases, but for certain values of A/D the average lift coefficient calculated can be quite different between these two methods. There is also a trade-off between modeling the physical system and facilitating numerical convergence of the response solution. This is the focal point of discussion in Chapter 6 of this thesis, where a more detailed discussion and some results are presented.

The lift coefficients obtained are used within the power-in regions for each mode, with the regions outside the modal power-in zones having hydrodynamic damping and no excitation force input. The local excitation force has the magnitude calculated from Eq. 2-16 with the lift coefficient from the lookup table in Fig. 2-2. The sign of the force is always assigned to be the same as the sign of the calculated mode shape at this location.

The following figures, Fig. 2-3 to 2-5, illustrate the difference in the VIV-response prediction when using the FRF analysis method or modal analysis. An estimate of the measured modal amplitudes is plotted as well, to show how the predictions for both methods compared to the actual measured responses. The predictions shown all use the optimized Strouhal number curve to be discussed in Chapter 3.

Comparison of FRF Simulator Results to Modal Analysis Simulator Results - May 4 0600h

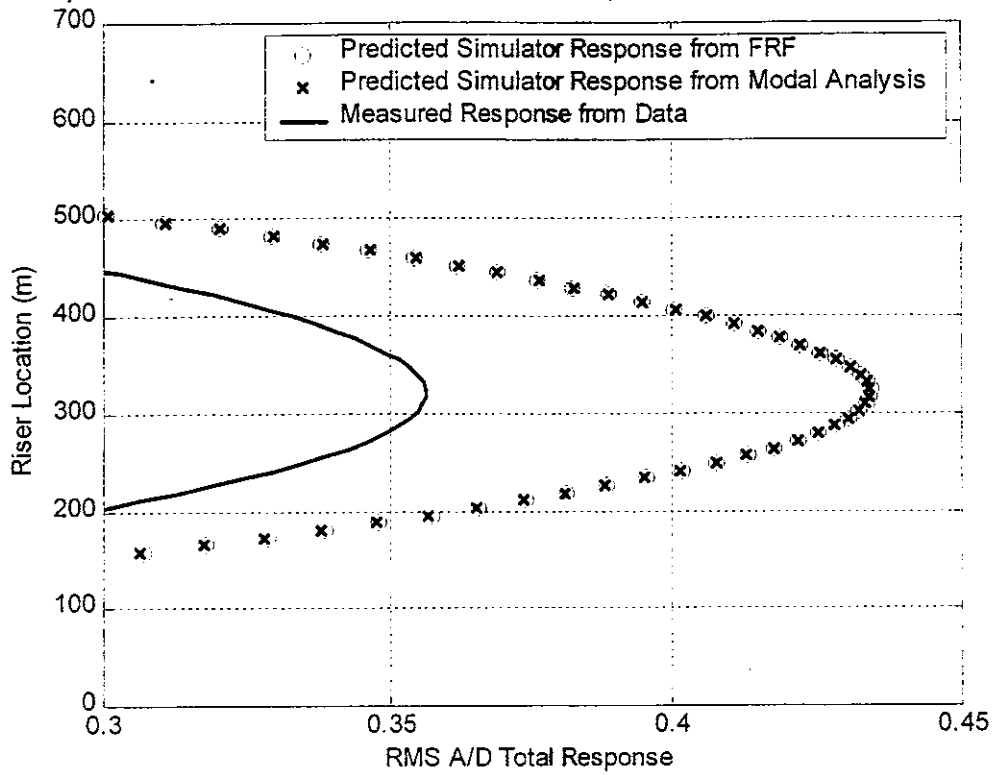


Figure 2-3: 1st-Mode-Dominant Case, Agreement Between FRF and Modal Analysis

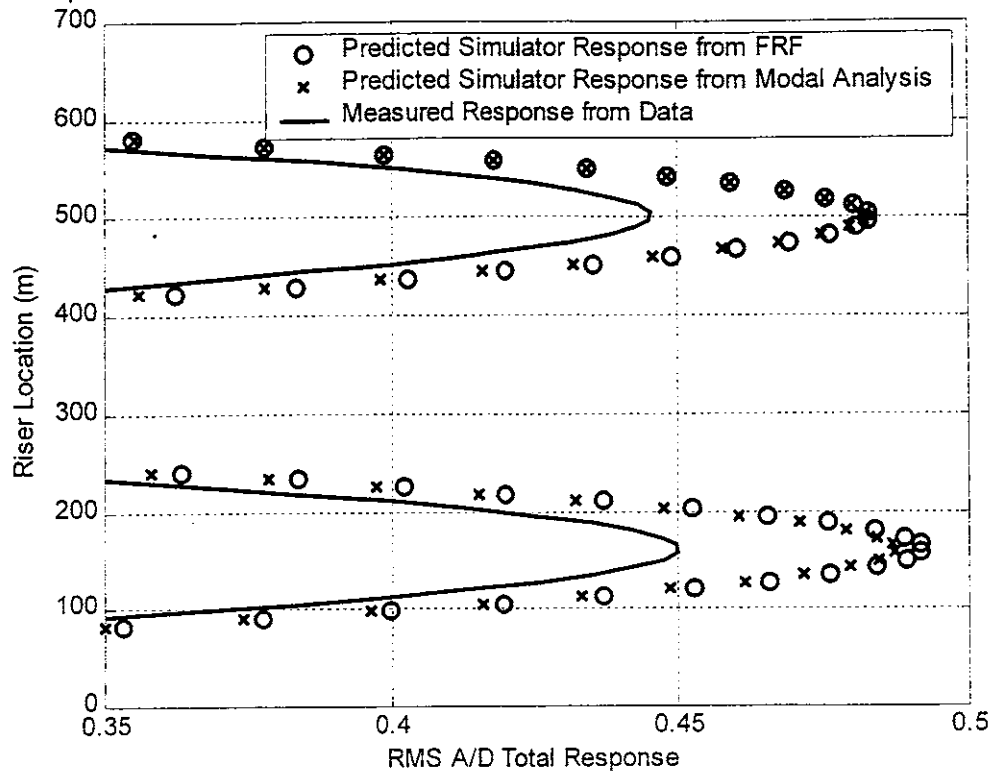


Figure 2-4: 2st-Mode-Dominant Case, Agreement Between FRF and Modal Analysis

Comparison of FRF Simulator Results to Modal Analysis Simulator Results - May 23 0800h




Figure 2-5: Case w/ “Worst-Case” Difference Between FRF and Modal Analysis

Of the hundreds of simulations carried out as research for this thesis, there was never a case in which the predicted maximum response from the simulator deviated more than about 5 percent between the FRF analysis and modal analysis. Fig. 2-5 was the case with maximum deviation between the two methods. The net result of this comparison is that a modal analysis approach to VIV-response prediction is as valid as a FRF analysis. The advantage to using modal analysis is that less calculation time is required. Even though the assumptions necessary for use of modal analysis are not completely satisfied, the results are acceptable.

Chapter 3

Reynolds-Number-Dependency of Strouhal Numbers

The relationship between flow velocity and excitation frequency is governed by the Strouhal number: [1]

$$f = \frac{St * V}{D} \quad (3-1)$$

where f is the local excitation frequency, St is the local Strouhal number, V is the local current velocity, and D is the local riser diameter. A variable Strouhal number implies an influence on the excitation frequencies of the drilling riser. The flow velocities corresponding to lock-in for a particular mode may be relatively higher or lower compared to neighboring modes depending on the local Strouhal number. The change in vortex-shedding frequency in the transcritical Reynolds number region is very significant.

In calculating the VIV-response of a riser many simulations make use of a single constant Strouhal number over the entire riser length. However, there are many cases in which the current velocity profile may simultaneously be in subcritical, transcritical, and / or supercritical Reynolds number flow regimes at different locations in the current profile. For cases in which the Strouhal number has significant variation according to Reynolds number, a single constant Strouhal number modeling technique may be inadequate.

Significant literature exists with measurements of Strouhal number as a function of Reynolds number for stationary cylinders, one particular example being some

experimental results from Achenbach and Heinecke. [10] Although this measured data does not exactly match the testing parameters of a sheared current profile in water on a moving drilling riser, this data is nevertheless useful as a baseline indicator of the range of Strouhal number values to expect for a given local Reynolds number. The Achenbach data is measured in air with a uniform velocity profile for a stationary cylinder with a constant diameter, but still provides insights as to the types of curves to be applied in the simulations explored in this thesis. Drilling risers are in water with a sheared velocity profile. Risers are not stationary and have variable diameter. Reynolds number is defined as:

$$\text{Re} = \frac{U \times D}{\nu} \quad (3 - 2)$$

where Re = Reynolds number, U is the current velocity, D is the riser diameter, and ν is the kinematic viscosity, assumed to be $1.56 \times 10^{-6} \text{ m}^2\text{s}^{-1}$ for salt water at 5 degrees Celsius. Assuming constant kinematic viscosity, Reynolds number varies locally along the riser as a function of local current velocity and riser diameter.

The largest variation in Strouhal number is observed for the smoothest cylinders, in which surface roughness is very small compared to the overall diameter of the cylinder. For such a case the Strouhal number may reach values as high as 0.50 in the transcritical Reynolds number flow regime. An example for a very smooth cylinder with negligible surface roughness is described in the Achenbach paper: [10]




Figure 3-1: Smooth Cylinder Strouhal Number as a Function of Reynolds Number

In this case the transcritical flow regime includes a Reynolds number range of about 3×10^5 to 3×10^6 . The transcritical region is indicated by high values of Strouhal number.

On the opposite end of the spectrum, a cylinder with a very high surface roughness exhibits almost no transcritical region at all; the Strouhal number in the subcritical region abruptly jumps to a slightly higher value typical of the supercritical flow region. An example from Achenbach:




Figure 3-2: Rough Cylinder Strouhal # as Function of Reynolds # – $k_s/D = 30 \times 10^{-4}$

where k_s is the surface roughness of the cylinder and D is the cylinder diameter. Note that the Strouhal number remains around 0.21 until a Reynolds number of about 1×10^5 , after which it attains a supercritical value of about 0.25.

From these two examples at the smooth and rough extremes, as well as measured data for cylinders of intermediate roughness, it is observed that the width of the transcritical region, the minimum Reynolds number corresponding to the transcritical region, and the peak Strouhal number within the transcritical region are all related to the surface roughness of the cylinder. The smoothest cylinders will have the largest transcritical regions, the highest minimum Reynolds number corresponding to the transcritical region, and the largest peak Strouhal number values. Very rough cylinders have practically no transcritical region; the transition from subcritical to supercritical flow regime occurs at a much lower value of Reynolds number than for smooth cylinders,

and the peak Strouhal number is relatively small and not very much greater than the subcritical value.

By comparing the Helland-Hansen riser data to simulation predictions, the Strouhal number to Reynolds number relationship that best matched the measured data fell between the smooth and rough curves. An intermediate roughness plot from Achenbach is shown for a surface roughness to diameter ratio of 75×10^{-5} :




Figure 3-3: Medium Roughness Cylinder St # as Function of Re # - $k_s/D = 7.5 \times 10^{-4}$

Although the data points within the transcritical region are somewhat scattered, most points follow one of three curves. The grouping of points near 0.50 closely resembles the curve for a very smooth cylinder (Fig. 3-1), the curve with the points slightly above 0.20 in the transcritical region resembles the curve for a very rough cylinder (Fig. 3-2), and the intermediate curve between the other two curves that has a peak Strouhal number value

of roughly 0.31 at Reynolds number $Re = 2.2 \times 10^5$. The Helland-Hansen riser response is predicted reasonably well by an intermediate curve as shown below. This curve was used as a starting point for the purpose of finding a more optimal one.

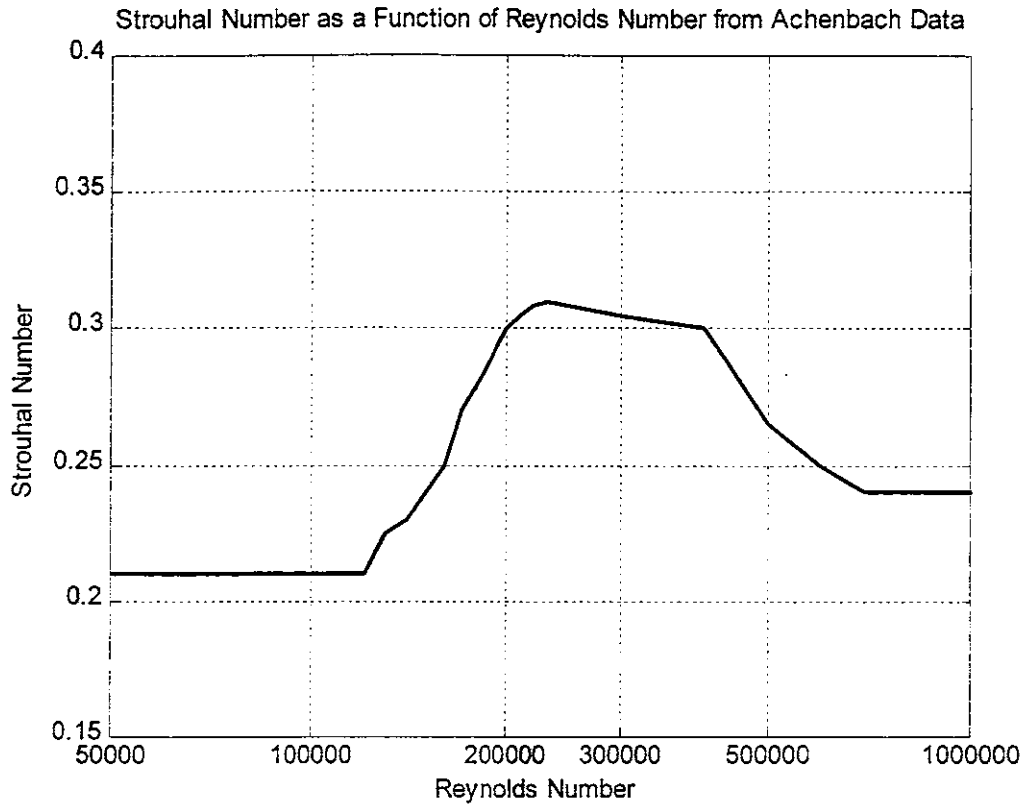


Figure 3-4: Basis Strouhal Number Curve Adapted from Achenbach Data

The VIV-response simulation program was run for this curve and for other curves with small permutations in the starting point, ending point, and point with maximum Strouhal number in the transcritical region. Two uninterrupted sequences of measurements, May 3-7 and June 5-7, served as particularly good measured data to compare with the simulator results because higher modes were present over the span of these days. Several variations on the “basis” Strouhal number vs. Reynolds number

curve were tried. Eventually a best-fit curve emerged, to be referred to as the “optimized” curve. This curve is shown in Fig. 3-5:

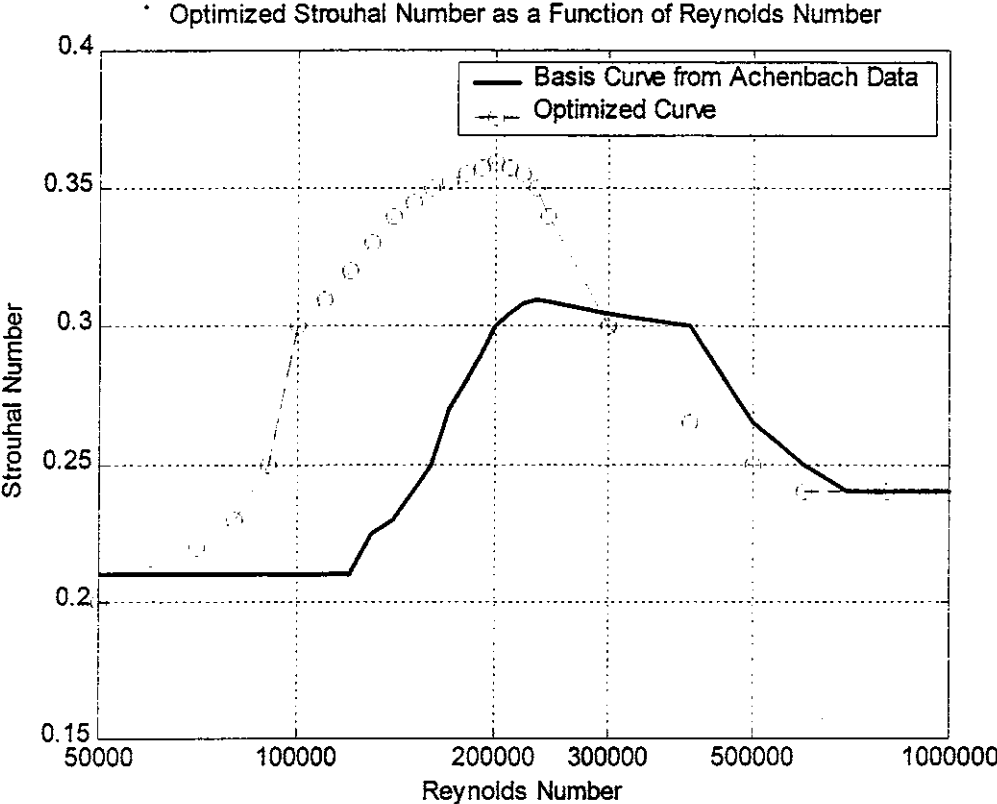


Figure 3-5: Optimized Strouhal Number Curve as a Function of Reynolds Number




Figure 3-6: First Mode Response Comparison Between Different St-Re Curves

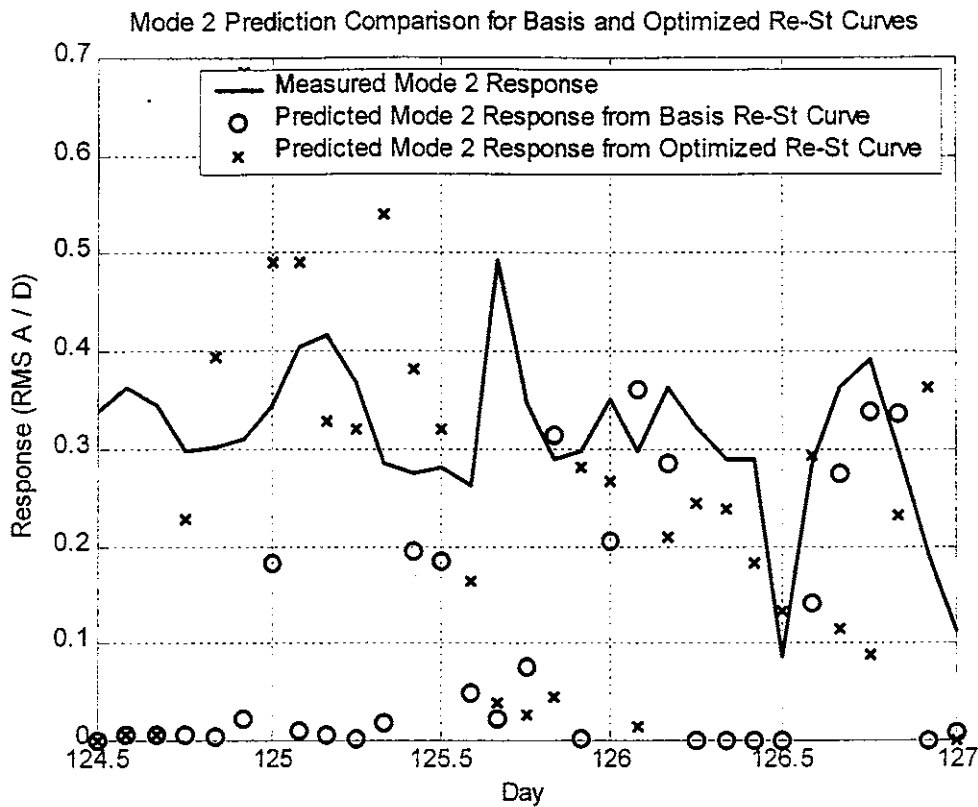


Figure 3-7: Second Mode Response Comparison Between Different St-Re Curves

Comparing the predicted modal responses in Fig. 3-6 and Fig. 3-7 using both the curve taken directly from the Achenbach reference and the optimized curve (Fig. 3-5), the optimized curve better matches the actual measured dominant mode responses for all but the first mode. For days 123 – 124.5, where the first mode is dominant, the optimized curve and the basis curve taken from Achenbach are equally good. The advantage to using the optimized curve is apparent in Fig. 3-7, where the norm of the difference between the predicted and measured mode 2 response is nearly 20% better using the optimized curve. Even casual observation of Fig. 3-7 reveals that the modal amplitude predictions using the optimized curve match many more cases with much better accuracy than the predictions using the basis curve do.

The optimized Strouhal-number-to-Reynolds-number dependency differs from the basis curve in that it has a slightly larger maximum Strouhal number (0.36 vs. 0.31) and begins and ends the elevated Strouhal number range at lower Reynolds number values. The slightly higher maximum Strouhal number is justified by the broad range of possible Strouhal numbers for a cylinder of intermediate roughness (see Fig. 3-3). The higher Strouhal numbers are also verified by independent simulations by Vikestad and Solaas [11], who calibrated the SHEAR7 program with the measured Helland-Hansen data and found that the Strouhal number range needed to match the measured data were $St \# = 0.30$ to 0.40 . These values are in the same range as the Strouhal numbers in this optimized curve relating Strouhal number to Reynolds number (peak $St \# = 0.36$). The shift to a lower starting and ending point for the transcritical region may be accounted for by cylinder motion. The basis curve is taken from data for a stationary cylinder. Turbulence may also contribute to the onset of the transcritical region at lower Reynolds numbers. The jump from a subcritical flow regime to a transcritical flow regime, and then from a transcritical flow regime to a supercritical flow regime, occurs at a lower Reynolds number when turbulence is increased, which may account for the shift evinced here. The recommended curve is shown in Fig. 3-5, where it can be seen compared to the basis curve taken from the Achenbach reference. All simulator response predictions, when not otherwise specified, make use of this curve to relate Strouhal number to Reynolds number.

Once the optimized Strouhal number to Reynolds number curve was obtained, cases were tested comparing predictions using this optimized curve to those using a constant Strouhal number. One of the most important criteria in judging comparisons is

determining which method correctly predicts the dominant responding mode. The simulator should then be able to conservatively estimate the amplitude of that mode's vibration, so that fatigue damage can be calculated. The following illustration compares a simulation run for a case with a constant Strouhal number of 0.25 to the same case using the optimized Strouhal-number-to-Reynolds-number curve. Both predictions are compared to the measured modal response for this May 5 0200h case.

Comparison of Response Predictions using Constant and Re-#-Dependent Strouhal Numbers

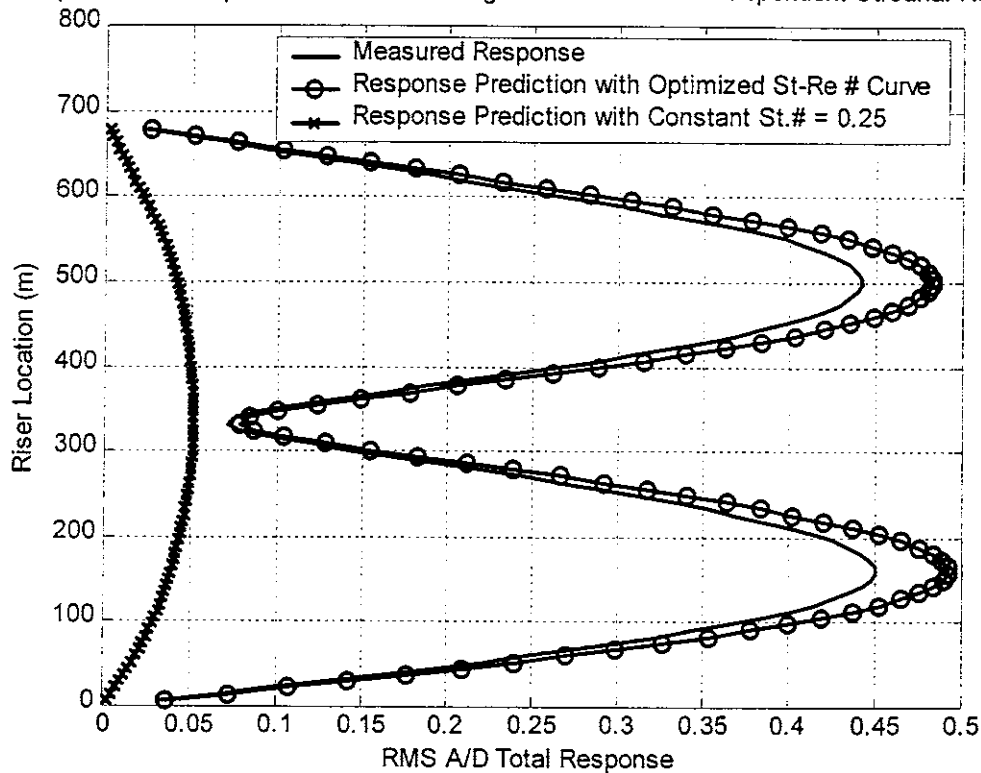


Figure 3-8: May 5 0200 - Illustrates Need for Reynolds-#-Dependent Strouhal Num.

As seen in Fig. 3-8, using a constant Strouhal number of 0.25, the first mode is incorrectly predicted as being dominant, and the response amplitude and fatigue damage

is underpredicted by an order of magnitude. Using the optimized Strouhal number to Reynolds number curve shown in Fig. 3-7, the second mode is correctly predicted as being dominant, and the amplitude of response very nearly matches the actual measured amplitude of response. The effect of the Strouhal number on modal power-in regions can be seen in Fig. 3-9 and Fig. 3-10, which are taken directly from the MATLAB VIV-response simulator and show the location and length of each power-in zone along the riser.

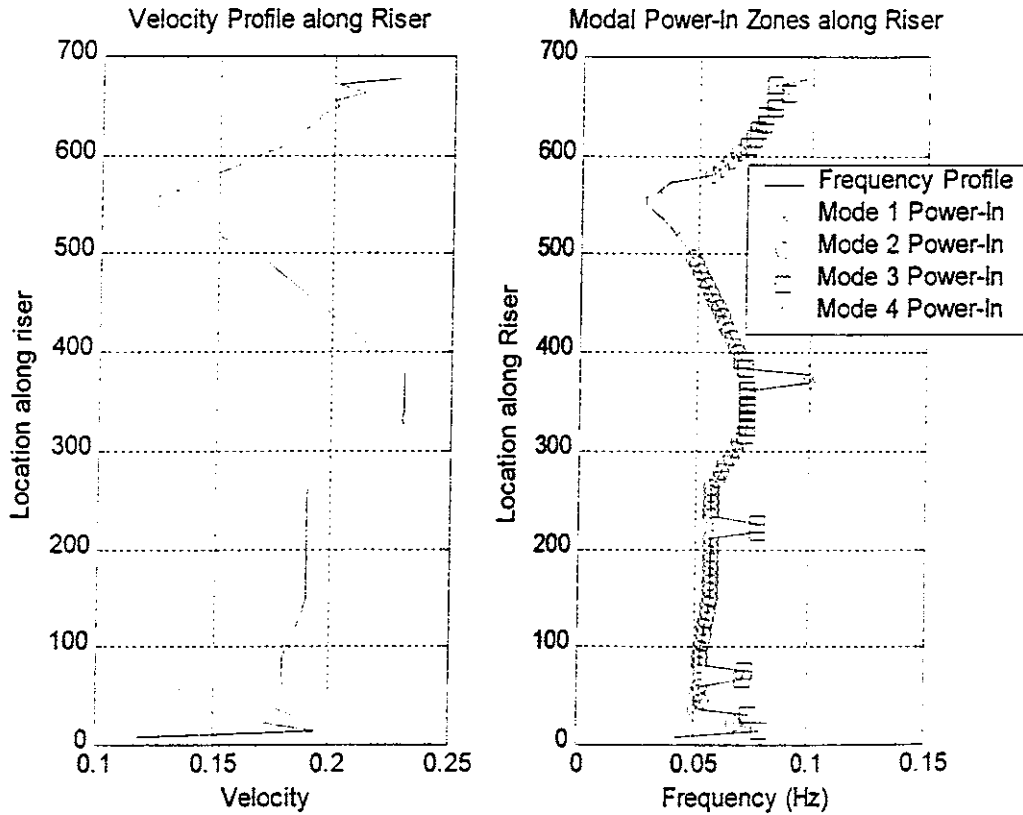


Figure 3-9: May 5 0200 – Power-in Locations from Simulator for Constant $St=0.25$

Fig. 3-9 shows the predicted power-in zone locations along the riser when the Strouhal number is a constant 0.25 all along the riser. The mode 2 power-in region is too small.

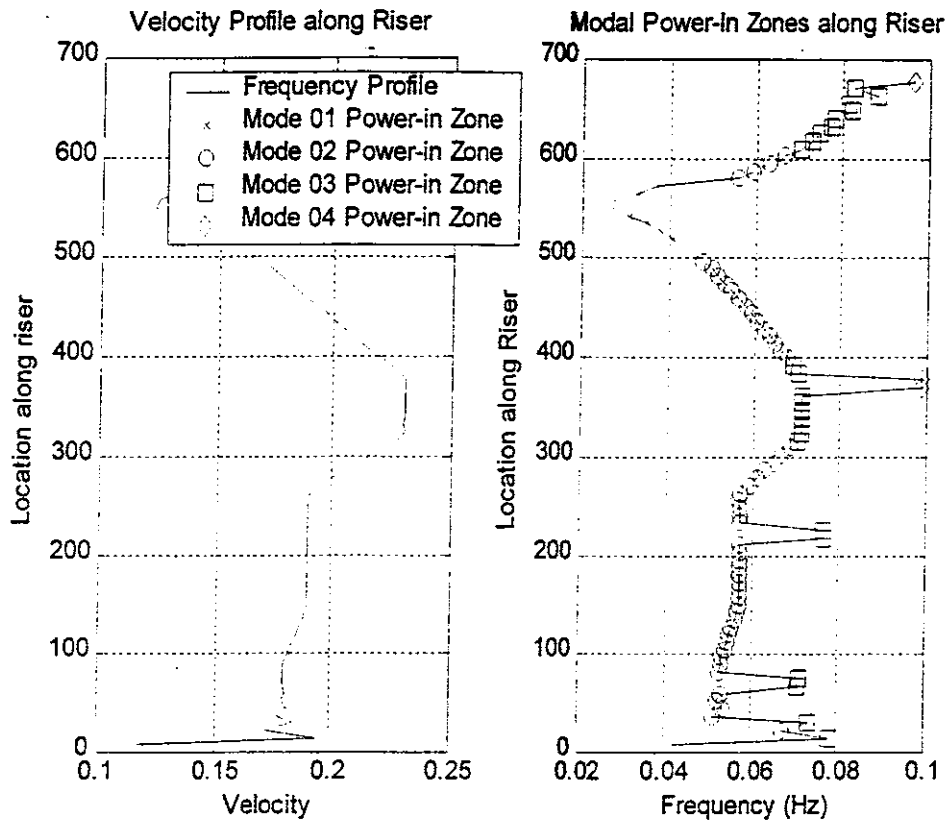


Figure 3-10: May 5 0200 – Power-In Locations for Optimized St #-Re # Curve

Figure 3-10 shows the simulator VIV-response prediction for the same case, but using the optimized Reynolds number versus Strouhal number curve shown in Fig. 3-7. The total length and location of the predicted power-in zone for mode 2 lead to a response prediction for mode 2 that very closely matches the measured data for this case.

Chapter 4

Resolving Conflicts Between Overlapping Power-In Regions

An important consideration when simulating vortex-induced vibration response is deciding which mode is excited at a location where more than one mode may theoretically be excited. Large reduced-velocity bandwidths and closely-spaced modes cause multiple modes to be predicted at a single location for a given current velocity. Since it is not physically realistic to have multiple frequency inputs at one given location where the current is assumed constant, a valid simulation must make an estimate as to which mode lays final claim to a certain location as a power-in point for that mode. This overlap conflict decision should be made in a conservative manner and calibrated with measured riser VIV response.

For the lowest numerical modes there is very little overlap present, since the power-in region is a certain percentage of the natural frequency of each mode and so the actual power-in frequency bandwidth is relatively small for modes 1-4. For these lowest-numbered modes, overlap is a rare occurrence, unless caused by a large difference in Strouhal number between two adjacent modes. As a result, there were few ideal cases in the Helland-Hansen data where there was significant power-in overlap and an opportunity to optimize an overlap resolution algorithm.

The highest amplitude of riser response, and the highest corresponding rate of fatigue damage, occurs when a single mode is dominant in the riser response. It is then very useful to know, based on a given set of riser properties and a given current velocity profile, whether single-mode dominance is likely to occur. An exact but time-consuming

process is to compute the response for each mode, given the maximum size of the power-in regions. The mode with the largest response is assigned the full power-in region.

Other modes have their power-in regions reduced and the response is calculated again.

A non-iterative method of resolving overlap makes use of a theoretical derivation by Vandiver in which the total response amplitude can be roughly predicted without actually running a full simulation. The following integral over the maximum power-in region for each mode is computed:

$$w(U, D, z, \omega_r) = \int U^3 D^2 \varphi dz \quad (4-1)$$

where each modal weight w is determined by integrating over the modal power-in region with the parameters U , local velocity, D , local riser diameter, and φ , mode shape. Modal weight w is a total rough estimate of the modal weight taking into account both lift force excitation and hydrodynamic damping, as taken from Vandiver. [12]

Single-mode dominance is indicated when the modal weight predicted by Eq. 4-1 for any single mode is much larger than the weight of any other mode. Velocity profiles which lead to single-mode dominance are usually large contributors to fatigue damage. An attempt was made to use Eq. 4-1 to develop a simple quantitative predictor of single mode dominance. For each mode determined to have a power-in region Eq. 4-1 was used. In each case the integral was conducted over the maximum possible power-in region for each mode. The two modes with the largest modal weight were then determined and the ratio taken between as the larger over the smaller value. The resulting value was plotted for many samples of data in chronological order, as shown in Fig. 4-1 and 4-2.

The simulation program was also run and the same ratio was found for the top two responding modes as found by running the simulation program. This result was also plotted. In most cases these two measures of modal dominance were in close qualitative agreement. They both indicated when the ratio was large or small at the same time. The next step was to compare the prediction to the actual data, to see if a single mode was indeed dominant. For this a new statistical tool was required.

Kurtosis, as explored by Vandiver [12], is a statistical process useful in determining the occurrence of single-mode-dominant events in measured data. Kurtosis is defined as:

$$kurtosis = \frac{\langle x^4 \rangle}{\langle x^2 \rangle^2} \quad (4-2)$$

where x is the local riser response and $\langle \rangle$ indicates a time-averaged value of the riser response as indicated in Eq. 2-7. The response, $x(t)$ is assumed to be a zero mean process. The kurtosis of a Gaussian random process is 3.0. The kurtosis of a single-frequency sinusoidal response is 1.5. Single-mode-dominant events can be quickly identified from the measured data by computing the kurtosis of a typical acceleration response of the riser for every data block (every two hours for the Helland-Hansen riser). When the kurtosis approaches 1.5 the response is dominated by a single sinusoidal component, which indicates lock-in.

Test cases were chosen from the measured Helland-Hansen data. The simple algorithm from Eq. 4-1 and the full response simulation were both used to construct the ratio of the top two responding modes. The result was plotted with the computed kurtosis

for a series of measured response events. The results are plotted in Fig. 4-1 and 4-2. When the ratios are high indicates a prediction of single mode dominance. If correct, these peaks should line up with valleys in the kurtosis plots, approaching the value of 1.5.

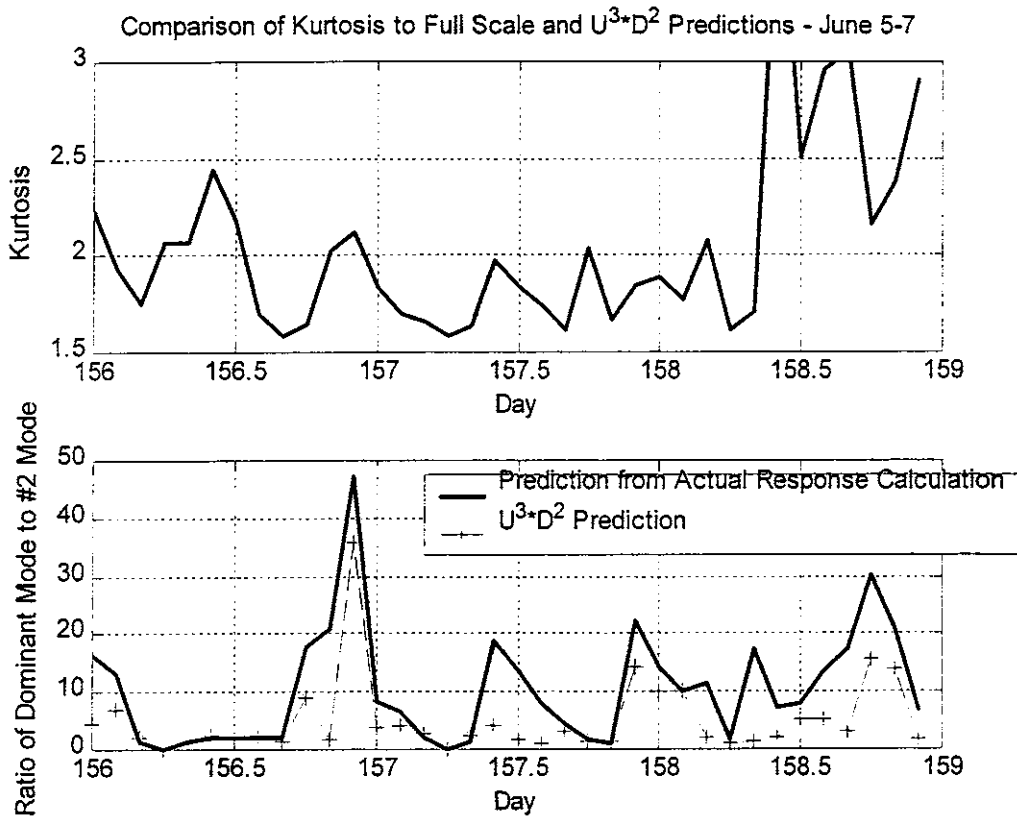


Figure 4-1: Comparison of $U^3 \cdot D^2$ Response Prediction to Full Response Prediction

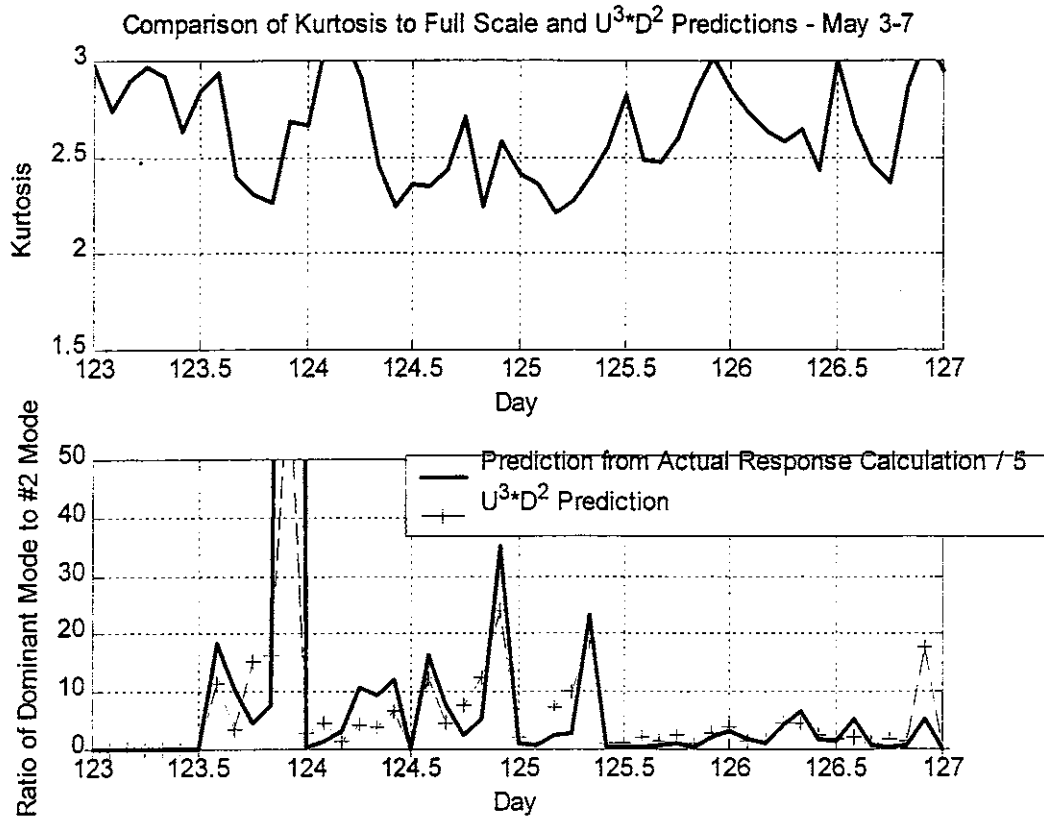


Figure 4-2: Comparison of $U^3 \cdot D^2$ Response Prediction to Full Response Prediction

Figure 4-1 shows very good agreement between the modal weight prediction and the calculation of the response ratio using the simulation program. During this period the kurtosis in the span of days from June 5-7 comes very close to 1.5, indicating pure sinusoidal behavior. In Figure 4-1 the predictors are correct less than half of the time. In general the number of predicted single mode events is correct, but not the time of occurrence.

Fig. 4-2 illustrates a span of days (May 3-7) in which the kurtosis does not indicate single-mode dominance as being prevalent. In Fig. 4-2 the method of predicting the dominant mode by using the rougher modal weight prediction underpredicts the ratio

of dominant to second-largest modal response by roughly a factor of 5. Other than this multiplication factor, the correlation between this rough prediction and the full-scale simulation result is good. Over this span of 5 days (May 3 – 7) the rougher method and the full-scale simulation follow the same trend, generally increasing and decreasing for the same time points. Again the valleys in the kurtosis curve and the peaks in the modal response ratio do not align more than 50% of the time.

It must be acknowledged that though the agreement between the single mode dominance predicted by the full scale simulation and the abbreviated method is good, neither single mode prediction is in good agreement with the kurtosis, which indicates single-mode dominance in the measured data. This implies that the model is not matching the real-world occurrences of single-mode dominance. One possible reason is that the predictors analyze the response of each mode independently and ignore the modal interdependence, leading to an inaccurate prediction of single-mode dominance. Another explanation is that the relationship between Strouhal number and Reynolds number has such a broad range of possibilities for a cylinder of medium smoothness (see Fig. 3-3) that the optimized curve found in Chapter 3 for the Helland-Hansen Strouhal number relationship may not be exactly accurate for each individual case tested.

Still, Fig. 4-1 would seem to indicate that for real-life single-mode-dominant cases, the rougher modal weight predictor can be a useful pre-processing tool in determining whether a single mode is likely to be dominant before performing a complete response prediction. But since the difference between the two predictions became greater by a factor of roughly 5 in the second span of days tested (Fig. 4-2), the good correlation between the preliminary and full-scale prediction may not hold for all cases. A more

thorough analysis of many more test cases may yield additional insight as to how to best predict modal dominance.

However, the U^3D^2 predictor is still a useful tool if used in a statistical way. If tested against a variety of profiles it will give an indication of the fraction of the profiles that will result in single mode dominance. Such an indicator is not reliable at precisely predicting single mode lock-in. This suggests that there is still much that is not adequately understood about vortex-induced vibration under field conditions with real risers.

Chapter 5

Convergence of Natural Frequencies and Added Mass Coefficients

One observed phenomenon in vortex-induced vibrations that is often not accounted for in many current simulators is the variance of natural frequency. Particularly for the lowest mode of the drilling riser, the natural frequency is observed to change up to 40%, as can be seen in the data obtained from research collaborators in Norway, in particular, from the research of Froydis Solaas and K. E. Kaasen. [3] Other low-numbered modes may not exhibit quite so drastic a variance, but they do exhibit enough variation to make modeling variable natural frequencies an important consideration in attempting response prediction. Fig 5-1 shows the measured natural frequencies of the first two modes as they vary over a span of several days from May 3 to May 7:




Figure 5-1: Measured Natural Frequencies from May 3 - 7

The underlying cause for the change in natural frequency is the variation of added mass coefficients. As mentioned briefly in Chapter 2, fluid added mass causes the observed natural frequencies in water and in air to be quite different. For cases in which most of the current acting on the riser is at a velocity lower than the velocity required to excite a particular mode, the added mass over the riser for that mode is close to a value of 1, implying that the riser in water vibrates as though it were twice as massive as a riser vibrating in air. But when the current velocity at a riser location becomes equal to or exceeds the velocity required to excite a particular mode, the equivalent added mass at that riser location may be substantially less than one, providing an explanation for the rise in natural frequency that is observed in the measurement data for increasing current

speed. The added mass at any one riser location varies for each mode; a particular location may have a velocity substantially less than the power-in velocity for one mode, but may have a velocity equal to or higher than the power-in velocity for another mode of interest. This creates cases in which a mode with a very low natural frequency has much less added mass and a high variation from its minimum observed frequency while a mode with a higher natural frequency sees no large reduction in added mass and has essentially no change from its minimum natural frequency.

In the response simulation program used in this thesis an attempt was made to model added mass variation as a function of reduced velocity distribution for each mode. In regions with reduced velocity less than that in the modal power-in region the added mass was set at 1.0. An added mass of 0.5 was used in the power-in region, and an added mass of 0 was used in the high-reduced velocity region. This yielded predicted natural frequencies that closely followed the actual measured variations in natural frequency. This was an iterative algorithm, because the power-in region depends on the natural frequency, which itself depends on the added mass. As an example of a typical iteration process to predict the added mass and natural frequencies of a drilling riser, Figs. 5-2 to 5-4 illustrate some steps in the iteration process for the current profile from May 5 0200 h (see Fig. 3-9 for current profile). The first step of the iteration process assumes an added mass value of 1 everywhere, yielding the minimum natural frequency, since the added mass at any location will have a maximum value of 1. Based on the natural frequency for mode 2 calculated with this default value of added mass, the mode 2 power-in zones appear as follows:

Initial Mode 2 Power-In Locations with Default Added Mass = 1 Everywhere - May 5 0200h

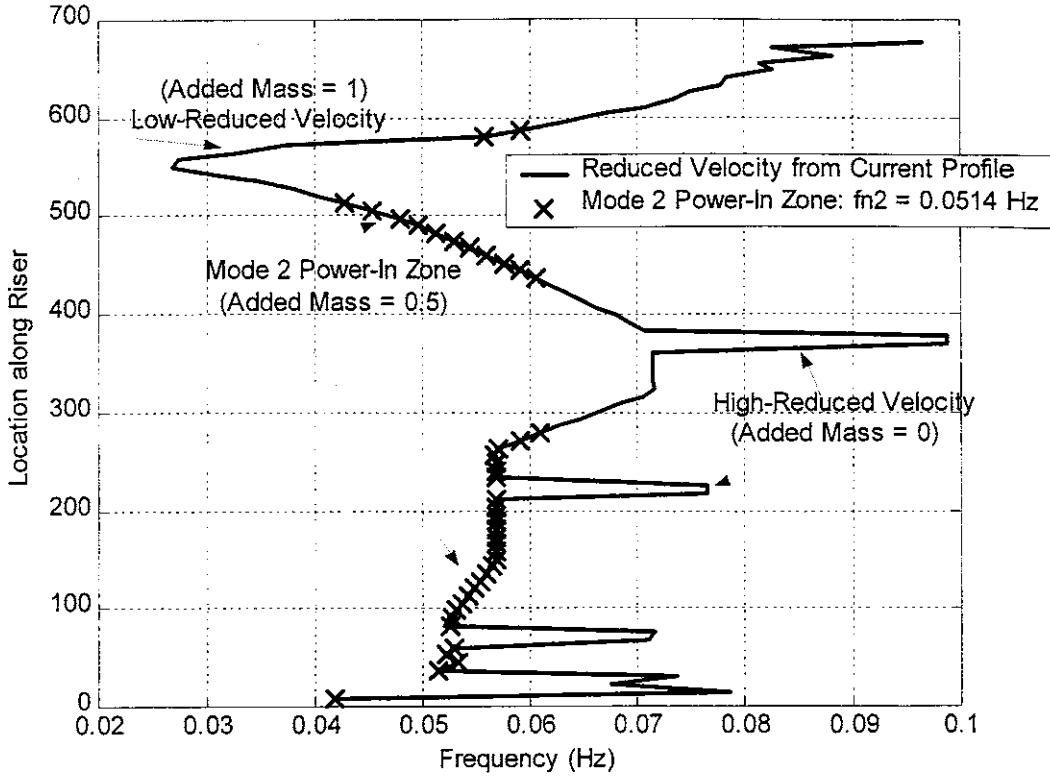


Figure 5-2: Initial Mode 2 Power-in Locations for Default Added Mass

As can be seen in Fig. 5-2, the riser is divided into three regions for a particular mode: a high-reduced velocity region, a power-in region, and a low-reduced velocity region. These are also the same regions also used to assign the local damping and to determine where the excitation input to the riser occurs for a specific mode. Based on the added mass model described, the small low-reduced velocity region keeps its added mass at 1, while the added mass is changed in the regions now indicated to be power-in or high-reduced velocity. This reassignment of added mass is recursively weighted so that only part of the areas that indicate a change of region between iterations are assigned a new added mass value. When this is not done and all of the areas that change from one

region to another are assigned their full new added mass values, the iteration process does not always converge to a final set of added mass and natural frequency values. After the first reassignment of added mass in this example, there are now areas in the power-in and high-reduced velocity regions of the riser with added mass less than one. When the natural frequency of mode 2 is recalculated, this decreased total mass causes a slight increase in the mode 2 natural frequency, so that in the second step of the iteration the new power-in zone corresponding to mode 2 with the altered added mass values and slightly higher natural frequency can be seen in Fig. 5-3:

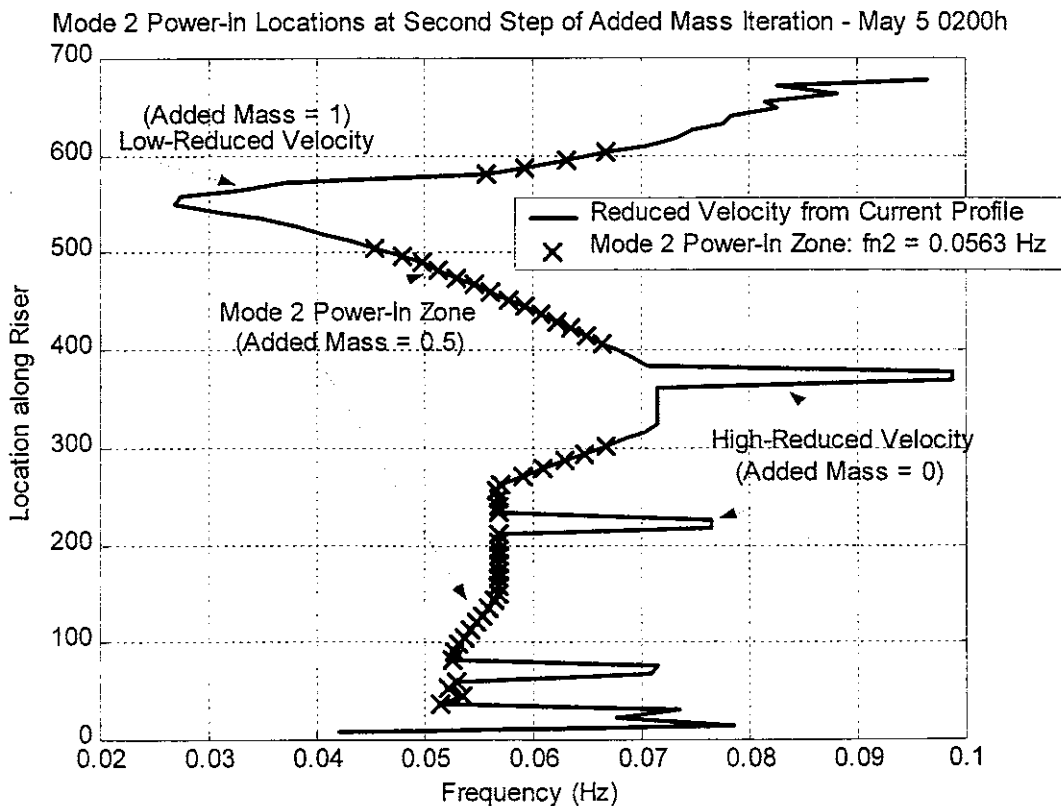


Figure 5-3: Second Step Mode 2 Power-in Locations for Iterative Added Mass

In this case the change in natural frequency and power-in location is not too large, and after only a couple of iteration steps the agreement between mode 2 natural frequency and power-in location is in good agreement. For the recursively assigned added mass, there will eventually be an iteration step, usually within a small number of steps (5 steps is usually sufficient), where the new recursively weighted added mass assignment results in so small a change that the new natural frequency calculated with the new added mass values is almost exactly the same as for the previous values of added mass. This means that the added mass values no longer need to be changed and the iteration can end with these converged values for added mass and natural frequency for that particular mode. Fig. 5-4 shows the converged values after 5 iteration steps for the case being demonstrated:

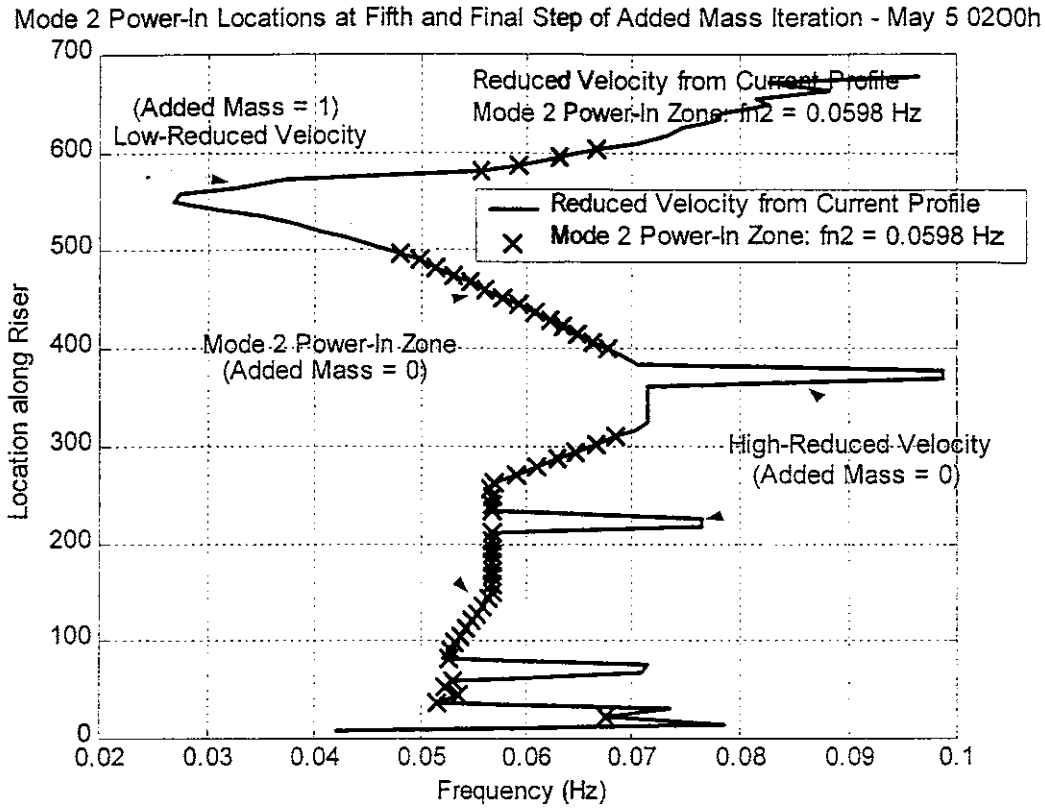


Figure 5-4: Converged Mode 2 Power-in Locations for Iterative Added Mass

Figure 5-5 illustrates the variation in natural frequencies for the measured Helland-Hansen data. The converged natural frequencies for modes 1 to 3 using the iterative method just described are compared to the measured natural frequencies for the span of 4 days May 3 – 7:

Comparison of Measured Natural Frequencies to Predicted Natural Frequencies - May 3-7

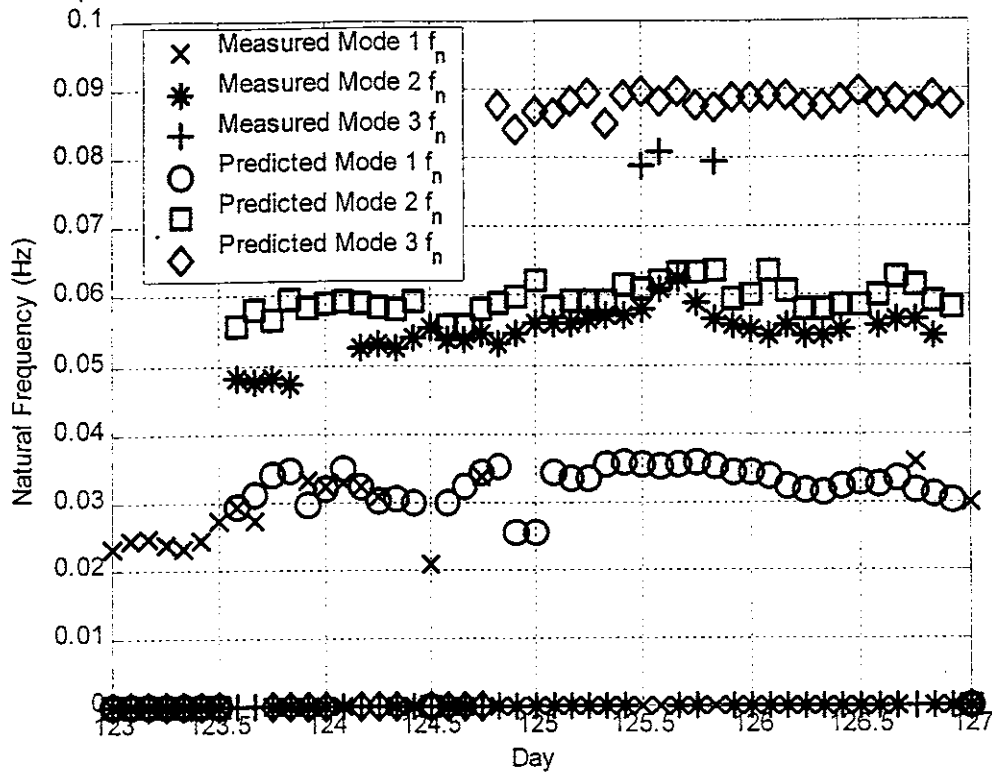


Figure 5-5: Comparison of Measured and Predicted Natural Frequencies May 3-7

As can be seen, even with the iterative procedure just described, there is still some variation between measured and predicted natural frequencies. However, the trend for increasing and decreasing natural frequency was close between the measured and predicted natural frequencies, so a further refinement of this method has promise of yielding even better results. Certainly, the predictions of natural frequency shown in Fig. 5-5 are better than they would be by merely setting the added mass of the riser to an assumed constant, in which case the natural frequency will never vary.

Chapter 6

Comparison of Response-Dependent Lift Coefficient Models

A last programming issue of note is the decision of which response amplitude should be used to assign the lift coefficient from the lookup table in Fig. 2-2. The lookup table was acquired from the Gopalkrishnan experimental measurements [9] with forced cylinder motion. In making a response prediction it is possible to apply these coefficients in a variety of ways within a power-in region. Since the local amplitude of response is dependent on mode shape, the lift coefficient will vary with location within a power-in region. An approximation to this is to apply everywhere in the power-in zone the same lift coefficient. It could be the lift coefficient found at the anti-node or it could be based on the average response amplitude found in the power-in region. The measured response was compared to predictions using these three methods. This was done in attempt to ascertain which methods were accurate and conservative for small, medium, and large response amplitudes.

Using a lift coefficient derived from the local modal A/D response along the riser makes the most logical sense, since this takes into account the mode shape when assigning lift coefficient. Fig. 6-1 displays the converged local lift coefficients along the riser for the dominant second mode in the simulated May 5 0200 h case.




Figure 6-1: Lift Coefficient as a Function of Max. Modal A/D for Different Methods

From this figure it can be seen that when using a locally-assigned lift coefficient, regions near modal nodes are not assigned the same lift coefficient as the regions near the anti-nodes, which are oscillating with much higher amplitude than the regions near the nodes. The result is that a locally assigned lift coefficient more realistically models the dependency on the mode shape of the local amplitude of response. The largest concern for using this method is that for a locally assigned lift coefficient, the response iteration may not converge due to the lack of a unique solution. Recall that an iteration is necessary, because the lift and damping are response-amplitude dependent. Fortunately, for all cases analyzed a single unique solution was found (see Fig. 6-2, 6-4), showing that

this most logical method of lift coefficient assignment is also numerically feasible and predicts a single unique response.

The algorithm currently used in SHEAR7 uses a lift coefficient which corresponds to that found at the anti-node of the mode. [6] This has the advantage of having a single value of lift coefficient everywhere in the power-in zone for the mode (see Fig. 6-1), guaranteeing numerical convergence of the response prediction. Also, for small amplitudes of modal response (less than about 0.3 RMS A/D) this method conservatively overpredicts the lift coefficient and corresponding excitation force, resulting in a conservatively overpredicted modal response amplitude. The significant weakness in assigning lift coefficient based on maximum modal amplitude is that as the RMS A/D becomes greater than 0.3, overpredicting the A/D response locally along the riser actually underpredicts the lift coefficient. The lift coefficient table in Fig. 2-2 shows that the assigned lift coefficient decreases once the absolute A/D becomes greater than a value about 0.5 (corresponding to an RMS A/D of about 0.3). So this method of assigning lift coefficient will overpredict response for very small response amplitudes, underpredict response for response amplitudes over 0.3 RMS A/D, and provide a predicted response amplitude in agreement with measure data only for response amplitudes in the neighborhood of 0.2 – 0.3 RMS A/D.

A good compromise between these two methods is to assign a constant lift coefficient everywhere along the modal power-in zone based on the average mode shape within the power-in zone. This has the advantage of guaranteed numerical convergence to a response prediction while not significantly overpredicting the mean local A/D response along the riser, because the maximum A/D is scaled down by the

mean modal amplitude value within the power-in zone. Lift coefficient as a function of A/D for the three methods are shown in Figures 6-2 and 6-4 for different current profiles. Comparison of response predictions using the various methods are shown in Figures 6-3 and 6-5.

Mode 2 Lift Coefficient as a Function of Response Amplitude - Large A/D Case - May 5 0200h

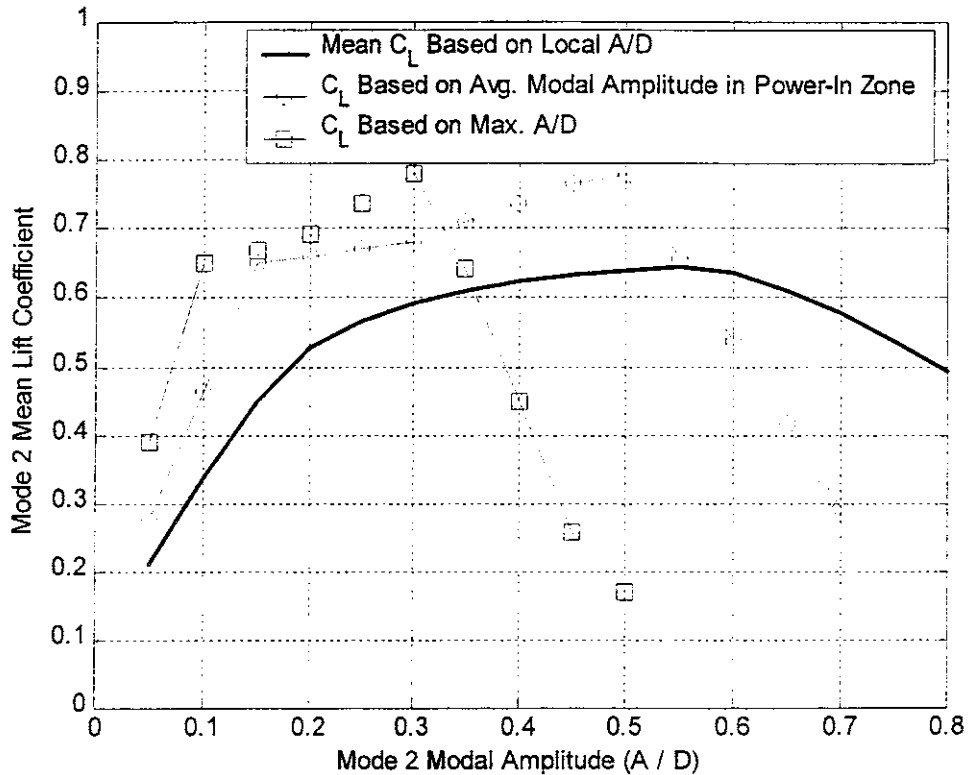


Figure 6-2: Lift Coefficient for Large A/D Response for Different C_L Methods

It can be seen that for the large A/D response case in Fig 6-2, the lift coefficient assignment based on maximum modal A/D drops off much more rapidly than the lift coefficient based on the other two methods. As a result, the method which assigns the local lift coefficient to be the lift coefficient at the maximum A/D location underpredicts the actual response, as can be seen in Fig. 6-3. Both the local lift coefficient assignment

and the lift coefficient assignment based on a mean modal amplitude within the modal power-in zone predict a response (in Fig. 6-3) that closely matches the measured response.

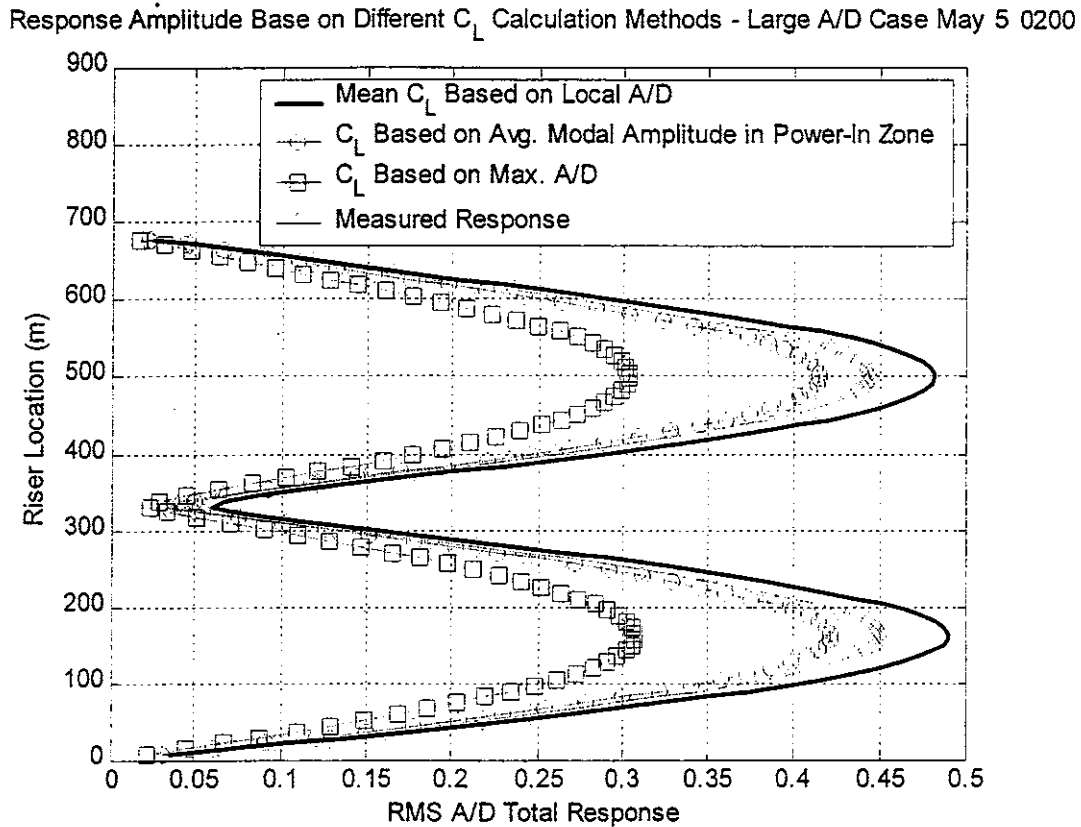


Figure 6-3: Response Predictions for Large A/D Case for Different C_L Methods

Fig. 6-3 clearly shows that a lift coefficient based on maximum modal amplitude is not satisfactory for a case with large response, since the large amplitude of motion causes the predicted lift coefficient to be much too small. For very small amplitude cases, the lift coefficient associated with the maximum modal amplitude will slightly overpredict the local lift coefficient. But such cases are fairly inconsequential since very

little response or fatigue occurs for such cases. For midrange response amplitude (somewhere between RMS A/D = 0.2 to 0.35) the mean lift coefficient and predicted response is similar using any of the three methods described. For this special region the lift coefficient (as seen in Fig. 2-2) is near the vicinity of its peak value and a difference in A/D does not produce a large difference in the lift coefficient. This can be observed in Fig. 6-4:

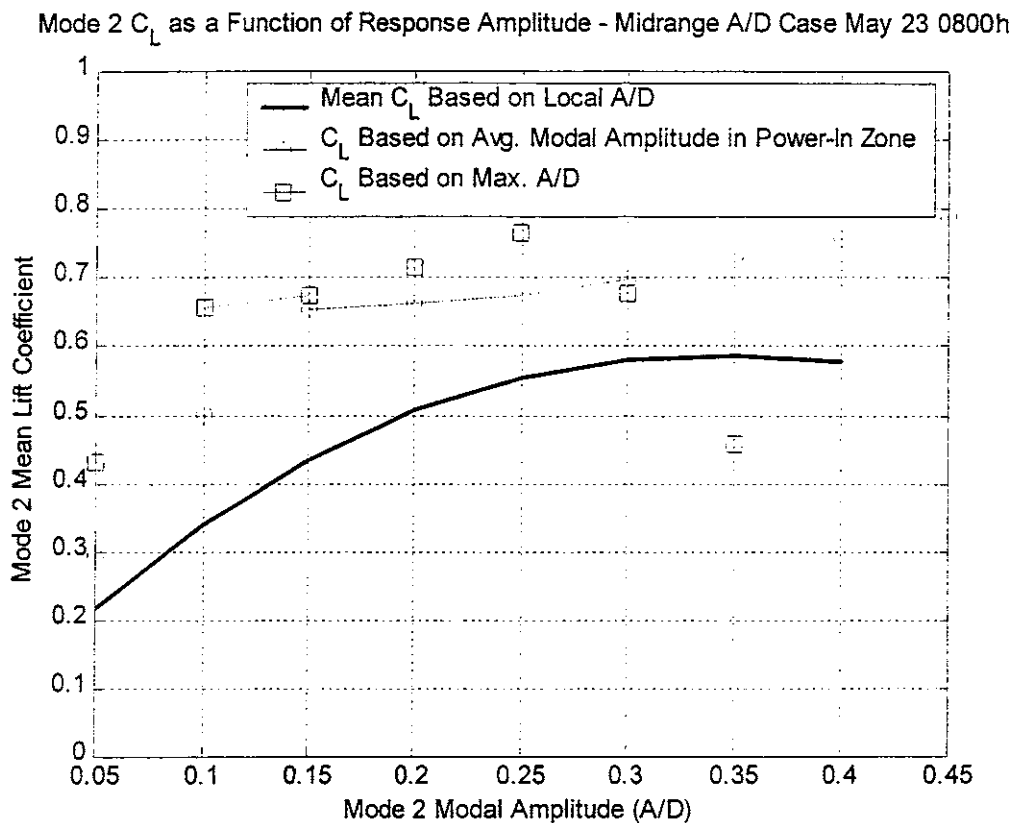


Figure 6-4: Lift Coefficient for Midrange A/D for Different C_L Methods

Comparing the mean lift coefficients for this midrange A/D case (Fig. 6-4) to the mean lift coefficients for a large A/D case (Fig. 6-2), it can be observed that the agreement between mean lift coefficients obtained by the different methods is better for

this midrange case than it is for the large A/D case. As explained, this is a result of the overprediction of the response by the method based on the maximum modal amplitude, which has little effect on a midrange A/D response lift coefficient prediction but begins to underpredict the lift coefficient for increasingly large A/D cases. The very similar response predictions for this midrange A/D case from the different lift coefficient assignment methods are shown in Fig. 6-5:

Response Amplitude based on Different C_L Calc. Methods - Midrange A/D Case May 23 0800

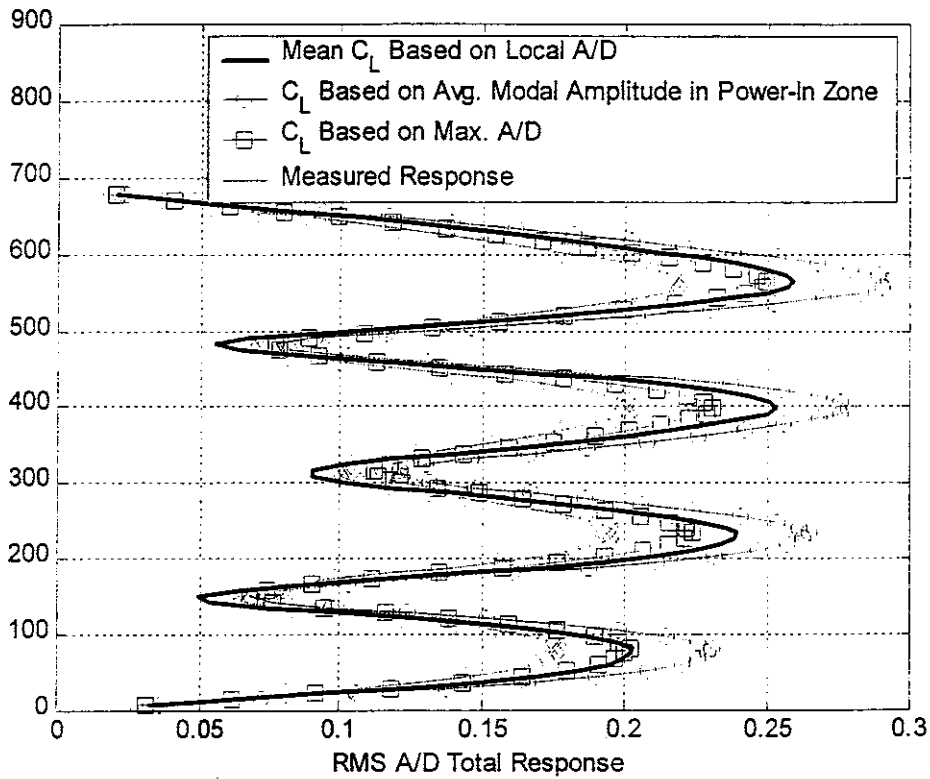


Figure 6-5: Response Predictions for Midrange A/D Case for Different C_L Methods

Chapter 7

Conclusion

Various simulation issues for predicting vortex-induced vibration in drilling risers were studied. Both programming issues and physical issues were tested and compared to full-scale measured data. The following are the main findings of this research:

- 1) The response prediction using modal analysis is nearly identical to the response prediction calculated from the frequency-response functions. This validates the use of modal analysis, the faster and more efficient method used to calculate VIV response.
- 2) Using a Strouhal number as a function of Reynolds number for a cylinder of a particular roughness more realistically simulates hydrodynamic effects on the riser response. This also eliminates the need to estimate a constant Strouhal number to be used in the simulation, which may not be conservative or accurate in predicting VIV response. However, the Strouhal number in the transcritical Reynolds number range is very sensitive to a variety of factors, such as roughness, turbulence, and cylinder motion. This prevents exact prediction of Strouhal number, and suggests that conservative practices in selecting Strouhal numbers should be used.
- 3) The use of a relatively simple formula may be useful in predicting single-mode dominance, the most important case in determining riser fatigue. This can be used to determine in advance which modes are likely to dominate before a full response prediction simulation is run.

- 4) The use of an iterative process for assigning added mass can be useful in predicting how the natural frequencies of the riser change with changing current velocity profiles. An iterative model for recursively assigning added mass for different reduced-velocity regions on the riser was able to yield quite close agreement between measured and predicted natural frequencies. However the lack of reliable measured tension data prevented an exact determination of an added mass iterative scheme.
- 5) Of the different possible methods of assigning lift coefficient within the power-in zone, assigning lift coefficient locally based on the mode shape is the most logical method. Comparisons between measured and predicted response using locally-assigned lift coefficients revealed good agreement.

The main focus of this thesis was to explore the programming and theoretical issues described above, compare them to full-scale measured data, and to seek out ways in which the simulation of VIV response could be improved. A working MATLAB simulator using much of the theory and models from the SHEAR7 VIV analysis program was created as a test bed for the many different parameters and methods explored in this thesis. The full-scale measurements from the Helland-Hansen riser were directly compared to all of the simulation results as a reference for calibration and justification of the results obtained and the methods explored in this thesis.

Areas for potential research related to this thesis include:

- 1) Close examination of measurements from other risers with different surface roughness and Reynolds number flow regimes to try to generalize

the optimized curve relating Strouhal number to Reynolds number to other drilling riser cases.

- 2) Further exploration of the simple model used to predict single-mode dominance.
- 3) Further study of the iterative model used for assigning added mass.

Comparison to measured natural frequencies in other riser data could help further develop the values used for added mass in the different reduced-velocity regions and help better predict changing riser natural frequencies.

Bibliography

1. Faltinsen, O.M. 1990. Sea Loads on Ships and Offshore Structures. Cambridge University Press, Cambridge, U.K.
2. Kaasen, K. E.; Norwegian Marine Technology Research Institute (MARINTEK); Lie, H.; MARINTEK (USA) Inc.; Solaas, F.; MARINTEK; And Vandiver, J. K., Massachusetts Institute of Technology. 2000. "Norwegian Deepwater Program: Analysis of Vortex-Induced Vibrations of Marine Risers Based on Full-Scale Measurements." Proc. Offshore Technology Conference 2000, Paper No. 11997, Houston, May 1-4, 2000.
3. Solaas, F. and Kaasen, K.E. 1999. "Norwegian Deepwater Programme Riser and Mooring Project – Improved VIV Predictions." Riser VIV Data Analysis, Helland-Hansen, Report MT51 F99-259 (513145.03.01), MARINTEK, Trondheim, Norway.
4. Meirovitch, Leonard. 1975. Elements of vibration analysis. Mc-Graw-Hill, New York.
5. Rao, S. S. 1995. Mechanical Vibrations. Addison-Wesley Publishing Company, Reading, Massachusetts.
6. Vandiver, J.K. and Li, L. 1997. "SHEAR7 Program Theoretical Manual." Department of Ocean Engineering, Massachusetts Institute of Technology, rev. Dec. 1997.
7. Venugopal, M. 1996. "Damping and Response Prediction of a Flexible Cylinder in a Current." Doctoral Dissertation, Massachusetts Institute of Technology, Department of Ocean Engineering, Supervisor: Prof. J. Kim Vandiver, Feb. 1996.
8. Vikestad, K.; Larsen, C. M.; and Vandiver, J. K. 2000. "Norwegian Deepwater Riser and Mooring: Damping of Vortex-Induced Vibrations." Proc. Offshore Technology Conference 2000, Paper No. 11998, Houston, May 1-4, 2000.
9. Gopalkrishnan, R. 1993. "Vortex-induced forces on oscillating bluff cylinders." Doctoral Dissertation, Massachusetts Institute of Technology, Department of Ocean Engineering, Supervisor: Prof. Michael S. Triantafyllou and Mark A. Grosenbaugh, 1993.
10. Achenbach, E. and Heinecke, E. 1981. "On Vortex Shedding from Smooth and Rough Cylinders in the Range of Reynolds Numbers 6×10^3 to 5×10^6 ." Journal of Fluid Mechanics, 1981, vol. 109, pp. 239-251.

11. Vikestad, K. and Solaas, F. 1999. "Calibration of SHEAR7 Using Full-Scale Riser Data from the Helland-Hansen Drilling Riser." Helland-Hansen, Report MT51 F99-0388 (513145.04.01), MARINTEK, Trondheim, Norway.
12. Vandiver, J. K. 2000. "Predicting Lock-In on Drilling Risers in Sheared Flows." FIV 2000 Conference, Lucerne, Switzerland, June 18-22, 2000.

Appendix A: Riser Properties

The Helland-Hansen riser, as seen in Fig. 1-1, was modeled as being 685 m long, with a structural diameter of 0.53 m, and consisting of 45 joints each 15.24 m in length. The following joints, listed in order from the bottom of the riser, were not covered with buoyancy: 1-2, 5, 15, 25, 39-45. For all other joints buoyant sections were added so that the effective diameter for these buoyant sections was 1.13 m. The equivalent mean riser density used in the simulations for this thesis was then 2877 kg / m for the regions without buoyancy and 1051 kg / m for the buoyant regions.

The top tension was modeled as 3.1 MN, though, as mentioned, the tension on the riser did in all likelihood change from day to day. Structural damping in the riser was modeled as 0.3%.

Appendix B: Reynolds Number Lift Modification Factor γ_L

The following lookup table (Fig. B-1) from the Gopalkrishnan experimental data [9] is used to find the value of Reynolds Number Lift Modification Factor γ_L in Eq. 2-17:




Figure B-1: Lookup Table for Reynolds Number Mod. Factor used in Eq. 2-17

Appendix C: Permutations of Strouhal Number Curve

The basis curve adapted from the Achenbach data [10], shown in Fig. 3-4, along with 4 permutations of this curve, were all used with the VIV-response simulator and compared with the measured modal amplitudes over 2 separate spans of data to determine the optimum relation between Strouhal number and Reynolds number to use for the simulator. The basic curve and the four other curves are shown in Fig. C-1:

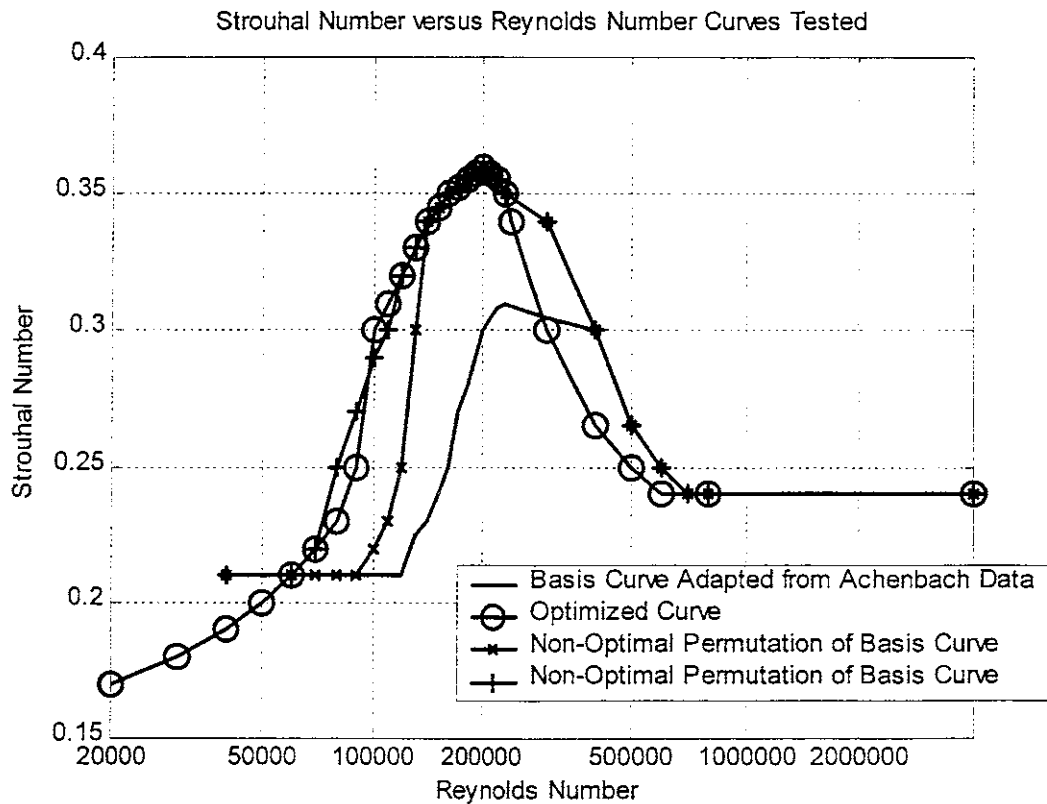


Figure C-1: Tested Permutations of Basis Curve Adapted From Achenbach Data