FLUID-STRUCTURE INTERACTION SIMULATIONS IN APPLICATIONS

FOR NUCLEAR ENGINEERING

A Dissertation

by

LANDON MITCHELL BROCKMEYER

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Yassin Hassan
Committee Members,	Elia Merzari
	Maria King
	Mark Kimber
Head of Department,	Yassin Hassan

December 2018

Major Subject: Nuclear Engineering

ABSTRACT

Nuclear reactors pump coolant through their core and heat exchangers at massive

mass flow rates to sustain energy production. These demanding requirements push

engineers to extreme levels of optimization to safely sustain the transfer of energy. High

flow rates introduce the possibility of flow-induced vibrations. Reactor core and heat-

exchanger/steam generator designs go through many stages of experimentation to ensure

that problematic flow-induced vibrations do not arise. Advances in computational

capabilities introduce the possibility of creating predictive simulations that accelerate the

iterative design process and replace expensive physical experiments. Simulation

methods for fluid-structure interactions are rapidly developing and undergoing extensive

verification and validation. Computational fluid dynamics code Nek5000 and

computational structural mechanics code Diablo have been coupled to create a highly

scalable, high-fidelity fluid-structure interaction code. A fully coupled model of

crossflow through a tube bundle has been simulated using the Nek5000-Diablo code for

validation purposes. Simulations at three velocities were performed to test the method's

capabilities of capturing the onset of large amplitude vibrations that occur at a critical

velocity for the tube bundle. The simulation results compared favorably to the

experiment on which it was based and gave further insight into the mechanisms behind

the vibrations. A 7-pin bundle of wire-wrapped fuel pins was simulated using Nek5000

and the forces exerted on the pins captured. A scheme was developed to synthesize force

histories of indefinite length replicating the Nek5000 force signals, forming a modified

one-directional coupling procedure. Multiple structure simulations were performed, observing the effects of pin-to-pin, and pin-to shroud contact scenarios on the resulting vibrations. The shroud was found to effectively limit vibrations to short wave-lengths on the order of $1/6^{th}$ helical pitch even when extensive gaps between pins formed in the reactor. Both the one-way and two-way coupling methods are successful in capturing the fluid and structure behaviors and provide a convenient method of analysis for these geometries.

ACKNOWLEDGEMENTS

CONTRIBUTORS AND FUNDING SOURCES

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER I

INTRODUCTION

Nuclear reactors represent feats of engineering for a myriad of disciplines.
Gigajoules of energy pass through the core every second for months at a time. The
demands on the coolant flow capacity have pushed thermal-hydraulics engineers to
impressive levels of design and optimization. The optimization process is never finished,
and the current fleet of reactors and next generation designs benefit from enhanced
understanding of the forces experienced. The extreme mass flow rates present in reactor
cores and heat-exchangers/steam-generators is conducive to the onset of problematic
flow induced vibrations (FIV). Heat-exchanger tube eruption and fuel pin failure are
problems that continue to exist in the nuclear industry. Flow-induced vibrations are
notoriously difficult to predict due to their extreme sensitivity to domain geometry and
flow condition. New geometries must undergo extensive experiment and testing in order
to ensure that vibrations will not arise. Otherwise, engineers must design within well
documented but tightly constrained design parameters. Modern advances in
computational capabilities and methodologies have recently reached the point where
coupled fluid-structure simulation are feasible for relatively large geometries. High
fidelity and highly scalable codes enable predictive simulations to assist with the design
and research process.

Chapters 3 and 4 of the present dissertation detail the use of coupled fluid-
structure modeling to gain insight and physical understanding of two applications of

1

engineering relevance. Chapter 3 details the fully 2-way coupled simulation of crossflow through a tube bundle, a flow geometry relevant to heat exchanger and steam-generator designs. In it a detailed verification and validation procedure is presented with the goal of adding credibility to the methodology and simulation software. The 3-D flow field and vibration measurements are studied to give insight into the nature of the onset of vibrations. Chapter 4 details the simulation of flow through a 7-pin wire-wrapped fuel pin bundle. The forces exerted on the pins are used to generate body forces for simulation of the wire-wrapped fuel pin structure. The structure simulation can give insights into the nature of the vibrations present that would be extremely difficult to measure experimentally. The purpose of these studies is not only to advance the understanding and research new aspects of these geometries, but also to contribute to the confidence in the methods presented.

Background

The fields of computational fluid dynamics (CFD) and fluid-structure interaction (FSI) have diverse methods and rich histories. The background section delves into the history, relevance, and basic methodology for CFD and FIV in general. Backgrounds into these fields are presented in order to provide necessary information and context for the results chapters.

Computational fluid dynamics encapsulates any fluid modeling done by calculation, typically by using the Navier-Stokes equations. The incompressible Navier Stokes equations for momentum is shown in equation 1.1.

$$\rho\left(\frac{\partial u}{\partial t} + u \cdot \nabla u\right) = -\nabla p + \nabla \cdot \tau + \rho f, \quad \nabla \cdot u = 0 \qquad \text{Eq.1.1}$$

CFD has been performed as early as the 1930s by hand, but today CFD is purely associated with the use of computers, in which case CFD really started in 1957 at Los Alamos National Laboratory[1]. Turbulence modeling is a branch of CFD that covers the simulation and modeling of turbulent flows. While this distinction may seem minor, the presence of turbulence presents a number of challenges to simulation, most notably the cascade of energy from large length scales to very small. The small lengthscales, called the viscous lengthscales, become increasingly small as the Reynolds number increases. For most applications direct simulation is not only prohibitively computationally expensive, but catastrophically so and will remain so for the foreseeable future. As a result, a wide range of numerical models have been developed to approximate turbulent flow.

RANS

The earliest efforts to model turbulence are still the most commonly used method in industry: the Reynolds Averaged Navier-Stokes (RANS) equations. RANS models are by far the least computationally expensive of the methods described here, and are the only method capable of modeling turbulent flow with Re > ~100,000 with rare

exceptions for LES. The RANS equations rely on separating the fluid flow into the mean

component and the fluctuating component in a technique called Reynolds

decomposition. This equation is then ensemble averaged yielding:

$$\rho\left(\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial \bar{x}_j}\right) = -\frac{\partial \bar{p}}{\partial x_i} + \frac{\partial}{\partial x_j}\left[\overline{\tau_{ij}} + \rho\langle u_i' u_j'\rangle\right] + \rho \bar{f}, \quad \frac{\partial \bar{u}_j}{\partial x_j} = 0 \qquad \text{Eq. 1.2}$$

where an overbar refers to the averaged quantity, and the prime refers to the fluctuating

component. These equations are not closed due to the unknown term, $\rho\langle u_i' u_j'\rangle$, with

dimensions of stress, and referred to as the Reynolds stress. There are two broad

methods for dealing with this lack of closure. Formulations such as the k-ε two-equation

model use the Boussinesq approximation, which assumes that the Reynolds stress is

proportional to the strain rate tensor, S_{ij}, using a so-called eddy-viscosity[2]. The k-ε

model solves PDEs for the transport of k, turbulent kinetic energy, and ε, dissipation

rate, which are used to formulate the local eddy viscosity. There are many formulations

that use the Boussinesq approximation which vary essentially in their calculation of the

eddy viscosity. These methods can and have been used very successfully for a variety of

flow domains. However they struggle with accuracy in situations where the Boussinesq

approximation is inaccurate, especially when the turbulence is highly non-isotropic.

The other common method of closure is to solve transport equations for the

Reynolds stresses[3]. There are five terms in the transport equation: production - the

transportation of energy from the mean flow to the fluctuating velocity, rapid pressure-

strain correlation - the redistribution of energy due to interaction between the fluctuating

velocity and the velocity gradient, slow pressure-strain correlation – the redistribution of

4

energy due to interaction between fluctuating velocity and is responsible for the return to isotropy, and dissipation – the dissipation of small eddies due to viscosity. These Reynolds stress equation models have the advantage that they do not rely on an isotropic eddy viscosity, thus allowing for anisotropy. Such models are better suited for certain flows, but come at the expense of increased computational cost due to the need to solve six transport equations

DNS

Direct numerical simulation (DNS) is the direct simulation of the Navier-Stokes equations down to the dissipative scales. As mentioned before, although this method is the most accurate, it is by far the most expensive, prohibitively so for most applications of interest. DNS is excellent for fundamental turbulence research due to its high level of accuracy, as long as proper attention is given to numerical dispersion. DNS is also a great tool for creating benchmarks, as it gives a complete 3-D image of the flow for a very fine spatial and temporal resolution, in addition to its physical accuracy. According to the Kolmogorov theory, the minimum criteria a DNS case must meet are the Kolmogorov length, velocity, and time scales (η, u_η, and τ_η respectively), which are each a function of kinematic viscosity, ν, and dissipation ϵ. These scales are formulated in equations 1.3, 1.4, and 1.5.

$$\eta \equiv \left(\frac{\nu^3}{\epsilon}\right)^{\frac{1}{4}} \qquad\qquad \text{Eq. 1.3}$$

$$u_\eta \equiv (\nu\epsilon)^{\frac{1}{4}} \qquad\qquad \text{Eq. 1.4}$$

$$\tau_\eta \equiv \left(\frac{\nu}{\epsilon}\right)^{\frac{1}{2}} \qquad\qquad \text{Eq. 1.5}$$

From these criteria alone, it is apparent that as viscosity decreases, the requirements for DNS grow rapidly.

LES

Large Eddy Simulation (LES) takes the middle ground as a compromise between speed and physical accuracy. Essentially LES directly solves the Navier-Stokes equations, like DNS, without resolving all the way down to the viscous lengthscales. According to the first Kolmogorov similarity hypotheses, at high Reynolds number, the local properties of small scales are isotropic and homogenous. Ideally, LES resolves the larger lengthscales, the inertial subrange, down to these small scales, which are predictable and easily modeled. This way the larger, energy carrying eddies which are affected by the geometry are directly simulated, and the interesting physics are preserved, without the strenuous requirements of the Kolmogorov scales, all with minimum modeling at only the most predictable scales. While LES is significantly less computationally expensive than DNS, current computational capabilities limit it to relatively small applications at reasonable Reynolds number (<100,000). Because LES simulations do not reach dissipation scales, energy must be artificially dissipated, or the

simulation may blow up. This can be done a number of ways, broadly either by constructing a sub-grid scale model, or by filtering out high frequencies. Commonly sub-grid scales are based on an eddy-viscosity[4]. Similar to RANS modeling, the turbulent effects of the sub-grid scales are assumed to behave as if there is an additional turbulent viscosity present. The filtering methods find a way to isolate the higher frequency oscillations in velocity and dampen them, thus adding an artificial energy dissipation[5]. Nek5000 uses a filtering approach that will be discussed in greater depth in the Numerical Methodology section.

FIV and FSI

The study of flow-induced vibrations, or FIV, spans many engineering fields including aerospace, mechanical, nuclear, biomedical, and petroleum. Typically, FIV are detrimental and the focus of study is to predict and prevent such vibrations from occurring. Vibrations can be prevented by adjusting certain geometric and flow parameters, adding flow altering or structure reinforcing components, or in some cases altering the design altogether. FIV can take many forms and there are dozens of identified mechanisms that can cause vibrations. Internal flow, external flow, two-phase flow, noise, temperature, pressure waves, and external vibrations can all affect FIV, especially when more than one is involved. The present study chiefly concerns FIV due to external flow, which largely divides into axial flow and cross flow. Axial flow can simply be thought of as flow that is parallel to a structure with a distinctly longer length scale, e.g. flow parallel to a long, cylindrical tube. Cross flow can be thought of as flow perpendicular to a structure. In the case of single-phase, steady flow, the primary sources

of vibrations for these arrangements are turbulent buffeting induced vibrations, vortex induced vibrations, and fluid-elastic instability induced vibrations. Additionally far-field effects like noise and external vibrations can be important as well, but the scope of this study pertains to the direct interaction of the fluid and structure present.

Turbulent buffeting induced vibrations are the most common, but also the smallest in amplitude of vibration. As the name implies, the energy for the vibration comes from turbulence in the fluid. Turbulent flow has random fluctuations in velocity and pressure, resulting in a randomly fluctuating force load on the structure due to pressure and viscous forces. This force load fluctuates over a continuous distribution of frequencies, and therefor has a non-zero power for the frequency corresponding to the structure's resonant frequency. This is true for every vibrational mode of the structure, but the first mode is dominant in most cases. The RMS amplitude of the vibrations corresponds to the magnitude of the power spectrum density, which increases with velocity of the flow. Turbulent buffeting does not result in a positive feedback effect like other vibrational mechanisms, and therefor results in comparatively smaller amplitude vibrations. Turbulent buffeting is relevant for both cross flow and axial flow.

Vortex induced vibrations (VIV) are likely the most well studied form of FIV. The quintessential example of VIV is the oscillation of a flexibly mounted cylinder in crossflow due to vortex shedding i.e. Kármán vortex streets. The flow forms a boundary layer around the front and top of the cylinder, but detaches as it rounds the back, forming a vortex. At sufficient Reynolds number, this results in the shedding of alternating vortices, causing an alternating pressure load on the cylinder. If the frequency of the

alternating pressure load synchronizes with the cylinder's natural frequency of vibration, oscillations can arise; this scenario is called lock-in. The vortex shedding frequency is inversely proportional to the flow velocity, meaning that the vortex shedding frequency will only be close to the structure's natural frequency for a small band of velocities[6]. Vortex induced vibrations can manifest for any geometry with vortex shedding, provided that the force exerted on the object is significant in comparison to the dissipation forces. Vortex induced vibrations are relevant for crossflow, but not for axial flow.

Fluid-elastic instability although more rare, is capable of producing devastating vibrations resulting in structural failure. Fluid-elastic instability can take many forms, but the common detail is that the fluid and structure interact in such a way that vibrations experience positive feedback e.g. galloping or fluttering. For example, galloping can take place for crossflow into the flat face of a half-cylinder. A small perturbation may force the half-cylinder to move upward slightly. The upward motion gives the flow a slight downward angle of attack on the half-cylinder, generating upward lift and increasing the displacement. Downward lift is generated when the half-cylinder moves back down. As the oscillations increase, the lift generated increases, resulting in a positive feedback loop creating larger and larger oscillations[7]. The large displacements can result in repeated stress on the vibrating component resulting in fatigue or failure. The vibrating structure may collide with nearby structures resulting in damage or failure. These vibrations can be very difficult to predict for complex geometries and they may arise suddenly with increase in velocity. While fluid-elastic instability can occur in both

9

axial and cross-flow, cross-flow is much more prone to its onset. It requires extreme flow conditions for fluid-elastic instability to arise in axial flow.

Vibrations and Nuclear Reactors

As mentioned before, most research into FIV is towards the means of prevention or limitation. The field of nuclear engineering is no exception. Most nuclear reactors use a liquid coolant to extract heat energy from the fuel rods, and carry it to a steam generator. The second loop, typically containing water, boils in the steam generator where the steam is used to turn turbines. Next generation designs often incorporate different coolants (sometimes gas), and occasionally a different thermodynamic cycle for electricity generation, but the commonality is coolant moving heat from the reactor to a heat exchanger. The energy scales involved with nuclear reactors, up to multiple GW of heat energy, mean that fluid system must undergo extreme conditions, including very high velocities and mass-flow rates. These high velocities coupled with the long thin fuel elements and heat exchanger tubes are at high risk for FIV. Any failure due to FIV within a nuclear reactor is potentially dangerous, and stopping the reactor for repairs is always very expensive.

Heat exchangers, including steam-generators, depending on their design, can experience both axial and crossflow, as well as in-between. Most consist of some version of the tube in shell design, which is good for high pressure applications. Essentially many thin tubes run through a large pressure vessel. Through the tubes flows the hot primary coolant, and around the tubes, inside the pressure vessel, the secondary coolant

is heated or boils. Smaller diameter tubes with thinner walls are beneficial for heat transfer, but more susceptible to vibration and breakage. Smaller vibrations from turbulent buffeting can be problematic, but the main concern is larger vibrations from vortices or fluid-elastic instability causing stress from bending or collision with one another. Tube failure is not common, but has been a problem for nuclear reactors in the past, sometimes because of an unexpected excitation from fluid-elastic instability[7].

Flow through the reactor itself is typically primarily axial flow, with some induced transverse flow. The fuel typically consists of long cylindrical pins. For this geometry, VIV is not a concern, and fluid-elastic instability initiates for velocities far greater than standard operating conditions. Turbulent buffeting vibrations, however, can be problematic. The long, cylindrical pins are held in a bundle by periodic spacer grids. These grids have springs that hold the pins in place while accommodating for expansion. Turbulent buffeting vibrations can result in rubbing and bouncing of the spring on the fuel cladding. Over the several years that the fuel pin spends in the reactor, this wear, called fretting, can cause enough damage to result in cladding failure. Fuel cladding failure results in excess fission fragments leaking into the primary coolant and is a breach in the first of three primary barriers between radio-active material and the environment. Next generation reactors exploring alternative fuel designs, such as wire-wrapped fuel pins, are susceptible to the same risks as traditional spacer grid style fuel pins.

Techniques for modeling FIV

The simulation of FIV using CFD has only recently become a topic of interest. The rapid increase in available computational power has made what was previously impossible now feasible. Previously, scaling experiments and analytical models were the primary methods for guiding design. Analytical models are effective for simple designs, but are limited. Experiments are very effective, but expensive and time consuming. Simulation, if proven to be capable, can help to identify problematic vibrations and allow for rapid iteration of design at relatively low expense.

One of the main challenges of simulating FIV is that a solver must be capable of simulating fluid and structural mechanics, as well as allow for both domains to have moving boundaries. Most existing CFD codes don't have these capabilities. To create a simulation method there are two broad options: build a solver from the ground up with the required capabilities, or take an existing CFD solver and a CSM solver and couple them.

Single code methods can take many forms. Most convenient are some multi-physics codes with built in coupling between physical models. There exist some generic PDF solvers using the finite element or finite volume methods which are capable of simulating a variety of physics, such as OpenFOAM [8]. These provide a shorter path to a working simulation than starting from the ground up. Implementation of a moving boundary and the complex boundary condition between fluid and solid is still required. A completely monolithic approach is possible by implementing a method such that the governing equations of the solid and fluid are dictated by a single set of equations. In

this method, the solver implicitly accounts for both sets of physics. This method has the distinct advantage of not necessitating iterations between the two sets of physics. This approach, however requires a new code to be written from the ground up. By doing so, one misses out on the advantage of using long verified methods and extensively validated codes. Due to the tight coupling of the fluid and structure physics, the one-equation approach can prove to be numerically stiff, leading to convergence issues.

A recent method developed specifically for the simulation of fluid-structure interactions is the immersed boundary method [9]. This method, as well as the multiple methods it has inspired, takes this monolithic approach. Its defining characteristic is that it avoids the moving mesh problem by solving the entire domain on a Cartesian grid. The boundary between fluid and solid is tracked separately from the grid and the finite volumes or elements nearby are given a source term to account for this boundary if it is nearby. These Cartesian methods show much promise, but are limited in the order of their accuracy and require more time to be developed and validated.

By coupling existing CFD and CSM codes one has the advantage of building on existing codes. There are a variety of high-quality CFD codes, each with unique advantages, so a developer can start with a CFD code that matches their desire for accuracy, ease of use, open-source, scalability, robustness, etc. In the present study, Nek5000 was chosen largely for its high-order accuracy and scalability. The details of which are discussed further in the numerical methodology section. Likewise the developer can choose a complimentary CSM code. The downside to this method is that

depending on the strength of the coupling required, communication and sub-iteration between the codes can come at a computational and time cost.

Coupling between CFD and CSM codes for FIV can roughly be broken into three catagories: strong 2-way coupling, weak 2-way coupling, and 1-way coupling. Strong 2-way coupling is the most robust and most accurate. The CFD and CSM codes iterate and communicate between eachother on every timestep in order to ensure that they both agree on the positions and velocities of the fluid and structure. Strong 2-way coupling is capable of solving for any kind of vibration, with the limiting factor being the meshes ability to handle the contortions.

Weak 2-way coupling is essentially strong two way coupling without the iteration. The fluid solves for the latest position of the structure, and informs the structure of the resulting forces, which moves accordingly. This can result in faster simulation, but is limited in its accuracy. Excessive vibrations or flows with relatively high-density fluid are unlikely to be solved accurately.

Finally, 1-way coupling is the weakest of the coupling methods, and makes the implicit assumption that the structure movement has no effect on the fluid flow. That is, the fluid communicates the forces exerted on the structure, but runs as a fluid only simulation. This method is very simple to implement, and does not necessitate direct communication between codes. The fluid case can be simulated and the force as a function of time collected for the boundary, then this data can be used as a boundary condition for the structure simulation. This technique is only appropriate for very small vibrations, where the stated assumption is reasonable, such as for turbulent buffeting

induced vibrations. The present study implements a strong two-way coupling approach

for a crossflow case, and a one-way coupling approach for an axial flow case.

CHAPTER II

NUMERICAL METHODOLOGY

Numerical simulation is a powerful tool for investigating engineering problems. Simulations are commonly used for rapid design iteration to complement experiments and prototypes. If proper care is taken to ensure physical accuracy, simulations can act as predictive tools. Simulations can act as numerical experiments, giving insights into a new system at a relatively lower cost than constructing a physical model. Before accepting the results of a simulation, proper background into the numerical methodology should be understood. The numerical methods used for the simulation are what give the results credibility. The following section describes the CFD and CSM codes used for the following chapters, including the key aspects that warranted their selection for each study. Following, the mathematics behind several of the methods used for analysis are described.

Nek5000

Nek5000 is primarily an LES/DNS CFD code developed at Argonne National Laboratory. Its most defining feature is its use of the spectral element method [10]. Nek5000 boasts high order accuracy and scalability to millions of ranks. Nek5000 has arbitrary Lagrangian-Eulerian capabilities, allowing for deforming and moving meshes, which is necessary for some modes of FIV. Nek5000 was selected for the LES of both the crossflow and the 7 pin wire-wrapped fuel pin geometry. Nek5000 was chosen for the crossflow domain mainly for its moving mesh capabilities, and its performance in

capturing transitioning flows. Nek5000 is beneficial for the 7pin geometry mainly for its scalability and computational efficiency. The curvilinear elements help to capture the intricate and difficult wire-wrap geometry, and are forgiving for the necessary deformed and high-aspect ratio elements.

The spectral element method is a high order variant of the finite element method. The mesh is composed of curvilinear hexahedral elements much like a finite element mesh might be. The solution within the elements, however, is described by nth order polynomials with $(n+1)^3$ collocation points distributed according to the Gauss-Lobatto quadrature. The advantage of the high order formulation is that the convergence is exponential in n, meaning that fewer gridpoints are required to resolve smaller wavelengths, which are implicitly present in turbulent and transitioning flow. The lower mesh count requires fewer computational resources, and allows for larger timestepping, compounding into a massive boost in scaling and efficiency.

Formulation

All flows simulated using Nek5000 in the present study feature incompressible fluid flow. The governing equations, incompressible Navier-Stokes, are conservation of momentum:

$$\rho \left(\frac{\partial u}{\partial t} + u \cdot \nabla u \right) = -\nabla p + \nabla \cdot \tau + \rho f \qquad \text{Eq 2.1}$$

and the continuity equation, a simplified form of mass conservation for incompressible flow:

$$\nabla \cdot u = 0 \qquad\qquad \text{Eq 2.2}$$

The flow for both the crossflow and wire-wrap bundle are isothermal and homogenous, meaning constant viscosity. Eq. 2.1 simplifies to:

$$\rho \left(\frac{\partial u}{\partial t} + u \cdot \nabla u \right) = -\nabla p + \mu \Delta u + \rho f \qquad\qquad \text{Eq 2.3}$$

Converting to the non-dimensional Navier-Stokes necessitate several non-dimensional variables:

$$u^* = \frac{u}{U} \qquad\qquad \text{Eq 2.4}$$

$$t^* = \frac{tU}{L} \qquad\qquad \text{Eq 2.5}$$

$$f^* = \frac{fL}{U^2} \qquad\qquad \text{Eq 2.6}$$

$$p^* = \frac{p}{\rho U^2} \qquad\qquad \text{Eq 2.7}$$

Which are non-dimensional velocity, non-dimensional time, non-dimensional body force and non-dimensional pressure respectively. The variable L is the characteristic lengthscale, and the variable U is the characteristic velocity. The non-dimensional Navier Stokes momentum and conservation equations are:

18

$$\frac{\partial u^*}{\partial t^*} + u^* \cdot \nabla u^* = -\nabla p^* + \frac{1}{Re}\Delta u^* + f^*$$

<div align="right">Eq 2 8</div>

$$\nabla \cdot u^* = 0$$

<div align="right">Eq 2.9</div>

Where:

$$Re = \frac{\rho U L}{\mu}$$

<div align="right">Eq 2.10</div>

Re is the Reynolds number of the flow. It is convenient to operate in a dimensionless

framework for consistency between cases. This makes tuning new cases and

troubleshooting easier. The equations are entirely characterized by the single non-

dimensional number Re. Variables are assumed to be in their non-dimensional form for

the remainder of this chapter.

ALE Formulation

Fluid -structure interactions with vibration amplitude appreciable to the fluid

flow necessitate an Arbitrary Lagrangian-Eulerian formulation, with the exception of

methods that use the immersed boundary method or one of its relatives. The Lagrangian

approach associates computational nodes with individual material particles. This is

useful for tracking interfaces and surfaces but is subject to large deformations. This

approach is used for solid mechanics but is not possible for fluid mechanics. The

Eulerian approach fixes nodes in space and tracks the continuum as it moves with

respect to the grid. This approach is typical of fluid mechanics. The Arbitrary

Lagrangian-Eulerian formulation seeks to combine the advantages of both of these
frameworks. Essentially nodes may move according to the continuum or remain in place.
Within the context of fluid simulations, this means that a portion of the mesh may
remain static, while another portion of the mesh adjusts according to a moving mesh.
The Lagrangian aspect operates to follow the same interface and accommodate the
motion by slightly readjusting nodes. Mesh displacements are limited by contortions of
the mesh. A more detailed description of the ALE method can be found in Donea et al.
[11]. The ALE formulation of the Navier-Stokes momentum equation involves adding a
mesh velocity to the convective term. Importantly, the time derivative of u is with
respect to the ALE coordinate system, not the Eulerian coordinate system.

$$\frac{\partial u}{\partial t} + (u - w) \cdot \nabla u = -\nabla p + \frac{1}{Re} \nabla \cdot \tau + \rho f \qquad \text{Eq 2.11}$$

$$\nabla \cdot u = 0 \qquad \text{Eq 2.12}$$

The mesh velocity, w, is introduced to the convective term. The viscous term is
returned to the stress tensor for generality. The non-dimensional form of τ does not
include the dynamic viscosity term. The Nek5000 formulation of the ALE method is
based on the formulation for spectral element method described by Ho [12]. The
complication arises from integrating the equation over an element using the spectral
element test functions. The test functions for velocity and pressure are v and q
respectively. Integrating equations 2.11 and 2.12 over the domain using (.,.) to represent
the inner product yields:

$$\frac{d}{dt}(v, u) = (\nabla \cdot v, p) - \frac{2}{Re}(\nabla v, \tau) + (v, \mathrm{NL}) + c(v, w, u) \qquad \text{Eq 2.13}$$

$$(\nabla \cdot \boldsymbol{u}, q) = 0 \qquad \text{Eq 2.14}$$

where NL is the non-linear convective term abbreviated for simplicity. The term $c(v, w, u)$ arrives when the time derivative is moved outside of the integration because the derivative is with respect to the ALE coordinate system. The new term is equated by:

$$c(v, w, u) := \int_{\Omega(t)} \sum_{i=1}^{3} \sum_{j=1}^{3} v_i \frac{\partial w_j u_i}{\partial x_j} dV \qquad \text{Eq 2.15}$$

Moving the time derivative to the outside of the integration allows for equation 2.13 to be easily converted to the desired numerical scheme.

LES Filtering

The defining factor of LES is that it does not resolve down to smallest scale vortices. Small scale vortices are responsible for the dissipation of energy from the system. While a given LES simulation may dissipate some energy, if sufficient dissipation isn't present, energy will build in the smallest resolved scales of motion giving non-physical results or even crashing the simulation. There are a variety of methods for modeling the dissipation of the subgrid scales. Commonly an eddy-viscosity model is used such as the Smagorinksy model[4].

The filtering in Nek5000 is accomplished by removing energy from the highest frequency modes within each element[5]. The principal can be illustrated on a one dimensional solution. In one-dimension, a given signal is represented by the summation of the contributions from an orthogonal set of N+1 test functions, where N is the order of the test function. The Gauss-Lobatto-Legenedre test functions do not increase in frequency with mode number, so energy cannot simply be removed from a single test function. The equivalent, though, is to interpolate the solution onto a set of (N-1) order polynomials, then back again to N order. Combining the interpolated solution with the original solution in a linear combination yields

$$S_f = \alpha \cdot S_{N-1} + (1 - \alpha) \cdot S_N \qquad \text{Eq 2.16}$$

where S_f, S_{N-1}, and S_N are the filtered, interpolated, and original solution respectively and α is the filtering percentage. By interpolating the solution to a lower order basis function, the highest frequency mode is removed. Linearly combining the interpolated solution with the original removes $\alpha*100$ % of the highest frequency. This process is performed at the end of every timestep with a user defined α removing energy from the smallest scales in order to stabilize the solution. This method essentially acts as a hyperviscosity filter. The baseline recommendation for Nek5000 simulations is to use ~7th order polynomials, and filter the top two modes at 5%.

Method of characteristics

The spacing between gridpoints in the SEM are not uniform within the elements. The grid points cluster near the outside edges. The smallest gap between points shrinks with N^2. This can have detrimental effects on the time-stepping when constrained by a CFL < 0.5, especially when using very high order elements. Nek5000 has implemented a time stepping scheme to get around the CFL limitation based on the method of characteristics. The method of characteristics is a modification of the typical backwards difference scheme for time-stepping. Applying a 2nd order backward difference scheme to the material derivative of velocity produces:

$$\frac{Du}{Dt} \rightarrow \frac{3u^n - 4u^{n-1} + u^{n-2}}{2\,\Delta t} \qquad \text{Eq 2.17}$$

where n is the number of the timestep. The method of characteristics applied to the material derivative of velocity produces:

$$\frac{Du}{Dt} \rightarrow \frac{3u^n - 4\tilde{u}^{n-1} + \tilde{u}^{n-2}}{2\,\Delta t} \qquad \text{Eq 2.18}$$

where $\tilde{u}$ is the velocity at the previous time step at the location from which the point originated before being convected forward, i.e.

$$\tilde{u}^{n-1}(x) = u(X^{n-1}(x), t^{n-1}) \qquad \text{Eq 2.19}$$

where X is position of the current particle before being convected to the current location. This time-stepping scheme does not have a strict CFL limit for stability. The limit to time stepping becomes accuracy of the solution. Typically CFL values of ~2 or 3 are used.

Diablo

Diablo is a multi-physics solver developed at Livermoore National Laboratory. Diablo uses an implicit scheme with finite elements in a Lagrangian framework[13]. The primary mechanics Diablo is designed for are nonlinear structural mechanics and heat transfer. As a multi-physics solver, it is constructed to allow for the addition of other continuum problems. The Diablo solver was formulated with scalability in mind.

The structural mechanics portion of the simulations presented in the following chapters are relatively simple to model, experiencing relatively small displacements and bending. While Diablo is a capable structural mechanics solver, the reason it was chosen for the FSI simulations has to do with its scalability and flexibility for modification. Fully coupled FSI simulations can be extremely computationally expensive, so a fast and scalable structure solution is the priority.

Two way coupling method

Nek5000 and Diablo have been fully coupled together for the purposes of simulating FIV and FSI cases. The codes complement eachother for their emphasis on accuracy and scalability. Additionally they are both coded in the language Fortran, Nek5000 in Fortran 77, and Diablo in Fortran 95. Both codes have existed long before coupling them together, so coupling of the codes requires modification to incorporate the communications between one-another.

Due to the level of accuracy and resolution needed to accurately simulate FSI, the communication between Nek5000 and Diablo is on a shared node basis, rather than interpolation across a communicator interface. Both Nek5000 and Diablo use hexahedral elements, making the interface between the two quadrilaterals. Communication takes place at these shared points, making up the communication interface. Due to a difference in surface description. For each Nek5000 surface which is contains 9 descriptive/communication points, there must be a corresponding 4 Diablo surfaces, which each contain 4 descriptive/communication points. The Fluid and solid meshes are typically constructed simultaneously in order to guarantee a conformal mapping. The solid mesh can perform a simple refinement routine resulting in the proper mapping between faces, then each mesh is exported for its respective solver. At simulation startup the codes identify the communicator points and locate their corresponding cross-code partner. Each processor maintains a map of the communicators generated before the first timestep.

The strong coupling between the structure motion and fluid flow make convergence for a given time-step more complicated than a simple solution cycle of both the fluid and solid codes. Stability can be challenging especially for cases with a high added mass from the fluid. In order to assure stability, an iterative method is necessary.

The scheme employed is an implicit partitioned predictor-corrector scheme like that described by Pereira [14]. At the beginning of a new time step, Nek5000 predicts the structure displacement:

$$p^{n+1} = p^n + \Delta t \left(\frac{3}{2} v^n - \frac{1}{2} v^{n-1} \right) \qquad \text{Eq 2.20}$$

where p is the location of the fluid structure interface, v is the velocity of the fluid structure interface, and n is the timestep count. Nek5000 then predicts the fluid structure interface velocity as:

$$v^{n+1} = \frac{1}{\Delta t} (p^{n+1} - p^n) \qquad \text{Eq 2.21}$$

Nek5000 performs a normal time-step iteration, and passes the force generated at the surface due to pressure and viscous forces across the communicator network to Diablo. Diablo performs a typical single time step calculation to determine the resulting displacement p^{n+1} and communicates them back to Nek5000. Nek5000 uses an under-relaxed estimate of the displacement from Diablo:

$$p^{i+1} = \alpha\, p^{Diablo} + (1 - \alpha)p^i \qquad\qquad \text{Eq 2.22}$$

where i is the iteration count and α is the under-relaxation parameter. Iteration continues until convergence is met such that:

$$\sqrt{\sum_N \frac{\|p^{i+1} - p^i\|^2}{\|p^{i+1}\|^2}} < \varepsilon \qquad\qquad \text{Eq 2.23}$$

where N is the number of communicators and ε is some tolerance. At which point the coupled simulation proceeds to the next timestep. This procedure is outlined by Merzari [15].

Analysis Methods

A number of the analysis methods warrant discussion into the procedure and mathematics behind them in order to better communicate the results.

Taylor microscale

The Taylor microscale is a length scale at which viscous effects in the flow have become significant. It is used as an indication of the approximate resolution at which energy dissipation is significant. The Taylor microscale definition is based on the longitudinal correlation which is given by:

$$R(x + r) = \langle u(x)u(x + r) \rangle \qquad \qquad \text{Eq 2.24}$$

These values are typically normalized to the value at R(x). The lateral two-point

correlation tells how effectively the fluctuations at some point x+r can be determined by

the fluctuations at the point x. This curve is well studied and used to derive both the

integral length scale and the Taylor microscale. The curve has the property that $R'(x) =$

0, and by definition $R''(x) = -\frac{1}{\lambda^2}$ where λ is the Taylor microscale [16]. The Taylor

microscale can be calculated from simulation results by recording three points close

enough together such that $R'(x) = 0$ and use them to calculate $R''(x)$ and thus the

Taylor microscale.

Power Spectral Density

The power spectral density (PSD) describes how the power of a signal is

distributed over frequency. The Fourier transform of a signal is:

$$\hat{x}(\omega) = \frac{1}{\sqrt{T}} \int_0^T x(t)e^{-i\omega t} \, dt \qquad \qquad \text{Eq 2.25}$$

The Fourier transform converts a given signal into frequency space. The power spectral

density is simply estimated as the square of the Fourier transform:

$$PSD(\omega) = |\hat{x}(\omega)|^2 \qquad \text{Eq 2.26}$$

This is true provided the signal is sufficiently long in duration.

Proper Orthogonal Decomposition

Proper orthogonal decomposition (POD) is a helpful method for identifying coherent structures within an oscillating signal based on their energy content. Given a set of snapshots depicting the field, POD determines a set of orthogonal modes that capture the most energy within a system.

A low order approximation of a field can be reconstructed with a linear combination of these modes guaranteeing that the maximum amount of energy possible is represented by this low order reconstruction regardless of the number of modes selected. In this way, the first several modes give a depiction of the highest energy structures that define the field. This gives a physical insight into the behavior of the field while helping to isolate the noise or randomness that might obfuscate.

The method presented here is referred to as the method of snapshots [17]. POD essentially constitutes solving an eigenvalue problem. For a given problem with N nodes and N_S snapshots, A matrix, R, is formed by:

$$R = S^T S \qquad \text{Eq 2.27}$$

where S is and N x N_S matrix consisting of the desired variable at every node n and every snapshot n_s. R is an N x N matrix. The next step is to solve the eigenvalue problem:

$$R\psi = \lambda\psi \qquad\qquad\qquad \text{Eq 2.28}$$

Yielding a diagonal matrix λ containing the eigenvalues and ψ an array of the

eigenvectors. With these, the proper orthogonal modes can be constructed as:

$$\phi_i = \frac{1}{\sqrt{\lambda_i}} S\psi_i \qquad\qquad\qquad \text{Eq 2.29}$$

where ϕ_i is the ith orthogonal mode.

The orthogonal modes are ranked in order to the energy contained. The amount

of energy contained within a given mode is proportional to its corresponding eigenvalue.

The percent of the energy described by a mode can be calculated as the modes

eigenvalue divide by the sum of all eigenvalues. The number of modes required to

capture 50 %, 90 %, or 99 % of the energy is a metric often used to give an idea of how

effective the POD is at reducing the field. For the structure simulations presented in

chapter 4, typically 3-6 modes capture 99 % of the energy, a comparatively low number

of modes.

CHAPTER III

CROSSFLOW THROUGH A TUBE BUNDLE FULLY COUPLED FSI[*]

Introduction

Crossflow through a tube bundle can result in complex interactions between fluid and structure. Experiments and simulations of single cylindrical pins experiencing crossflow are common and well understood. The mechanisms that activate vibrations are relatively intuitive and make for an excellent baseline for understanding and studying flow induced vibrations. The case of a tube bundle, while studied for its engineering applications, also make for an interesting study on emerging complex behaviors from simple changes. When placed in a bundle, the mechanisms of vibration become less intuitive, and highly sensitive to the geometry and arrangement of the tubes. While the nature of the vibrations has been effectively studied using experiments in the past, simulation allows for high fidelity 3-D imaging of the flow for a better understanding of the mechanics behind the vibrations. FIV comes at a high computational cost due to the necessity of LES to capture the periodic nature of the flow, and the full, two-way coupling of fluid and solid physics necessary. Crossflow thus makes for an excellent candidate for a validation study for the coupled Nek5000-Diablo FIV code, providing high scalability and fidelity necessary.

31

Flow induced vibrations can take many forms depending on flow regime, material properties, geometry, and a variety of additional factors. Crossflow through a tube bundle can be characterized as a bluff body interaction. In this family of FIV the fluid impacts a body, creating drag and lift forces which cause the structure to flex. The moving structure affects the fluid domain and pressure field, and simultaneously the structure experiences a restoring force. The interaction and feedback between the changing fluid and structure domain can result in behaviors as simple as a sinusoidal vibration, or as complex as chaotic fluttering.

A variety of researchers have put extensive effort into characterizing flow behaviors possible for simple geometries including the flexibly mounted cylinder. Research approaches such as experiment, to semi-empirical models, to numerical models have been used to study this flow [18]. Modern efforts include the use of Direct Numerical Simulation (DNS) and Large Eddy Simulation (LES) to give high fidelity results with physically accurate simulation [19]–[21]. The complexity of the flow compounds when additional bluff bodies are added to the domain, adding additional degrees of freedom and modes of interaction between bodies. These interaction mechanisms in their simplest form are illustrated by the research conducted on two cylinder domains[20], [22]. This complexity compounds as more tubes are added to form an array or bundle.

Flow induced vibrations from crossflow through tube bundles has been studied in the past primarily be means of experimentation. There has been a plethora of experiments with varying tube diameters, bundle size, arrangement, restoration force,

32

among other key parameters. Experiments of this type were primarily carried out in the 1980s and 1990s. Many of these experiments are detailed and compared by Price [23]. Recently computational fluid dynamics methods have been utilized to study FSI of tube bundles. Schroder utilized k-ω and k-ε RANS models to model the flow induced vibrations and compare to experiment [24]. The results showed that while RANS could predict key parameters such as vortex shedding frequency with low error, it poorly modeled the flow separation that occurs on the downwind side of the pins. For a higher fidelity and more physically accurate look, Shinde utilized LES to model FSI for crossflow [25]. The results showed excellent agreement with experiment.

The experiment on which the simulations are based and validated against were performed by Weaver[26]. The geometry consists of crossflow through an in-line square lattice of pins. Visualization of the flow through the flexibly mounted bundle is given for Reynolds numbers from 110 to 1100. Three distinct vibration regimes are present for the given geometry and flow rates. Between Reynolds numbers of 30 to 150, any vibration that may exist are insignificant. Between Reynolds numbers of 150 to 3875, the vortices in the flow become unstable, and the turbulence produces small amplitude vibrations in the pins. At Re = ~3875, a threshold is breached, and vortex shedding induced vibration increases rapidly with flow rate causing column-synchronized streamwise oscillations. At Re = ~9040, fluid-elastic instability becomes dominant, creating unstable oscillations in both the streamwise and transverse directions. This experiment was chosen for simulation due to its:

1. Detailed characterization of laminar and low Reynolds flow, including distinct changes in behavior at specific velocities.

2. Small domain with well documented setup including natural frequency and logarithmic decrement.

3. Successful capture of three vibrational modes of interest: turbulence induced vibrations, vortex shedding induced vibrations, and fluid-elastic instability.

With successful validation, the coupled code may be used to further study crossflow through a tube bundle with increased confidence that the results are physically accurate.

Methodology

The numerical methodology behind Nek5000 and Diablo were largely discussed in the numerical methodology chapter, importantly including the Nek5000-Diablo coupling scheme and arbitrary Eulerian-Lagrangian formulation. Details specific to the present case are presented and discussed here.

Domain

The geometry simulated is based on the Weaver experiment[26], with a number of simplifications in order to minimize the computational expense while maintaining the physical relevance. Figure 1 depicts the fluid and solid domains, annotating the boundary conditions. The fluid is water at room temperature, and the tubes are
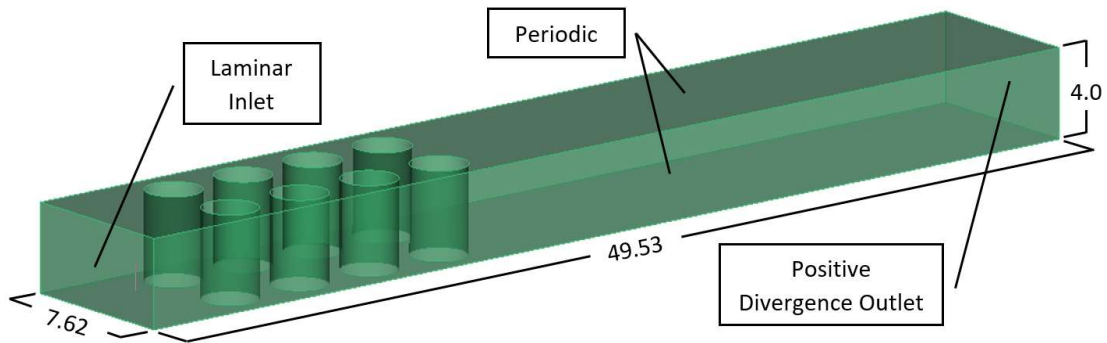
Figure 1. Simulation domain with labeled boundary conditions. Dimensions are in cm.

solid acrylic rods with diameter 2.54 cm. The simulation fluid is given a density of 1.0 g/cm^3 and viscosity of 1E-6 m^2/s. The acrylic has a density of 1.18 g/cm^3. The elastic modulus is set to the arbitrary large value of 2E13 Pascal and a Poisson's ratio of 0.49. The elastic properties of the acrylic rods are assumed to be unimportant, as the vibration movement is due to the spring mounting, not the bending of the tubes.

The Weaver experiment conditioned the inlet such that turbulence intensity < 0.5 %, and velocity variation outside the boundary layer < 1 %, so the simulation inlet condition is a flat, laminar velocity. The inlet is set 4.45 cm upstream of bundle. The outlet enforces the condition div U > 0, preventing backflow with minimal disturbance to the flow. The outlet is set eight pitches downstream of the bundle exit to isolate the wake vortices from outlet effects.

The top and bottom, and left and right sides are connected with periodic boundary conditions. The width and depth were chosen in order to capture flow structures on the order of the tube to tube pitch. It is possible that larger fluid structures are not captured in the simulation. It is also possible that flow structures are artificially

constrained to wavelengths according to the width and depth of the bounding surfaces. It is assumed that the dominant flow structures are those passing between tubes, and that structure behavior in the tube axial direction will have minimal effect on the streamwise and transverse vibrations. However, larger scale simulations may be helpful to verify this assumption. The simulation fluid domain with dimensions and boundary conditions listed is shown in figure 1.

A major concession for using the top and bottom periodic boundary conditions must be acknowledged. The experiment domain pins are approximately 12 diameters in length, and end at their spring mounting at a wall at their top and bottom. The boundary layer and pin edge effects assuredly have some effect on the pin vibration behavior. These effects are not captured in any way by the simulation, and it is difficult to estimate the impact of this simplification. This assumption will be considered when analyzing the results of the simulation and comparing to the experiment.

Connecting the fluid domain to the cylindrical tubes is an FSI-interface condition. As discussed in the numerical methodology section, at this interface the fluid and structure finite elements faces line up, with four diablo finite elements for each Nek5000 spectral element. Nek5000 simply treats this as a boundary condition containing mesh velocity information. Diablo treats this as an external force load at each point.

The experiment configuration, depicted in figure 2, contains many additional tubes, some flexibly mounted, some rigidly mounted. The tubes in the experiment are 30 cm long. The structural domain consists of eight cylinders, highlighted in green in Figure

36

2, with height 4 cm and diameter 2.54 cm, which fit into the fluid domain shown in

figure 1. In the experiment, the tubes are flexibly mounted at their tops and bottom. The

simulation, to mimic the flexible mountings, has four springs mounted on the center of

each tube as depicted in figure 3. The springs do not interact with the fluid or structure,

except to apply a restorative force on the pins. The springs are made to be 10 cm long to

ensure that y direction springs apply negligible force for x-displacements, and vice-

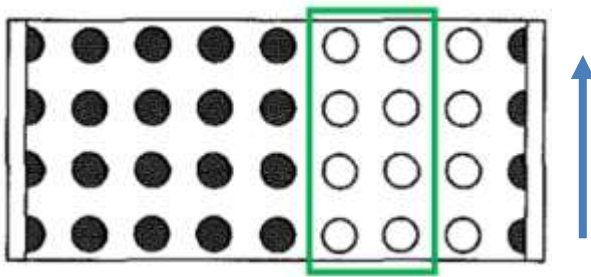versa. The non-tube side of the spring is stationary. The spring constant was chosen to

Figure 2. Top down view of experiment bundle geometry. The green rectangle highlights the portion of the experiment domain captured by the simulation domain. White tubes are elastically mounted and black tubes are rigid.
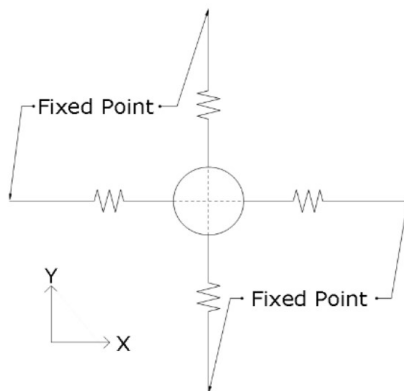

Figure 3. Spring mounts on the simulation pin in order to apply a restorative force equivalent to the experiment.

match the natural frequency in air of 25.5 Hz, as found in the experiment. The tubes are constrained to move only in the x and y direction, and with no rotation over any axis. The tubes are assumed to experience negligible bending and are assumed to remain oriented straight up and down, i.e. no tilting.

Fluid Mesh

The fluid mesh, depicted in figure 4, consists of 29,760 elements. The crossflow portion is 20 elements across. The entire domain is 10 elements thick. Following a brief outlet length from the bundle, the mesh coarsens toward the outlet. Vortices form in the wake of the pins, which may impact the flow in the bundle; so, for a short distance they are well resolved. The coarser mesh following may not resolve the vortices, but rather serves as a buffer for any non-physical outlet effects. Near the tube surface the mesh is finest in order to capture the boundary layer. A y+ value of less than one is achieved for 58% of all solid faces, and less than two for 98%.

Each element is further discretized according to the order of polynomial chosen. A polynomial order of 7 is used for the bulk of analysis, making for a total 15.2 million GLL points. A coarser mesh with polynomial order 9 is also used for verification purposes, containing 29.8 million collocation points. The distribution of GLL points within each element for polynomial order 7 is depicted in figure 4.
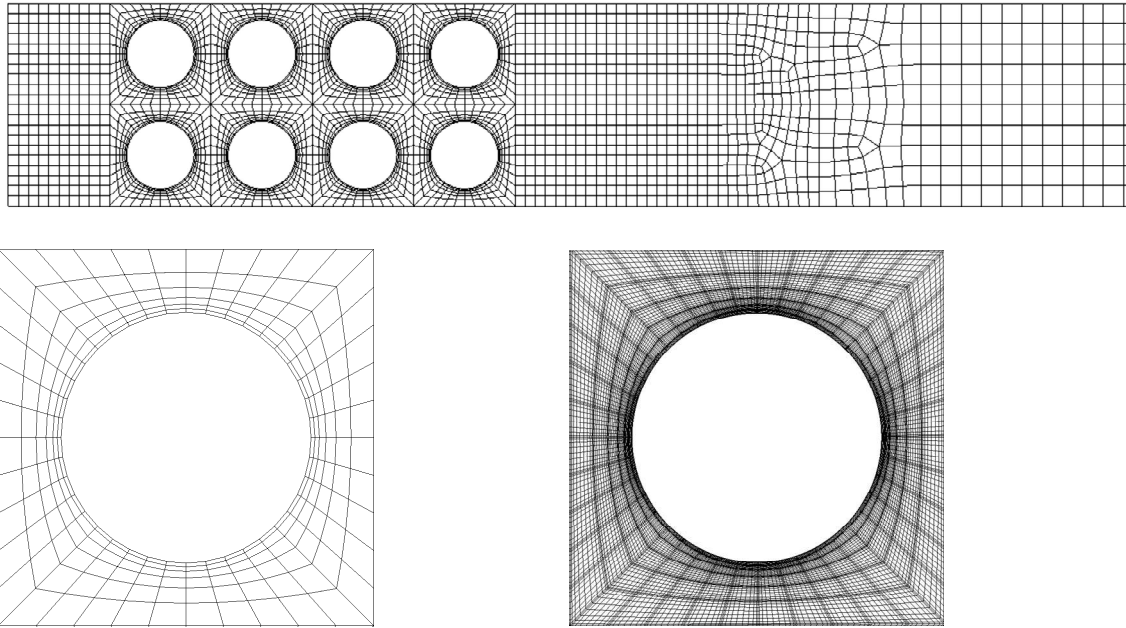
Figure 4. Top: View of entire fluid mesh. Left: Mesh distribution around a tube. Right: Mesh around a tube showing the collocation points for the 7th order spectral elements.
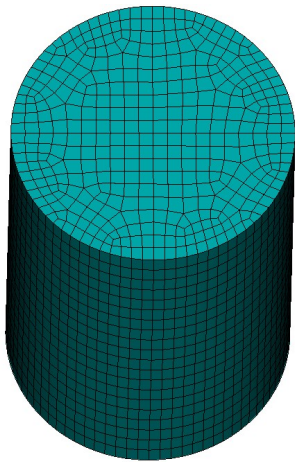

Figure 5. Mesh of a single cylinder of the solid mesh. The curved surface has four faces for every Nek5000 face.

Solid Mesh

The solid mesh consists of 8 identical cylinders which fit inside of the fluid mesh. For each fluid element face at the interface, there are 4 diablo elements. The solid mesh is 80 elements in circumference, and 20 elements in height containing a total of 8180 elements per cylinder. A depiction of a single cylinder of the solid mesh is given in figure 5.

Low Velocity Case

In addition to the study capturing the onset of large amplitude vibrations, the results of a preliminary examination of the onset of turbulent induced vibrations is also presented. Fundamentally, the boundary conditions and mesh construction are similar but differ due to the change in assumptions. The mesh is depicted in figure 6 and contains some minor differences from the previously shown mesh, as less effort was put into reiteration for optimization. Most importantly, due to the relatively minor presence of turbulence due to the lower velocity, the mesh is a single element thick, and uses reflective boundary conditions on the top and bottom, rather than periodic.
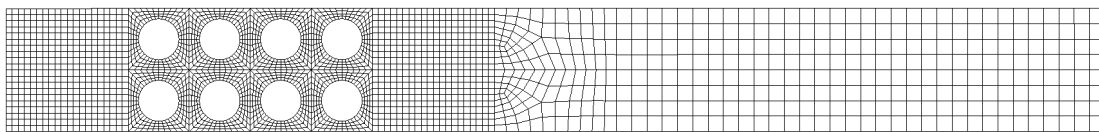

Figure 6. Mesh for low velocity, turbulent buffeting case.

Verification

The primary assumption when using LES to simulate flow is that the domain is refined to the degree necessary to capture the larger, anisotropic, energy-carrying eddies. However, the Kolmogorov theorems are specifically applicable to highly turbulent flow. The flow in the present domain contains portions of turbulence, but also portions that are laminar and transitioning. While mesh refinement can be used to demonstrate mesh independence, by the nature of LES the physics simulated changes with the mesh discretization. A fully resolved LES mesh becomes DNS, defeating the purpose of LES. The approach to verifying the present LES mesh thus consists of two portions. First the flow is simulated using the desired mesh, as well as a finer mesh. Velocity profiles of key areas are sampled and compared to determine if any macroscopic differences appear from the change in mesh size. Second the mesh is constructed in such a way to meet certain scale criteria. For the free stream, the Taylor microscale is computed. The Taylor microscale does not refer to a specific eddy size like the Kolmogorov scale, but rather
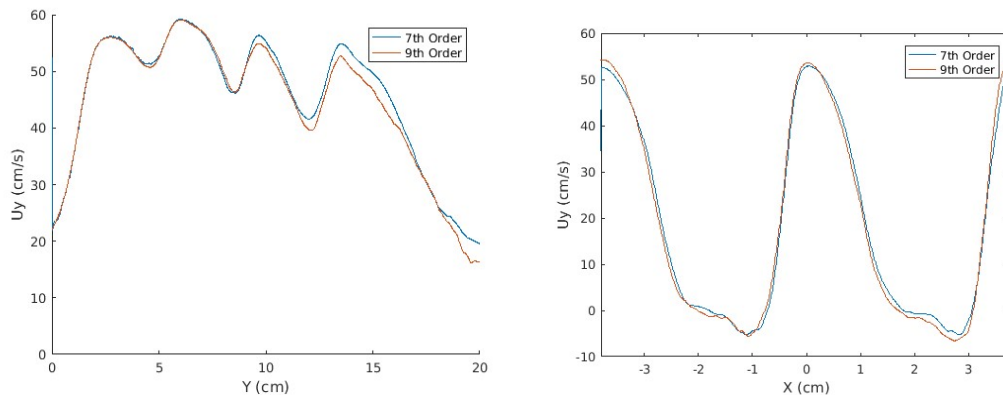

Figure 7. Left: Streamwise velocity as a function of y down the center of the bundle for the 7th and 9th order meshes. Right Streamwise velocity as a function of x across the bundle between the 2nd and 3rd row for the 7th and 9th order meshes.

41

acts as an average of the dissipative scales. Having a mesh size smaller than the Taylor microscale indicates that the LES will resolve some portion of the energy dissipation. The method of computing the Taylor microscale is discussed in the Numerical Methods chapter. Additionally, near the tube surfaces a y+ value of ~1 is ensured in order to resolve the turbulent boundary layer.

Velocity line comparisons are shown for the 7^{th} and 9^{th} order fluid only simulations in figure 7. The lines compared are depicted in the geometry in figure 8. The 7^{th} and 9^{th} order mesh LES results compare well. The velocity plot following the streamwise direction shows indistinguishable error for the laminar inlet and transition to turbulence. Visible difference appears as the flow passes between the third and fourth row. Differences here are < 5%. The flow at this position is highly oscillatory. Discrepancy may be caused by a difference in averaging period. The velocity plot comparison across the bundle also shows excellent agreement with no notable discrepancy.
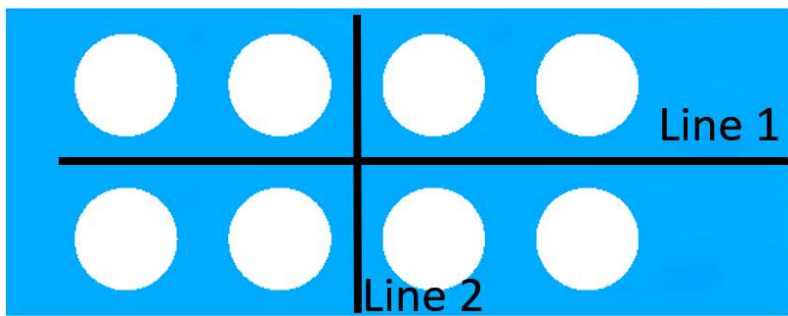

Figure 8. Position of lines used for velocity comparison and Taylor microscale calculation. The direction of flow is from left to right.
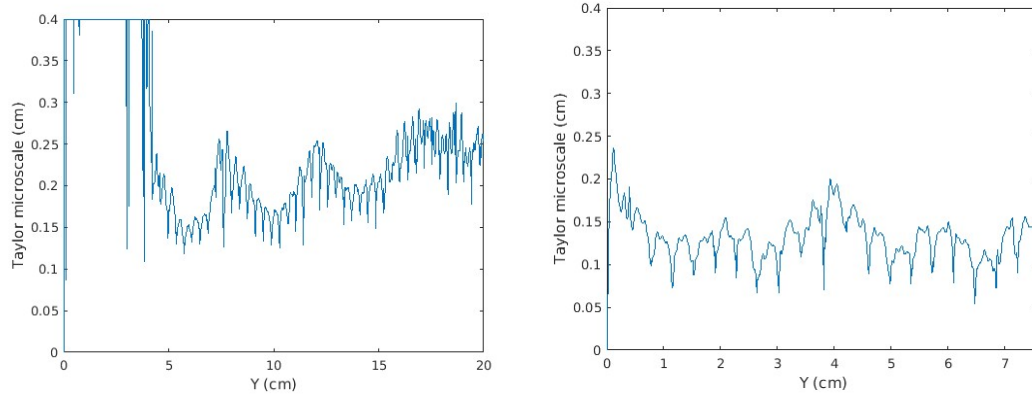
Figure 9. Left: Taylor microscale as a function of position down the center of the bundle. Right: Taylor microscale as a function of position across the bundle between the 2nd and 3rd rows.

On the same segments that the average velocities are measured, figure 9 shows the Taylor microscale as a function of position. The Taylor microscale is calculated for the streamwise direction in both cases.

The flow is laminar for the inlet of the bundle and as it passes between the first row. As a result the Taylor microscale loses its meaning as there is no dissipation. The Taylor microscale then shows a pattern of decreasing as it passes between pins, then increasing as the flow expands. As the space contracts, the velocity increases, increasing the local Reynolds number, resulting in smaller eddies. The Taylor microscale plot across the bundle shows a pattern of dropping in the wake of the pins and rising in the jets in between. The turbulence in the wake of the pins requires higher resolution than the turbulence in the jets between pins. In either case, the Taylor microscale largely stays larger than 0.1 cm, and never dips below 0.05 cm. The coarsest portion of the mesh has

eight GLL points over 0.38 cm, indicating that the mesh is finer than the Taylor microscale, so significant dissipation should be resolved.

In order to verify that the springs acting as a restorative force are operating properly, a simulation with Diablo uncoupled was run. A single cylinder while mounted on the springs was perturbed by a small force then allowed to oscillate freely. There is no fluid interaction, so the pin is effectively oscillating in a vacuum, allowing it to oscillate at its natural frequency. The experiment gives the pins' natural frequency in air to be 25.5 Hz. The added mass by air for a cylinder is approximately the product of the air density and the volume of the cylinder. The added mass by air adds approximately 0.1% of the cylinder mass and can be assumed to be negligible. The pin was displaced 0.02 cm and allowed to oscillate for 10 seconds. The displacement followed an expected sinusoidal curve. A total of 255 oscillation were completed over the course of 10 seconds for a natural frequency of 25.5 Hz as expected.

Validation

The primary motivation of this crossflow study is to validate the Nek5000-Diablo coupled methodology's ability to capture the onset of large amplitude vibrations. Large amplitude vibrations are the result of a positive feedback between the fluid and structure that maintain energy into the vibration frequencies despite the natural dissipation of energy. Thus, capturing the onset of this behavior demonstrates not only the simulations ability to capture the specific mechanism responsible for vibration, but that the fluid and structure are interacting properly to capture the shift. Essentially, the onset is a more

scrutinizing test. Adding to the difficulty, the fluid simulation involves laminar, transitioning, and turbulent flow. Transitioning flow is notoriously difficult to capture.

The validation encompasses multiple regimes of the experiment. First particle imaging snapshots of the experiment are compared with instantaneous vector fields from the simulation. The snapshots encompass the transition from purely laminar flow to the turbulent flow present at the onset point. These snapshots were selected to show how the flow behavior evolves as turbulence emerges. This serves as a background to understanding the behavior of fluid in crossflow.

In order to capture the onset of the high amplitude vibrations, simulations were performed at three velocities surrounding this onset point. The vibration amplitude response for the streamwise and transverse directions as a function of velocity serves as the primary comparison between the simulation and experiment. Additionally, frequency spectra are compared between the simulation and experiment.

Development of Turbulence

The publication by Weaver contains flow visualizations for average inlet flow velocities of 0.43 cm/s, 0.59 cm/s, 1.9 cm/s, and 12 cm/s. Each image was specifically chosen to illustrate a marked change in behavior. The experimental snapshots are compared with corresponding instantaneous flow fields from the simulations in figure 10. For each flow velocity, the behaviors in simulation and experiment are described.
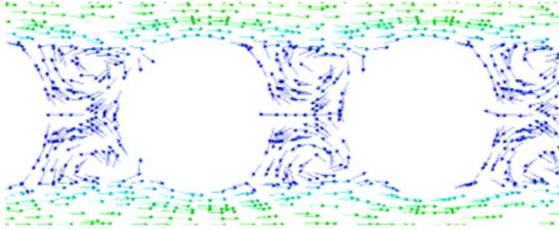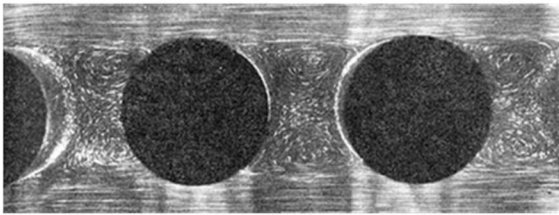
0.43 cm/s

For low velocities, the flow through the channels formed by the pins is laminar. The wakes behind the pins form symmetric, static vortices. The flow shows no change as it passes through the bundle. No significant difference between simulation and experiment is observable.
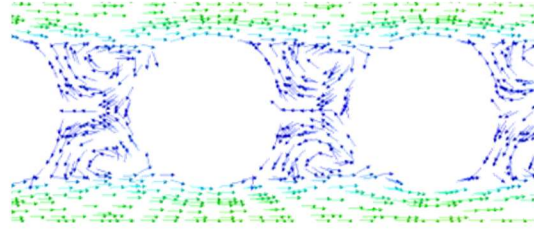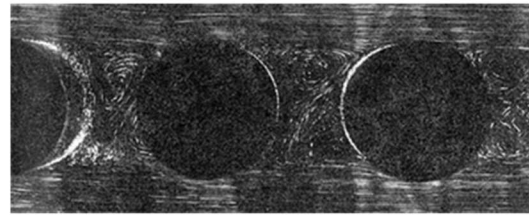
0.59 cm/s

For slightly increased flow velocity, the flow through the channels remains laminar. The difference the slight increase in velocity makes is the introduction of instability in the vortex behavior. The columns of high-velocity flow entrain the vortices, perturbing both, and causing asymmetric, alternating vortices. These perturbations are not enough to cause significant pressure disturbances and vibrations but are a visible development. The asymmetries in the vortices are apparent in the vector plot for the simulation but less easily observed in the experiment imaging.

1.9 cm/s

In the experiment, the increase in velocity causes the instabilities in the vortices to break down into turbulent action. The turbulent behavior interferes with the columns of high velocity flow, causing instabilities. The result is a cascade of increasingly perturbed flow through the bundle. The simulation indicates some coherent structures in the wake of the first two pins which does not seem to appear in the experiment.

0.43 cm/s

0.59 cm/s
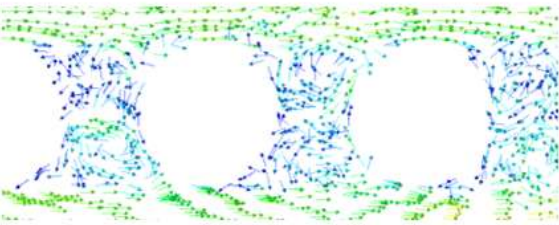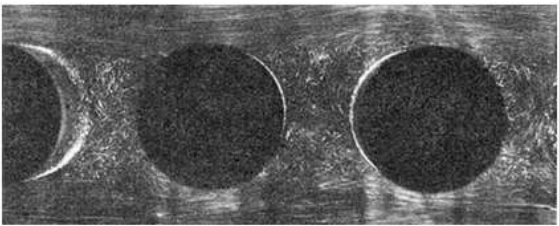
1.9 cm/s
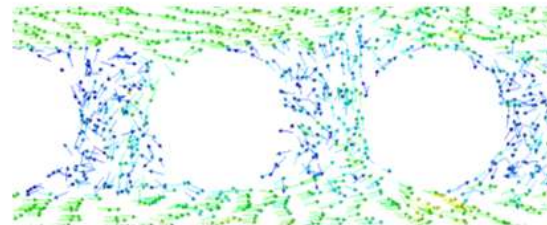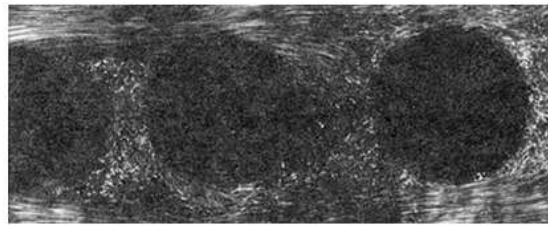
12 cm/s

Figure 10. Visual comparison of experiment snapshots with simulation vector diagrams.

12 cm/s

In the case of high-velocity flow, the previously observed turbulent wakes have become fully developed turbulent action and bleed heavily into the main column flow. The flow is laminar at the inlet to the bundle, but it steadily devolves into increased turbulent action as it passes between the second and third rows of pins. In both the experiment and simulation, the vortices can be seen to push out into the main flow, skewing the flow to an angle. For this velocity the experiment exhibits very small amplitude vibrations that are not present in the fluid-only simulation

For the preliminary geometry, a single case of fluid-structure interaction was simulated for comparison with the experiment. The flow velocity simulated was 6.8 cm/s. This puts the flow in range for small but measurable vibrations. The graphical comparison can be seen in figure 11 alongside the higher velocity cases. This case establishes that the simulation can capture purely turbulent buffeting vibrations before proceeding onto the transition to larger vibrations.

Onset of large vibrations

The publication by Weaver gives the RMS amplitude of vibration for pins in the 2^{nd} and 3^{rd} row as a function of inlet velocity. Plots are given for both streamwise and transverse vibrations. The inlet velocity spans from 6 to 50 cm/s. The curve resembles that of a single cylinder. At low velocity turbulent buffeting induced vibrations are present with very low amplitude. At some velocity, streamwise vibrations increase dramatically, peak, then begin to decrease, forming a hump. This is characteristic of

vortex induced vibrations. That is, there is some oscillating fluid behavior whose frequency is dictated by the inlet velocity. At velocities for which the oscillations coincide near or with the structure's natural frequency, the vibrations gain amplitude, causing a peak. As velocity increases further this resonance decreases and amplitude falls. Finally, fluid elastic vibrations begin, and amplitude grows indefinitely with velocity. The initial onset of large magnitude vibrations is of great interest to designs including tube bundles.

Three velocities were chosen for the coupled simulation to straddle the onset velocity. The first velocity chosen is just below the onset of the large vibrations, where the amplitude remains in line with that expected of turbulent buffeting. The second is shortly after the onset, where the amplitude is still small, but obviously not due to turbulent buffeting alone. Finally, the third velocity is chosen such that the vibrations are large, approximately half of the local peak in amplitude. These velocities are 15.3 cm/s, 18.1 cm/s, and 20.8 cm/s respectively.

Figure 11 gives the comparison between the simulation and experiment for the three velocities chosen. Additionally, for context, all the data from the experiment and the datapoint for the low Re simulation are included. The data shown is for the streamwise and transverse RMS amplitude for pins in the second and third row.
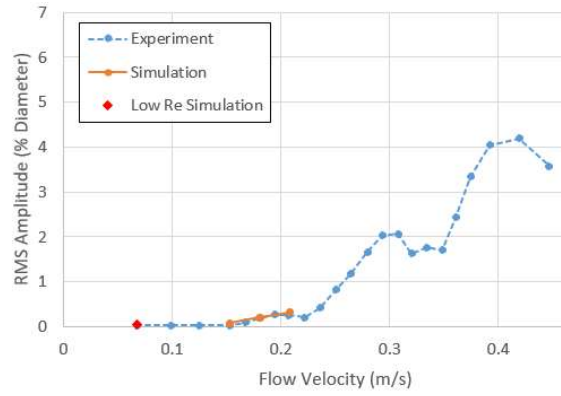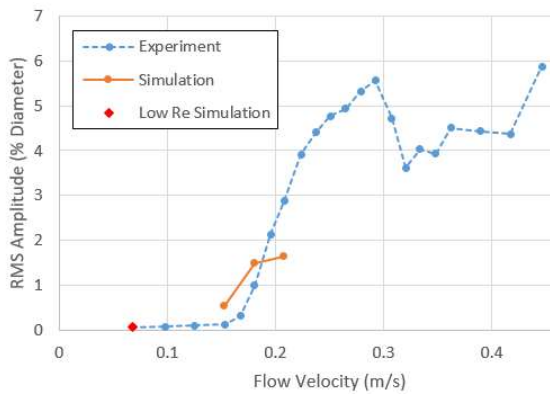
a. Streamwise RMS amplitude for 2nd row

b. Transverse RMS amplitude for 2nd row
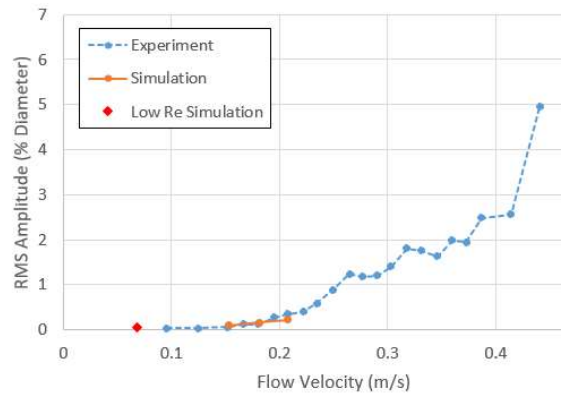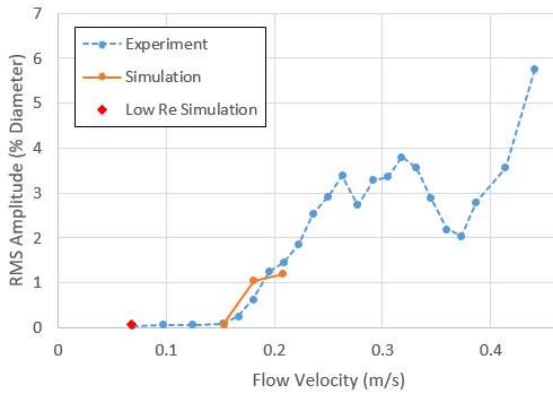
c. Streamwise RMS amplitude for 3rd row

d. Transverse RMS amplitude for 3rd row

Figure 11. RMS amplitude comparison plots for experiment and simulation.

The streamwise vibrations in the second row (Figure 11a.) increase the most with velocity. The simulation predicts a sharp increase in vibration amplitude over the course of the three velocities. However, the simulation increase is not as pronounced as the experiment. The simulation predicts a higher amplitude for the lowest velocity data point, which apparently is not experiencing turbulent buffeting alone. The middle velocity data point for the second row is proportionally higher than the first point. The highest velocity data point, however, does not increase proportionally, but rather shows a marginal increase in amplitude over the second point.

The streamwise vibrations for the 3^{rd} row tube show better agreement between simulation and experiment. At the lowest velocity the simulation shows small vibrations, meaning that the onset of large vibrations has not been reached. The second velocity, like the second row, shows substantially greater amplitude than experiment. The final velocity, again like the second row, shows only a marginal increase in amplitude. While the onset point seems to have already been reached for the pin in the second row, this is not true for the third. In the experiment, the onset of large amplitude vibrations is slower, and slightly delayed for the third row.

The transverse vibrations for both the 2^{nd} and 3^{rd} rows show little change with the velocity increase. In each case, the increase in amplitude, while greater than expected for turbulent buffeting is marginal compared to the streamwise vibrations. The simulations show the same behavior as the experiment with minute discrepancies.

The frequency spectra for the simulations and experiments are shown in figure 12. The simulation spectra are shown for 15.3 cm/s, 18.1 cm/s, and 20.8 cm/s, the same as the

vibration amplitude data points. The experiment spectra are shown for 15 cm/s and 19 cm/s. The velocities do not match perfectly, but the behavior is consistent. Due to the relatively short simulated time, the frequency spectral analysis must use relatively course frequency stepping: 0.9639, 1.1484, and 1.2994 Hz respectively. As such, the peaks are not well defined to specific frequencies, but the strongly peaked results make identification easy. Additionally, as the pin oscillations are nearly sinusoidal for most cases, calculation of the frequency based on wavelength is possible.

The 15.3 cm/s simulation shows a dispersed streamwise frequency spectra with a strong peak around 15 Hz. Though not as distinct, ~12 Hz and ~18 Hz both show small peaks. Comparatively, the experiment streamwise spectra for 15 cm/s show distinct peaks at 15.7 and 18.6 Hz. The 18.1 cm/s and 20.8 cm/s simulation streamwise spectra both show distinct single peaks at around 15 Hz, matching well with the distinct single peak at 15.7 Hz for the experiment at 19 cm/s. By measuring the wavelength of the oscillations of the two higher velocity cases, the peak is determined to be at 14.9 Hz. The transverse spectra in each case, like the amplitudes, show little of interest. Small peaks at 14.9 Hz can be found for the simulations and 15.7 for the experiment.
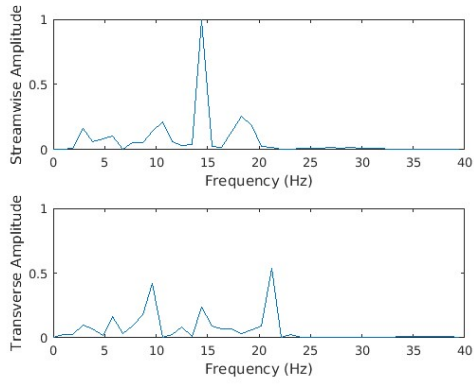
a. Simulation frequency spectra, 15.3 cm/s
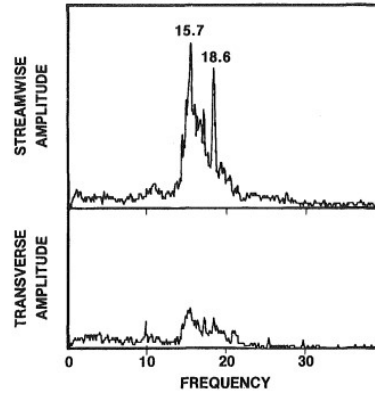
b. Experiment frequency spectra, 15 cm/s
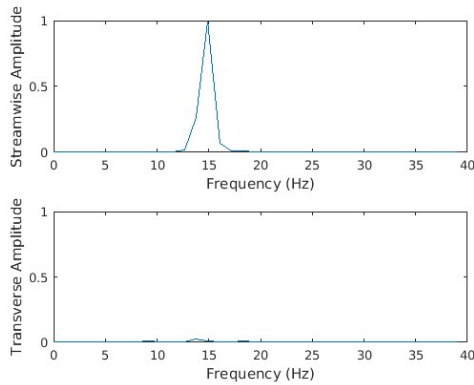
c. Simulation frequency spectra, 18.1 cm/s
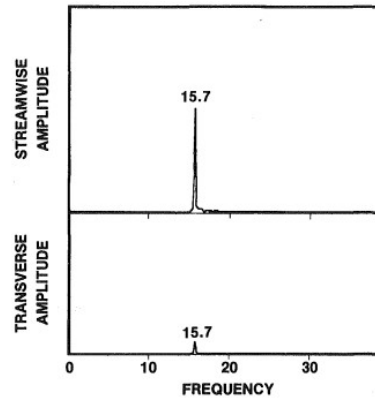
d. Experiment frequency spectra, 19 cm/s
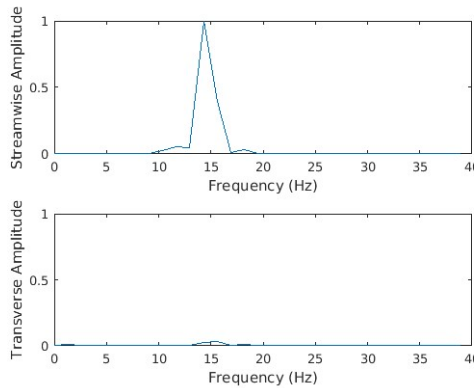
e. Simulation frequency spectra, 20.8 cm/s

Figure 12. Frequency spectra for simulation and experiment.

The sharp contrast between the lowest velocity spectra and the two higher velocity spectra is indicative of the change from primarily turbulent buffeting induced vibrations, to the onset of the flow structure based vibrations. Turbulent buffeting vibrations should show a peak at the structure's resonant frequency, but incoherence otherwise, as observed for both experiment and simulation. As the periodic fluid structure imparts more energy into the motion, the amplitude increases, and the near sinusoidal motion dominates the frequency spectra, as the tube essentially becomes a simple harmonic oscillator.

There are many possibilities as to why, but the resonant frequency of the tubes in the simulation is slightly lower than the resonant frequency in the experiment. Likely sources of this discrepancy stem from the assumptions made for domain simplification. While the experiment maintains a laminar inlet, there are necessarily boundary layers at the top and bottom of the experiment chamber where the velocity profile is not flat. The effect of the boundary layer at the top and bottom of the vibrating tubes is not present in the simulation. Likewise, edge effects at the top and bottom of the tubes in the experiment are not captured. Another possible source of the discrepancy is the assumptions made about the restorative force. The simulation assumes that the restorative force behaves like a spring, i.e. the force is proportional to displacement. No logarithmic decrement is included. The behavior of the mounting for the experiment may behave in a more complicated manor.

The early onset of the increased amplitude vibrations in the simulations (as seen for streamwise vibrations in the 2nd row) and the lower dominant frequency of the

experiment are likely related. If the resonant frequency is lower, then the flow structure responsible for the onset of larger vibrations will activate the resonant frequency at a lower velocity. As previously discussed, vibrations related to oscillating flow structures exhibit a peak of amplitude at a certain velocity which corresponds to the resonant frequency. If the resonant frequency is lower, then it will be activated by a lower velocity. Therefor the onset velocity will be proportionally lower. This may explain the seeming early onset of large vibrations found in the simulation.

Analysis

The position that the center of the pins trace is helpful for understanding the nature of the pin vibrations. The amplitude and frequency data suggest that the pins oscillate primarily in the streamwise direction and that the streamwise and transverse oscillations are synchronized. Details like the angle of vibration, and relationship with neighboring pins can point to flow mechanisms for further study. The path traced by all pins for the cases straddling the onset point are depicted in figure 13. The displacements are amplified by a factor of 15 for visualization purposes. The neutral positions of the pins are depicted as for context and clarification by a black dot. The flow is in the positive y direction.

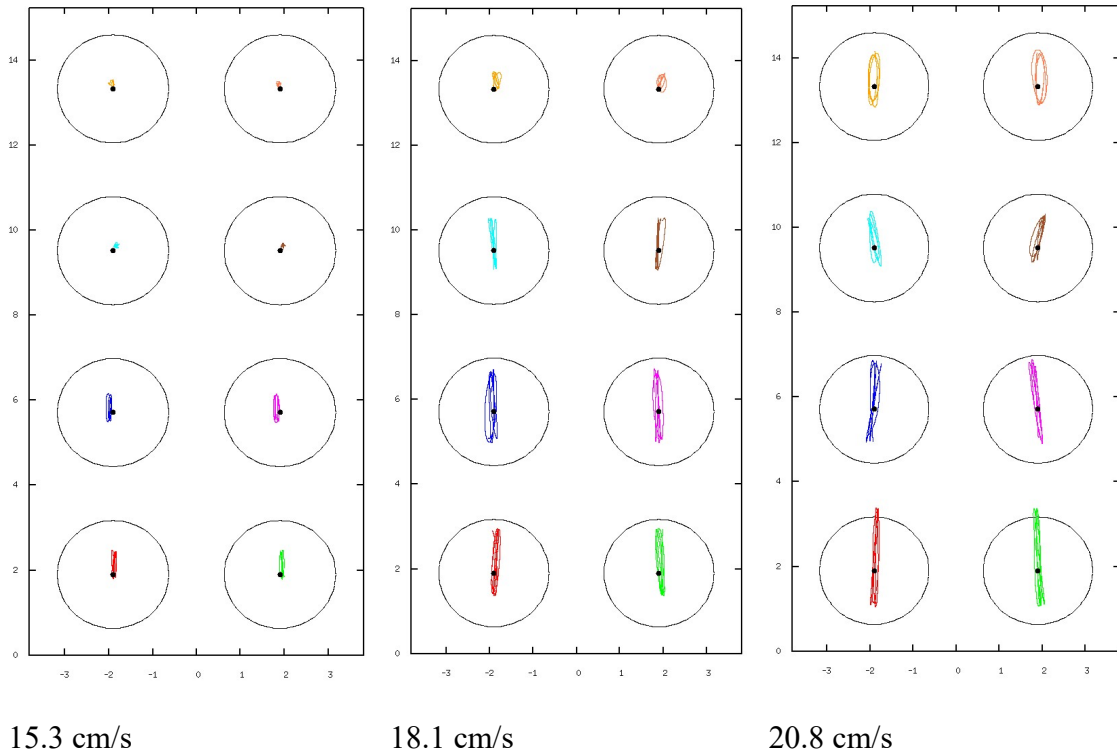15.3 cm/s   18.1 cm/s   20.8 cm/s

Figure 13. Position traced by pin centers, amplified by a factor of 15. The black circles represent the neutral positions of the pins, and the black dot the neutral center.

In each case, the pins settle into a new average center position when the fluid flow is present. A displacement in the y direction is expected regardless of the velocity of the flow, as naturally there will be pressure drop across the pins. This displacement in the y direction does not vary significantly across the different velocities. The velocity differs by around ~10% between cases, and the pressure drop is caused primarily by the drag, which does not change significantly at the onset of the larger vibrations. Displacement in the x-direction, however does give information on the flow structure and vibrations. For the lower velocity case, below the onset velocity. The pin centers have arranged themselves in an alternating pattern as a function of row number. The first row settles to the left, second to the right and so forth. The third and fourth row of pins experience very small oscillations compared to this displacement, indicating that a steady force, which is not related to the oscillations, is responsible. The second row would likely behave the same if not for the premature onset as discussed previously. This steady force is indicative of a non-symmetric flow structure. The pattern of the displacement alternating between rows indicates that this non-symmetric flow structure switches back and forth as it passes through the bundle. The higher velocity cases show that as the oscillations increase, the displacement in the x direction diminishes, as the displacement passes through the original neutral center. The middle velocity fourth row pins' behavior supports this observation. While the rest of the bundle has switched to the larger vibrations and the diminished x displacement, the fourth row has smaller vibrations and a distinct x displacement. The fourth row of the higher velocity case,

57

however, shows the fourth row as it developed larger amplitude vibrations has returned to the diminished x-displacement.

While the x-displacement appears to diminish as the vibration amplitude increases, this transverse behavior appears to be replaced by a slight angle of vibration. The periodic structure applying the force must have some asymmetrical component. This is hinted at when looking at the time averaged velocity profiles of the fluid only simulations (Figure 7) as they show an asymmetrical distribution. The lower velocity case shows no recognizable angle in their oscillations, even for the prematurely activated second row. The degree of the angle of vibration appears to grow as the amplitude increases. This may explain why the transverse amplitude increases later than the streamwise amplitude in the experiment results presented in figure 11 previously.

As the velocity increases, it appears as if the onset of larger vibrations develops in the pins in a sequence from front to back. At the lower velocity, only the first two pins show significant amplitudes. For the middle velocity, the first three rows show large vibrations. Finally, at the highest velocity, all four rows show large vibrations. The flow pattern responsible for the vibrations appears to develop from the inlet toward the outlet.

Finally, the fourth row vibrations, only significant for the highest velocity, show characteristically different behavior than the first three rows. The transverse vibrations are noticeably larger and more consistent than the previous three rows. This difference is likely caused by nature of being the last row. The wake behind the fourth row is unobstructed, allowing for traditional vortex shedding to form, likely contributing to the transverse oscillations one would expect from vortex shedding.

15.3 cm/s   18.1 cm/s   20.8 cm/s
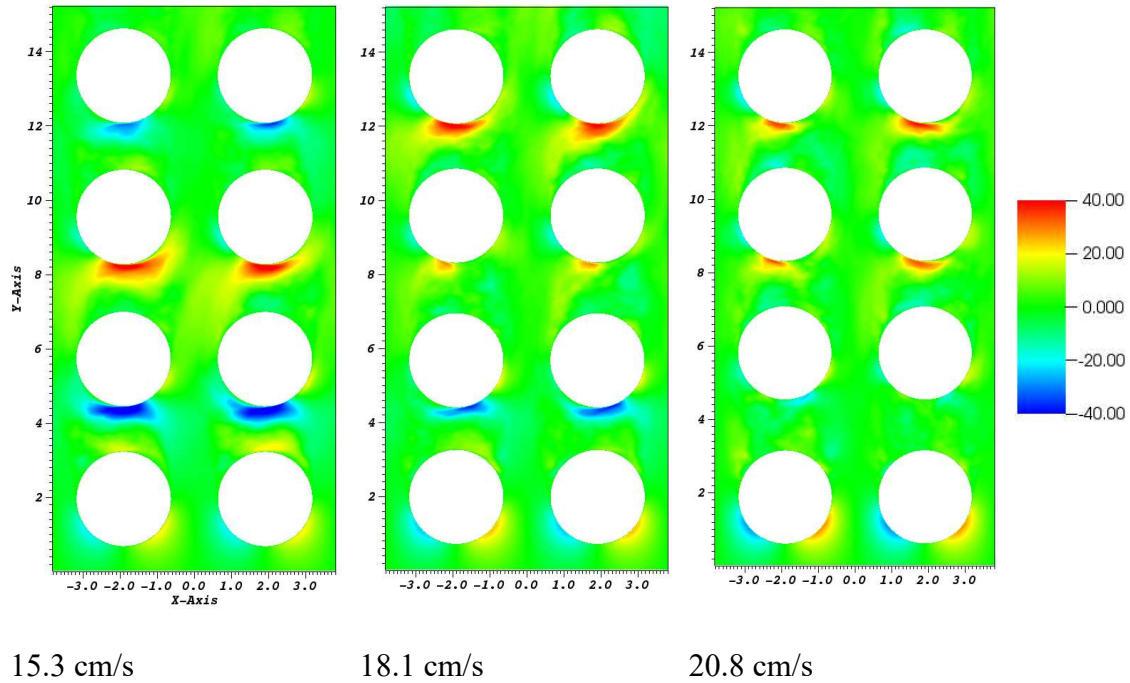
Figure 14. Transverse velocity for each simulation case. Units are in cm/s.

The time averaged transverse velocity gives insight into the flow structure development as velocity increases. The time averages presented have not reached statistical stationarity. Due to the high computational expense of the FIV coupling, reaching statistical stationarity was not feasible for the present study. Time averages should be observed with this in mind and used only for qualitative analysis. The time-averaged x-velocity is shown for each velocity in figure 14. The most apparent development is the decrease in time-averaged x-velocity directly beneath the pins as velocity increases. The lowest velocity case has a persistent stream passing beneath the pins in the second and third row, seemingly snaking through the geometry. As the flow velocity increases, these permanent fixtures disappear, implying that the direction of

59

flow directly beneath the pins oscillates rather than stays fixed. This behavior explains

the consistent x-displacement for the lowest velocity case. The transverse flow behavior

of the higher velocity cases likely has a more oscillatory behavior that is not captured by

the averaged velocity.

CHAPTER IV

7-PIN WIRE-WRAPPED FUEL BUNDLE ONE-WAY COUPLING[†]

Introduction

Wire-wrapped fuel pins, unlike more common bare fuel pins, generate a comparatively strong transverse flow in proportion to the helical pitch of the wire. Transverse flow promotes heat transfer and inter-channel mixing, meaning cooler hot-spots. However, tied to this secondary flow is a non-uniform and fluctuating pressure distribution of greater magnitude than for bare pins. The more intricate pressure field gives rise to greater bending moments, as well as concern for the generation of vibrations. For the thin fuel pins, bending can be a structural design limitation. Regarding vibrations for wire-wrapped fuel pins, which use wires to maintain spacing between pins, the non-uniform pressure field and lack of spacer grids gives reason to believe that vibrations may be of greater concern than for bare rods. Identifying the magnitude of these concerns, as well as the mechanisms behind them, is necessary for the future design of reactors using wire-wrapped fuel pin bundles.

In general, axial flow is less prone to vibration than crossflow. The only significant source of vibration for flow conditions that would reasonably exist in a nuclear reactor are turbulent buffeting induced vibrations. These vibrations are generally

---

[†] A portion of this chapter is reprinted with permission from "Simulation of Fluid-Structure Interaction of Crossflow Through a Tube Bundle and Experimental Validation," by L. Brockmeyer, J. Solberg, E. Merzari, Y. Hassan. 2017. *Proceedings of ASME 2017 FEDSM*, Copyright 2017 by ASME.

small and grow steadily with Reynolds number. There is no concern for a resonance emerging at an unknown flow condition. Small vibrations allow for the assumption that the structural motion does not affect the fluid flow. Therefore, full, two-way coupling is not necessary. A fluid flow simulation can be run to gather the force applied on the surface, then the surface force can be applied to a structural simulation as a boundary condition. This one-way coupling is dramatically less computationally expensive than two-way coupling.

Wire-wrapped fuel pins have been of interest for sodium fast reactors for several decades. Early on, studies were performed experimentally, and design was aided by subchannel codes. With the increase of available computational power, computational fluid dynamics codes have been used to study the flow field in ways that are difficult to observe experimentally. Efforts have been primarily focused on determining the effects of changes or simplifications of the geometry and determining the effectiveness of modelling. Hamman and Berry[27] and Raza and Kim[28] used RANS modelling to observe the effects of wire shape, finding that the structure of the flow remains the same, but parameters such as pressure drop and turbulent kinetic energy are sensitive to the changes. Merzari[29] using RANS and LES determined that the geometry at the base of the wire (i.e. the connection between pin and wire) has little effect on the flow field. Pointer[30] and Brockmeyer[31] used RANS and LES to study the effect of the bundle size on the flow behavior, finding the flow behavior to fundamentally change from smaller to larger bundles.

Fischer was the first to model the wire-wrapped fuel rod bundle using Large

Eddy Simulation (LES), providing more detailed flow resolution for a 7 pin bundle[32].

Fischer's LES results agree with the RANS modeling of Ahmad & Kim. Direct

Numerical Simulation (DNS) of the geometries of interest is not possible at the time of

writing due to the large computational expense. However, Ranjan performed DNS on a

simplified geometry to observe in great detail the flow behavior as it passes over the

wire [33]. The geometry consists of a pin attached to a wall in a channel with angled

flow. The results indicate that a separation region develops on the leeward side of the

wire. This region shows large deviations from the law of the wall behavior. Thus,

models that utilize wall functions may not accurately model the flow near the wire.

While many studies observe the pressure field as a means of analysis, few

specifically take a structural point of view. Raza and Kim[28] observed, like many

others to follow, a large radial pressure gradient across a horizontal slice of 7 and 19 pin

bundles. Also observed is the surface pressure field on center and edge pins of 7 and 19

pin bundles. Findings include that the pressure is lowest where the wire blocks the gap

between neighboring subchannels. The researchers noted that the pressure distributions

on the edge fuel pins were qualitatively and quantitatively different than for the center

fuel pins, while the pressures for the respective 7 and 19 pin bundles are qualitatively

similar. Jeong [34] took a detailed look at the pressure distribution on the wire

specifically, comparing the driving force on the wire in the x, y, and z direction as a

function of height for a center and an outside pin. The authors emphasized that the closer

the pins are to the shroud, the greater magnitude driving force in all three directions the

wires experienced. The corner pin experienced forces up to 263% of those in the center subchannel.

The primary vibrations of concern for wire-wrapped fuel pins are expected to be those from turbulent buffeting, as is true for bare rods with spacer grids [35], [36]. The present study seeks to extend from observing the pressure effects on walls, to using the forces from the fluid simulation as the forces imposed on a computational structure simulation. This coupling allows for the direct analysis of the turbulent buffeting vibrations. This process entails simulating a 7 pin bundle using LES as implemented by Nek5000, similar to that done by Merzari [37]. The force data is collected from the simulation then used to synthesize longer signal histories to be used as body forces exerted on a structure simulation of the seven-pin bundle. The structure simulation is repeated several times using the same input data while changing how the pins are coupled with one another. Comparison between cases helps to inform on the nature of the energy transfer between pins in contact, as well as the effect on vibrations of gaps between pins. This FIV simulation is the first to be applied to the wire-wrap geometry.

Methodology

The implemented one-way coupling method uses Nek5000 and Diablo in their original forms, with no modification to their underlying code for communication. The numerical method chapter largely covers Nek5000 and Diablo's numerical methods. The present LES simulation implements the Pn-Pn pressure scheme, and uses the method of characteristics, allowing for CFL values ~2. The flow is turbulent with a Reynolds

64
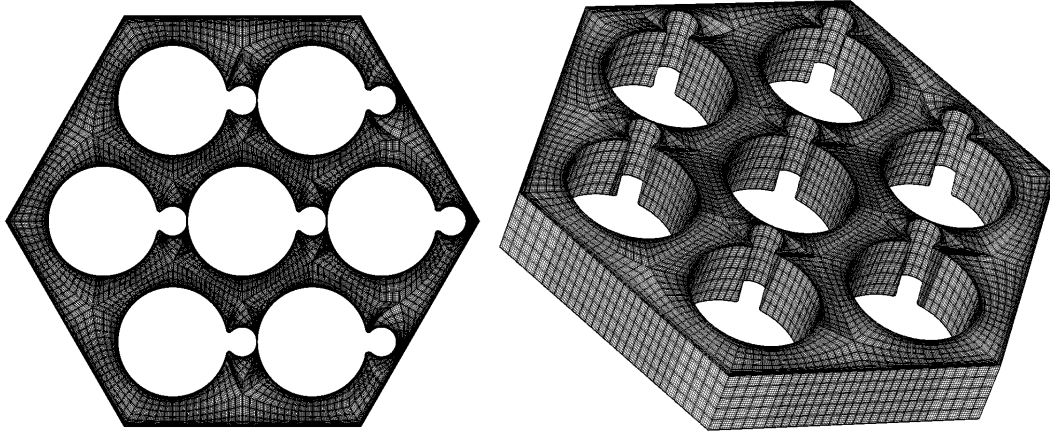
Figure 15. Fluid domain and LES mesh. The domain extends 26.5 cm, or 40.45 diameters.

Table 1 Geometric Parameter Values

Parameter	Length (mm)
Helical Pitch (P_H)	265
Pitch (P)	8.35
Pin Diameter (D)	6.55
Wire Diameter (D_W)	1.75
Fillet Diameter (D_f)	0.35

number of 26,300 according to hydraulic diameter as a lengthscale. The fluid simulation

is run for a single flow-through time of the full helical pitch after reaching convergence

to fully developed flow.

Domain

      The fluid and structure domain compose a 7-pin wire-wrapped fuel pin bundle. A

depiction of the fluid and structure domain are shown in figures 15 and 16 respectively.

The parameters defining the geometry are listed in Table 1. While geometrically they

complement each other, the fluid domain is shorter than the structure domain at one

65

Figure 16. Structure domain and mesh.

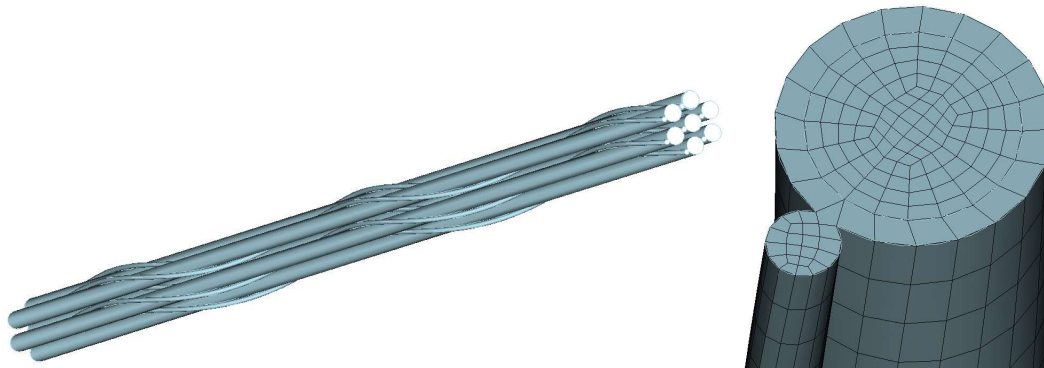pitch and 2.64 pitches respectively. The structure length was chosen to match that of an experiment performed by De Pauw [38]. The fluid domain is limited to one pitch in order to minimize computational cost. The signal syntheses scheme used allows the single fluid pitch to inform force inputs for the entire structure body. The fluid, lead-bismuth eutectic (LBE), is simulated at 573 K, has a density of 10337 kg/m^3, and a dynamic viscosity of 1.717E-3 Pa·s. The structure, stainless steel (316L) has a density of 799 kg/m^3, a Young's modulus of 177 GPa, and a Poisson's ratio of 0.25.

The mesh for wire-wrapped fuel pins can be challenging to generate due to the repeated instance of point contact of a wire to the neighboring pin, as well as line contact of the wire to its corresponding pin. The wire is attached to the pin using fillets in order to avoid meshing the small space where the pin touches the wire. According to Merzari[29] geometric simplifications near this contact point have negligible effect on the flow field. The wire is slightly smaller than the gap between pins, so the wire leaves a small gap as it passes a neighboring pin. As long as the velocity in this small gap is approximately zero,

this simplification is not expected to have an effect on the flow field. These or similar geometric compromises are ubiquitous and necessary in wire-wrap bundle modeling.

The fluid domain boundary conditions consist of walls for the surface of the wire-wrappers and pins, as well as around the outside for a hexagonal shroud. The inlet and outlet are periodic boundary conditions, implicitly implementing fully developed flow. Wire-wrapped bundles restore to fully developed flow in a shorter span than a similar bare bundle, conservatively requiring a full helical pitch. However, for the inlet of a fuel bundle this is a poor assumption. As the fuel pins are mounted at the inlet, but not on the outlet, they essentially act as cantilevered beams. The vibrations are expected to increase in amplitude with distance from the inlet. The fully-developed assumption is valid for the portion of the domain expected to have the most problematic vibrations.

Fluid Mesh

The fluid mesh, depicted in figure 15 is made of 1.296 million elements. Being spectral elements, they are further discretized according to the polynomial order. Using 7^{th} order polynomials brings the number of collocation points to 663 million. The element distribution near the pins was selected to ensure that the growth ratio is less than 1.5 while making the elements near the boundary as thin as possible in order to resolve the boundary layer. The y+ value near the wall is less than 2 for 99% of wall faces, and less than 1 for 68% of wall faces.

Structure Mesh

The structure mesh, depicted in figure 16, is identical for each of the 7 pins. The hexagonal shroud surrounding the bundle is not included in the mesh and is not simulated. Each fuel pin contains 55100 elements and 61722 nodes, for a total of 385700 elements and 432054 nodes. The wire and pin are considered a single piece for the simulation and are meshed together. The mesh consists of the same cross-section of elements repeating and rotating down the axis of the pin. A total of 380 layers of 145 elements each.

Coupling

The method of converting the fluid simulation results into inputs for the structure simulation is more involved than simply extracting the fluid data. While the CFD simulation is computationally expensive, the structure simulation is rather inexpensive. Additionally, for the best statistical results involving vibrations, it is better to have longer simulated time for the structure than is feasible for the fluid simulation. A method has been developed to replicate the fluctuating force data for longer time periods while maintaining the same power spectrum. Proper coherence of signals is also maintained.

The synthesized time history is generated using the power spectrum distribution (PSD) of the signal to be mimicked. A PSD is calculated as the square of the Fourier transform of a signal. It is an estimate of how the power of a signal is distributed over frequency. Any number and any length of synthesized histories can be generated by adhering to this PSD. However, in addition to adhering to the PSD, signals with

coherence in the original measurement should have coherence in the synthesized histories. Synthesizing an artificial time history involves taking the Fourier transform of this PSD. As there are an infinite number of signals that can result in a given PSD, there is no one unique signal to generate. As such, a random white noise function is applied to the phase information before the inverse Fourier transform is called. The process for generating two signals with proper coherence involves using a portion of the same white noise, in proportion to the coherence between the signals, as the input for the inverse Fourier transform.

While generating a time history for every point on the structure mesh would be optimal, the generation, storage, and implementation of the vast amount of data required is not feasible. Every time history read into Diablo comes at a time and memory cost. Additionally, orchestrating the coherence between each point is not realistic. The wire-wrapped fuel pins largely operate like cantilever beams and express their movement overwhelmingly in their first few modes. The force input data can be distributed on a very coarse grid and the first few modes will be unaffected. Additionally, the pins only interact with neighboring pins once every $1/6^{th}$ of a helical pitch, so the force distribution within a $1/6^{th}$ helical pitch will not affect any interaction with neighboring pins. Taking advantage of this, the synthetic force histories are generated for the each $1/6^{th}$ pitch of each pin for both the x and the y direction. Generating forces for 7 pins, each with 16 segments in two directions makes for a total of 224 synthetic histories. The histories generated are five seconds in length.

In this case one full helical pitch of fluid flow is simulated. Due to the rotational repetition in the bundle, this means that each distinct $1/6^{th}$ helical pitch appears six times in the geometry, but with shifted orientation. This repetition is exploited by generating the PSD for a given segment six times. PSDs can be averaged together to improve their statistical convergence.

A graphic demonstrating the coupling procedure from simulating the fluid to simulating the structure is depicted in figure 17.

Pin interconnection

The solid mesh of the wire-wrapped pins leaves a small space between the wire and neighboring pins as they pass by, meaning that neighboring pins do not have a connection. How the pins interact and connect with one another is a substantial aspect of the fluid-structure interaction. The pins will transmit energy and dampen one another if in contact. Similarly contact with the surrounding shroud can affect the pin vibrations. The nature of these connections is difficult to measure from experiment, and the force of connection can vary wildly for different portions of the bundle. Especially as the fuel swells, pins will exert different forces on one another and even lose contact with neighboring pins. This inconsistency in the connection and coupling between pin-to-pin and pin-to-shroud contact presents a difficult modeling dilemma. A number of
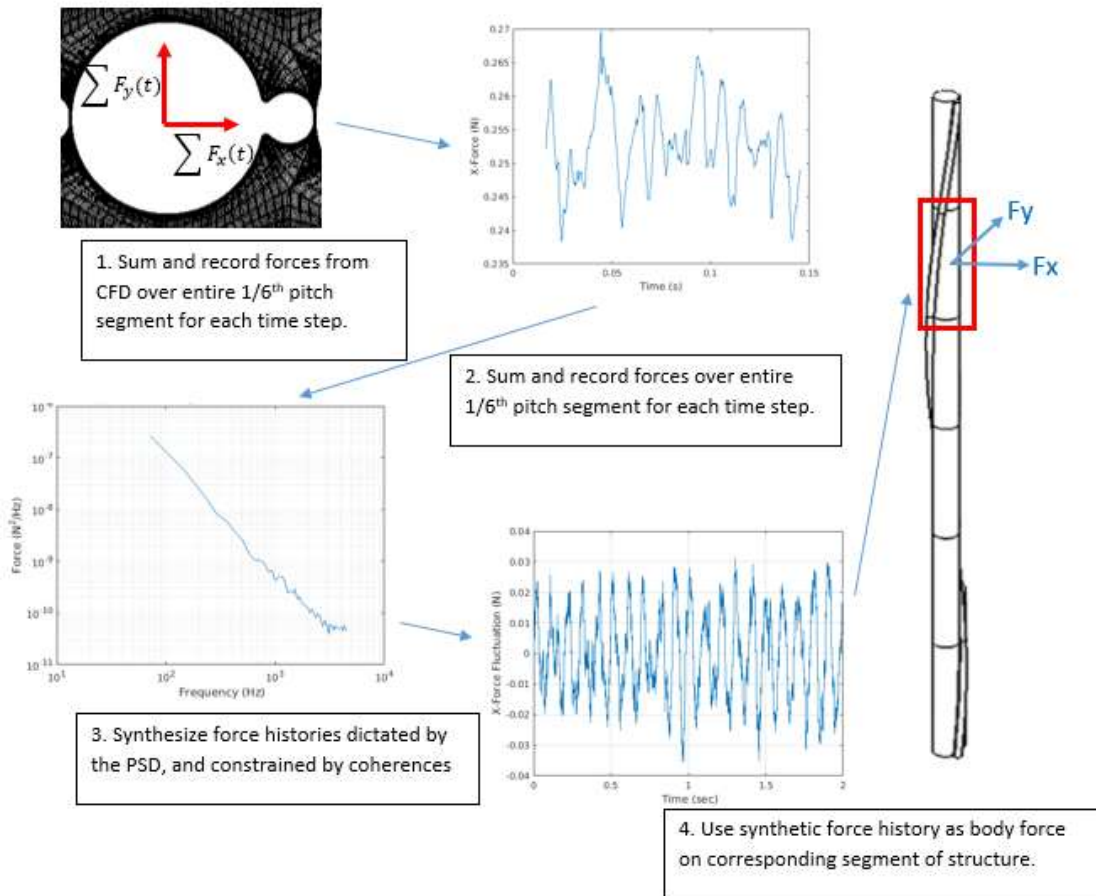
Figure 17. Method for coupling fluid and structure. Steps 1-4 summarize the process of gathering the force data, calculating the PSD, synthesizing time histories, then implementing them as body forces for the structure simulation.

constraints are tested to simulate the coupling between the solid elements. The effects of the different constraints are compared in the results.

For the majority of the simulation presented, neighboring pins at the point when the wire passes between them, are assumed to be in contact with one another. This contact is assumed to have enough normal force such that the tip of the wire in contact with the neighboring pin cannot translate with respect to the nearest node of the neighboring pin. The wire can however rotate about this connection point about any axis. This form of contact is modeled by connecting the pin and wire with a rigid beam extending from the wire, as depicted in figure 18. This beam does not appear or interact in any way with the fluid flow, it acts as a constraint between two nodes on separate pins.

The connecting beam is constrained to only rotate on the end connecting the to the neighboring pin. This connection essentially extends the wire to meet the neighbor and forces a rolling contact. This is the only constraint used for inter-pin connections,
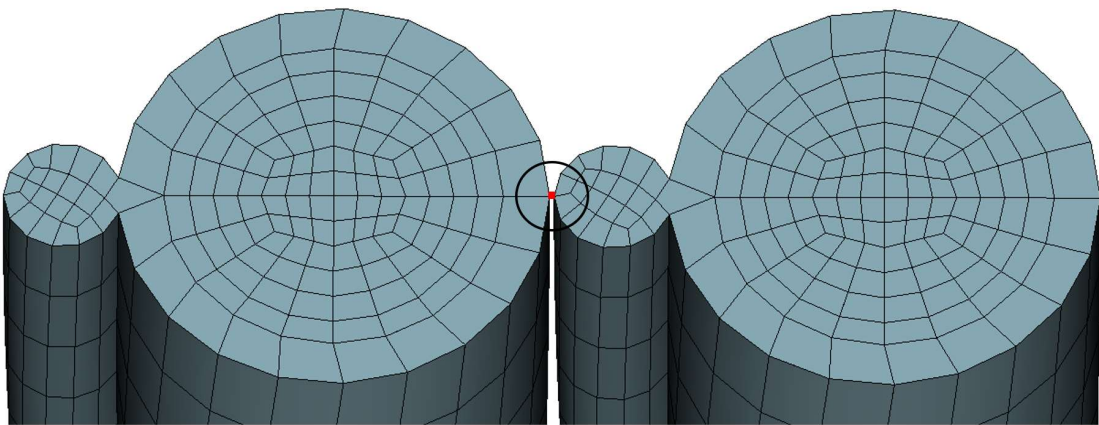

Figure 18. Contact between neighboring pins as modeled by a rigid beam connecting the nearest nodes. The beam allows for rotation about the node on the pin (left), but not about the wire (right).

other than no connection. A wire-wrapped pin will have 12 such connections in a single helical pitch, 6 from its wire passing neighbors, and 6 from neighbor wire passing itself.

The connection between the pins and the outside shroud is just as important to the behavior of the pins. The pin behavior as a whole is constrained by how forgiving the outside shroud is. The shroud, depending on the construction can itself deform and couple its motions with the pins. For the present simulation the shroud is assumed to be rigid and stationary. Two models for the interaction between the wire and wall were used. The first operates in the same way as the inter-pin coupling. A short segment protrudes from the tip of the wire where it is constrained against translation on the other end. The segment is still permitted to rotate around this point. The second model contains a longer rigid beam, two diameters long, connecting the center of the pin to a point outside of the shroud. The beam perpendicularly intersects the shroud as it passes through. The long beam permits the pin to translate approximately parallel to the wall, while not allowing it to translate to or away from the wall, making for a more forgiving connection as if pushed together by a weaker force allowing for slipping. The effects of these connection models are compared against unconstrained pins in order to observe the dampening and energy transfer effects of the connections between pins.

Verification of LES

The verification process for the LES for the seven pin bundle is based primarily on comparing the mesh to the scales of motion, namely the y+ for near the wall, and the Taylor microscale for the free stream. The mesh for the present study is informed by

similar wire wrapped meshes generated using the same method [32], [37], [39]. The present mesh has the greatest mesh density of any of these previous studies. The greater mesh density is primarily driven by the y+ value, as well as the goal of gathering a higher fidelity force load distribution for the coupling with Diablo.

A snapshot of the y+ distribution at an instant is depicted in figure 19. The y+ value is less than 1 for 66.3% of all surface collocation points, and < 2 for 99.6% of all collocation points. Figure 19 indicates that particularly high y+ areas tend to be concentrate on strips of wall facing the subchannels. The majority of y+ greater than two exists on the outside shroud, where wire-wrapped fuel pin bundles experience their highest flow rate and where the channels are the largest. The center subchannel
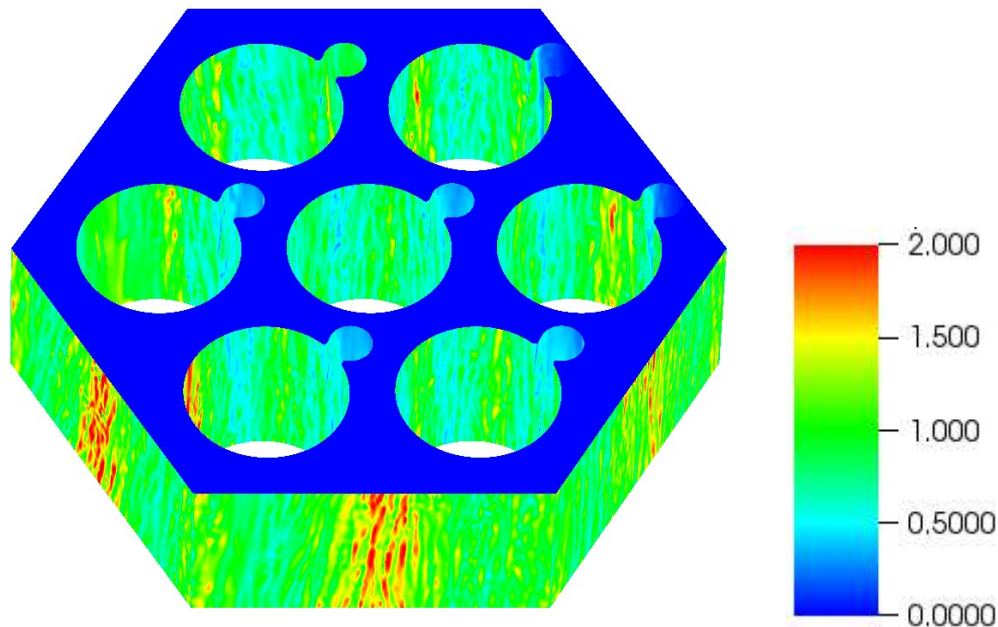

Figure 19. Snapshot of y+ distribution for a slab of the bundle one pin diameter in thickness.

experiences considerably lower y+ values than the outer subchannels. Boundary layer

definition of the surrounding shroud is not a priority for the present study, and the y+

value on the pins is ~1 for the majority of the geometry.

The Taylor microscale does not refer to a specific eddy size like the Kolmogorov

scale, but rather acts as an average of the dissipative scales. Having a mesh size smaller

than the Taylor microscale indicates that the LES has entered the dissipative portion the

the turbulent cascade. This metric does not guarantee that all scales of motion that will

significantly affect the flow behavior have been resolved. It does offer a reasonable

target for mesh size when the Kolmogorov length scale is not feasibly obtainable. The

method of computing the Taylor microscale is discussed in the Numerical Methods

chapter. A plot of the Taylor microscale in a center subchannel as a function of heigh for

a single helical pitch is depicted in figure 20, and across a gap between two pins in

Figure 20. Taylor microscale as a function of height through the center of a subchannel. The graph is interrupted by a passing pin at ~13 cm and ~26 cm.

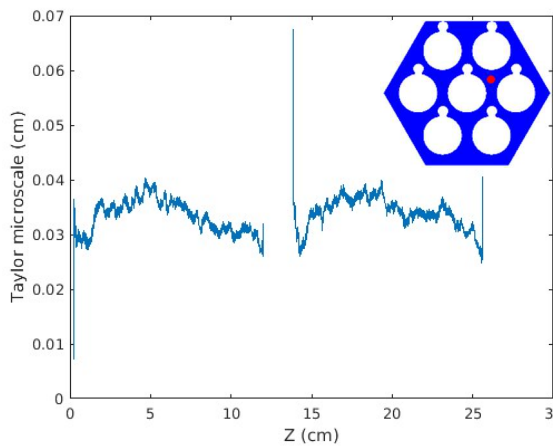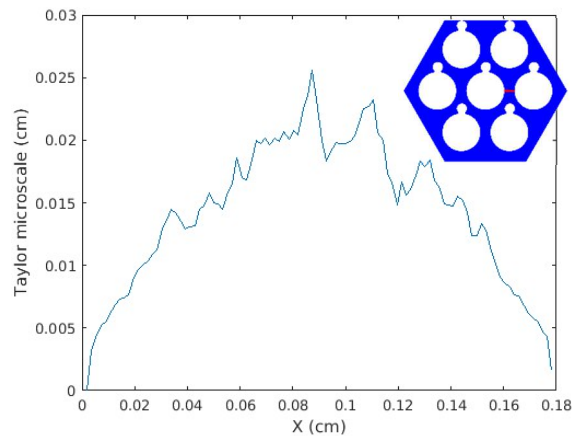**Figure 21. Taylor microscale as a function of height through the center of a subchannel. The graph is interrupted by a passing pin at ~13 cm and ~26 cm.**

figure 21. The location of these measurements are highlighted in figures 20 and 21 as well.

Concerning the subchannel centered plot, the Taylor microscale maintains a consistent value of ~0.03 – 0.035 cm for the majority of the subchannel, swelling when the pin passes through other subchannels. The plot is interrupted twice as the pin passes through the subchannel at $\sim Z = 13$, and $\sim Z = 26$. As the line approaches the pin, the Taylor microscale loses meaning as the boundary layer laminarizes, increasing dramatically. The plot across the gap between two pins shows a similar curve to what on may expect for channel flow. The Taylor microscale peaks at ~0.02 cm in the center, then falls down to 0 at the walls. The microscale is lower in between the bundles than in the center of the subchannels. A benefit of the meshing method selected, is that the mesh naturally becomes finer in compact areas such as the gap between pins. The mesh is additionally arranged to become finer as it approaches the wall in order to capture the boundary layer. The axial spacing, which is the coarsest, remains the same. The mesh elements in the axial direction are 0.073 cm long, making for an average spacing between collocation points of 0.01 cm. Thus, for the freestream flow, the mesh is finer than the Taylor microscale, meeting the stated criteria.

Verification of Diablo solution

The vibration of the fuel pins is compared for both sensitivity to time-stepping and mesh count. The computational cost of the structure solution is comparatively low, so there is little reason to limit either the timestep or the mesh count, so both have been
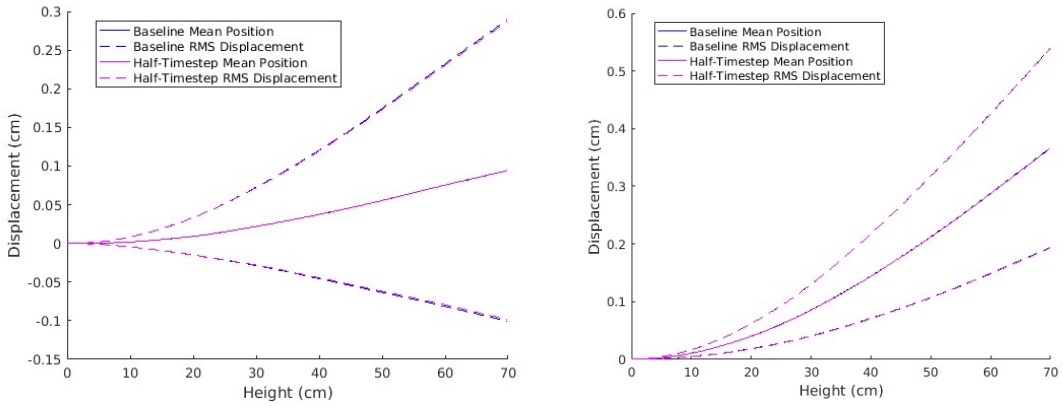
76

Figure 22. Mean position and RMS amplitude for baseline and half timestep Diablo simulations. Left: x-displacement, Right: y-displacement.
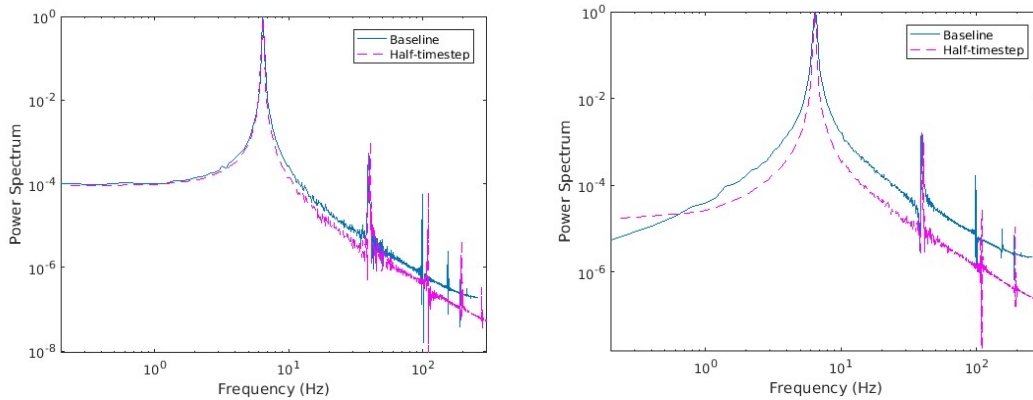
Figure 23. Power spectrum for vibration of the tip of the vibrating pin. The three largest peaks are at 5.4 Hz, 33.2 Hz, and ~100 Hz. Left: spectrum for x-displacement, Right: spectrum for y-displacement.

chosen conservatively. Regardless, verification is performed to ensure the pin vibration

is not affected by either.

Vibrations of the center pin undergoing forces generated from the fluid

simulation are simulated for two different time-steps. The pin is simulated without

coupling to neighboring pins in order to maximize displacement and velocities. The

vibration amplitude and frequency power spectra of the fuel pin are compared. The first and baseline simulation is performed using a time-step of 1.907E-3 seconds, and the second simulation using a time step of half, 9.535E-4 seconds. Mean and rms vibrations from the displacement are plotted for both the baseline and half timestep simulations in figure 22 for the x and y directions. The frequency spectra for the x and y displacement at the tip of the pin for both cases are depicted in figure 23.

The displacement results have less than 1% difference for any point. The power spectra show excellent agreement for the two largest, lowest frequency peaks, 5.4 Hz and 33.2 Hz. At the higher frequencies there is minor discrepancy in the frequency of a peak ~100 Hz. By far the dominant frequencies are in excellent agreement between the two simulations.
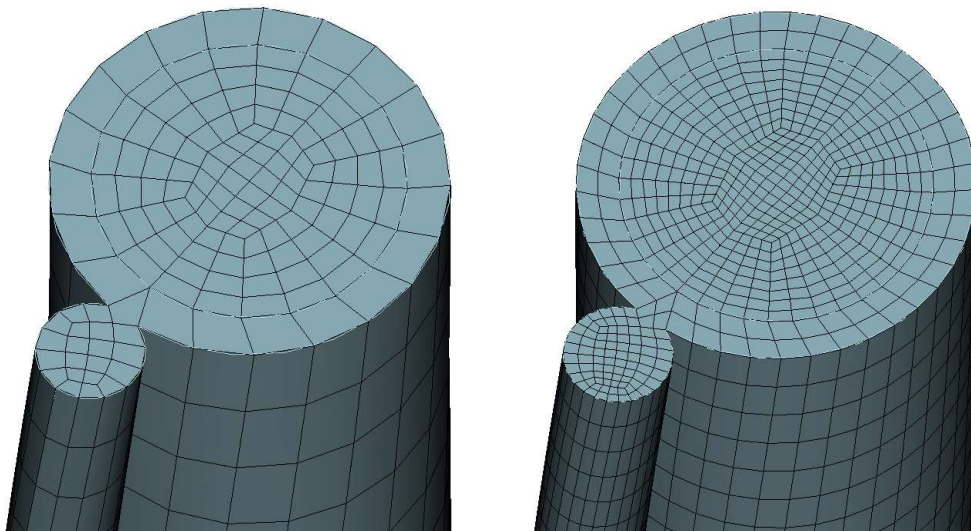

Figure 24. Samples of meshes used for structure mesh verification. Left: baseline mesh, Right: Refined mesh.

A similar analysis was performed to establish independence of the mesh size for the structure. The finer mesh shared identical construction with a simple refinement that splits each element into 8 equal sized elements, making for a total of 3085600 elements and 3456432 nodes. Samples of the two meshes are shown in Figure 24. The same comparisons for displacement and frequency spectra were performed for the two simulations, which revealed imperceptible differences between the two. The plots for the refined mesh are identical to those shown for the baseline mesh in figures 22 and 23.

Results

A convenient property of performing simulations is that the same simulation is repeatability. A simulation can be repeated with identical initial and boundary conditions isolating individual permutations, and any differences are necessarily a result of the premutation. For each of structure response simulations, identical fluid forces are exerted on the structure as a function of time. Simulations with varying inter-pin contact methods are carried out for comparison. The first section explores the effect of the inter-pin connection, and pin-shroud connection. The second section explores the effect of the presence of gaps, or missing connections between pins on the nature of the vibrations. The primary methods of comparison include: rms displacement measurement, proper orthogonal decomposition, power spectral density of motion, and centerline tracing as a function of height. Focus is given to the center pin. While the outside pins exhibit interesting behavior the center pin is the closer approximation to what can be expected for a pin to experience in a larger bundle.

Inter-pin and Pin-Shroud Contact

This section compares the vibrational response of the pins under three scenarios. The first scenario has the pins completely independent of contact with one another. Essentially the pins act as cantilevers, as the only restrictive force they experience is the connection at the bottom. This acts as a baseline, revealing how the unconstrained pin is naturally inclined to respond to the fluid forces. The second scenario has the pins interconnected to one another, but with no connection to the shroud. This helps to determine how the pins interact with one another, and how their motions amplify or cancel one-another. The third scenario has the pins connected to eachother, as well as to the shroud with the short beam connection. This scenario, in comparison with the others, shows how the shroud dampening affects the nature of the vibrations.

The frequency spectra for each of the three cases is given in figure 25. The four highest modes of each are depicted in figure 26. The mean displacement and RMS amplitude in the x and y directions for the three cases are shown in figure 27.

The no contact model experiences considerable larger displacements than either of the other cases. It's expected that neighboring pins would absorb energy and prevent movement, so this result is expected. The four most prominent orthogonal modes very closely resemble those that would be expected for a vibrating pin, essentially a first, second, third and fourth order polynomial. However 96% of the motion can be described in the first mode alone, 99% with the first two modes. The frequency spectra is the same as that used for the verification. There is one dominant frequency at ~5 Hz, and another
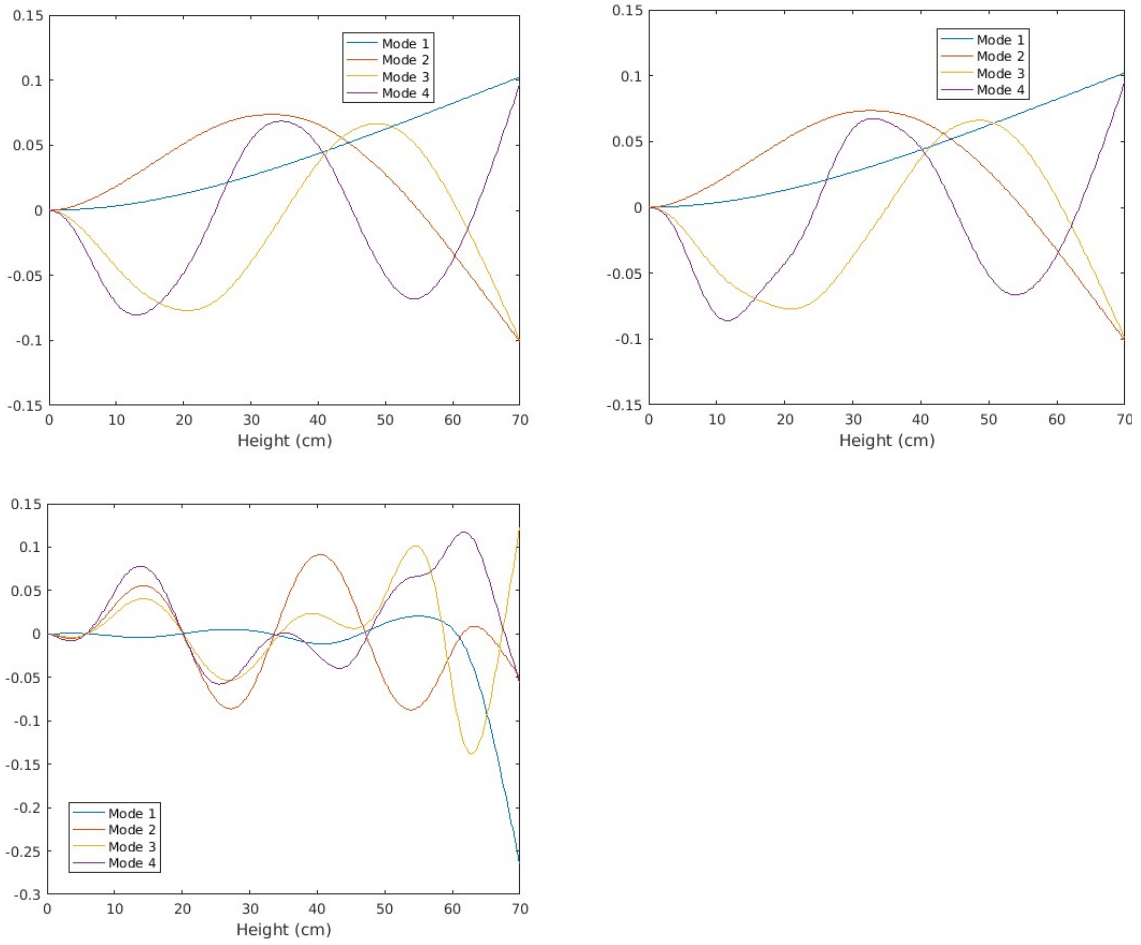
Figure 25. Four highest energy modes for, Top left: no contact, Top right: inter-pin contact, Bottom: inter-pin and shroud contact.
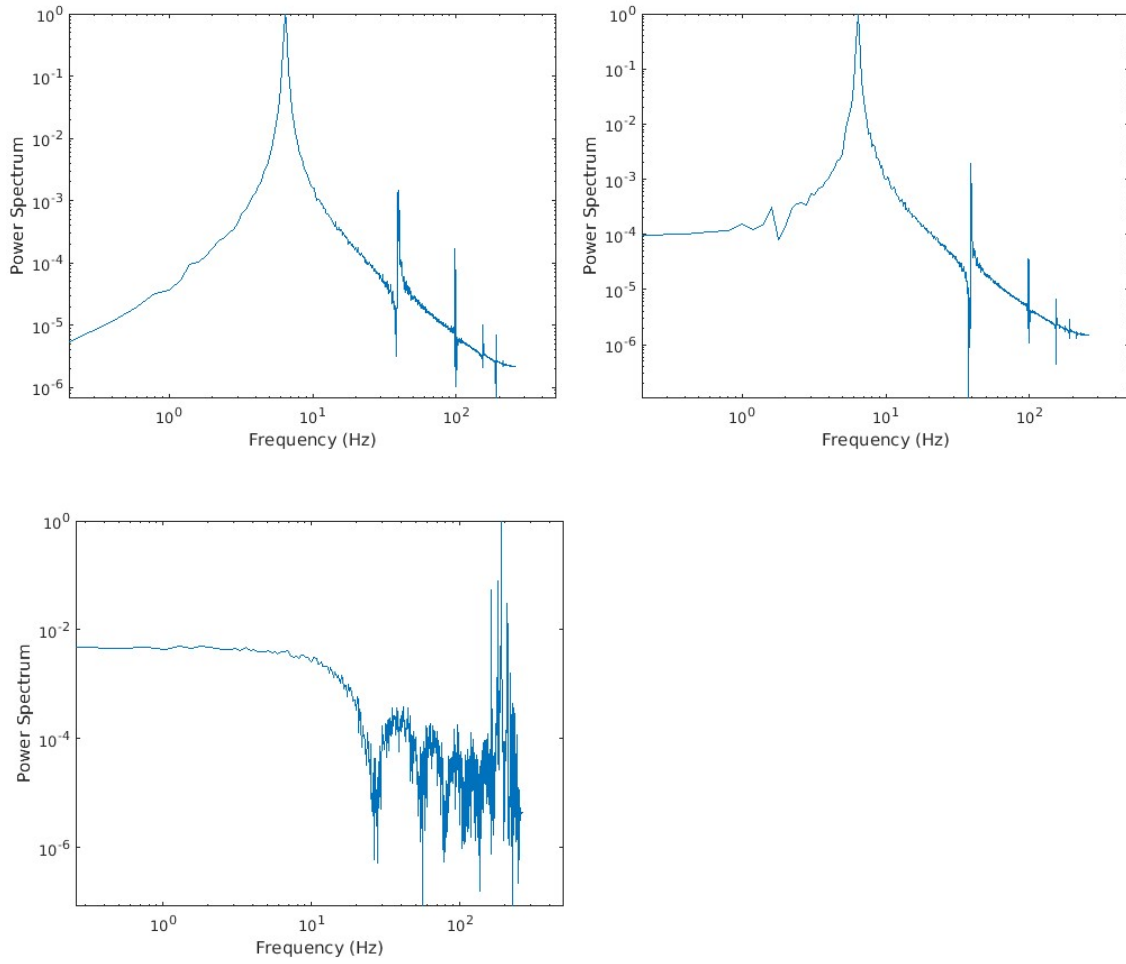
Figure 26. Frequency spectra for y-displacement Top left: no contact, Top right: inter-pin contact, Bottom: inter-pin and shroud contact.
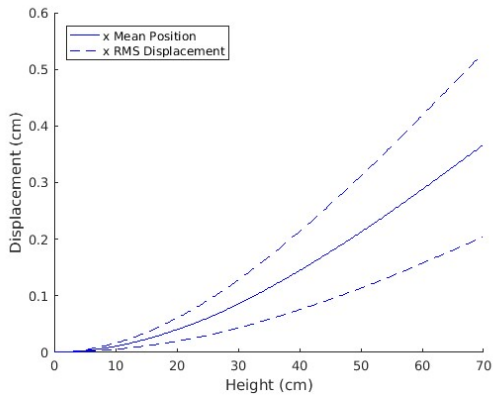
a. No contact, x displacement
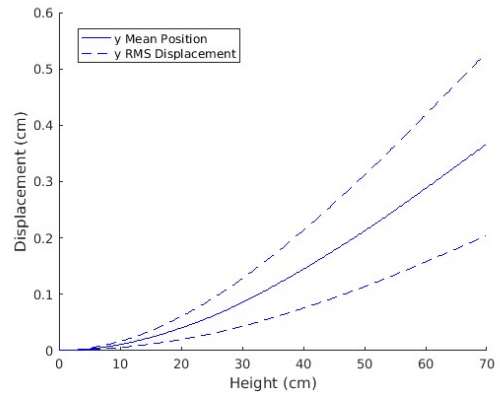
b. No contact, y displacement

c. Inter-pin contact, x displacement

d. Inter-pin contact y-displacement

e. Inter-pin and shroud contact, x displacement

f. Inter-pin and shroud contact, y displacement

Figure 27. Mean position and displacement for 3 contact scenarios.

83

peak at ~33 Hz. The matching x and y displacements indicate that the pin vibrates in a near circular motion.
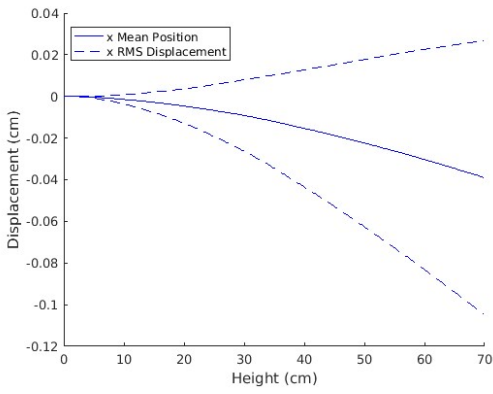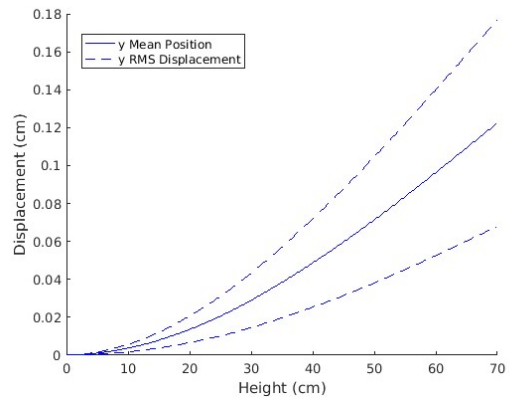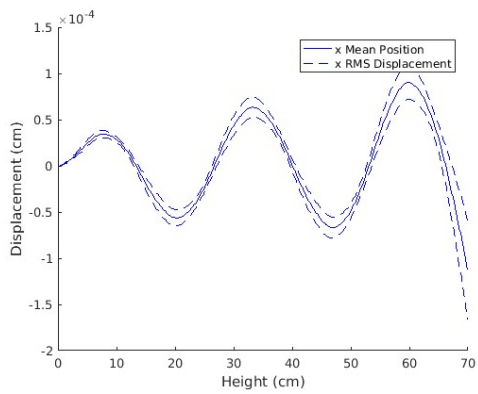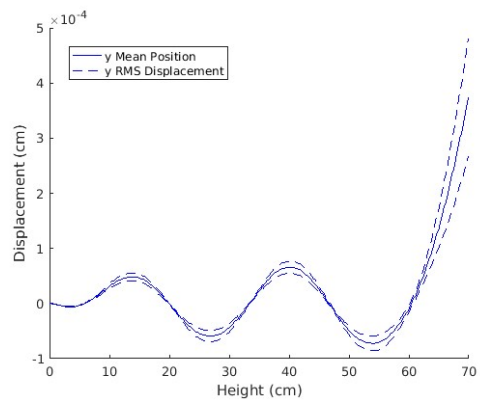
Comparing the inter-pin connected data to that of the unconnected pin reveals a mitigating impact. Most notably, although the frequency and modes appear very similar to the unconnected case, the inter-pin case has significantly smaller amplitude vibrations, less than 1/3rd. The mean position of the x component has switched from bending in the positive x direction to the negative. This is an example of the natural motions of neighboring pins cancelling each other out, and mitigating motion. Coupled together, the entire bundle acts as a single cantilever beam. The constraint of a neighboring pin can be seen to impact the fourth mode of motion. Although, similar to the unconnected pin, 94% of the motion can be described by the first mode.

The fully constrained bundle shows completely different behavior than the previous two cases. The pin motions are two orders of magnitude smaller, the modes are more complex, and the frequency shows no strong peaks. The bundle behavior has become fundamentally different, essentially changing from having a single constraint at the bottom, to a network of constraints. Now that the neighboring pins are not free to bend, only shorter wavelength, higher order oscillations with nodal points on the pin connections are permitted. The top four modes show the consistency of these nodal points. By far the largest displacement comes from the motion of the tip of the pin, as for the last portion it is unconstrained. The mean position of the pin bows out between points of contact with neighboring pins. Unlike the previous cases, the displacement does not grow with the length of the bundle, but oscillates consistently with the pitch.

The frequency spectrum is flat except for very high frequencies ~200 Hz, which are less of a concern. The first mode describes about 63% of the motion and the second 16%, meaning more modes are present and the motion is more complex.

This comparison emphasizes the role of the surrounding shroud. The impact of the shroud is significantly larger than that of the inter-pin connections. However, this is likely a consequence of assuming a rigid shroud. The inter-pin connection tends to cancel out motions between pins by coupling their motions to act as a single body, while the surrounding shroud tends to act as a periodic constraint, forcing only high frequency vibrations with shorter wavelengths to be possible.

Another pin-shroud connection is presented to observe how the pin motion changes with a less constraining shroud connection. The longer-beam connection, discussed in the methods section, allows for the pin to move parallel to the shroud, but not perpendicular. Displacements and RMS amplitude, mode decomposition, and frequency spectra are shown for this shroud constraint in figure 28.

The mean displacement of the pin for the more forgiving shroud connection is about the same as for the previous shroud connection. However the amplitude of vibrations is significantly greater. The top portion of the pin especially has much greater amplitude vibrations. The top four modes appear to be less tightly constrained about the nodal points that were prevalent in the previous case. Nodal points are still present, but their locations are less rigid, as the outside pin is permitted to translate in the z-direction. The frequency spectra while more flat is largely the same, as it is dominated by the high
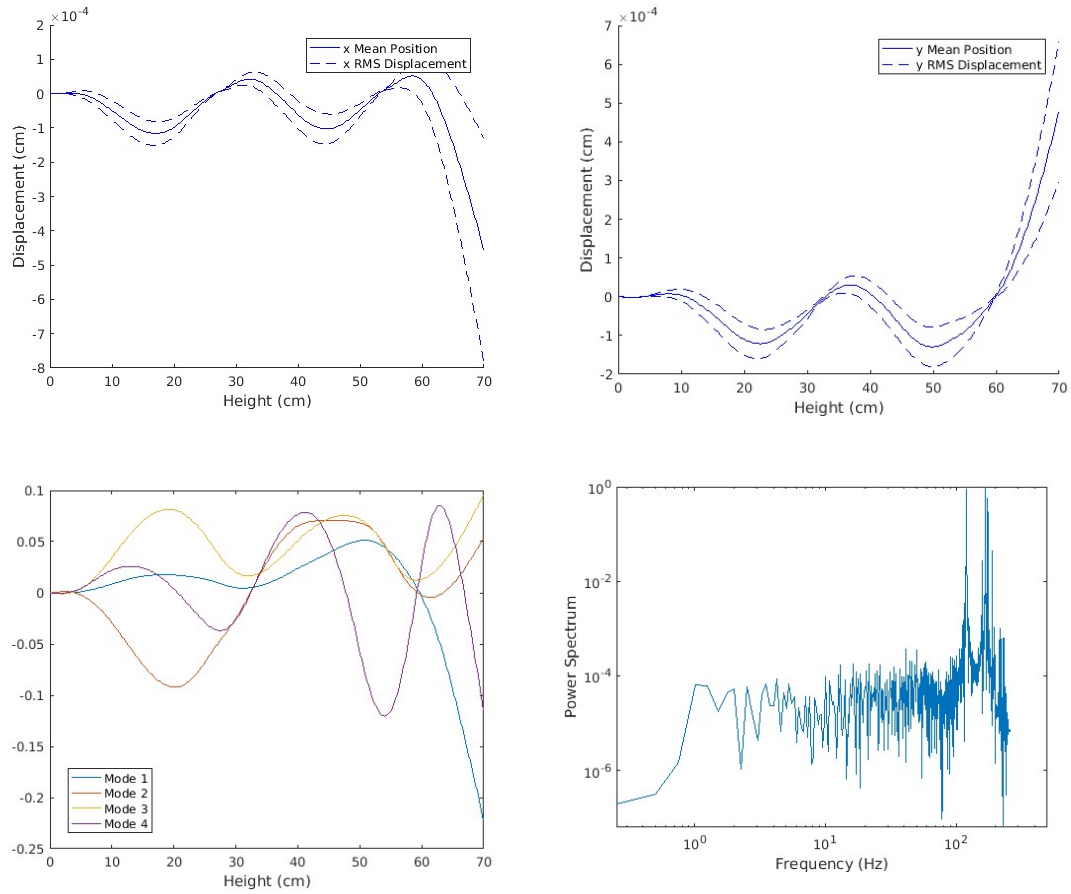
Figure 28. Results for long beam shroud connection. Top Left: x displacement, Top Right: y displacement, Bottom Left: Top four modes of y-motion. Bottom Right: y motion frequency spectra.

frequency spikes ~100-200 Hz. The first two modes describe 50% and 23% of the motion respectively.

Despite the greater freedom of wavelengths and nodal positions, the pin is largely constrained to high frequency, low wavelength vibrations. Regardless of how the outside pins are connected to the shroud, as long as radial motion is constrained, the center pin is forced from 6 angles to have periodic nodal points. Unless gaps form or the shroud deforms, the center pin cannot vibrate at lower frequencies.

Presence of Gaps

The surrounding shroud forces the pin vibrations to short wavelengths that must comply to certain nodes, thus restricting vibrations in the pin and avoiding frequencies that might stoke the pins natural frequency. However, it is possible and likely that as the pins begin to swell inside of a reactor that the pin to pin contact will become less consistent and gaps may form. These gaps remove constraints from the pin and open the possibility to longer wavelength vibrations.

Three gap patterns have been devised for testing the effects of gaps on the center pin vibrations. Gaps are represented by the lack of a connection between pins. A pin with a gap with a neighboring pin is no longer constrained and may move too or away from the neighboring pin. The first case is a bundle such that the neighbors on the immediate left and right of the center pin have gaps with the center pin for the entire length of the bundle. The second case is a bundle where the top 1/2 pitch of the center pin has no contact with neighboring pins. Finally, the third case is a bundle where the

87

middle 1/2 pitch is surrounded by gaps, approximately from 29 -42 cm height. For each of these cases the x and y displacement and RMS amplitude, the top four modes and frequency spectra depicted in figures 13, 14, and 15 for the first, second and third gap cases respectively. In each case pin-shroud contact is modeled using the short beam method.

The case of no contact with pins in the positive or negative x direction (figure 29) determines how the other pins compensate for the center pin's relative lack of constraint. Despite the decrease in number of contact points, there is little effect on the behavior of the center pin vibrations. The displacement and RMS amplitude have the same magnitude as the full contact model, with the exception of the slightly greater displacement and vibration at the edge of the pin for both the x and y direction. The first and second modes contain 78% and 10% of the motion respectively. The first mode pertains exclusively to the motion at the tip of the pin. The frequency power spectrum again shows the higher frequencies, 100-200 Hz, peak. These being associated with the vibration of the tip of the pin. The modes and frequency spectra for the x direction do not show any exceptional behavior.

Despite the loss of $1/3^{rd}$ of all of the constraints for the center pin, the center pin vibrations behave largely the same as the fully constrained version. The other four directions of contact sufficiently suppress motion in the x direction that motion in the x-plane cannot develop.

Figure 29. Results for no contact with left or right pins. Top Left: x displacement, Top Right: y displacement, Bottom Left: Top four modes of y-motion. Bottom Right: y motion frequency spectra.

Figure 30. Results for no contact around the top 1/2 pitch. Top Left: x displacement, Top Right: y displacement, Bottom Left: Top four modes of y-motion. Bottom Right: y motion frequency spectra.
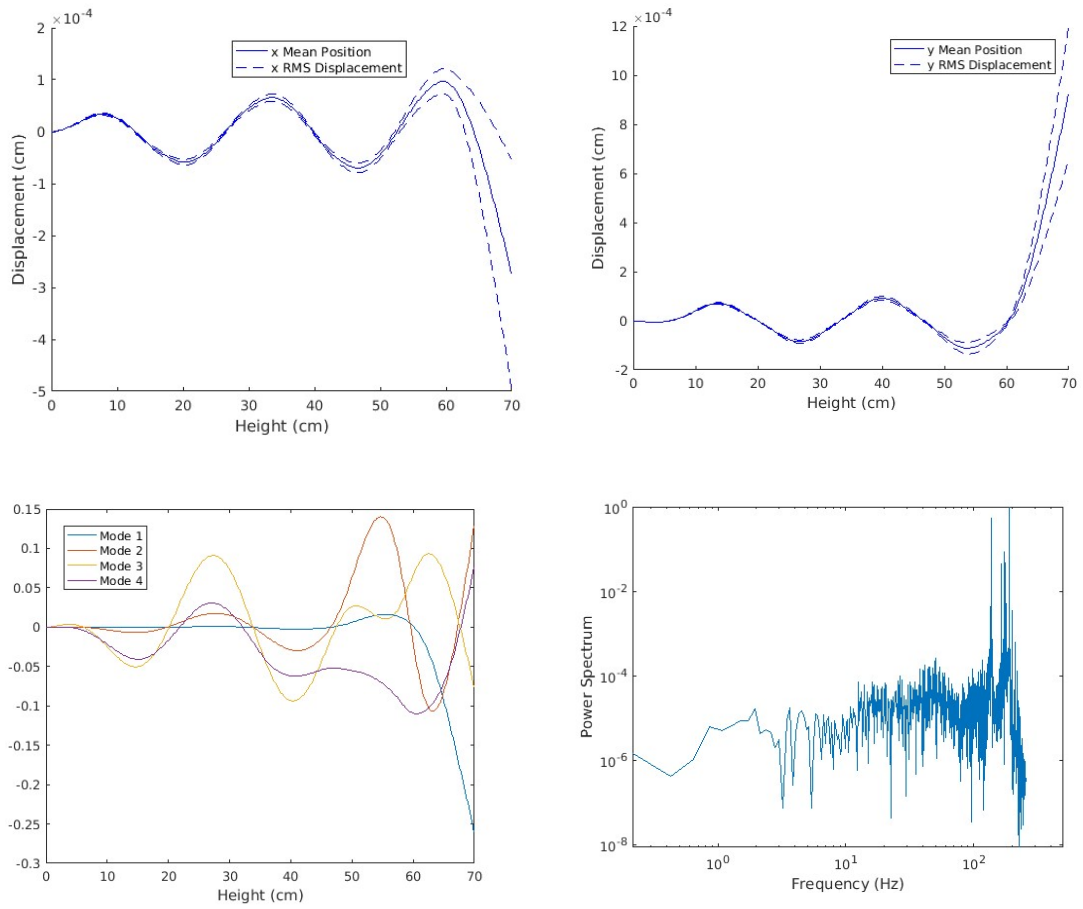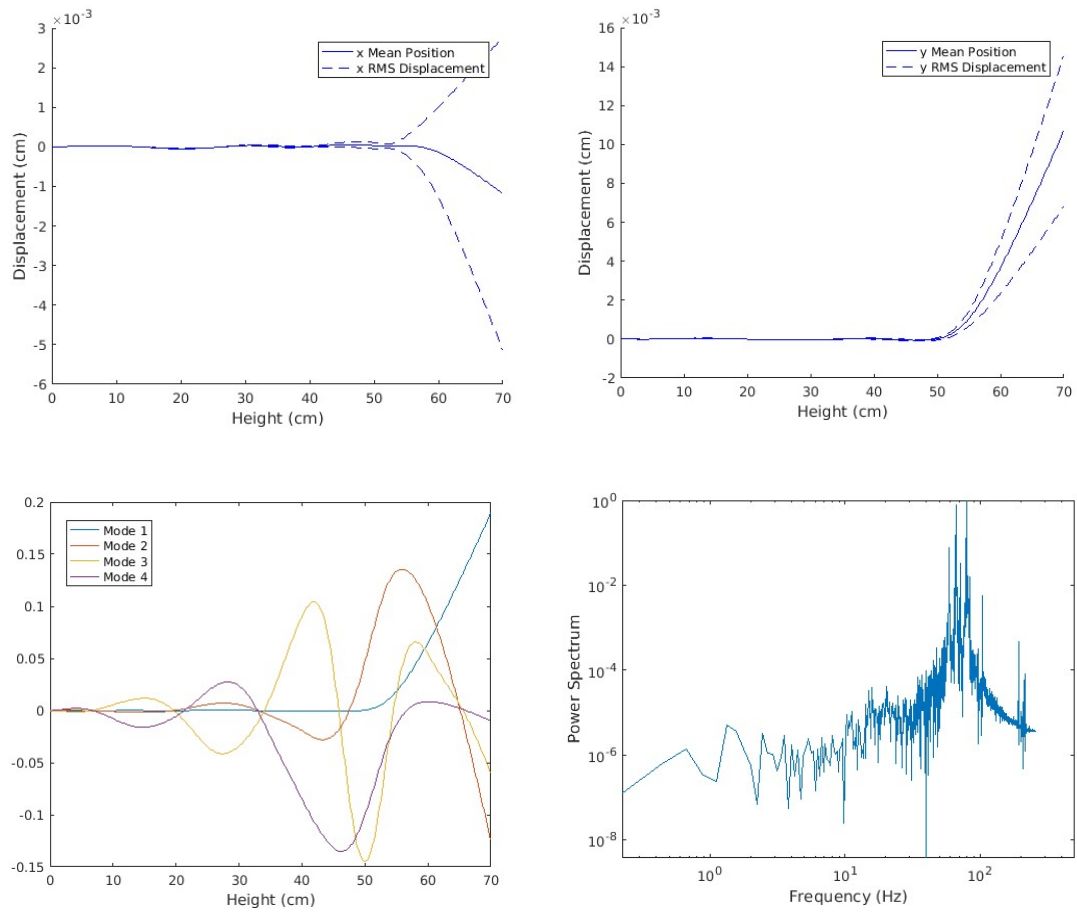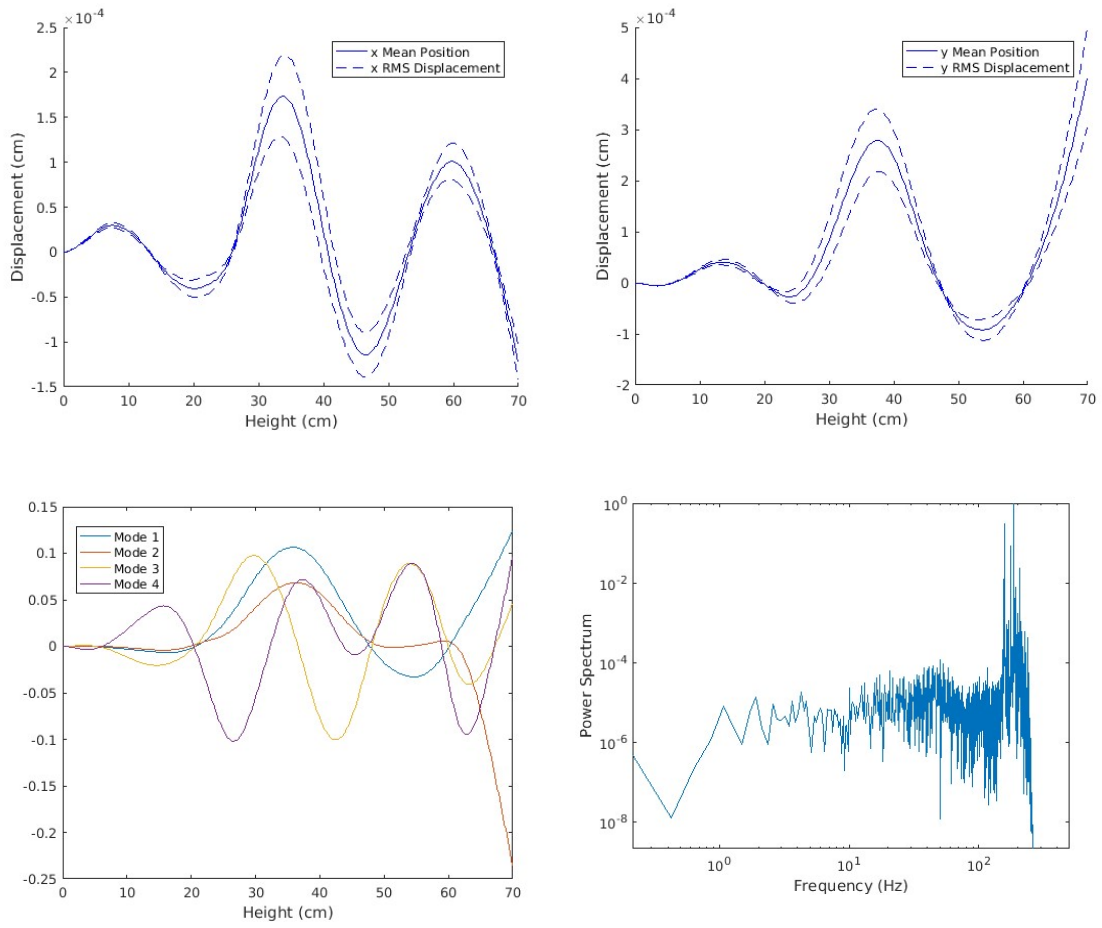
Figure 31. Results for no contact around the middle 1/2 pitch. Top Left: x displacement, Top Right: y displacement, Bottom Left: Top four modes of y-motion. Bottom Right: y motion frequency spectra.

The case of no contact around the top 1/2 pitch of the center pin (figure 30) emphasizes the largest vibrations present in the fully constrained model. The tip of the pin behaves as a short cantilever, after it experiences its last contact point. By extending this length of no contact, a significant length of the pin experiences amplitudes up to 25 times larger than the fully constrained model. The vast majority of the motion, 93%, is described by the first mode, which exclusively pertains to the vibrating pin tip. The longer vibrating segment results in a significant drop in peak frequency to ~60-90 Hz. While the vibrations in this case are significantly lower than the simulations without shroud contact, the vibration amplitudes and frequencies in this case are approaching concern.

The case for no contact around the middle 1/2 pitch of the center pin (figure 31) explores how vibrations manifest in pockets of no contact. The mean position and rms amplitude increase sharply for both the x and y directions in this pocket of no contact. The greater displacement extends further outside of this pocket as well, increasing the vibration amplitude slightly all across the domain. The vibration in this segment are still less than those of the pin tip. The pin tip on a fully constrained bundle is more of a threat than a 1/2 segment with no contact. As two vibrations are dominant, the top two modes are fairly split at 52% and 25%. The frequency spectra shows the same high frequency peaks as discussed previously

CHAPTER V

CONCLUSIONS

A fully coupled FSI simulation was developed for crossflow through a tube bundle. The simulation is based on and modeled after an experiment for validation purposes. Efforts to verify the meshes ability to resolve the relevant physics included comparison to a finer mesh, and calculation of the Taylor microscale. Comparisons between the simulation and experiment snapshots show that the simulation captures the transition from laminar to turbulent flow and the flow mechanics that change with it. The central validation effort concerned the onset of large amplitude vibrations at a critical velocity. Three velocities were chosen for simulation straddling this onset point. The results show that the streamwise vibrations increase dramatically around this point. However, it appears that in the simulation the 2^{nd} row of pins prematurely begins to increase in amplitude. Also observed is that the simulation natural frequency is slightly lower than the experiment's natural frequency: 15 Hz as opposed to 15.7 Hz. These observations point to the possibility that the domain simplifications made for the simulation in pursuit of feasible computation times resulted in the slight decrease in natural frequency of the pin bundle. This decrease in natural frequency leads to the premature onset of large amplitude vibrations observed. This aside, 3^{rd} row streamwise and both 2^{nd} and 3^{rd} row transverse simulation amplitudes were in line with the experiment. Further analysis into the vibration shape and averaged velocity analysis pointed to some fundamental changes in the flow structure from below the onset point and above. Below the onset point, a persistent transverse flow is present just beneath

93

each row of pins past the first. This transverse flow alternates directions at every row. As the larger amplitude vibrations increase, this persistent channel disappears and is replaced by some oscillating flow. The flow structure responsible for the shift in vibrations appears to work its way from front to back through the bundle.

A scheme was developed to couple a fluid simulation from Nek5000 to a structure simulation using Diablo. The one way coupling allows for inexpensive and statistically significant FSI computations of vibrations due to turbulent buffeting through a wire-wrapped fuel pin bundle. The scheme involves gathering the force exerted on the wire-wrapped pins from Nek5000, calculating the power spectrum density of the force exerted, synthesizing long force histories that satisfy the power spectrum, as well as coherence with related forces, then using the force histories as inputs for a structure simulation in Diablo. Both the fluid and solid components went through a verification process. Multiple wire-wrapped fuel pin solid simulations were run to implement and compare a variety of contact and constraint scenarios. Inter-pin connectivity was found to result in neighboring pins opposing and cancelling motion, resulting in a dampening effect. The entire bundle acts as a single beam. The surrounding shroud, however fundamentally changes the vibration characteristics. The rigid shell forces the outside pins to constrain the center pin. Every point of contact with a neighboring pin becomes a forced nodal point for any vibrations present. This forces only short-wavelength vibrations to be able to exist, which is against the long, slender beams inclination to low mode vibrations. In reality the shroud may couple with the bundle vibrations or deform in ways not predictable when assuming a rigid shroud. Gaps in the bundle, even for large

94

segments do little to encourage vibrations. Removing all connections to the center pin in the x-direction has little effect on the vibrations. Only the tip of the pin vibrates with any meaningful amplitude, which is compounded if the top of the bundle has looser connections between pins. In general, the hexagonal array ensures that even systematic gaps don't allow the pins free motion. The effect of coupling of the bundle with the shroud is a significant unknown.

These two coupling studies and their respective schemes showcase the capabilities of highly scalable, high-fidelity fluid-structure interaction simulations and their practical applications for studying engineering challenges for nuclear power reactors. The advancement of these methods contributes to our ability to design, maintain, and optimize existing and future nuclear reactors. Flow-induced vibrations have historically been difficult to predict and design around, but with persistent computational advancements like that described here, engineers can design novel heat exchanger and reactor core concepts that boast superior thermal characteristics. The technology advancements throughout the field of nuclear engineering promise the safe operation of existing reactors, and the tremendous advancement of new reactor generations.

REFERENCES

[1]   N. L. Johnson, "The Legacy and Future of CFD At Los Alamos," *Proc. 1996 Can. CFD Conf.*, vol. 836, pp. 1–22, 1996.

[2]   B. E. Launder and D. B. Spalding, "The numerical computation of turbulent flows," *Comput. Methods Appl. Mech. Eng.*, vol. 3, no. 2, pp. 269–289, 1974.

[3]   B. E. Launder, G. J. Reece, and W. Rodi, "Progress in the Development of a Reynolds-Stress Turbulence Closure," *J. Fluid Mech.*, vol. 3, pp. 537–566, 1975.

[4]   J. SMAGORINSKY, "General Circulation Experiments With the Primitive Equations," *Mon. Weather Rev.*, vol. 91, no. 3, pp. 99–164, 1963.

[5]   P. Fischer and J. Mullen, "Filter-based stabilization of spectral element methods," *Comptes Rendus l'Académie des Sci. - Ser. I - Math.*, vol. 332, no. 3, pp. 265–270, 2001.

[6]   C. H. K. Williamson and R. Govardhan, "Vortex-Induced Vibrations," *Annu. Rev. Fluid Mech.*, vol. 36, no. 1, pp. 413–455, 2004.

[7]   S. Kaneko, F. Inada, T. Nakamura, and M. Kato, *Flow Induced Vibrations*. Elsevier, 2008.

[8]   H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, "A Tensorial Approach to Computational Continuum Mechanics using Object-Oriented Techniques," *Comput. Phys.*, vol. 12, 1998.

[9]   C. S. Peskin, "The immersed boundary method," *Acta Numer.*, vol. 11, pp. 479–517, 2002.

[10]   M. . Deville, P. F. Fischer, and E. H. Mund, *High-Order Methods for Incompressible Fluid Flow*. Campbridge University Press, 2002.

[11]   J. Donea, A. Huerta, J. Ponthot, and A. Rodr, "Arbitrary Lagrangian – Eulerian Methods," *Encycl. Comput. Mech.*, pp. 1–25, 1999.

[12]   H. Lee-Wing and A. T. Patera, "A Legendre spectral element method for simulation of unsteady incompressible viscous free-surface flows," *Comput. Methods Appl. Mech. Eng.*, vol. 80, no. 1–3, pp. 355–366, 1990.

[13]   J. M. Solberg *et al.*, "Diablo: A Parallel, Implicit Multi-physics Finite Element Code for Engineering Analysis User Manual," 2015.

[14]   J. P. Gomes, S. Yigit, H. Lienhart, and M. Schäfer, "Experimental and numerical study on a laminar fluid-structure interaction reference test case," *J. Fluids Struct.*, vol. 27, no. 1, pp. 43–61, 2011.

[15]   E. Merzari, P. Fischer, and J. Solberg, "A High-Fidelity Approach for the Simulation of Flow-Induced Vibration," in *Proceedings of the ASME 2016 Fluids Engineering Summer Meeting*.

[16]   S. B. Pope, *Turbulent Flows*. Campbridge, UK: Campbridge University Press, 2000.

[17]   L. Sirovich, "Turbulence and the Dynamics of Coherent Structures," *Q. Appl. Math.*, vol. XLV, 1987.

[18]   R. D. Gabbai and H. Benaroya, "An overview of modeling and experiments of vortex-induced vibration of circular cylinders," *J. Sound Vib.*, vol. 282, no. 3–5, pp. 575–616, 2005.

[19] S. Dong and G. E. Karniadakis, "DNS of flow past a stationary and oscillating cylinder at Re = 10000," *J. Fluids Struct.*, vol. 20, no. 4 SPEC. ISS., pp. 519–531, 2005.

[20] H. Al-Jamal and C. Dalton, "Vortex induced vibrations using Large Eddy Simulation at a moderate Reynolds number," *J. Fluids Struct.*, vol. 19, no. 1, pp. 73–92, 2004.

[21] G. V. Papaioannou, D. K. P. Yue, M. S. Triantafyllou, and G. E. Karniadakis, "On the effect of spacing on the vortex-induced vibrations of two tandem cylinders," *J. Fluids Struct.*, vol. 24, no. 6, pp. 833–854, 2008.

[22] G. Palau-Salvador, T. Stoesser, and W. Rodi, "LES of the flow around two cylinders in tandem," *J. Fluids Struct.*, vol. 24, no. 8, pp. 1304–1312, 2008.

[23] S. J. Price, "A Review of Theoretical Models for Fluidelastic Instability of Cylinder Arrays in Cross-Flow," *Journal of Fluids and Structures*, vol. 9, no. 5. pp. 463–518, 1995.

[24] K. Schröder and H. Gelbe, "Two- and three-dimensional CFD-simulation of flow-induced vibration excitation in tube bundles," *Chem. Eng. Process. Process Intensif.*, vol. 38, no. 4–6, pp. 621–629, 1999.

[25] V. Shinde, E. Longatte, and F. Baj, "large eddy simulation of fluid-elastic instability in square normal cylinder array," in *Proceedings Of The Asme Pressure Vessels And Piping Conference*, 2015, vol. 5.

[26] D. S. Weaver and A. Abd-Rabbo, "A Flow Visualization Study of a Square Array of Tubes in Water Crossflow," *ASME*, vol. 107, no. September 1985, pp. 354–362, 1985.

[27] K. D. Hamman and R. A. Berry, "A CFD simulation process for fast reactor fuel assemblies," *Nucl. Eng. Des.*, vol. 240, no. 9, pp. 2304–2312, 2010.

[28] W. Raza and K. Y. Kim, "Effects of wire-spacer shape in LMR on thermal-hydraulic performance," *Nucl. Eng. Des.*, vol. 238, no. 10, pp. 2678–2683, 2008.

[29] E. Merzari, W. D. Pointer, J. G. Smith, A. Tentner, and P. Fischer, "Numerical simulation of the flow in wire-wrapped pin bundles: Effect of pin-wire contact modeling," *Nucl. Eng. Des.*, vol. 253, pp. 374–386, 2012.

[30] W. D. Pointer, J. Smith, A. Siegel, and P. Fischer, "RANS Simulations of Turbulent Diffusion in Wire-Wrapped Sodium Fast Reactor Fuel Assemblies," *Int. Conf. Fast React. Relat. Fuel Cycles Challenges Oppor. FR09*, pp. 249–250, 2009.

[31] L. Brockmeyer, L. B. Carasik, E. Merzari, and Y. Hassan, "Numerical simulations for determination of minimum representative bundle size in wire wrapped tube bundles," *Nucl. Eng. Des.*, vol. 322, 2017.

[32] P. Fischer, J. Lottes, A. Siegel, and G. Palmiotti, "large eddy simulation of wire-wrapped fuel pins I : Hydrodynamics in a Periodic Array," 2007.

[33] R. Ranjan, C. Pantano, and P. Fischer, *Direct simulation of turbulent swept flow over a wire in a channel*, vol. 54, no. 21–22. 2011.

[34] J.-H. Jeong, J. Yoo, K.-L. Lee, and K.-S. Ha, "Three-dimensional flow phenomena in a wire-wrapped 37-pin fuel bundle for SFR," *Nucl. Eng. Technol.*, vol. 47, no. 5, pp. 523–533, 2015.

[35]   P. J. Blau, "A multi-stage wear model for grid-to-rod fretting of nuclear fuel rods," *Wear*, vol. 313, no. 1–2. pp. 89–96, 2014.

[36]   K. T. Kim, "The study on grid-to-rod fretting wear models for PWR fuel," *Nuclear Engineering and Design*, vol. 239, no. 12. pp. 2820–2824, 2009.

[37]   E. Merzari *et al.*, "Benchmark exercise for fluid flow simulations in a liquid metal fast reactor fuel assembly," *Nucl. Eng. Des.*, vol. 298, no. March 2011, pp. 218–228, 2016.

[38]   B. De Pauw *et al.*, "Vibration monitoring using fiber optic sensors in a lead-bismuth eutectic cooled nuclear fuel assembly," *Sensors (Switzerland)*, vol. 16, no. 4, p. 571, 2016.

[39]   L. Brockmeyer, E. Merzari, A. Obabko, and Y. Hassan, "analysis of pressure field on wire-wrapped pin bundle surface for concerns of FSI," in *Proceedings of the ASME 2018 5th Joint US-European Fluids Engineering Summer Conference*.