

An Information-Theoretic Exploration of Multi-Body Fluid-Structure Interactions

by

Benjamin Pocock

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Mechanical and Mechatronics Engineering

Waterloo, Ontario, Canada, 2019

© Benjamin Pocock 2019

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Fluid-structure interactions are highly complex and difficult to resolve using classical physics tools. Recently, information theory has been proposed as an alternative for exploring these often highly non-linear and complex multi-physics problems. By viewing the fluid-structure system as a flow of information, valuable knowledge can be gained without needing to resolve the details of the interaction. Here, a combination of simulations and experiments are used to assess how the behaviour of the fluid-structure system can be recast into an information-theoretic framework. The proposed information-theoretic tool is transfer entropy, which takes time series of some aspect of the system behaviour, for example structural displacements, and infers directed casual relationships between the components. Furthermore, due to dependence on the lag parameter, transfer entropy can infer information about the critical time scales of the system. The first experiments consist of two structures, constrained to allow only unidirectional communication and designed to have enough known properties to validate the information-theoretic analysis. Additional experiments are performed that add a third structure, thereby increasing the number of communication pathways. This work shows that, while transfer entropy is model free, the order of operations and interpretation of the results are significantly impacted by the characteristics of the input data. As a result, it is imperative to have a clear research goal and an understanding of the critical behaviours of the system. With that, it is relatively inexpensive to gain valuable information from even the most complex fluid-structure systems. Overall, transfer entropy is proven to be a useful tool for the analysis of fluid-structure interactions, provided it is not naively applied and the data exhibits some degree of randomness.

Acknowledgements

I owe the completion of this work to many people. Without the support of those around me, I would never have made it through.

I would like to thank my supervisor and friend, Professor Sean Peterson, for his support at every step of this process. He is the one who brought me to Canada and he has always been there since to remind me not to “lose the forest for the trees.” I would also like to thank all of my colleagues in the fluids research group and the Second Heart research group. They set such high standards for the quality of work and level of effort, it was an inspiration to be around.

I owe special thanks to Michael Dickey who has kept me sane by listening to my rambling and providing his unwavering support. Also to Bjorn Dawson and Erik Marble who helped me get away from my desk and explore Canada.

Finally, thank you to my family, even though we were on different sides of the planet, knowing I had their support allowed me to take the chances that led me to where I am.

Table of Contents

| | |
|--|-------------|
| List of Tables | viii |
| List of Figures | x |
| Nomenclature | xiv |
| 1 Introduction | 1 |
| 1.1 Study Objectives and Thesis Overview | 2 |
| 2 Background | 4 |
| 2.1 Variables and Probability | 4 |
| 2.2 Information Theory Primer | 7 |
| 2.2.1 Information Content | 8 |
| 2.2.2 Shannon Entropy | 9 |
| 2.2.3 Conditional Entropy | 11 |
| 2.3 Transfer Entropy | 12 |
| 2.3.1 Communication Lag | 14 |
| 2.3.2 Sampling Rate | 18 |
| 2.3.3 Symbolisation | 22 |
| 2.3.4 Conditional Transfer Entropy | 23 |
| 2.3.5 History Length | 26 |
| 2.3.6 Equivalence to Granger Causality | 26 |
| 2.3.7 Calculating TE | 27 |
| 2.3.8 Expansion for Conditional TE | 29 |
| 2.3.9 Transfer Entropy Summary | 29 |
| 2.4 Fluid Dynamics | 30 |
| 2.4.1 Flow Around a Cylinder | 30 |
| 2.4.2 Flow Around an Airfoil | 34 |
| 2.4.3 Fluid Structure Interactions | 35 |

| | | |
|----------|--|-----------|
| 3 | Experimental Setup | 38 |
| 3.1 | Flow Facility | 38 |
| 3.2 | Models | 42 |
| 3.2.1 | Cylinder Model | 42 |
| 3.2.2 | Airfoil Models | 44 |
| 3.2.3 | Model Placement | 49 |
| 3.3 | Drivers and Electronics | 53 |
| 3.4 | Data Capture | 57 |
| 4 | Results of Simulated Experiments | 59 |
| 4.1 | Validation of TE Algorithm | 60 |
| 4.2 | Analysis of Calculation Parameters | 61 |
| 4.3 | Simulation Conclusions | 68 |
| 5 | Methodology and Results of Two Body Experiments | 70 |
| 5.1 | Preliminary Analysis | 70 |
| 5.2 | Methodology | 72 |
| 5.2.1 | Captured Data | 72 |
| 5.2.2 | Symbolization | 74 |
| 5.2.3 | Determine Test Length | 76 |
| 5.3 | Results | 77 |
| 5.3.1 | Down-Sample First Method | 77 |
| 5.3.2 | Symbolize First Method | 81 |
| 5.3.3 | Alternate Symbols | 89 |
| 5.3.4 | Check for Self-Feedback | 89 |
| 5.3.5 | Null Distribution | 91 |
| 5.4 | Single Pathway Conclusions | 92 |
| 5.5 | Cross-Stream Offset Experiment | 94 |
| 6 | Results of Three Body Experiments | 95 |
| 6.1 | Simulations | 95 |
| 6.1.1 | Mutual Driver | 96 |
| 6.1.2 | Relayed Information | 100 |
| 6.1.3 | Relayed Information With Active Relay | 102 |
| 6.1.4 | Multiple Pathways | 106 |
| 6.1.5 | Simulation Conclusions | 111 |
| 6.2 | Experiments | 111 |

| | | |
|----------|---|------------|
| 6.2.1 | Preliminary Analysis | 111 |
| 6.2.2 | Conditional TE | 114 |
| 6.2.3 | Path Significance | 117 |
| 6.2.4 | Hypothesis Testing Approach | 118 |
| 6.2.5 | Conclusions | 119 |
| 7 | Conclusions and Recommendations | 121 |
| 8 | Key Ideas | 124 |
| | References | 127 |
| A | Statistical Significance | 133 |
| A.0.1 | Bootstrapping | 134 |
| A.0.2 | Statistical Comparison to Zero | 136 |
| B | Derivation of Probabilistic TE | 137 |
| B.1 | Pairwise TE | 137 |
| B.2 | Conditional TE | 139 |
| C | Comparison to ARIMA Residuals Method | 140 |
| D | Additional Figures | 172 |
| E | Analysis of Pathways: Three Body Experiments | 182 |

List of Tables

- 3.1 Cylinder properties 43
- 3.2 Test conditions. 44
- 3.3 Mass ratios for various airfoil designs 48
- 3.4 Variables used in the voltage model, Eq. 3.3 56

- 5.1 Case numbers and parameters 72

- 6.1 Significant delays identified from pairwise TE 107
- 6.2 Delays related to indirect communication 109
- 6.3 Summary of delays for three body experiments. 115
- 6.4 Summary of delays and TE measurements 118

- A.1 Repeated Measures of TE 134
- A.2 Statistical Significance Comparison; Bootstrapping or no Bootstrapping . . 136

- C.1 Test Cases 143
- C.2 Results of an exhaustive study for case 1, U with parameter ranges $p = 1, 2, \dots, 10$; $q = 1, 2, \dots, 10$; $\lambda = -2.00, -1.75, \dots, 2.00$ 151
- C.3 Parameter estimates and standard errors for ARMA(9,6) and MA(1) models with $\lambda = 0$ 153
- C.4 Coarse exhaustive study for case 1, D 154
- C.5 Refined exhaustive study for case 1, D 154
- C.6 MLEs and SEs for an ARMA(7,8) model with $\lambda 0.5$ for case 1, D 155
- C.8 Refined constrained model AICs 156
- C.7 AICs for select constrained models fitted to case 1, D 156
- C.9 MLEs and SEs for the chosen model for C1, D 157
- C.10 Models selected by the coarse exhaustive study for cases 2 to 6. Note that the active airfoil cases do not meet the normal residuals condition 158
- C.11 Summary of significant lags predicted by the two methods ordered by significance. 165
- C.12 Parameter estimates and SEs for Case 2 168

| | |
|---|-----|
| C.13 Parameter estimates and SEs for Case 3 | 169 |
| C.14 Parameter estimates and SEs for Case 4 | 169 |
| C.15 Parameter estimates and SEs for Case 5 | 170 |
| C.16 Parameter estimates and SEs for Case 6 | 171 |
| E.1 Possible relay pathways. | 184 |

List of Figures

- 2.1 Intuitive example of causality. 8
- 2.2 Relationship between $H(X)$ and $p(x)$ for cardinality 2 11
- 2.3 $TE(\delta)$ for various two variable systems; adapted from Wibral *et al.* [66] 17
- 2.4 Example of how different sampling harmonics affect data representation 20
- 2.5 An example of down-sampling impulse data 21
- 2.6 Example of binning and symbolisation of data 23
- 2.7 Intuitive three variable system. 25
- 2.8 Cylinder wake. 32
- 2.9 Coordinate system and forces acting on the airfoil. 35
- 2.10 A schematic of a pivoted plate in a uniform flow with a passing parallel vortex 36

- 3.1 Hydraulic schematic of the test facility. 39
- 3.2 A rendered image of the flow facility. 39
- 3.3 Flow facility drawing 40
- 3.4 Isometric view of the test section with an example model installed. 41
- 3.5 Model of the plates which form the top of the test section. 42
- 3.6 Isometric and Section views of Cylinder Assembly. 43
- 3.7 Airfoil design drawing. 45
- 3.8 Airfoil pitching moment as a function of pitch angle 47
- 3.9 Renders of the lower airfoil section to highlight the design features. 49
- 3.10 Model Locations: In-line 50
- 3.11 Experiment Schematic 51
- 3.12 Model mounting location options: Array 52
- 3.13 Model driver render 54
- 3.14 Model driver subassembly render 54
- 3.15 Schematic of new control system 56
- 3.16 Voltage required by the stepper motor as a function of the rotational speed. 57
- 3.17 Motor mount rendering. 57
- 3.18 A rendering of the camera position and optical rails 58

| | | |
|------|---|----|
| 4.1 | A sample of the time series used in this simulation. | 60 |
| 4.2 | Comparison to analytical results | 61 |
| 4.3 | Scatter of TE results for simulated experiments | 62 |
| 4.4 | Scatter of TE_{net} results for simulated experiment | 63 |
| 4.5 | Shannon entropy versus Δ for Synthetic Data | 64 |
| 4.6 | TE results for select Δ 's. | 65 |
| 4.10 | Identified critical sample period using data driven methods | 65 |
| 4.11 | Identified critical communication delay using data driven methods | 65 |
| 4.7 | TE contour for $c = 0.2$ | 66 |
| 4.8 | TE contour for $c = 0.2$ with the entropy filter applied | 67 |
| 4.9 | Final results of $TE_{X \rightarrow Y}$ using data driven methods to select Δ and δ parameters. | 68 |
| 5.1 | Sample of raw data for single airfoil experiments | 73 |
| 5.2 | Sample of cylinder symbolization. | 74 |
| 5.3 | Sample of airfoil symbolization. | 75 |
| 5.4 | Test length convergence | 77 |
| 5.5 | TE results for case 1 | 78 |
| 5.6 | $ TE_{\text{net}} $ for all 15 cases calculated using the down-sample first method | 79 |
| 5.7 | Critical lag values from sample first method. | 80 |
| 5.8 | Entropy of variables for test case 1 using the down-sample first method. | 81 |
| 5.9 | Entropy of case 1 using the symbolize first method | 82 |
| 5.10 | TE results for case 1 | 83 |
| 5.11 | $ TE_{\text{net}} $ for all 15 cases calculated using the symbolise first method | 84 |
| 5.12 | TE of case 1 as a function of lag for a set sample period, $\Delta = T_D$ | 86 |
| 5.13 | Dominant communication delays of cases 1 to 12 with linear fit. The different colours indicate either the first or second peak. | 87 |
| 5.14 | $TE_{C \rightarrow A}$ as a function of separation between structures. | 88 |
| 5.15 | $TE_{C \rightarrow A}$ as a function of separation between structures. | 88 |
| 5.16 | Alternate symbol comparison. | 89 |
| 5.17 | TE results along self feedback pathways. | 90 |
| 5.18 | $TE_{C \rightarrow A}$ for case 1 with indicators of communication delays predicted from the self feedback results. | 91 |
| 5.19 | Mean of the null distribution for case 7 | 92 |
| 5.20 | TE vs cross-stream offset. | 94 |
| 6.1 | Communication pathways for mutually driven X and Y | 98 |

| | | |
|------|--|-----|
| 6.2 | TE results for simulation of three variables with Z mutually driving X and Y | 99 |
| 6.3 | Communication pathways for a system where information generated at Z is relayed to Y by X . The path $X \rightarrow Y$ is drawn starting at the end of the $Z \rightarrow Y$ to symbolize the idea that no information is generated at X | 101 |
| 6.4 | Pairwise TE results. For clarity, components with no measurable TE are not shown. | 102 |
| 6.5 | Conditional TE results for simulated system with three variables with a relayed pathway. | 102 |
| 6.6 | Communication pathways for a system where information generated at Z is relayed to Y by X | 104 |
| 6.7 | Pairwise TE results. For clarity, components with no measurable TE are not shown. | 104 |
| 6.8 | Conditional TE results for simulated system with three variables with information generation at Z and X | 105 |
| 6.9 | Identified communication pathways for the system described in Fig. 6.6 | 105 |
| 6.10 | Communication pathways for a system where information generated at Z is relayed to Y by X and also directly. | 106 |
| 6.11 | Pairwise TE results. For clarity, components with no measurable TE are not shown. | 107 |
| 6.12 | Conditional TE results for simulation of three variables with both mutual driver and cascade effects. | 108 |
| 6.13 | Identified communication pathways for the system described by Fig. 6.10. | 108 |
| 6.14 | Identified communication pathways for the system described by Fig. 6.10. | 110 |
| 6.15 | Relative placements of structures. | 112 |
| 6.16 | Pairwise TE results for Test T14 | 113 |
| 6.17 | Feedback results for Test T14 | 113 |
| 6.18 | Dominatant communication delays | 114 |
| 6.19 | Causality map | 116 |
| 6.20 | Conditional TE results for three structure experiment. | 116 |
| 6.21 | Conditional TE results for three structure experiment. | 117 |
| A.1 | An overview of the case resampling bootstrapping method | 135 |
| C.1 | Schematic of airfoil placement in experimental setup. | 142 |
| C.2 | The raw observation series for case 1, U | 144 |
| C.3 | Sample ACF, PACF and data distribution for case 1, U | 145 |
| C.4 | The raw observation series for case 1, D | 146 |
| C.5 | Sample ACF, PACF and data distribution for case 1, D | 147 |

| | | |
|------|--|-----|
| C.6 | Sample ACF, PACF and histogram of transformed data with $\lambda = 0$ for case 1, U | 152 |
| C.7 | Residual ACF and histogram for case 1, D after applying a Box-cox transformation $\lambda = 0.50$ and constrained ARMA(7,8) model. | 158 |
| C.8 | Overlaying the time series of case 1, U and case 1, D. | 159 |
| C.9 | a. Cross correlation of the residuals of case 1, U and case 1, D. b. Transfer Entropy of case 1 | 161 |
| C.10 | a. Cross correlation of the residuals of case 2. b. Transfer Entropy of case 2 | 162 |
| C.11 | a. Cross correlation of the residuals of case 3. b. Transfer Entropy of case 3 | 162 |
| C.12 | a. Cross correlation of the residuals of case 4. b. Transfer Entropy of case 4 | 163 |
| C.13 | a. Cross correlation of the residuals of case 5. b. Transfer Entropy of case 5 | 163 |
| C.14 | a. Cross correlation of the residuals of case 6. b. Transfer Entropy of case 6 | 164 |
| C.15 | A comparison of the affect of the two most likely U models on the RCCF for case 1 | 167 |
| D.1 | Electronics Schematic: cRio. | 172 |
| D.2 | Entropy as a function of sampling period for single airfoil experiments; down-sample first method | 173 |
| D.3 | $TE_{C \rightarrow A}$ for all 15 cases of the single airfoil experiments, calculated using the down-sample first method. | 174 |
| D.4 | $TE_{A \rightarrow C}$ for all 15 cases of the single airfoil experiments, calculated using the down-sample first method. | 175 |
| D.5 | Entropy as a function of sampling period for single airfoil experiments; symbolise first method | 176 |
| D.6 | $TE_{C \rightarrow A}$ for all 15 cases of the single airfoil experiments, calculated using the symbolise first method. | 177 |
| D.7 | $TE_{A \rightarrow C}$ for all 15 cases of the single airfoil experiments, calculated using the symbolise first method. | 178 |
| D.8 | TE as a function of δ , for a single Δ | 179 |
| D.9 | Self feedback of A (orange line) and C (blue line) compared to the null threshold (black dashed line). Data is presented for a single Δ and a range of δ for all 15 cases. | 180 |
| D.10 | $TE_{C \rightarrow A}$ for all cases with indicators of communication delays predicted from the self feedback results. | 181 |
| E.1 | Map of identified communication paths | 186 |
| E.2 | Simplified map of identified communication paths | 186 |
| E.3 | Simplified map of identified communication paths | 187 |

Nomenclature

Dimensionless Quantities

| Symbol | Description | Definition |
|--------|-----------------|---------------------|
| Re | Reynolds number | $\frac{U_0 c}{\nu}$ |
| St | Strouhal Number | $\frac{f_f D}{U}$ |

Symbols

| Symbol | Description | Units |
|---------------------|--|-------|
| B | Buoyancy force | N |
| Δ | Sample period | s |
| Δ' | Sample period as number of indices | - |
| δ' | Primary lag parameter as number of indices | - |
| δ_{XY} | Primary lag parameter from source variable X to target variable Y | s |
| $\hat{\delta}_{XY}$ | Dominant communication delay from variable X to Y , identified by local maxima of $TE_{X \rightarrow Y}$ | s |
| η | Information content | bits |
| f_f | Fundamental shedding frequency | Hz |
| $H(\dots)$ | Shannon entropy of the variable listed in brackets | bits |
| l | Airfoil chord length | m |
| Ω_X | Sample space of variable X | - |
| $\omega_{X,i}$ | An event. See Section 2.1 for description of subscripts | N/A |

| | | |
|--------------------------|--|-------------------|
| $Pr(\dots)$ | Probability of event listed in brackets | - |
| s | Separation distance between structures | m |
| T_0 | Event period | s |
| T_A | Acquisition period | s |
| T_C | Convective time scale between structures | s |
| T_D | Decision period | s |
| $TE_{X \rightarrow Y}$ | Transfer entropy from source variable X to target variable Y | bits |
| $TE_{X \rightarrow Y Z}$ | Conditional transfer entropy from source variable X to target variable Y , conditioned on variable Z | bits |
| T_S | Shedding Period | s |
| $TE_{X \rightarrow Y}$ | Transfer entropy from X to Y | bits |
| U_0 | Free-stream velocity | m s^{-1} |
| V | Volume | m^3 |
| W | Weight force | N |
| ζ_{ZY} | Conditional Lag parameter from conditioning variable Z to target variable Y | s |

Chapter 1

Introduction

Fluid-structure interactions (FSI) are prominent in many engineering and scientific fields. They are typically highly complex, involving multi-disciplinary physics and non-linear dynamics [25, 46]. This makes it exceedingly difficult in most cases to disentangle and understand the underlying physics. Few FSI systems allow for analytical solution, those that do are only made possible by making significant simplifying assumptions. As such, it is more common to apply numerical methods, which must resolve the physics of the structural dynamics, throughout the fluid and at the interface. Time dependent and non-periodic interactions are particularly expensive to solve due to the large number of variables. These problems lend themselves to data driven methods which can focus on the behaviours of interest without having to resolve all of the details of the interactions.

Data driven time-series analysis can be used to focus analyses by identifying cause and effect relationships. This has long been performed through correlation of time-series. However, correlation is rarely a suitable indicator of causality. Additionally, most methods for doing this require the fitting of a model, which is expensive and often ill-suited to non-linear, non-normally distributed data [24, 31, 61]. Information theory can provide a model free method for dealing with non-linear data [8]. Information is a contextual measure based on the likelihood of a variable being in a specific state [13]. Measuring the mutual information of two variables, that is the amount of non-unique information, is analogous to correlation, but, being symmetric, cannot be used to differentiate the direction of any implied interaction [42]. Neither correlation method is equivalent with causality, so another measurement is sought.

Transfer entropy (TE) is an asymmetric, model free, non-linear measure that has been used in a range of fields to identify directed causal relationships as a function of the

communication delay [9, 32, 41, 59, 63, 64, 66, 67]. Coupled with some knowledge of the system, this can be used to identify the communication pathways and other time-scale information. It is based on the information-theoretic property, Shannon entropy. This property relies on uncertainty about the state of a variable, meaning TE is only applicable for non-deterministic systems.

One of the key challenges of applying transfer entropy to fluids data is the structure of the data itself. TE was developed for data with a more natural information-theoretic interpretation, such as text characters. In these data, the temporal evolution of an event is unimportant and the information is naturally discretised. With the data common to fluid dynamics, the temporal dimension of an event is not always easily defined and is susceptible to change. For example, a pressure blast with a very short interval may cause an oscillatory motion response of a nearby structure. In this scenario, a short message causes a long response, something that is uncommon in information-theory. For this reason, a clear hypothesis and a reasonable understanding of the information-theoretic concepts must be obtained before the method of recasting fluid data in an information-theoretic framework can be developed.

Zhang *et al.* [74] took the first steps toward analysing the interactions of fish with their surrounding fluid using transfer entropy (TE). Their motivation is to disentangle the fish's other senses from the fluid interaction, allowing for a better focused FSI investigation. They use a simplified experimental setup to confirm that TE is an appropriate tool for identifying causal relationships between two structures that are communicating via a fluid, using only the time-series of the structure displacements. The following work builds upon the framework set out by Zhang *et al.* [74] by adding a third structure to the analysis, thereby increasing the possible interactions. Additionally, this work begins to explore what other knowledge can be obtained by measuring TE, as well as the broader applicability of TE to other FSI problems. The focus here is systems with dynamic structures, using the structural displacements as the only input. However, it is noted that it is likely possible to extend the method to other information, such as force or direct fluid measurements.

1.1 Study Objectives and Thesis Overview

The objective of this work is to explore what knowledge can be obtained about complex FSI systems using information theoretic tools. Specifically, it will be focused on the application of transfer entropy to two and three body FSI problems.

To assess the applicability of transfer entropy to FSI problems, one must first develop an understanding of the nuances of its application. It is acknowledged that most readers will have limited prior knowledge of the general concepts of information theory, so this document begins with a focused introduction to information theory (Section 2.1, Section 2.2) before discussing the particulars of transfer entropy specifically (Section 2.3). This introduction draws from published literature as well as worked experiments to demonstrate some of the intricacies and, where possible, create an intuitive understanding of the concepts. Section 2.4 gives a very brief overview of the applicable fluid phenomena to accompany the experiment design (Chapter 3) and to set expectations of the experiment results to compare to the information-theoretic analysis.

The methodology for measuring transfer entropy from the type of data available in FSI systems is not fully developed. This work uses simulations (Chapter 4) and two body experiments (Chapter 5) to compare different methods suggested in the literature and to fill in the gaps. The methodology is then expanded to three body experiments (Chapter 6) to explore the use of conditional transfer entropy.

Chapter 2

Background

2.1 Variables and Probability

This section is designed to define some of the concepts that will provide a foundation for discussing information theory and to introduce the notation that will be used for the remainder of this document. For general introductions to probability and statistics, the reader is directed to the works of Bossomaier *et al.* [8], Dekking *et al.* [16], Dixon [17] and Schay [49].

In the broadest sense a variable is any value that can be manipulated. Here we will be talking about discrete random variables, denoted by an italicised capital letter, which record a number of outcomes. When referring to an arbitrary outcome in a variable, the lower case letter will be used. A subscript value can be added to refer to a specific outcome. For example, a coin is flipped five times giving the results X such that $X = [H, T, T, H, T]$, where H is used to indicate heads was showing and T to indicate tails. The fourth outcome can be referred to as x_4 . The number of outcomes or the length of X is denoted $|X|$. In general:

$$X = [x_1, x_2, x_3, \dots, x_N]$$

Many of the variables used here will actually be time series. The only difference is that the order of the outcomes is related to the time at which the outcome was observed. In this sense it can be said that a time series variable contains time information as well as the observation. For example, if X contained observations of the position of the sun in the sky taken every hour, with the first observation (x_1) being taken at midday, then observation x_3 would contain the position of the sun at 2 pm. When it is important to differentiate a

variable as a time series, a subscript t will be added. Time shifts, which will arise in later discussions, are treated naturally via index shifting, such as

$$\begin{aligned} X_t &= [x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9] \\ X_{t-1} &= [-, x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8] \end{aligned}$$

The values that an outcome can take are called events. In our example series there are 5 outcomes and 2 events, heads or tails. The collection of events is called the sample space which has the general form:

$$\Omega_X = \{\omega_1, \omega_2, \omega_3, \dots, \omega_M\}$$

where Ω_X is the sample space of X containing all of the possible events, ω . Specific events are referred to using a subscript. The outcomes of the variable must take one of the values of the sample space, so for the example $\Omega_X = \{H, T\}$.

For much of the work here, multiple variables will be operated on at a time. In these cases the sample space contains all possible combinations of the states in the variables. For an example using two coins, X and Y , the individual sample spaces are $\Omega_X = \Omega_Y = \{H, T\}$. For both coins the sample space is $\Omega_{XY} = \Omega_X \times \Omega_Y = \{(H, H), (H, T), (T, H), (T, T)\}$ where each event now includes two states; for example $\omega_2 = (H, T)$. The order of the values inside the event is very important as they relate to a specific variable and cannot be mixed. As such the general notation is

$$\Omega_{XY} = \{(\omega_{X,1}, \omega_{Y,1}), (\omega_{X,2}, \omega_{Y,2}), \dots, (\omega_{X,M}, \omega_{Y,M})\} \quad (2.1)$$

where the subscript of ω denotes both the variable and the event to which it belongs. For countable variables, such as those dealt with here, the probability that a variable takes on the value of a specific event can be estimated by the number of times that event occurs as an outcome.

$$\Pr(X = \omega_i) = \frac{N(X = \omega_i)}{|X|} \quad (2.2)$$

where $N()$ is the number of occurrences of the event inside the bracket and $|X|$ is the length of X . This should be read as the probability that the event ω_i occurs in X is equal to the number of times that event occurs divided by the number of outcomes. In the single coin example variable, the event H occurs twice, $N(\omega_1) = 2$, in 5 outcomes, $|X| = 5$. Therefore:

$$\Pr(X = H) = \frac{2}{5}$$

Notice that this does not exactly match the known outcome, $\Pr(X = H) = 0.5$,

because the sample length is so short. To achieve the expected outcome requires $|X| \rightarrow \infty$.

When dealing with more than one variable it is sometimes desirable to calculate joint or conditional probabilities which account for all variables. Joint probability is the probability of two events occurring, written as $\Pr(X = \omega_X, Y = \omega_Y)$, read as the probability that event ω_X occurs in X and event ω_Y occurs in Y . When the two variables are independent, the joint probability is the product of the two individual probabilities. In the two coin example, the probability of both coins showing heads¹ is $\Pr(X = H, Y = H) = \Pr(X = H) \Pr(Y = H) = \frac{1}{4}$.

Conditional probability is the probability of an event given that another event is occurring. Written as $\Pr(X = \omega_X | Y = \omega_Y)$, read as the probability that event ω_X occurs in variable X given that ω_Y has occurred in Y . For example, imagine the two coins are glued together so that when one shows heads, they both show heads. The sample space is still the same but there is no chance of one coin showing heads at the same time as the other shows tails, $\Pr(X = H | Y = T) = \Pr(X = T | Y = H) = 0$. The remaining probabilities must be $\Pr(X = H | Y = H) = \Pr(X = T | Y = T) = 1$. The joint probabilities can then be calculated using Eq. 2.3 [16].

$$\Pr(X = \omega_X, Y = \omega_Y) = \Pr(X = \omega_X | Y = \omega_Y) \Pr(Y = \omega_Y) \quad (2.3)$$

Notice that in this system, with the coins glued together, the joint probability of both coins showing heads is $\Pr(X = H, Y = H) = 0.5$. This differs from the previous case because they are no longer independent.

In the following work, the solutions are not so intuitive. To explain all of these concepts in a more relevant way we will use a less tired example. Someone wants to know how the probability of seeing a kangaroo changes depending on the time of day so they make three recordings per day, one in the morning, one in the afternoon and one in the evening. The results are stored in variable K with the events defined as:

$$k_t = \begin{cases} 1, \text{Kangaroo was seen during time } t \\ 0, \text{Kangaroo was not seen during time } t \end{cases}$$

¹Using the analytical results $\Pr(X = H) = \Pr(Y = H) = 0.5$

The time of day is stored in another variable D which follows the rules

$$d_t = \begin{cases} \text{M, Observation } t \text{ was taken in the morning} \\ \text{A, Observation } t \text{ was taken in the afternoon} \\ \text{E, Observation } t \text{ was taken in the evening} \end{cases}$$

The sample spaces of the variables individually are $\Omega_K = \{0, 1\}$ and $\Omega_D = \{\text{M,A,E}\}$. The combined sample space $\Omega = \{(0,\text{M}), (0,\text{A}), (0,\text{E}), (1,\text{M}), (1,\text{A}), (1,\text{E})\}$. Hopefully it is clear that the variables K and D are linked by the time of the observation so an outcome of K needs to be compared to its companion in D . To calculate any of the probabilities (basic, joint or conditional), the number of occurrences of the events must be found by counting. Consider the example results from the first 9 observations:

$$\begin{aligned} K &= [1, 0, 0, 0, 0, 1, 0, 0, 1] \\ D &= [\text{M}, \text{A}, \text{E}, \text{M}, \text{A}, \text{E}, \text{M}, \text{A}, \text{E}] \end{aligned}$$

The event ‘a kangaroo is seen,’ occurs three times in nine observations, $\Pr(K = 1) = \frac{1}{3}$. The event (1,M) occurs once in nine observations so the joint probability of that event is $\Pr(K = 1, D = \text{M}) = \frac{1}{9}$. Similarly the conditional probability is $\Pr(K = 1|D = \text{M}) = \frac{1}{3}$ because from the 3 samples where $D = \text{M}$, $K = 1$ occurs once.

2.2 Information Theory Primer

At first it can be difficult to visualise how anything can be learned by looking at systems as information events which we have defined as contextual and devoid of meaning. The following example is intended to demonstrate how powerful this is in an intuitive way. Figure 2.1 shows two arbitrary time series. It does not matter what the different states are or what is being measured, only how they compare within the system. Without any math, it can be seen that it is likely that variable X is causing a response in Y at a delay (lag) of 19 time steps. This is deduced from the following facts:

1. The behaviour (or pattern) of X and Y are the same.
2. Changes in X occur before changes in Y , so Y is not causing X .
3. The lag between an event occurring in X to it occurring in Y is constant.

This is simply lagged correlation, which is not equivalent with causality, but it is an intuitive, information-theoretic, data driven measure. Of course, this is a highly simplified




Figure 2.1: Time series for variables X and Y to intuitively demonstrate how causality can be inferred from time series.

case but the idea that something can be learned about a system simply by looking at the time series of some behaviour, without even knowing what the system is, is central to the field of information theory. Suppose that system was the displacement of two structures in a fluid. Without resolving the fluid flow, it is possible to infer that they are interacting and to estimate the time it takes for the communication to occur. In the following sections these concepts will be formalised.

2.2.1 Information Content

In information theory, information and meaning are entirely separate ideas [51]. We could send two emails, the first saying "Hello", the second saying "eIHol". The information of both is identical but the first has meaning and the second is gibberish. Mathematically information content, or simply information, is a measure of the likelihood of an event occurring. It is a property of an event defined as [28].

$$\eta(\omega) = -\log_a(\Pr(\omega)) \quad (2.4)$$

where a defines the units of the result. When information is calculated by a natural logarithm, the units are nats. For all of the work here $a = 2$ and information is measured in bits.

A more colloquial understanding can be obtained from thinking about it as a measure of how significant some information is. Using the kangaroo example from earlier, let us say that the kangaroo spotter spent 10 years recording whether or not they saw a kangaroo, but they were taking the observations in downtown Toronto. Intuitively we would say that seeing a kangaroo in Toronto was far more significant information than seeing one in Australia because it is less likely or because it occurs less often; this also works out to be mathematically true. Assume in those 10 years, 10959 observations are made and a kangaroo was spotted once. From Eq. 2.2, $\Pr(K_{\text{Toronto}} = 0) = \frac{10958}{10959}$, $\Pr(K_{\text{Toronto}} = 1) = \frac{1}{10959}$ so the information generated by that event (from Eq. 2.4) is $\eta(K_{\text{Toronto}} = 1) \approx 13$ bits. By comparison, seeing one in outback Australia is fairly likely; assume $\Pr(K_{\text{Australia}} = 1) = \frac{1}{2}$, then $\eta(K_{\text{Australia}} = 1) = 1$ bit. This demonstrates how the measure of information is related to uncertainty of the outcome, and it shows that the same event can have different information content in a different system; in Australia the event ‘a kangaroo is seen’ contains 1 bit of information, in Toronto it contains 13 bits.

The idea that the same event in two different data sets can have entirely different information contents is an important property to understand. Formally, it is stated that information is a contextual measure meaning the measured value only has meaning in the context in which it was calculated. This is not very intuitive, since we are used to dealing with measures like energy where 1 J in this system is the same as 1 J in that system. The difference being that the value of information is measured against the data set in which it exists.

It will soon be shown that this property will propagate through to transfer entropy. The consequence that has the biggest effect on this work is that it becomes very difficult to compare the results obtained from different data sets. We never actually compare information content so that discussion is saved for later.

2.2.2 Shannon Entropy

This section starts with an important distinction; Shannon entropy is not thermodynamic entropy.² The moniker Shannon entropy is often used when clarity is needed but for

²Claude Shannon formulated Shannon entropy in 1948 [51], some hundred years after the idea of thermodynamic entropy became popular. Tribus & McIrvine [60] detail an interview with Shannon about why he would call his tool the same name as something already in use. He explains that he had originally called it uncertainty but that name was already in use in statistical mechanics. The suggestion of entropy came from John von Neumann who asserted that "no one knows what entropy really is, so in a debate you will always have the advantage."

this document it is sufficiently clear to simply call it entropy. Entropy is the amount of uncertainty contained in a variable, that is how uncertain the observer is about what state the variable will be in. It is described mathematically as the expected value of the information content of all of the possible events.

$$H(X) = \langle \eta(\omega) \rangle \quad (2.5)$$

Substituting Eq. 2.4 into Eq. 2.5:

$$H(X) = - \sum_{\omega \in \Omega} \Pr(X = \omega) \log_2 (\Pr(X = \omega)) \quad (2.6)$$

where X is the variable and the summation occurs over all possible states ω contained in the sample space Ω . This sounds confusing but is actually fairly straightforward. Let us continue the kangaroo analogy from before. We created a variable K containing kangaroo sightings in downtown Toronto and we are reasonably certain that no kangaroos will be sighted. Intuitively, if entropy is a measure of uncertainty then entropy of this data set should be low. Mathematically, this is proven from the probabilities of seeing/ not seeing a kangaroo; $\Pr(K_{Toronto} = 0) \rightarrow 1$, $\Pr(K_{Toronto} = 1) \rightarrow 0$. Therefore from Eq. 2.6:

$$H(K_{Toronto}) = -0 \log_2 0 - 1 \log_2 1 \rightarrow 0$$

Repeating this experiment in the Australian outback where the probability of seeing a kangaroo is about 0.5, we are very uncertain if we will see a kangaroo on any given day: the entropy is high.

$$H(K_{Australia}) = -0.5 \log_2 0.5 - 0.5 \log_2 0.5 = 1$$

To show that the result is only dependant upon the likelihood of an event, not the value it takes, the experiment is repeated again at a kangaroo sanctuary where we are almost certain of seeing a kangaroo; $\Pr(K_{Sanctuary} = 0) \rightarrow 0$, $\Pr(K_{Sanctuary} = 1) \rightarrow 1$.

$$H(K_{Sanctuary}) = -1 \log_2 1 - 0 \log_2 0 \rightarrow 0$$

It is important to understand that an outcome is most uncertain when the probability of any outcome is uniform, i.e when $\Pr = 1/|\Omega|$. For the case where cardinality is 2 (K can only be one of two values), uncertainty is maximum when the probability of either event

is 0.5 as shown in Fig. 2.2.




Figure 2.2: The relationship between $H(X)$ and $p(x)$ for variables of cardinality 2.

2.2.3 Conditional Entropy

In a manner analogous to conditional probability, conditional entropy accounts for another *conditioning variable*. That is, the conditional entropy of X given Y ($H(X|Y)$) is the uncertainty about the state of X when the state of Y is known [28]. The derivation of the probabilistic form (Eq. 2.7) is given as part of the derivation of transfer entropy in Appendix B.1 and in Cover & Thomas [13].

$$H(X|Y) = - \sum_{x,y \in \Omega} \Pr(X = x, Y = y) \log_2 \Pr(X = x|Y = y) \quad (2.7)$$

where x and y would be more correctly written as ω_x and ω_y based on the previous discussion but since the two statements are technically equivalent and the single letter is less cumbersome, that is the common notation.

To demonstrate the use of conditional entropy, the kangaroo example is once again employed. In Section 2.1, observations were made at three different times of day for three days. Substituting the joint and conditional probabilities into Eq. 2.7 gives a conditional entropy of $H(K|D) = 1.14$ bits. Without some reference or comparison this number is meaningless because the magnitude is contextual. A valid comparison would be to $H(K)$;

if the conditional entropy and the unconditioned entropy are the same then no change in uncertainty was created by the knowledge about the conditioning variable. In other words, the time of day has no effect on kangaroo sightings if $H(K|D) = H(K)$. For the example data $H(K) = 2.20$ bits so uncertainty about K is reduced by knowing D and it can be said that the time of day does affect kangaroo sightings. It cannot yet be said if this change is statistically significant but that will be covered in Appendix A.

Conditional entropy can similarly be used to measure the dependence of the current state of a variable on its previous state by replacing the conditioning variable with a lagged version of the main variable.

$$H(X_t|X_{t-1}) = - \sum_{x, \dot{x} \in \Omega} \Pr(X_t = x, X_{t-1} = \dot{x}) \log_2 \Pr(X_t = x|X_{t-1} = \dot{x}) \quad (2.8)$$

where the dot over \dot{x} indicates that it relates to the lagged variable X_{t-1} . The time series notation is the same as that in Section 2.1. The example data is not long enough to be useful here but it is not difficult to imagine that this form could be used to see if the uncertainty of seeing a kangaroo today would be changed by knowing if one was seen yesterday.

All of the concepts covered so far can be applied to the example from the start of this section. Looking only at variable Y in Fig. 2.1 we would be very uncertain about what state Y will be in at time step 101 (y_{101}), that is, the entropy of Y ($H(Y)$) is high. However, we can see that Y seems to be following X with a lag of 19 so considering both variables, we would be reasonably certain that the state of y_{101} will be the same as $x_{82} = 0$. This means the conditional entropy of Y given X ($H(Y|X)$) is low. More technically, the conditional entropy of Y given X delayed by 19 time steps ($H(Y_t|X_{t-19})$) is low. Entropy measured at other time lags will be higher because the correlation will be less; this is a property that will be discussed further in Section 2.3.1.

2.3 Transfer Entropy

This section contains a review of the theoretical aspects of transfer entropy. Some of the content is obtained from published literature, as cited, but much of it is based on logical extensions of the general concepts or the results of some experimentation, as supported by worked examples or thought experiments; it is designed to form the foundation of the discussions of experimental results.

Transfer entropy (TE) is an information-theoretic measurement which can be used to infer the existence of directed causal relationships between pairs of variables. Since its inception by Schreiber [50] in 2000, it has been used to identify functional circuits in brains [64, 66, 67], infer cardio-respiratory interaction [63], analyse climate change [63], financial markets [32] and animal behaviour [9, 41, 59]. Parametric methods for achieving similar results are also widely used but they are typically time consuming to apply and are often sensitive to non-normal distributions of the input data. These points are discussed in detail in Appendix C. A popular data driven method for inferring causality is to measure mutual information. However, this method is symmetrical so the direction of the causal relationship is indistinguishable. Additionally, Paluš *et al.* [42] demonstrated that mutual information is not equivalent to causality. Transfer entropy has the benefit of being both data driven, asymmetric and equivalent with Granger causality (see Section 2.3.6).

In the discussion of conditional entropy it was shown that the raw value had little physical meaning; it is the amount of uncertainty about one variable given knowledge about another. However, the difference between the entropy of X and the conditional entropy of X given Y results in the reduction of uncertainty in X given the influence of Y . This is called mutual information, I , and is analogous to, but not equivalent to, causality.

$$I(X : Y) = H(X) - H(X|Y) \quad (2.9)$$

In some cases this measure can be used to infer causality but it does not take into account the shared history of the two variables, limiting its applicability. Good examples of this are provided in Bossomaier *et al.* [8] and Kaiser & Schreiber [29]. By adding the condition of previous states³ to the difference between entropy and conditional entropy, the formal description of transfer entropy is obtained.

$$\text{TE}_{Y \rightarrow X} = H(X_t|X_{t-1}) - H(X_t|X_{t-1}, Y_{t-\delta}) \quad (2.10)$$

The notation here $\text{TE}_{Y \rightarrow X}$ is read as the transfer entropy from Y to X . The first variable is called the source variable, the second is the target variable. The parameter δ is the lag as shown in Fig. 2.1 and discussed in Section 2.3.1. As previously mentioned, the measure is asymmetric so this is a measure of the communication from Y to X not from X to Y .

³See Section 2.3.5 for a discussion of the history length. The method here accounts for only one previous state.

Measurements in the other direction are performed by swapping the variables:

$$\text{TE}_{X \rightarrow Y} = H(Y_t|Y_{t-1}) - H(Y_t|Y_{t-1}, X_{t-\delta}) \quad (2.11)$$

These can be interpreted as the reduction of uncertainty about the state of the target variable caused by knowledge about the state of the source variable, while also considering the shared history. Due to the equivalence of TE with Granger causality, it can be said that statistically significant TE values indicate the presence of a significant causal relationship. The asymmetry of TE allows for the direction of the relationship to be inferred.

2.3.1 Communication Lag

In Eq. 2.10, it is seen that three variables, $(X_t, X_{t-1}, Y_{t-\delta})$, are required to calculate $\text{TE}_{Y \rightarrow X}$. The process of time shifting variables is demonstrated in Section 2.1. Once shifted, the start of the series are trimmed to the same length such that the first time index is $\tau = \delta + 1$, as shown by the boxed section below (assuming $\delta = 3$).

$$\begin{array}{rcl} X_t & = & [x_1, \quad x_2, \quad x_3, \quad \boxed{x_4, \quad x_5, \quad x_6, \quad x_7, \quad x_8, \quad x_9}] \\ X_{t-1} & = & [-, \quad x_1, \quad x_2, \quad \boxed{x_3, \quad x_4, \quad x_5, \quad x_6, \quad x_7, \quad x_8}] \\ Y_{t-3} & = & [-, \quad -, \quad -, \quad \boxed{y_1, \quad y_2, \quad y_3, \quad y_4, \quad y_5, \quad y_6}] \end{array}$$

So far the lag parameter δ has been treated as a simple quantity corresponding to the time it takes for information created at the source to be transferred to the target as shown in Fig. 2.1. This definition is sufficient for the original purpose of TE, where one bit of information fit inside one time instance, but for systems where an event spans a finite time, a different analysis is required.⁴ In reality, what is obtained from Eq. 2.10 and Eq. 2.11 are the TE values at the defined lag. Generally, it will not be known what the communication time is, so there is no way to know what the lag parameter should be. It is often sufficient to simply determine whether TE is significant at any δ , thereby identifying the existence of a directed causal relationship. Wibral *et al.* [66] and Zhang *et al.* [74] infer additional information about the system by estimating the communication time from the dependence of TE on the lag parameter; although their interpretation methods differ.

⁴In the original definition, the system was entirely information. Using a binary example, either a 0 or 1 would be sent by the source at every time instance and the information received by the target would also be in the form 0 or 1. In a physical system it is possible for the information to transform as it is transmitted. For example, a step impulse to the end of a simple cantilever beam insights an oscillatory response.

Zhang *et al.* [74] selects the communication delay, $\hat{\delta}$, by identifying the δ which maximises $|\text{TE}_{\text{net}}|$ (Eq. 2.12). Logically, this is mostly applicable to simple systems where communication occurs in one direction because, if the two communication periods are similar, the two peaks will cancel each other out.

Wibral *et al.* [66] uses a similar method but uses the maximums of the TE components to identify $\hat{\delta}$; i.e. $\hat{\delta}_{XY}$ is the lag parameter that maximises $\text{TE}_{X \rightarrow Y}$. This allows for bidirectional communication but is more susceptible to spurious results.

$$\text{TE}_{\text{net}} = \text{TE}_{X \rightarrow Y} - \text{TE}_{Y \rightarrow X} \quad (2.12)$$

Wibral *et al.* [66] supported their method by simulating a range of systems with two variables, exploring the various possible information pathways, as shown in Fig. 2.3. They correctly predict the communication times of those systems by selecting the local peaks of each TE component independently. In all cases the local peaks of $\text{TE}(\delta)$ match the actual communication times, confirming that this method is viable for a range of systems. However, they approach the analysis with *a priori* knowledge of the systems behaviour. For example in case b they conclude that the existence of multiple local peaks corresponding to the multiple communication times is proof of their method. In case a there are also multiple peaks ($\delta = 14, 20, 23$) but the side peaks are ignored as noise because they knew that there was only one communication path. Going into the analysis blind could potentially lead to the conclusion that case a has multiple communication paths. It is not possible to state whether those peaks would be removed by using TE_{net} because some of the data $\text{TE}_{Y \rightarrow X}$ is missing.

The broad spectrum communication, case c can be defined either by multiple communication paths that are very close together or as an inconsistent communication time over a range of lags. These two definitions are almost the same in the discrete domain. In this case Fig. 2.3 c shows that the $\text{TE}(\delta)$ spectrum has a broad peak (compared to a) covering all of the communication times. From this it is reasonable to say that communication occurs over some range but it becomes very difficult to select a single TE value to use for comparison with other systems.

Cases with a feedback loop as shown in Fig. 2.3 e create an interesting effect. Looking only at the components $\text{TE}_{X \rightarrow Y}$ and $\text{TE}_{Y \rightarrow X}$ in Fig. 2.3 e it is seen that the communication path $X \rightarrow Y$ is identified as before but additional peaks are seen at $\delta_{X \rightarrow Y} \pm \delta_{X \rightarrow X}$. This could easily be misinterpreted as additional paths from $X \rightarrow Y$ but Wibral *et al.* [66] offers a useful check for feedback. Figure 2.3 f shows the transfer entropy from X along the

path $X \rightarrow X$ ($TE_{X \rightarrow X}$). This component has a clear peak at the feedback communication time indicating that this check should be performed before concluding that a system has multiple communication paths. One factor not discussed by Wibral *et al.* is the effect of periodicity in the system. When the source variable is periodic, $TE(\delta)$ is also periodic⁵ and it becomes impossible to correlate a response to a specific causal event. The only conclusion that can be drawn in this case is that the actual communication time is some integer multiple of δ_{peak} .

⁵If both variables are periodic, TE is zero as the system is deterministic.

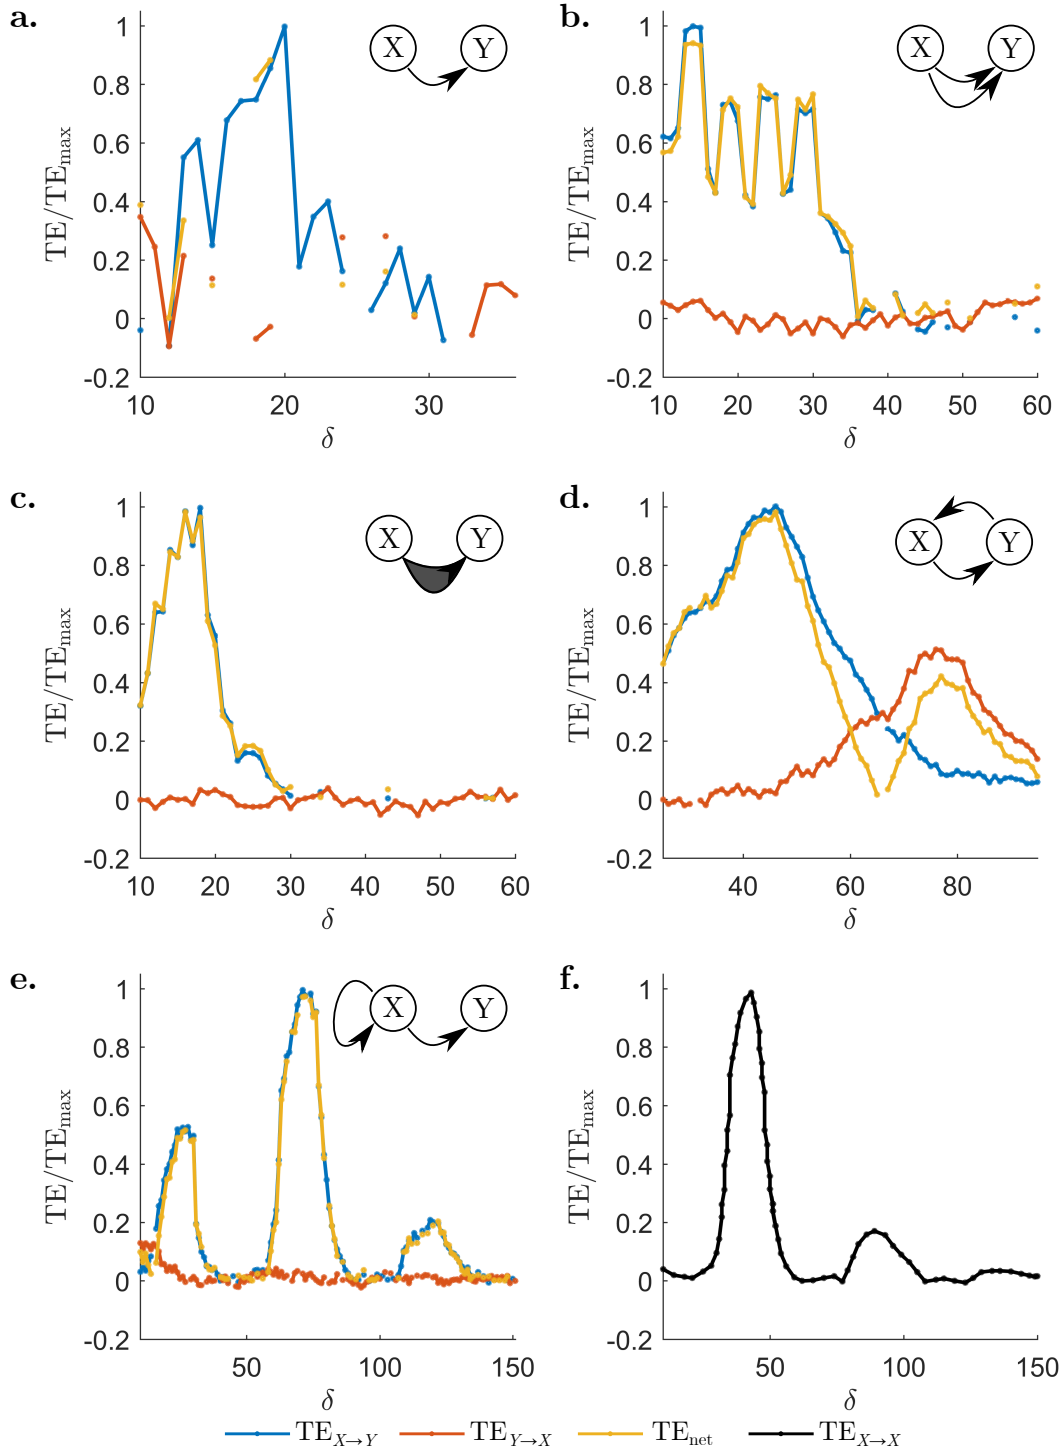


Figure 2.3: $TE(\delta)$ for various two variable systems; adapted from Wibral *et al.* [66].
a. Single pathway, single delay; $\delta' = 20$. b. Single pathway, multiple delays; $\delta' = 15, 20, 25, 30, 35$. c. Single pathway, broadband delay; $\delta' = (18, 22)$. d. Bi-directional communication; $\delta'_{XY} = 45, \delta'_{YX} = 75$. e, f. Single delay from X to Y plus self feedback of X ; $\delta'_{XY} = 75, \delta'_{YX} = 45$.

2.3.2 Sampling Rate

In most cases the input data will be discrete. The rate of discretisation (the sampling frequency) has a significant effect on the detectability⁶ of TE so it must be chosen carefully. In many other applications it is generally accepted that faster sampling will give a better approximation of the continuous process being sampled so it is best to sample as fast as possible. When calculating TE this is not necessarily the case, there is an ideal sampling period which is unlikely to be known *a priori*. The selection of the correct sampling period is often glossed over in literature pertaining to experimental applications. In many cases the ideal sampling period is related to, but not necessarily equal to, the event period, that being the period at which events occur in the system [4, 50]. Sometimes this is an obvious quantity such as for a square wave periodic signal which changes state every half period (see Fig. 2.4). In other cases it may not be so clear, such as when the event timings are chaotic. Zhang *et al.* [74] perform tests to identify where the results are relatively insensitive to changes in the sampling period. Other texts do not detail how they chose the sampling period. The minimum sensitivity method worked well but it is not suitable when different cases with different event timings are being compared.

Paluš & Vejmelka [43] describe how oversampling does not improve accuracy but does increase susceptibility to noise. This is because noise often has a much lower period than the event period so low sampling rates can act to smooth out data. In addition to the sensitivity to noise, oversampling causes a significant reduction in the magnitude of TE. This affects both components ($TE_{X \rightarrow Y}$, $TE_{Y \rightarrow X}$) so it is not necessarily an issue but it increases the potency of rounding errors and can affect comparisons if not properly accounted for. The reduction in magnitude is caused by the low number of events compared to the total number of data points. In other words, a dataset that is sampled at the same rate as events occur will see an event at each data point while the same dataset sampled ten times faster will see an event every tenth data point. To understand the full implication of this requires a quick departure from the main point. Looking at Eq. 2.10 it is seen that the maximum possible value of TE is $H(X_t|X_{t-1})$, when the term $H(X_t|X_{t-1}, Y_{t-\delta})$ is zero. Remembering that conditioning cannot increase entropy it is clear that the upper limit of TE is the entropy of the target variable i.e. the maximum of $TE_{Y \rightarrow X}$ is $H(X_t)$. It has already been shown in Section 2.2.2 that entropy is maximised when the probability of each state is equal. Over-sampling skews the probability thereby reducing the entropy and causing a decrease in the maximum magnitude of TE. This is investigated in Chapter 4 and Chapter 5. At

⁶Detectability errors are discussed below.

the other end of the scale, under-sampling causes events to be missed thereby skewing the results.

Although not specific to TE, Barnett & Seth [5] found that the detectability of causality had an oscillatory relationship with the sampling period. The harmonic of sampling period (Δ) with the event period (T_0) is crucial because of the way it affects the fidelity of the down-sampled data. Using the square wave from Fig. 2.4 as an example, when the sampling frequency is an even harmonic ($\Delta = \frac{T_0}{2}, \frac{T_0}{4}, \frac{T_0}{6} \dots$), the system is faithfully replicated. When the harmony is slightly off the system will usually be correctly replicated but occasionally there will be a skip in the data. When the sampling frequency is an odd harmonic, the system is completely skewed. In Fig. 2.4 the even-harmonically sampled data has two data points per peak/trough while the odd-harmonic sample has two data points in the trough and one at the peak giving the impression that the duty cycle is around 0.66 rather than 0.5. The off-harmonic sample usually has one data point per peak/trough but occasionally adds in a second, thereby changing the periodicity of the signal. It follows logically that the entropy of the down-sampled data will be different at sampling periods that cause the data to be misrepresented. A perfect sample of X will have half of the values in the high state and half in the low state, therefore entropy equals 1 bit ($H = -0.5 \log_2 0.5 - 0.5 \log_2 0.5 = 1$ bit). Compare that to the odd harmonic sampling where the duty cycle is 0.66, then $H = -0.66 \log_2 0.66 - 0.34 \log_2 0.34 = 0.92$ bits.

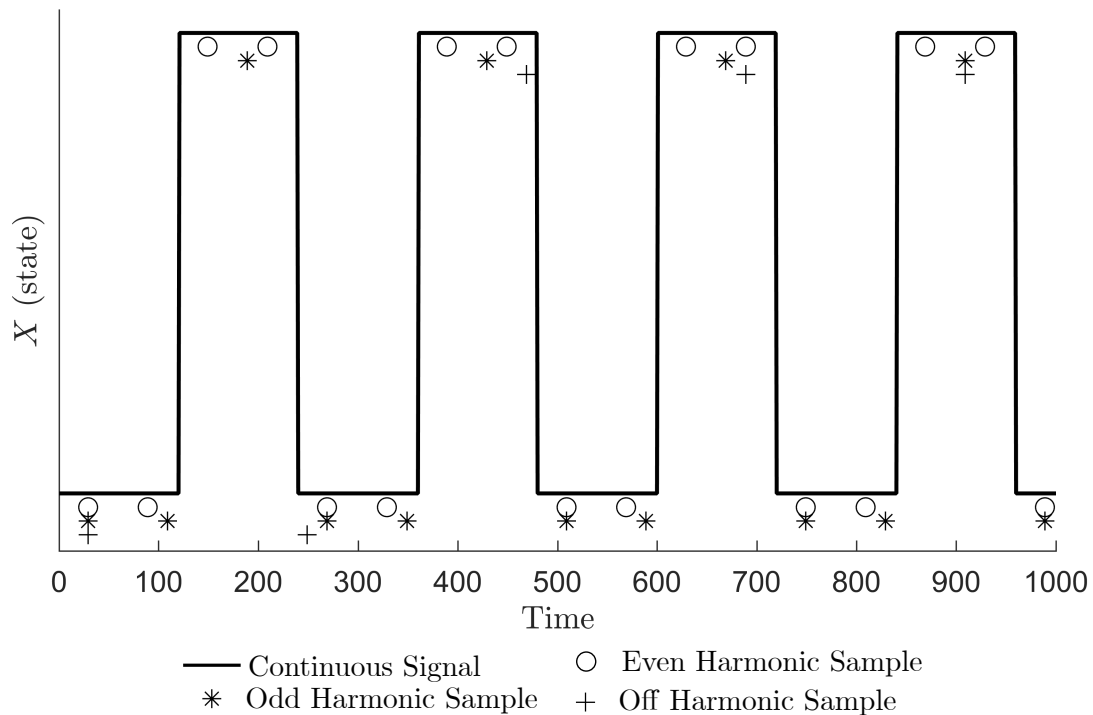


Figure 2.4: Example of how different sampling harmonics change how data is represented. The samples are shifted in the y direction for clarity.

In practice, it is best to acquire data at a higher rate and down-sample it to the desired sampling period. This can be done by simply picking every i th value to reduce the data from the acquisition period (T_A) to the sample period (Δ).⁷ The reason for acquiring data at a higher rate will be demonstrated here: It was previously noted that sub-sampling had the potential to miss events when the event impulse was significantly less than the time between events, this is demonstrated in Fig. 2.5, where the first down-sample captures two of the three peaks and the second only picks up one. The only difference between the two is the time index at which the down-sampling process started.

⁷The down-sampling rate is defined from the ratio of the T_A to δ .




Figure 2.5: An example of how down-sampling impulse behaviour can cause events to be missed. The down-sampled data is shifted vertically for clarity.

To overcome the differences in the data caused by the starting point Zhang *et al.* [74] starts the down-sampling at index 1, then down-samples again starting at index 2 and so on up to Δ/T_A , ensuring all acquired data is accounted for in one of the down-sampled series. Then TE is calculated for each down-sample and the mean of all the results is given as the correct result. In Section 4.2 this method is tested and it is shown that the measured mean does match the analytically calculated value, thereby validating the method. It should be obvious that having to calculate TE Δ/T_A times per lag and sampling period drastically increases the computation time so it is still advisable to choose T_A to balance the resolution of results and computation time. To give a typical example, optical displacement measurements can easily be acquired at 50 Hz ($T_A = 0.02$ s). For a process with an event period of 2 s, TE would be calculated at $\Delta = 2$ s, 100 times per lag. Given that the actual event period and lag values are unlikely to be known, it can be seen that the problem can quickly become computationally intense.

2.3.3 Symbolisation

The cardinality, $|\Omega|$, of the variables has a significant impact on computational intensity,; explored in Section 2.3.7, and the required test length; these factors are improved by lower cardinality. Remembering that the cardinality is defined as the number of unique values that can occur in a variable, it can be seen that for most measurements this value will be very high. Even a modest 12 bit data acquisition system (DAQ) can represent 4096 unique values. The cardinality can be reduced by binning the data (Fig. 2.6 b.), effectively lowering the spatial resolution but this still often has relatively high cardinality. It is often better to remember that events are information and therefore independent of their original meaning. As such they can be further abstracted to indicate behaviours like moving/ not moving or moving clockwise/ counter-clockwise (Fig. 2.6 c and d). This process, introduced by Staniek & Lehnertz [54], is called symbolisation and is very common for empirical measurements of TE. If done correctly, the reduction in spatial resolution of the variables is achieved by removing unimportant information, leaving only the behaviours of interest. Symbolising is done by creating a set of rules to relate the data to a behaviour similar to those used in Section 2.1. For the moving/ not moving symbols, the rules are based on a comparison of the current state to the previous state:

$$x_t = \begin{cases} 1, & x_t \neq x_{t-1} \\ 0, & x_t = x_{t-1} \end{cases}$$




Figure 2.6: Examples of binning and symbolisation of data. a. shows a typical measured displacement with a spatial resolution of 0.005 resulting in a cardinality of $|\Omega| = 400$. b. Data is binned to a spatial resolution of 0.25. c. Data is symbolised based on the direction of movement with symbols for moving up, moving down and not moving. d. Data is symbolised based on whether the object is moving or not.

2.3.4 Conditional Transfer Entropy

Conditional transfer entropy considers how additional variables affect the source and target variables already discussed. The additional variables are called conditioning variables because they are used to pre-condition the data to remove potential mutual influences or cascade effects [8]. In effect, conditional TE measures the dependence of the interaction between the source and target variables on the conditioning variable. As such, some work has been done in recent years to use conditional TE to map causality networks in multivariate systems [2, 48, 53].⁸ The mathematical description of conditional TE defined in Eq. 2.13 is a combination of the definitions in Bossomaier *et al.* [8] and Aste & Di Matteo [3].

$$\text{TE}_{X \rightarrow Y|Z} = H(Y_t|Y_{t-1}, Z_{t-\zeta_{ZY}}) - H(Y_t|Y_{t-1}, X_{t-\delta_{XY}}, Z_{t-\zeta_{ZY}}) \quad (2.13)$$

⁸The method proposed here differs slightly from those in the cited literature to better match the approach of Zhang *et al.* [74]. Here, instead of calculating conditional TE with a large history length, the data is down-sampled and single lags are applied to the source and conditioning variables.

The notation follows the conventions already set out and the variables can be replaced in the exact same way as they were between Eq. 2.10 and Eq. 2.11 to calculate the other TE components. The extra variable also comes with an extra time lag [3] to account for the time it takes for information to be transmitted from Z to Y as shown in Fig. 2.7. Note that when the delay is from the conditioning variable, the symbol ζ is used in place of δ to make future comparisons easier. The equation can be extended to include as many conditioning variables as necessary but is shown here with only three variables.

A simple thought experiment will help clarify the physical interpretation. In Fig. 2.7 a wave drives two floating structures which are placed far enough apart that there is no possible way they could be communicating. From the earlier discussions it should be clear that the expected TE measured between X and Y is zero but looking only at the time series X and Y (Fig. 2.7 b.) it would be easy to conclude that they are actually communicating with a delay equal to $\hat{\delta}_{ZY} - \hat{\delta}_{ZX}$. Only by considering a measurement of the wave (Z) can it be seen that they are actually mutually driven. In terms of transfer entropy, a significant peak in $\text{TE}_{X \rightarrow Y}$ would be measured at a delay $\delta_{XY} = 25$, the temporal separation of the events at X and Y respectively. Evaluating conditional TE with a conditioning lag equal to the actual communication delay from $Z \rightarrow Y$ ($\delta_{ZY} = 60$) will condition out the effect of Z on Y . i.e. $\text{TE}_{X \rightarrow Y}(\delta = 25) \neq 0$ and $\text{TE}_{X \rightarrow Y|Z}(\delta = 25, \zeta = \delta_{ZY} = 60) = 0$.





Figure 2.7: Intuitive example of how conditional transfer entropy can be used to remove mutual driver effects. In this case both X and Y are driven by waves Z but have no interaction themselves. Note that Z is a measurement of the wave height at that point not a structure. The dashed arrow is used to show that the calculation would be performed along that path but no interaction exists.

The other potential compound when considering systems of more than two bodies is

cascade effects,⁹ that is when Z communicates with X and X relays that information on to Y . An example of this would be if X was driven by waves Z as before but now Y is not driven by the waves, it is mechanically coupled to X . The time series would be qualitatively similar to that in Fig. 2.7 despite the communication path being entirely different. In this case, the false communication path is $Z \rightarrow Y$ with a delay equal to $\hat{\delta}_{ZX} + \hat{\delta}_{XY}$. Evaluating conditional entropy along this path, conditioned at lag $\zeta = \hat{\delta}_{XY}$ gives $\text{TE}_{Z \rightarrow Y|X}(\delta = 60, \zeta = 25) = 0$.¹⁰ For comparison, in this system $\text{TE}_{X \rightarrow Y|Z}(\delta = 25, \zeta = 60) \neq 0$.

2.3.5 History Length

In Section 2.2.3 it was shown that the current state of a variable may be affected by previous states. The equations so far offered have all only accounted for the effect of the last state (the effect of X_{t-1} on X_t) but the historical effects may be more far reaching. To include this longer history time, the variables can be replaced with Takens embedding vectors [8, 50], which contain all of the previous states from times $t - k + 1$ up to and including t , such that

$$X_t^{(k)} = (x_t, x_{t-1}, \dots, x_{t-k+1})$$

where k is the history length. Computational intensity is significantly increased with increasing k , so this parameter should be balanced. For Markovian systems, the current state is dependent upon the previous M states, where M is the order of the Markov chain. Schreiber [50] proved that ideally the history length (k) should encompass all of those states and only those states; therefore $k = M$. For non-Markovian systems, the ideal history length is infinite [30].

2.3.6 Equivalence to Granger Causality

The concept of causality versus correlation is ubiquitous in modern science. The most commonly used definition of causality in time-series analysis is Granger causality. The distinction ‘Granger causality’ rather than simply ‘causality’ exists to address the short comings of the measure. There are a number of other definitions of causality which are discussed in Chicharro & Ledberg [11] and Pearl [44]. However, Barnett *et al.* [4] proved

⁹Cascade and relay will be used interchangeably.

¹⁰It should be noted that this relies on some randomness in the interactions as is the case for all TE measurements. The applicable subtleties of this sort of analysis are not widely documented so they are investigated further in Section 6.1.

that Granger causality and TE are equivalent measures, so that is the causality definition that will be used in this work.

Granger [22] summarises the conditions for his definition as “1. The cause occurs before the effect; and 2. The cause contains information about the effect that is unique, and is in no other variable.” The first condition is a formal restriction on the communication lag being negative. The second explains the necessity of some degree of randomness in the time-series. For example, if variable Y was perfectly coupled to X and X contained a periodic signal, no Granger causality could be inferred because there is no unique information in either. However, if the coupling was imperfect or X was not periodic, it is possible to detect Granger causality. The formal definition was first published in 1969 by Granger [23]. In that work, a parametric prediction model is used to predict the state of x_t using all past values of X ($P(x_t|X)$ ¹¹). Then the state of x_t is predicted using all past values of X and all past values of Y ($P(x_t|X, Y)$). It is determined that if the variance of the prediction error is decreased by the inclusion of Y in the prediction model, Y can be said to cause X . The degree of causality is then determined by the coherence of the cross power spectra.

Transfer entropy was introduced as a measure of the reduction in uncertainty of X_t given the past of X and the past of Y compared to the uncertainty of X_t given only X . Barnett *et al.* [4] equates the ideas of ‘reducing uncertainty’ and ‘prediction,’ then goes on to prove the two concepts are analytically identical for normally distributed time-series. That work does not comment on how the two may diverge with non-normal data but it does state that non-zero TE will still imply non-zero Granger causality.

2.3.7 Calculating TE

The TE equation was previously defined as a difference of entropies, which is the most intuitive formulation, but it is not the most practical. A number of TE estimators are outlined by Bossomaier *et al.* [8], but when the input data is discrete and countable, Eq. 2.11 can be recast as a collection of probabilities as derived in Appendix B.1, following the work of Porfiri [45].

$$\text{TE}_{X \rightarrow Y} = \sum_{y, \hat{y}, x \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}) \log_2 \left(\frac{\Pr(Y_t, Y_{t-1}, X_{t-\delta}) \Pr(Y_{t-1})}{\Pr(Y_{t-1}, X_{t-\delta}) \Pr(Y_t, Y_{t-1})} \right) \quad (2.14)$$

¹¹The notation is adapted from [23] for consistency

where y, \dot{y}, x are the possible states of the variables $Y_t, Y_{t-1}, X_{t-\delta}$. The over dot on the y is used to indicate that it is related to the lagged variable Y_{t-1} . It would be more consistent to use the notation $\omega_{Y_t}, \omega_{Y_{t-1}}, \omega_{X_{t-\delta}} \in \Omega$ but this is overly cumbersome. Using a two symbol example, each of the variables can take on a value of 0 or 1. Given that there are three variables $Y_t, Y_{t-1}, X_{t-\delta}$, each with two possible states, there are eight possible combinations so the sample space is

$$\Omega = \{(0, 0, 0), (0, 0, 1), (0, 1, 0), (0, 1, 1), (1, 0, 0), (1, 0, 1), (1, 1, 0), (1, 1, 1)\}$$

Each term of Eqn. 2.14 corresponds to one of these combinations of states such that

$$\begin{aligned} \text{TE}_{X \rightarrow Y} &= \Pr(Y_t = 0, Y_{t-1} = 0, X_{t-\delta} = 0) \log_2 \left(\frac{\Pr(Y_t = 0, Y_{t-1} = 0, X_{t-\delta} = 0) \Pr(Y_{t-1} = 0)}{\Pr(Y_{t-1} = 0, X_{t-\delta} = 0) \Pr(Y_t = 0, Y_{t-1} = 0)} \right) \\ &+ \Pr(Y_t = 0, Y_{t-1} = 0, X_{t-\delta} = 1) \log_2 \left(\frac{\Pr(Y_t = 0, Y_{t-1} = 0, X_{t-\delta} = 1) \Pr(Y_{t-1} = 0)}{\Pr(Y_{t-1} = 0, X_{t-\delta} = 1) \Pr(Y_t = 0, Y_{t-1} = 0)} \right) \\ &\dots \\ &+ \Pr(Y_t = 1, Y_{t-1} = 1, X_{t-\delta} = 1) \log_2 \left(\frac{\Pr(Y_t = 1, Y_{t-1} = 1, X_{t-\delta} = 1) \Pr(Y_{t-1} = 1)}{\Pr(Y_{t-1} = 1, X_{t-\delta} = 1) \Pr(Y_t = 1, Y_{t-1} = 1)} \right) \end{aligned}$$

The effect of cardinality on calculation intensity is evident here. For a cardinality of two for each variable there are eight terms. If all of the states of a 12 bit DAQ were used, there would be 69×10^9 terms.

The probabilities are calculated from their frequency of occurrence (Eq. 2.2) by counting how often a particular combination occurs and dividing by the total number of samples. For example take the time series

$$\begin{aligned} Y_t &= 1, 1, 0, 1, 0, 1 \\ Y_{t-1} &= 1, 1, 1, 0, 1, 0 \\ X_{t-\delta} &= 0, 1, 1, 0, 0, 0 \end{aligned}$$

The combination $(Y_t, Y_{t-1}, X_{t-\delta}) = (0, 0, 0)$ does not occur so the probability $\Pr(Y_t = 0, Y_{t-1} = 0, X_{t-\delta} = 0) = 0$. However the event $Y_{t-1} = 0$ occurs twice in a series of six data

points so $\Pr(Y_{t-1}) = \frac{2}{6} = 0.333$. Following this through Eq. 2.14 becomes

$$\begin{aligned} \text{TE}_{X \rightarrow Y} &= 0.000 \log_2 \left(\frac{0.000 \times 0.333}{0.333 \times 0.000} \right) + 0.000 \log_2 \left(\frac{0.000 \times 0.333}{0.000 \times 0.333} \right) \\ &\quad + 0.167 \log_2 \left(\frac{0.167 \times 0.667}{0.333 \times 0.333} \right) + 0.167 \log_2 \left(\frac{0.167 \times 0.667}{0.333 \times 0.333} \right) \\ &\quad + 0.333 \log_2 \left(\frac{0.333 \times 0.333}{0.333 \times 0.333} \right) + 0.000 \log_2 \left(\frac{0.000 \times 0.333}{0.000 \times 0.333} \right) \\ &\quad + 0.167 \log_2 \left(\frac{0.167 \times 0.667}{0.333 \times 0.333} \right) + 0.167 \log_2 \left(\frac{0.167 \times 0.667}{0.333 \times 0.333} \right) \\ \text{TE}_{X \rightarrow Y} &= 0.004 \text{ bits} \end{aligned}$$

A number of terms take the form $0 \log_2 \frac{0}{0}$, L'Hospital's rule can be used to show that these terms equal zero. Typically, this process will need to be repeated many times to obtain results over a range of Δ and δ .

2.3.8 Expansion for Conditional TE

The calculation of conditional transfer entropy is not significantly more complicated than the standard form although the computational intensity does increase. The probabilistic formulation of the conditional TE equation (Eq. 2.15) is derived in Appendix B.1.

$$\begin{aligned} \text{TE}_{X \rightarrow Y|Z} &= \\ &\sum_{y, \hat{y}, x, z \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta_{XY}}, Z_{t-\zeta_{ZY}}) \log_2 \left(\frac{\Pr(Y_t, Y_{t-1}, X_{t-\delta_{XY}}, Z_{t-\zeta_{ZY}}) \Pr(Y_{t-1}, Z_{t-\zeta_{ZY}})}{\Pr(Y_{t-1}, X_{t-\delta}, Z_{t-\zeta_{ZY}}) \Pr(Y_t, Y_{t-1}, Z_{t-\zeta_{ZY}})} \right) \end{aligned} \quad (2.15)$$

2.3.9 Transfer Entropy Summary

In the preceding sections, it was demonstrated that TE can be used to infer directed causality between two variables, using only time series data. One important limitation is that there must exist a significant amount of randomness in the system, either by stochastic action of the source variable or imperfect coupling of the two variables.

The process for calculating TE for a single set of parameters Δ , δ and ζ requires the following operations:

- Collect data at a higher rate than events occur in the system.

- Down-sample the data to the desired, often unknown, sampling period, Δ . This is done Δ/T_A times to ensure all data is accounted for.
- Lag the variables as necessary.
- Symbolize the data as necessary, making an effort to minimise the cardinality while retaining the important behaviour of the system.

In most cases it will be necessary to repeat this process over a range of Δ , δ and ζ where appropriate.

2.4 Fluid Dynamics

This section is an overview of some of the fluid dynamic concepts that will provide a foundation for the design of experiments. The main focus is on systems where one or more dynamic structures are interacting with an external fluid. The study of these systems is a topic with broad engineering applications including energy harvesting [18, 33] and flight control [26, 56, 72], as well as scientific applications such as the study of animal propulsion and flocking behaviour [45, 72, 74]. However, these systems are often poorly understood because the large number of parameters and significant non-linearity make it difficult to apply many of the classical analysis tools [25]. The following reviews are aimed to identify the significant parameters for the purpose of designing the experiments and to support the claim that parametric analysis tools are difficult to apply to fluid-structure interaction problems.

2.4.1 Flow Around a Cylinder

It is well established that bluff bodies often shed vortices in their wake. For static, uniform, infinitely long cylinders, this occurs for Reynolds numbers¹² between 50 and 2×10^5 [69]. The characteristics of the cylinder wake and its vortices are heavily dependent upon the Reynolds number, aspect ratio and surface roughness of the cylinder, incoming turbulence intensity and flow blockage [37], [38], [73], as well as by movement of the cylinder.¹³ For the purposes of this study, only a high level understanding of the process is required because the system being analysed is inherently chaotic and many of these factors are beyond experimental control.

¹²based upon cylinder diameter, d and free stream velocity, U .

¹³Numerous references throughout the text.

Jeon & Gharib [27] summarises the process of a cylinder shedding vortices as having three stages; vortex growth, saturation and shedding. Initially, vorticity is generated at the cylinder’s surface due to pressure gradients tangential to the surface. Vorticity can only be generated at boundaries of the fluid domain for homogeneous fluids [36], so the surface must be the only source of vorticity. The vorticity is transmitted to the near wake by the viscous action in the shear layers. Instabilities in the separated shear layer cause it to roll up into wake vortices. The reader is directed to the review papers of Roshko [47] and Williamson [69] for in-depth discussions of these vortex structures and the mechanisms which shape them. The constant influx of vorticity causes the vortex structures to grow until saturation occurs. The vortex saturates at uniform intervals referred to as the critical formation time,¹⁴ which, for static cylinders, is equal to the inverse of the Strouhal number, St . Once a vortex has saturated, it separates from the structure and ‘pinches off,’ thereby preventing the influx of additional vorticity, and is advected away. The periodicity of this process results in a vortex street, a series of vortices travelling downstream with a uniform pattern as shown in Fig. 2.8 a. For the Reynolds numbers studied here ($Re \sim 10^3$), the static cylinder is in the shear layer transition regime [69], defined by an increase in three-dimensional structures. Fig. 2.8 shows how the cylinder wake becomes less ordered as the Reynolds number increases but periodicity is still evident. The Strouhal number, St , is relatively insensitive to Re at this range with a value $St \sim 0.2$. The shedding period is then defined as;

$$T_s = \frac{St U}{d} \quad (2.16)$$

The particulars of the saturation mechanisms, times and strengths, as well as the instabilities, are thoroughly discussed in the review papers of Roshko [47] and Williamson [69].

Rotating cylinders have been broadly studied for a wide range of applications. One of the earliest studies being that of Heinrich Gustav Magnus who studied the deflection of cylindrical missiles caused by their axial rotation [57]. Investigations into the wake of rotating cylinders began with Ludwig Prandtl in 1925, inspiring a broad investigation of constantly rotating cylinders, such as those by Glauert [21], Mittal & Kumar [34], Moore [35] and Wood [70]. The exact results of these studies vary significantly, especially in the low Reynolds number regimes, most likely due to the high level of sensitivity to experimental conditions. For the present study, it is sufficient to surmise that cylinder rotation causes an increase in circulation. At high cylinder surface speed to free stream velocity, ratios, V_s/U the stagnation points move off of the cylinder surface, the flow becomes stable and the

¹⁴a non-dimensional form of the shedding period.




Figure 2.8: Wake of cylinders with different Reynolds numbers. a, $Re=140$, b, $Re=10^4$. Modified from Van Dyke [62]

circulation becomes trapped inside closed streamlines around the cylinder. Typically it is said that the flow is stable for ratios greater than 2 but Mittal & Kumar [34] demonstrates that this stability does not persist for all cylinder speed ratios.

Better approximations of the flow studied here are rotationally oscillating cylinders, studied by Filler *et al.* [20], Morton [36], Okajima *et al.* [40], Tokumaru & Dimotakis [58] and Wu *et al.* [71], or impulsively started rotating/translating cylinders as in Chew *et al.* [10] and Coutanceau & Menard [12]. These studies show that oscillating a cylinder at or near the fundamental shedding frequency causes the shedding cycle to lock-in to the oscillation frequency. Additionally, Tokumaru & Dimotakis [58] show that very high frequency, large magnitude oscillations can almost eliminate shedding. Morton [36] studied the vortex formation process for an accelerating cylinder rotation, finding that acceleration significantly shortened the critical formation time due to the additional vorticity flux. It was also noted that higher rates of acceleration led to earlier separation of the first vortex. Similar results were seen by Chew *et al.* [10] and Coutanceau & Menard [12] in studies of cylinders which were impulsively started into a rotation and translation motion, although they attribute this starting vortex to the translation, analogous with the phenomenon seen when an airfoil is accelerated. The pertinent result of all these studies is that the timing of vortex shedding can be affected by rotations of the cylinder and the effect is proportional to the acceleration rate and the magnitude of the rotation.

No studies have been found which deal with short, random, rotational impulses of cylinders. However, it is evident that cylinder rotation causes an increase in vorticity compared to the static cylinder case. It follows that when the cylinder stops rotating, this vorticity must either be dissipated by interactions at the cylinder surface or shed, the latter being the most intuitive result. This indicates that a rotational impulse of the cylinder is a reasonable way to generate information in the fluid. The sensitivity of the results to experimental conditions suggests that the test parameters (rotation speed, etc) should be selected empirically.

Once vorticities are shed they are advected downstream forming a Karman vortex street. Viscous forces in the fluid act to diffuse the vorticity, eventually resulting in a turbulent, non-periodic wake. In the applicable Reynolds number ranges and with a relatively low turbulence free-stream, it was found that the wake would devolve to fully turbulent at around 40 to 50 diameters downstream [7]. This will be affected by the wake impinging upon solid structures as it will in this study but it demonstrates the upper limit of usable stream-wise separations between the cylinder and secondary structures.

2.4.2 Flow Around an Airfoil

This section will discuss the basics of flow around an airfoil. This review is highly focused as airfoil dynamics are not central to this study. For this study, the most important property is the pitching moment, described below. For a broad discussion of aerodynamics, the reader is referred to Anderson Jr [1]. The applicable concepts will be introduced using the example of an isolated, static airfoil in uniform flow before moving on to the more complex case of a pitching airfoil. The field of airfoil aerodynamics is very broad so the following conditions are imposed to limit the scope of this review based on the experiments to be conducted.

1. This study will deal solely with thin airfoils, defined as having a chord length, l , much greater than the maximum thickness, t_0 .
2. The airfoils will be symmetric about the chord line.
3. The airfoils will be span-wise uniform.
4. The Reynolds number will be low, on the order of 10^3 .

In a general sense, airfoils cause a deflection of the incident air, creating pressure at the surface. This pressure results in a net force acting on the body which is commonly separated into the components of lift and drag, which are aligned normal to and along the fluid velocity vector, U , respectively, as shown in Fig. 2.9. In addition to the pressure induced force, there is a viscous force which is often not considered as it becomes negligible at high Reynolds numbers, with the exception of low lift conditions when the pressure forces are low. The forces on the body also cause a pitching moment. For the described class of airfoil, the center of pressure and aerodynamic center are both located at the quarter chord point,¹⁵ although this is not true for all airfoils. Additionally, the moment about the quarter chord point is, theoretically, zero. Viscous effects reduce the validity of these assumptions but they serve as a reasonable starting point for this study given the lack of information available for airfoils in this Reynolds number range.

It is well established that rapid pitching of an airfoil causes a delay in stall and increases the forces and moments acting upon an airfoil [72]. Yu *et al.* [72] uses a flat plate with $Re = 1.45 \times 10^4$ to demonstrate that the rapid pitching motion causes an increase in circulation which increases the moment about the quarter chord by up to a factor of 5. They also show that the minimum absolute moment ($|C_M|$) occurs when the the pivot is placed somewhere between the quarter and half chord.

¹⁵One quarter of the way from the leading edge to the trailing edge, following the chord.

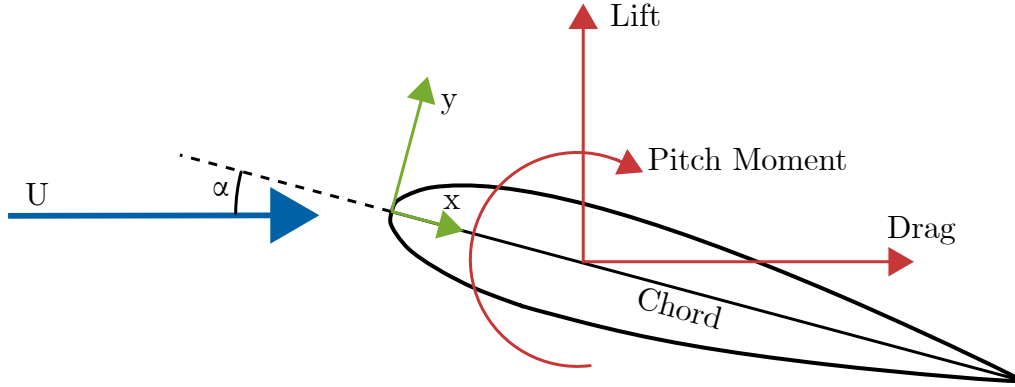


Figure 2.9: Coordinate system and forces acting on the airfoil.

2.4.3 Fluid Structure Interactions

Broadly speaking, in most fluid-structure interactions, the boundary conditions of the fluid domain change due to motion of the structure, while the structure responds, at least in part, to the imposed fluid loading. Typically, it is not possible to formulate an analytical solution for such systems [25], due to the complex and often highly non-linear nature of the interaction. In some cases it is made possible by making significant assumptions. For example, Falnes [18] developed an analytical solution for energy extraction from an array of wave energy converters by assuming the structures are infinitely small and the array is infinitely large, thereby making use of the asymptotic nature of the solution; similar models exist for wind farms, such as that of Meyers & Meneveau [33]. It is more common to analyse them experimentally or numerically, although the large number of parameters that typically exist make this difficult [25]. The following section reviews some of the specific FSI systems of interest to this study; for a more complete review, readers are directed to the review papers of Hou *et al.* [25] and Rockwell [46].

Of interest to this study is the response of a pivoted plate to a passing parallel vortex as shown in Fig. 2.10. This system depends upon many parameters, so it has not been studied exactly as it will be employed here. Therefore, similar systems will be discussed here to highlight these parameters and help set expectations for the experiments; this is not intended to be a complete review of the field. Walker [65] demonstrated that a vortex passing an infinite, fixed, flat plate generates an unsteady surface pressure, that is affected by vortex strength, Γ , and passing distance, ξ . From classical mechanics, this would result in an unsteady pitching moment on the plate. The response of a pivoted

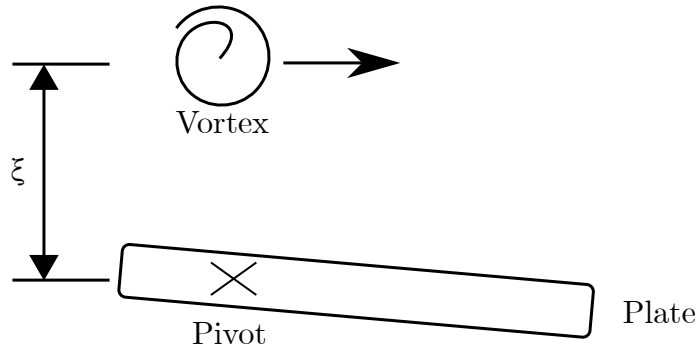


Figure 2.10: A schematic of a pivoted plate in a uniform flow with a passing parallel vortex

plate would be affected by the location of the pivot as well as the damping, restoring and inertial forces. The dynamics of the flow around the plate as it starts to move becomes increasingly complex. Studies on airfoils driven in an oscillatory pitching motion, such as that of Widnall & Wolf [68], confirm that the lift remains unsteady even if the airfoil is pitching. That study found that unsteady lift was generated for passing distance to chord length ratios greater than two across a range of loading parameters. The motion of the structure being affected by the passing vortex adds more parameters to the system. The most important for this work is the arrival time of the vortex relative to the structure's motion, which was shown to significantly affect the dynamics of the fluid. For example, vortices in a quiescent fluid may induce a leading edge vortex or not, depending upon the phase of oscillation when the vortex arrives at the structure [46].

In cases where the vortex passes very close to, or impinges upon, the structure, the fluid dynamics are dominated by distortion of the vortex. This has a significant effect on the response of the structure as well as on the wake and the response of subsequent structures [46].

From this discussion, it is clear that the parameters that affect the studied system are vortex strength, vorticity distribution, passing distance, pivot location and forces, structure mass and geometry (in part due to the added mass) and the phase. To counter this the models used to analyse these systems are typically inviscid and neglect the solid body forces [6, 26, 56, 65, 68]. Extending these parametric models to the free response of the structure is, therefore, not trivial. This is especially true when the vortices are generated stochastically, as they will be in the following experiments. For that purpose,

it is sufficient to surmise that a vortex passing a pivoted plate or airfoil is likely to induce a pitching moment. This moment will lead to motion described by classical mechanics providing the resisting forces are sufficiently low. Additionally, the response is unlikely to be accurately modelled using any parametric models, so another analysis method should be applied.

Chapter 3

Experimental Setup

3.1 Flow Facility

The experiments performed here placed an array of structures into a uniform fluid flow. The facility is a constant pressure, closed water channel designed and documented by Sommer [52]. Water is stored in two, 250 L reservoirs below the channel, and pumped¹ up to the inlet settling chamber (see Fig. 3.1) with an estimated maximum volumetric flow rate of 0.0019 m³/s. The settling chamber has two purposes, to dampen the turbulence of the incoming fluid and to maintain constant pressure at the channel inlet. Water enters via a perforated pipe designed to slow the fluid and fills up to the level of the top of the spillway wall. The free surface ensures a constant pressure at the inlet to the water channel. The return of the spillway feeds back into the reservoir. The settling tank and water channel are connected by three evenly spaced vertical pipes. These pipes extend through to the bottom of the water channel and are perforated to encourage two dimensional flow in the channel [52]. The water channel (Fig. 3.2) consists of:

- A pre-conditioning volume, where fluid enters.
- The conditioning element; a hexagonal aluminium honeycomb detailed in Sommer [52].
- A 3.7:1 two dimensional contraction.
- The test section, discussed below.
- The exit region where fluid is conditioned for a smooth exit.

¹Pump: WEG, model: 00156ES3EB56C

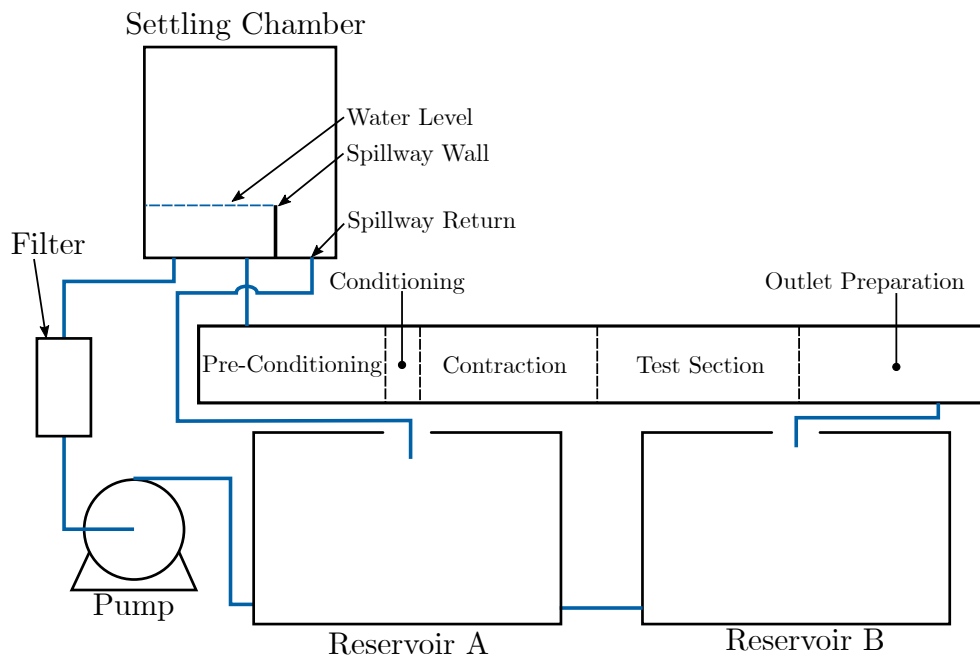


Figure 3.1: Hydraulic schematic of the test facility.




Figure 3.2: A rendered image of the flow facility with no model installed. Structural components have been removed to highlight the functional components.




Figure 3.3: Flow facility drawing. Blue arrows indicate water flow into/ out of the facility. The red line highlights the footprint of the water channel.

The test section is 0.60 m long, 0.27 m wide and 0.16 m deep with a removable top and optical access available through the three other sides (see Fig. 3.4). Models are attached to the top plate and a camera is placed underneath to capture the movements of the models.




Figure 3.4: Isometric view of the test section with an example model installed.

In all cases the working fluid was water with a free-stream velocity of 0.04 m/s and a pressure head of approximately 0.3 m as measured at the start of the test section, mid-depth.

The test section roof was modified from that described by Sommer [52] to make the channel more adaptable to different experiments. It consists of a series of interlocking plates, shown in Fig. 3.5, which can easily be exchanged to suit the experiments. In this case the interface is made of three plates; an insert plate supports the two interface plates, one which mounts a single cylinder model and one with a series of threaded holes for mounting airfoils, see Section 3.2 for more information about the models. The plates are designed to maintain a flush inner surface and are symmetrical to allow the plates to be installed in any orientation, increasing the flexibility of the system. The two interface plates are easily removed to reconfigure the models or provide access to the inside of the channel.




Figure 3.5: Model of the plates which form the top of the test section.

3.2 Models

3.2.1 Cylinder Model

The two types of models used are a circular cylinder and an airfoil. For many of the experiments the cylinder is driven to rotate about its axis by a motor located outside of the channel. As such, it is required to pass through the interfacing plates while maintaining a water tight seal without hindering the rotational freedom. The cylinder is constructed from two concentric cylinders (Fig. 3.6), the outer cylinder has a diameter of $d = 44$ mm and is in contact with the flow, while the inner cylinder serves as an axle, supported by a bearing assembly outside the channel. It passes through an oil seal mounted in the interface plate, as shown in Fig. 3.5. The fluid pressure on the seal (McMaster Carr PN. 5154T841) is approximately 2200 Pa, about 3% of the seal’s pressure rating. A plate mounted in the underside of the cylinder contains the alignment markings used to optically track the cylinder’s rotation. The marking plate is printed from PLA (Polylactic Acid) plastic; it is designed such that when the counter-sunk center bolt is tightened, the plate expands to create a friction connection with the cylinder. Wells printed into the underside are filled with green acrylic paint to create the markings.

Properties of this cylinder are provided in Table 3.1, whilst test conditions are provided in Table 3.2.




Figure 3.6: Isometric and Section views of Cylinder Assembly.

Table 3.1: Properties of the cylinder. Assuming free-stream of water at 0.04 m/s where appropriate.

| Property | Symbol | Value | Unit |
|----------------|-----------|-------|------|
| Length | L | 0.160 | m |
| Diameter | d | 0.044 | m |
| Aspect Ratio | AR | 3.63 | - |
| Blockage Ratio | β_R | 16.3 | % |

Table 3.2: Test conditions.

| Property | Symbol | Value | Unit |
|----------------------------|------------|--------------|------|
| Free Stream Velocity | U | 0.04 | m/s |
| Reynolds Number (Cylinder) | Re_{Cyl} | 1760 | - |
| Reynolds Number (Airfoil) | Re_{AF} | 2400 | - |
| Strouhal Number | St | $\sim 0.2^*$ | - |
| Shedding Frequency** | f_f | 0.18 | Hz |

* [19], [69] Excluding the effect of the body/bodies in the cylinder wake. This serves only as a guide.

** Calculated from St estimate. This is the expected shedding frequency from the cylinder without rotation or interference from the other bodies.

3.2.2 Airfoil Models

The airfoil profiles are NACA 0018 with the exception of a small provision in the lower span to accommodate a bearing as shown in Fig. 3.7. The airfoil are printed from PLA plastic with a wall thickness of 1 mm and coated in epoxy resin (Smooth-On XTC-3D), which seals the parts and allows them to be sanded smooth. Due to printer constraints, the airfoils are printed in two pieces and connected end to end with steel connector pins. A solid shaft (brass or stainless steel) screws into the interface plate then passes through the span of the airfoil, attaching to the bearing, allowing the airfoils to pitch freely about their quarter-chord.

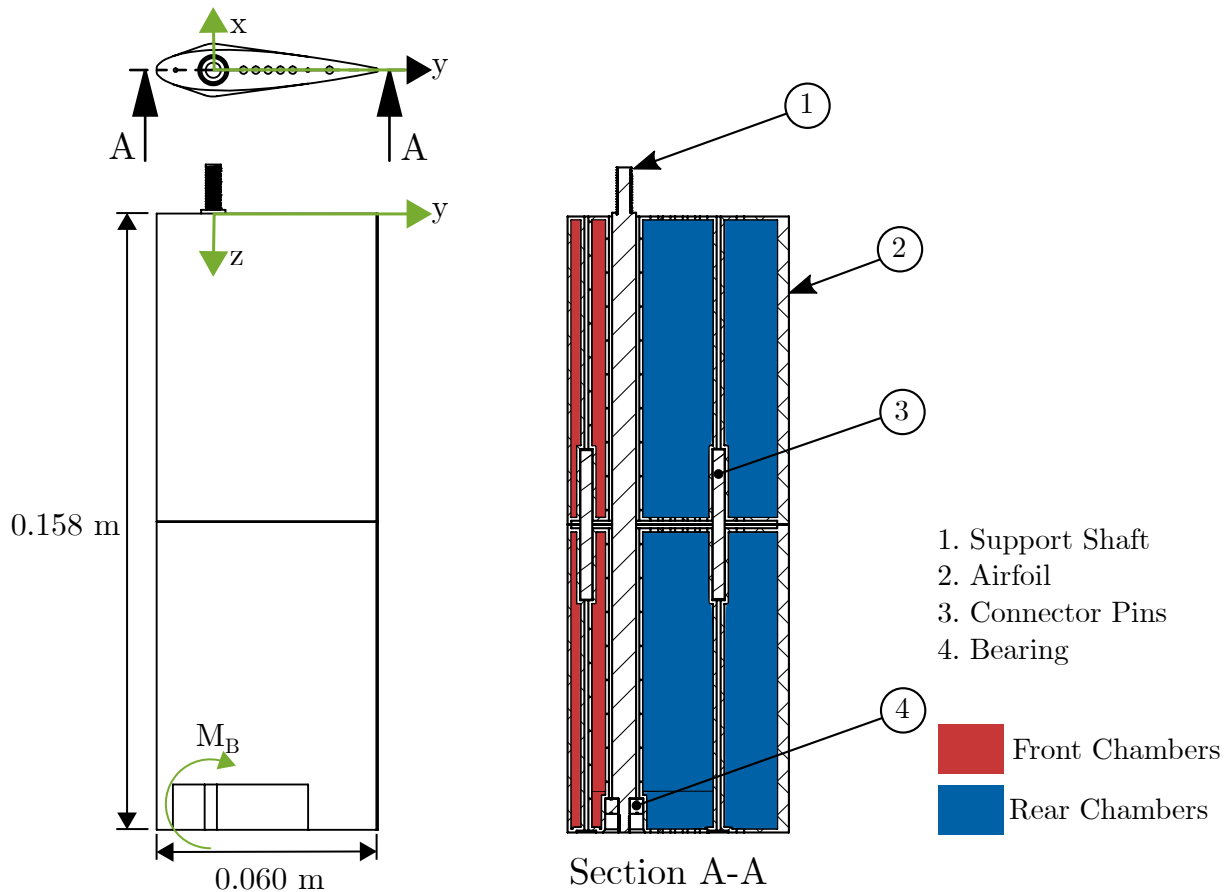


Figure 3.7: Drawing of the airfoil assembly. The airfoil’s coordinate system is shown by the green arrows.

The system is designed to minimise resistance to rotation about the pivot, thereby maximising the airfoil’s response. The following factors were considered for this purpose:

1. Pitching moment due to free stream. i.e. Restoring moment.
2. Weight or buoyancy driven rotation.
3. Bearing friction.
4. Mass ratio of moving components.

From the discussion in Section 2.4.2, the minimum pitching moment is expected to occur between the quarter and half chord points. For that reason and because it is the thickest point, the pivot is placed at the quarter chord.

One of the most critical design factors was the balance of the weight and buoyancy vectors. The main reason for this is the tilt of the channel. The channel is tilted by 5 degrees (inlet up) to allow it to bleed as it fills. This would cause a rotational moment

about the pivot for pitching angles other than zero if the vectors are not properly balanced. This is easily imagined, if we take the airfoil and pivot in air, where buoyancy is negligible, and rotate it so the pivot shaft is horizontal, the airfoil would want to fall into a position where the center of mass is directly below the pivot. This happens inside the channel due to the table tilt, albeit to a less extreme degree. The moment about the pivot is calculated as

$$M_z = \hat{\mathbf{z}} \cdot (\mathbf{R}\mathbf{W} \times \mathbf{C}_G + \mathbf{R}\mathbf{B} \times \mathbf{C}_B) \quad (3.1)$$

where \mathbf{W} is the weight vector and \mathbf{B} is the buoyancy vector, both aligned with gravity. They are projected onto the reference frame of the airfoil by multiplying by the rotation matrix \mathbf{R} , Eq. 3.2; note that for brevity Eq. 3.2 has been simplified using the knowledge that the table roll angle is zero. \mathbf{C}_G and \mathbf{C}_B are the displacement vectors from the origin to the center of gravity and center of buoyancy, respectively.

$$\mathbf{R} = \begin{vmatrix} \cos(\alpha) \cos(\beta) & \cos(\beta) \sin(\alpha) & -\sin(\beta) \\ -\sin(\alpha) & \cos(\alpha) & 0 \\ \cos(\alpha) \sin(\beta) & \sin(\alpha) \sin(\beta) & \cos(\beta) \end{vmatrix} \quad (3.2)$$

where α (Fig. 3.11) is the pitching angle and β is the inlet up table tilt.

The most obvious solution is to balance $|\mathbf{W}| = |\mathbf{B}|$ and set the centers of gravity and buoyancy to be coincident but this is difficult to achieve exactly. In an effort to achieve this balance, three designs were analysed; one fully solid part and three hollow parts with various degrees of flooding. The no flooding model is fully sealed, trapping air in all chambers, the partially flooded model has vent holes (see Fig. 3.9) to flood the rear chambers with water and the fully flooded model has vents to flood all of the chambers. The different chambers of the flooded models are shown in Fig. 3.7, with colour coding to show the partial flooding condition.

The solid lines in Fig. 3.8 show that \mathbf{W} and \mathbf{B} cause a small moment about the pivot (the z axis). The no flooding design creates moments in the opposite direction to the other designs because it is the only positively buoyant design. Little difference is seen in the moment created about the pivot axis by the other three designs. Note that these results are hydrostatic and are calculated for a table tilt $\beta = 5$ deg. The effects of moving fluid are not considered here because the point of this analysis is to minimise the forces other than those related to the flow.




Figure 3.8: Pitching moment about the pivot (solid line) and magnitude of the moment on the bearing (dashed line) as a function of pitch angle for four different airfoil designs submerged in quiescent water. The channel is tilted 5 deg inlet up.

The bearing used in the pivot is a nylon bearing with glass balls,² chosen for its high corrosion resistance and low stiction. The nylon race is very susceptible to distortion so the tolerance of the bearing housing must be very low. To achieve this from the 3D printed part, the housing recess was printed slightly oversize, filled with epoxy then reamed to the correct diameter. This bearing is also sensitive to out of axis torsion. Shakedown testing proved that having bearings at the top and bottom of the airfoil caused high resistance due to low accuracy of the bearing alignment inherent to the chosen manufacturing process. It was determined that good results could be obtained by installing a single bearing in the bottom of the airfoil if the out of axis torsion was minimised by balancing the moments caused by the weight and buoyancy forces. As before the moment on the bearing (M_B , Fig. 3.7) is calculated for the hydrostatic case and the four designs are compared in Fig. 3.8. The dotted lines show that the best results are obtained by the partial or fully flooded designs.

The clearance around the shaft also acts to minimise the axial alignment error. The shaft has a radius of 6.35 mm in a 6.80 mm hole and the bearing is 156 mm from the upper opening. Therefore the maximum alignment error is $\tan^{-1}(0.225/156) = 0.8$ deg. Tighter tolerances are not feasible given the accuracy of the printing process used.

²Canada Bearings part number POM623

The ideal mass ratio of the airfoil for maximum amplitude response is 1. For a first approximation the mass ratio is simply calculated as the ratio of the mass of fluid displaced by the static airfoil to airfoil mass; $m^* = m/V_D\rho_f$. Added mass is ignored for simplicity but it is acknowledged that this will act to increase the mass ratio, so a value less than one is better than a value greater than one. For airfoils with flooding, the mass of the fluid inside the airfoil is included in the airfoil mass because it can be considered trapped. The mass ratio's for the four designs are listed in Table 3.3.

Table 3.3: Mass ratios for various airfoil designs where $m^* = m/V_D\rho_f$

| Design | m^* |
|-------------------|-------|
| Solid | 1.20 |
| No flooding | 0.54 |
| Fully Flooded | 1.12 |
| Partially Flooded | 1.02 |

Given the above analysis, the selected design is the hollow airfoil with 1 mm thick walls. The compartments at the leading edge are sealed and the tail compartments are allowed to flood, see Fig. 3.7. Flooding is achieved through vents drilled in the top and bottom of each airfoil section as highlighted in Fig. 3.9. The recess between the connected parts is also allowed to flood to ensure no air is trapped in the lower section. Other features of the airfoil, shown in Fig. 3.9, are:

- The holes which accept the connector pins are reinforced by a 2.2 mm rod printed with the part, this connects the outer face to the inner structure.
- The mating faces are recessed to ensure the outer face has no gap.
- Wells are printed in the bottom face of the airfoil which are filled with green acrylic paint to allow for optical tracking from below.




Figure 3.9: Renders of the lower airfoil section to highlight the design features.

3.2.3 Model Placement

Two different sets of interface plates were made for this facility. The single airfoil experiments in Chapter 5, where the airfoil was placed directly downstream of the cylinder used one style, here referred to as the in-line setup. All other experiments used the plates shown in all of the above renders, here call the array setup.

In-line Setup

The interface plate for the in-line setup is shown in Fig. 3.10. It has a single plate for mounting the cylinder and airfoil, with all positions in-line with the free-stream. The possible separation distances are 0.052 m to 0.156 m with a resolution of 0.026 m. This was originally designed with a larger cylinder in mind but the experiments were performed with a 0.044 m diameter cylinder which leads to the inconvenient non-dimensional separations of $s/d = 1.18, 1.77, 2.36, 2.95, 3.55$. This was corrected for the second setup.




Figure 3.10: Model mounting locations for in-line configuration, as viewed from below.

Array Setup

The cylinder is mounted to a separate plate, with its location fixed at 205 mm from the beginning of the test section and on the cross-stream mid-line. Airfoils can be installed into any of the threaded holes in the second interface plate. This plate has symmetrical mounting holes allowing it to be installed in either of the configurations shown in Fig. 3.12. The result is that the airfoils can be installed in any position from $s/d = 2.5$ to 4.0 downstream of the cylinder with a resolution of $0.5d$ and at any cross-stream location $p/d = \pm 1.5$ from the mid-line with a resolution of $0.25d$. The coordinate systems and nomenclature are shown in Fig. 3.11.




Figure 3.11: Schematic showing the coordinate system as viewed from below and the dimension convention.

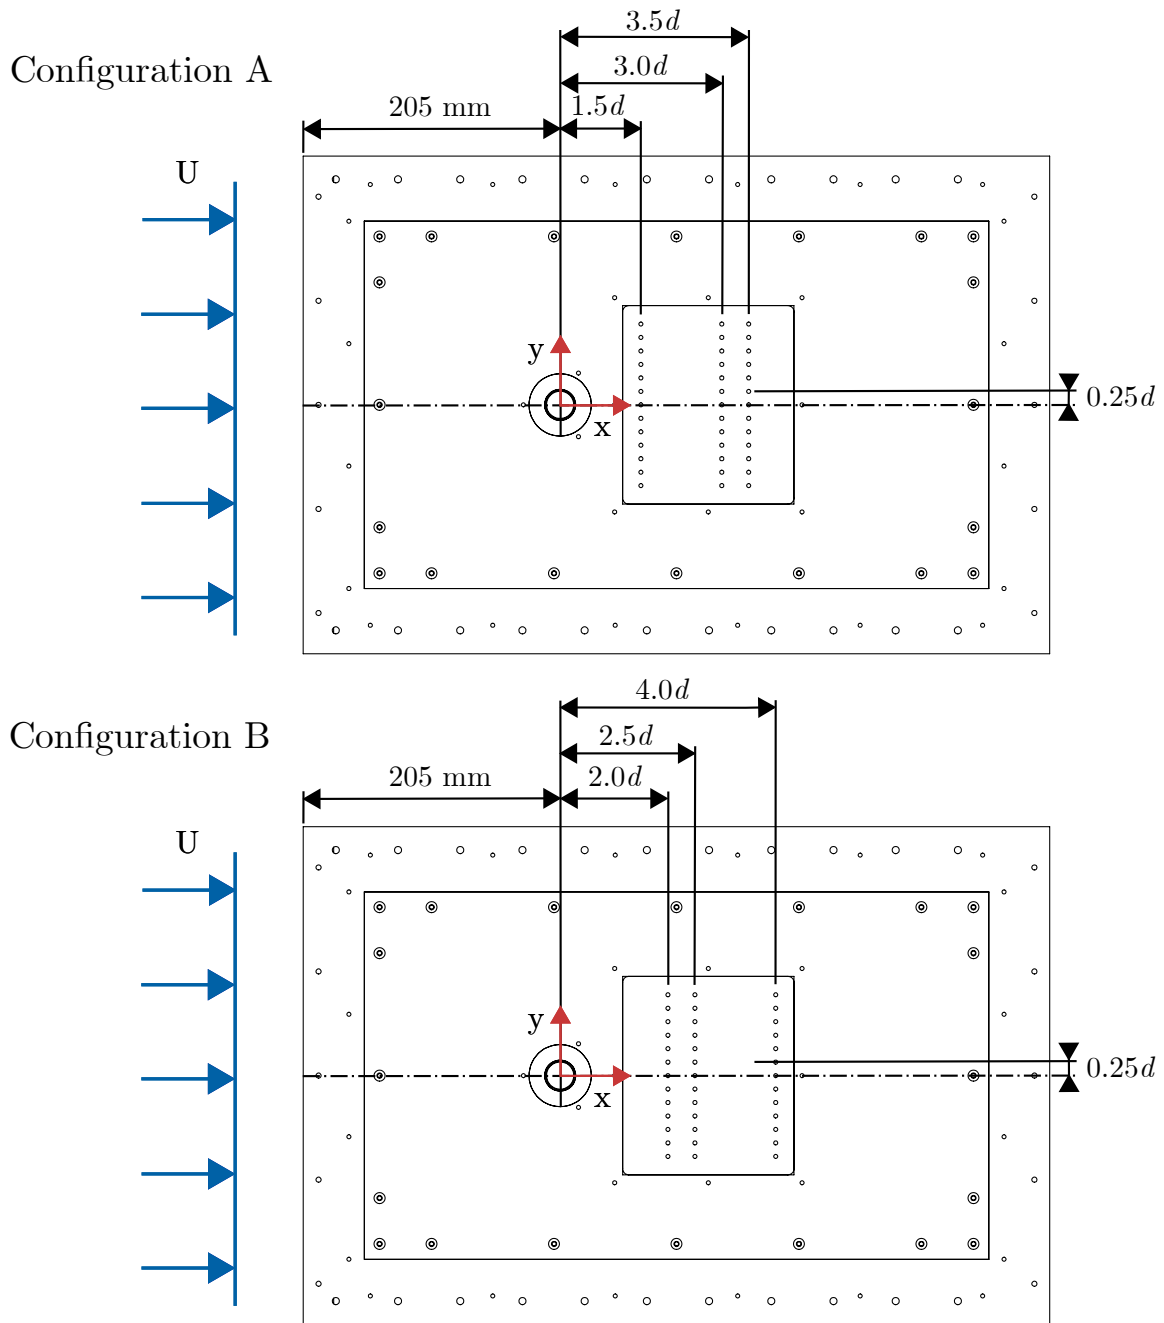


Figure 3.12: Schematic of possible model mounting locations as viewed from below. The two different configurations are achieved by rotating the insert plate.

3.3 Drivers and Electronics

This section details the electro-mechanical systems used to drive the cylinder. The desired movement is a rotational impulse at effectively random intervals. The characteristics of the cylinder rotation were guided by the literature reviewed in Section 2.4.1,³ but were ultimately determined experimentally to incite a measurable response of the airfoil. The chosen impulse is a movement of 0.75 revolutions in 0.3 s which gives a surface speed of 0.345 m/s, approximately nine times the free stream velocity.

cRio Based Control System

The first driver, used for the experiments in Chapter 5, utilised existing components. The facility previously had two linear stages (Danaher 2RB12 with Kollmorgen T23 steppers) with rotation heads (CTP12 steppers) as detailed in Sommer [52]. This was more degrees of freedom than required for this work but in an effort to maintain cross functionality, the drivers were built into a universal frame with improved bearing support as show in Fig. 3.13. The sub assembly of each stage is shown in Fig. 3.14. The vertical tube contains bearings and retaining collets to ensure high axial stiffness of the models. The full assembly is easily attached to alignment dowels in the interface plates.

³In general, a larger amplitude rotation with higher acceleration should incite a stronger response but there are exceptions to this. The lack of applicable data led to the decision to choose the values empirically.




Figure 3.13: Rendered image of the model driver assembly. The system consists of a frame and two movement stages, each with a linear motion and a rotation axis.




Figure 3.14: Rendered image of one of the movement stages. The linear stage is driven by the T23 stepper motor. The rotation axis is driven by the CTP12 stepper motor through a 10:1, right angle gearbox and is supported by the bearing tube. Movements are shown by blue arrows.

The motors are controlled by a National Instruments (NI) cRio 9074 with NI-9512 motor

driver interfaces and Kollmorgen P70530-SDN stepper drivers.⁴ Originally the control system was sitting on top of the water channel. For safety, the electronics were built into a separate cart before testing began.

The control software was programmed in NI LabView. The software makes a random decision whether or not to impulse the cylinder at predefined decision period (T_D). The randomness of the decision is sufficient for the system to be considered stochastic (see Chapter 5), despite the regularity of the decision period. If the decision to impulse is made, a command is sent to the driver interface and the cylinder is driven by one of the CTP12 stepper motors. This system suffered from robustness issues and ultimately the cRio control unit failed. Rather than replace the unit, a new system was built using open-source components.

Arduino Based Control System

The new control system was designed to drive a single CTP12 stepper motor, limiting the functionality to a single degree of freedom. The micro-controller is programmed to make random decisions of whether to impulse the cylinder or not, at regular intervals, similar to the LabView program used previously. The decision period, probability of a positive decision and the rotation direction can be changed with simple buttons and are displayed on a small LCD screen. When the decision to move is made, the micro-controller sends a ‘run’ command over the serial interface to the motor driver, see Fig. 3.15. The motor driver converts the micro-controller’s command into a series of electrical impulses, powering each of the phases of the stepper motor in the correct order, with precise timing. It has many settings to tune the response of the motor, as detailed in [55]. A number of these are required to be accurately set for this system to work due to the relatively high power output of the motor. Charging the coils of the stepper motor generates a back electromotive force (BEMF) which opposes the incoming electric current [55]. To overcome this, the voltage into the motor must increase with the motor speed. The voltage model employed by this stepper driver is given in Eq. 3.3 with variables defined in Table 3.4. With the exception of the electric constant, these values are from the motor’s datasheet [15].

$$V = \begin{cases} R_m |I_{PH}| + \frac{k_e}{4} \omega, & \omega \leq \frac{4R_m}{2\pi L_m} \\ R_m |I_{PH}| + \frac{k_e T_m}{2\pi L_m} + \frac{2\pi L_m |I_{PH}| + k_e}{4} \left(\omega - \frac{4R_m}{2\pi L_m} \right), & \omega > \frac{4R_m}{2\pi L_m} \end{cases} \quad (3.3)$$

⁴The control system schematic, Fig. D.1, shows the layout of this system.

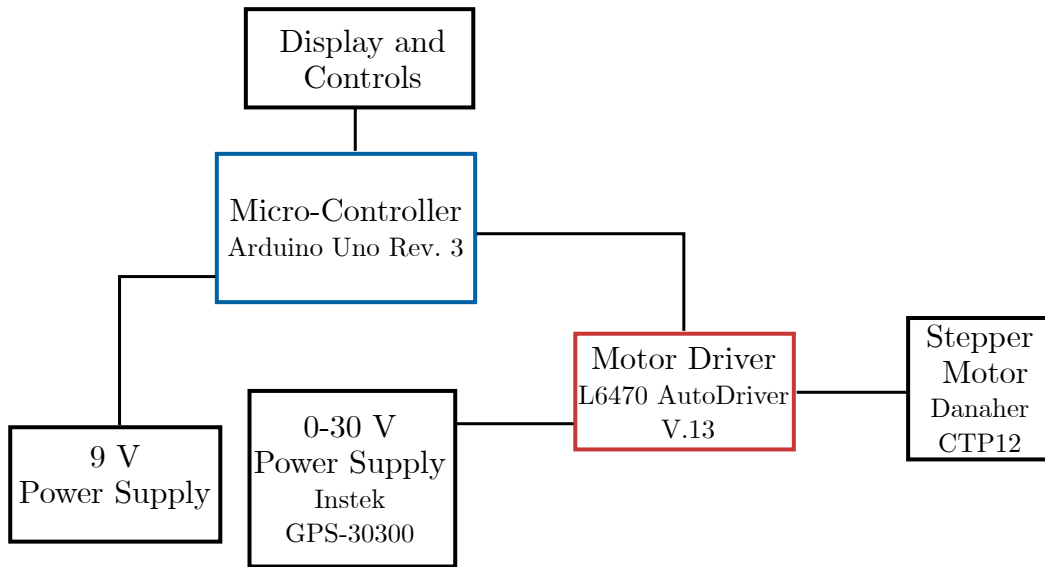


Figure 3.15: Schematic of new control system

Table 3.4: Variables used in the voltage model, Eq. 3.3

| Variable | Definition | Value | Unit |
|----------|----------------------------|-------|-----------------|
| I_{PH} | Phase current | 1.0 | A |
| k_e | Electric constant of motor | 0.01* | $\frac{V}{rad}$ |
| L_m | Phase impedance | 0.012 | H |
| R_m | Phase resistance | 6.5 | Ω |

* Measured

This results in the model shown in Fig. 3.16. The maximum supply voltage for this stepper driver is 45 V, limiting the speed to 2100 steps/s (10.5 rev/s).⁵ To achieve the desired cylinder speed of 2.5 rev/s, the gearbox must be removed and the stepper directly connected to the cylinder. Removing the gearbox results in a significant reduction in available torque but as long as good alignment of the motor, cylinder and cylinder seal is maintained, it is sufficient. This required a new mount to made, shown in Fig. 3.17, which is connected directly to the cylinder mount and these pieces are manufactured with tight tolerances.

⁵The motor driver is able to generate steps much faster than this but the voltage required to turn the motor at higher speeds is not available. The previous system was able to supply over 100 V to the motors.




Figure 3.16: Voltage required by the stepper motor as a function of the rotational speed.




Figure 3.17: Rendered image of the direct motor mount assembly.

3.4 Data Capture

The only data recorded for these experiments is the structure displacements as a function of time. This is done by capturing video from below the structures (Fig. 3.18) and optically

tracking green markers attached to them. The camera is a Nikon D7200⁶ recording at 50 Hz. It is mounted on optical rails built into the frame of the flow facility allowing it to film the structures through the window in the floor of the flow channel. The advantages of this method are that it is cheap, it is easy to expand to many structures without needing a new sensor per structure and the data is always synchronised.

The camera is controlled using the open-source software DigiCamControl. The software for tracking the structures was written in Matlab. At the start of the program, the user is presented with a sample frame, they draw boxes around each structure then the software outputs the angle of each structure relative to the image for the full data set. This does not need to be oriented with the flow or corrected for bias because the data will be symbolised based on relative angles (Section 5.2.2).



Figure 3.18: A rendering of the camera position and optical rails

⁶with lens: Nikon AF-S Nikkor 18-70mm 1:3.5-4.5G ED, set to 18 mm focal length, $f/4.5$

Chapter 4

Results of Simulated Experiments

This study is designed to validate the Matlab code used to calculate transfer entropy, written by this author. Results from simulation data will be compared to an analytical solution. The rules used to generate the data for this experiment are guided by those used by Bossomaier *et al.* [8] to formulate the analytical solution below. The system contains two variables (X, Y) each with two possible states ($-1, 1$). X is deterministic and has no dependence on Y . Every time step the value of X flips creating a periodic signal with a period of 2 time steps, t . Y has a probability of following X such that $\Pr(y_t = x_{t-1}) = \frac{1+c}{2}$ where c ranges from -1 to 1. X is independent of its own past so $\Pr(y_t = -x_{t-1}) = \frac{1-c}{2}$. The analytic solution is given as [8].

$$\begin{aligned} \text{TE}_{X \rightarrow Y} = & 0.5[(1+c)\log_2(1+c) + (1-c)\log_2(1-c)] \\ & - 0.5[(1+c^2)\log_2(1+c^2) + (1-c^2)\log_2(1-c^2)] \end{aligned} \quad (4.1)$$

$$\text{TE}_{Y \rightarrow X} = 0 \quad (4.2)$$

To allow for an investigation into the effects of sampling period and the lag parameter, the data is reshaped into a square wave. To maintain the rules of the system, the period of oscillation of X is $24t$ and the communication delay $X \rightarrow Y$ is $12t$ as shown in Fig. 4.1.

Before analysing the calculation method, it is worth considering the analytical result in Fig. 4.2. From the rules of the system, it appears that the ‘strength’ of the causal link, $X \rightarrow Y$, is proportional to the absolute value of c . At $c = 0$, the probability that Y takes the value X is equally likely as Y taking the value of $-X$. At high values of $|c|$, Y has a high probability of following one of X or $-X$. However, in Fig. 4.2, it is seen that this is not the case. This difference comes from the difference between Granger causality and

the intuitive understanding of causality. As discussed in Section 2.3.6, TE is equivalent to Granger causality which is predicated upon the existence of unique information in the two considered variables. It is apparent that for $c = 1$, $Y = X$ so there is no uniqueness and Granger causality is zero, despite the fact that Y is caused by X .

4.1 Validation of TE Algorithm

Results calculated from a 5×10^6 data point synthetic data series deviate from the analytic solution with an RMS error of only 0.17×10^{-3} bits, thereby validating the code. The results are too close for the difference to be clearly shown in Fig. 4.2.




Figure 4.1: A sample of the time series used in this simulation.

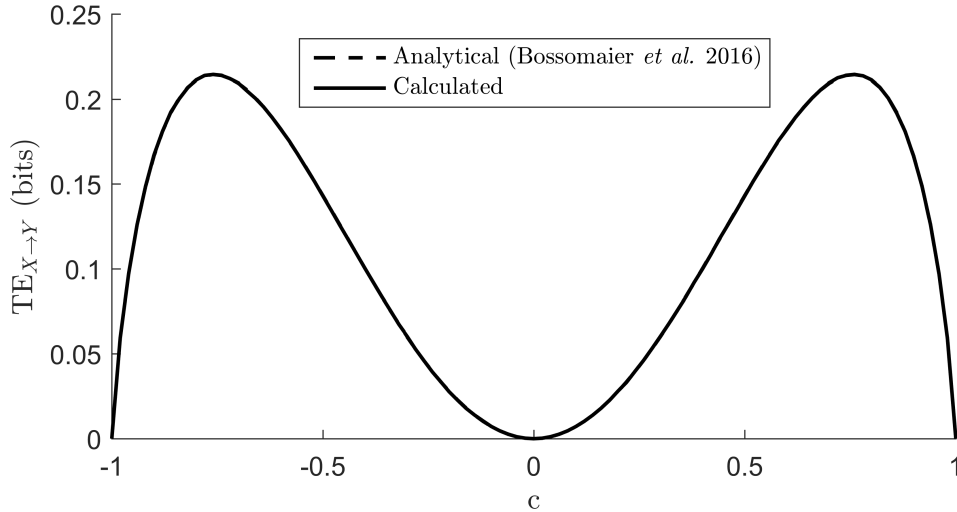


Figure 4.2: A comparison of calculated and analytic solutions to validate the algorithm used to calculate transfer entropy. $TE_{Y \rightarrow X}$ is not shown because it is exactly 0 for both methods. RMS error = 0.17×10^{-3} bits

4.2 Analysis of Calculation Parameters

As described in Section 2.3, calculating transfer entropy requires two parameters Δ and δ to be selected. The values of these parameters are reported in the literature to be related to physical characteristics of the system as described in Chapter 2. In short, Δ relates to the time period between information events and δ is the amount of time it takes for information to pass from the target to the source. The parameters can be selected using *a priori* knowledge of the system as they were in the preceding section. However, since the purpose of using this tool is to learn information about the system, data driven methods for selecting the parameters are required. This section details a simple method for achieving this and validates the relationships between the parameters and the physical properties.

In the described system, the communication lag is known to be $\hat{\delta} = 12t$ and the period of oscillation is $24t$ which means a change occurs every $12t$; this should be the critical Δ . The expectation is that, at these parameters, $TE_{X \rightarrow Y}$ and $TE_{Y \rightarrow X}$ will match the analytical solution and $|TE_{\text{net}}|$ will be maximised. To test this, the data is down-sampled for a range of Δ from 1 to 24. Similarly, TE is calculated for a range of δ from 1 to 24. This results in two, two-dimensional arrays $TE_{X \rightarrow Y}$ and $TE_{Y \rightarrow X}$ per c value. All of the results are plotted in Fig. 4.3 which, through its wide spread and large number of

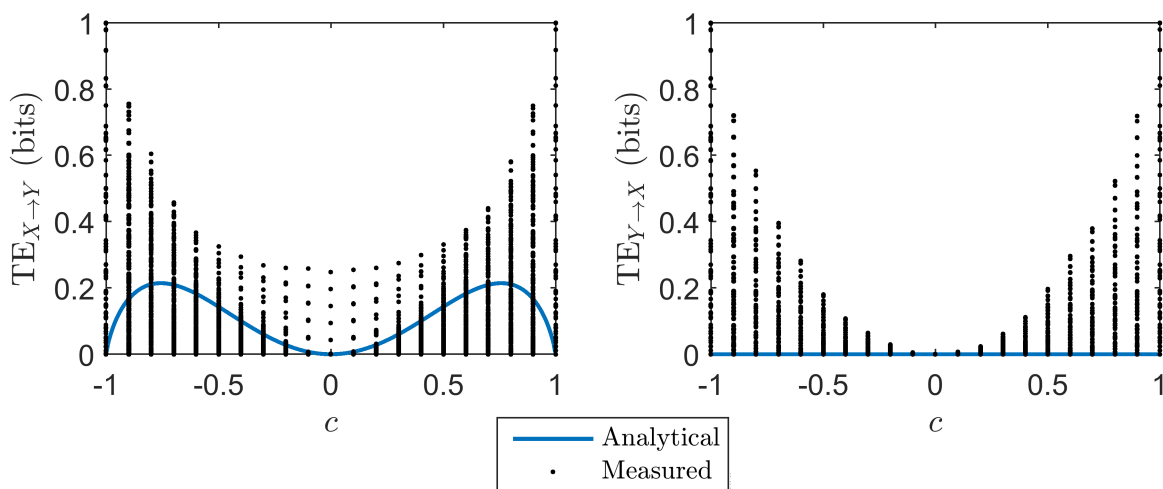


Figure 4.3: Scatter of TE results for parameters $\Delta = (1, 24)$ and $\delta = (1, 24)$ compared to analytical results. Note that the analytical result for $\text{TE}_{Y \rightarrow X} = 0$.

spurious values,¹ demonstrates the need to correctly select the calculation parameters Δ and δ . Fig. 4.4 demonstrates that $|\text{TE}_{\text{net}}|$ is less chaotic and tends to conform better to the analytical results which lends credit to the idea that this is a useful tool for parameter selection.

To simplify the discussion, the case of $c = 0.2$ will be looked at in isolation. Fig. 4.7 shows the results for all Δ and δ parameters tested, for each TE component ($\text{TE}_{X \rightarrow Y}$, top left, $\text{TE}_{Y \rightarrow X}$, top right and $|\text{TE}_{\text{net}}|$, bottom). The component $\text{TE}_{X \rightarrow Y}$ is dominated by the regions about $\delta = 4, 16$, $\Delta = 8$ which is not the expected result² but it does match Fig. 4.3 where there are many values greater than the analytical solution. It can be seen in Fig. 4.5 that this spike at $\Delta = 8$ is also seen in the entropy calculations, $H(X)$ and $H(Y)$, which is propagating through to the TE calculation. This one case is not a good demonstration of the principle but it will be demonstrated in Chapter 5 that filtering out sampling periods based on the entropy is a robust, data driven method for avoiding spurious results. In Fig. 4.8 this filter is applied by setting all values corresponding to $\Delta = 8, 16, 24$ to 0. In Fig. 4.3, these filtered values are highlighted red. Applying this filter allows for a better analysis of the patterns, revealing there are additional spikes in the data (Fig. 4.8 at $\Delta = 4, 6, 12, 18$) that were not detected by the entropy. The spikes occur when the sampling period is a harmony of the event period (T_0) similar to the demonstration in

¹In this context 'spurious' is used to describe values greater than the analytical result. This is because it is asserted that a false positive detection of causality (or a high measurement of TE) is incorrect, where a low measurement is a detectability issue.

²The expected result is $\Delta = 12$, $\delta = 12$.




Figure 4.4: Measured TE_{net} compared to analytical results for simulated experiment. Filtered Δ values are those that are identified as causing erroneous measurements of entropy.

Fig. 2.4. The periodicity of this signal increases the harmonic content such that anomalies are created at $\Delta = iT_0/j$ where i and j are any positive integer. The power of the harmonic is affected by the magnitudes of i and j as well as the number of modes in which it occurs. To avoid these anomalies, Zhang *et al.* [74] calculated TE over a range of Δ , then identified where the results are relatively insensitive. However, that step is not necessary here and it will be demonstrated in Chapter 5 that filtering using the entropy is sufficient for that experimental data.

On a macroscopic level there is a clear pattern in both the Δ and δ directions. The pattern along δ has a period of $12t$ and is expected due the periodicity of X which, as explained in Section 2.3.1, makes it impossible to distinguish which cycle is causing the response. The oscillations with respect to Δ are variations in the detectability as detailed by Barnett & Seth [5]. While it is not expressly stated, the results of Barnett & Seth [5] suggest that Δ can be selected by which value causes TE to be maximised. This does not work at the component level ($TE_{X \rightarrow Y}$) for this data but it does work using TE_{net} , i.e. TE_{net} is maximum at $\Delta = T_0 = 12t$. Interestingly, at this sampling period the periodicity hides all of the lag information and $TE_{X \rightarrow Y}$ is uniform with respect to δ as shown in Fig. 4.6. However, the TE measurement does match the analytical result. Using a slightly different sampling period, the communication delay can be accurately detected as the value that




Figure 4.5: The relationship between sampling period and entropy for the synthetic data at $c = 0.8$

maximises $\text{TE}_{X \rightarrow Y}$ as is done by Wibral *et al.* [66]. This is evidenced by the peak at $\delta = 12$ in $\text{TE}_{X \rightarrow Y}(\Delta = 7)$ in Fig. 4.6, but the measurement of TE is incorrect. The figure also shows that δ can only be selected from the maximum of TE_{net} , as advised by Zhang *et al.* [74], if Δ is correctly selected first. Here, the slight shift in Δ had a minimal effect but for $\Delta = 7$, the communication lag is incorrectly identified as $9t$. This indicates that the sampling period should be selected before analysing communication delays or both should be identified together.

Applying the principles just discussed, the results for the full range of c are truncated to those sampling periods which cause $|\text{TE}_{\text{net}}|$ to be maximised. In all cases, the results are independent of δ and match the analytical result as shown in Fig. 4.9. Note the slight increase in error compared to Fig. 4.2 is caused by the reduction in the sample length to 5×10^4 data points which was done to reduce the computation time. Fig. 4.10 shows that in all cases except $c = 0$, the selected sampling period is $\Delta = 12$. For the case $c = 0$, the TE magnitude is zero so the selection is sensitive to noise. To estimate the communication delay, the sampling period is shifted by 1 from that just identified.³ i.e. if $\Delta = 12$ is identified as the correct sampling period, $\Delta = 11$ is used to estimate the communication lag. To remove the effect of periodicity, the range of δ is reduced to 1 to 23 otherwise the result was as likely to be 0 or 24 as it was 12. As seen in Fig. 4.11, the correct lag, $\delta = 12$

³This is only necessary because of the periodicity of X .




Figure 4.6: TE results for select Δ 's. Analytical result is $\text{TE}_{X \rightarrow Y} = 0.0279$ bits.

is identified in all cases except $c = 0$ as before.




Figure 4.10: Identified critical sample period using data driven methods

Figure 4.11: Identified critical communication delay using data driven methods




Figure 4.7: Contours of transfer entropy for all calculated Δ and δ . $TE_{X \rightarrow Y}$ at top left, $TE_{Y \rightarrow X}$ at top right and $|TE_{\text{net}}|$ at the bottom.




Figure 4.8: Contours of transfer entropy with respect to Δ and δ with the entropy filter applied. $TE_{X \rightarrow Y}$ at top left, $TE_{Y \rightarrow X}$ at top right and $|TE_{\text{net}}|$ at the bottom.




Figure 4.9: Final results of $TE_{X \rightarrow Y}$ using data driven methods to select Δ and δ parameters.

4.3 Simulation Conclusions

By comparing synthetic data to an analytical result, it has been that the numerical method introduced in Section 2.3 is equivalent to the analytical method and the code used is correct. Additionally, It has been confirmed that there is are direct relationships between the parameters Δ and δ and the system properties event timing and communication delay, respectively. Data driven methods for identifying the correct values of these parameters, without *a priori* knowledge of the system, have been introduced. Some of the nuances of interpreting the data have also been discussed; particularly in relation to periodic variables. This experiment also highlighted the difference between causality and Granger causality, seen as detectability issues for c values that resulted in a lack of unique events in the two variables, $c \rightarrow \pm 1$.

In summary:

- In most cases, it was possible to identify the causal relationship $X \rightarrow Y$ using transfer entropy.
- The calculation parameter Δ was shown to be related to the physical time at which events occur. In this case how often the symbol X changes.
- Some values of Δ caused the data to be misrepresented. These values show up as a spike in entropy and thus can be filtered out.

- For this data, it was possible to select Δ using $|\text{TE}_{\text{net}}|$ once the data has been filtered using the entropy results.
- The calculation parameter δ was shown to be related to the time it takes for the information to get from the source to the target.
- The communication delay could be found from $\text{TE}_{X \rightarrow Y}$, as suggested by Wibral *et al.* [66], once an appropriate Δ was selected.
- For this data, it was not appropriate to select the communication delay based on $|\text{TE}_{\text{net}}|$, as had been suggested by Zhang *et al.* [74], due to the periodicity of X .

Chapter 5

Methodology and Results of Two Body Experiments

5.1 Preliminary Analysis

The purpose of these experiments is to validate the hypothesis that transfer entropy can be used to infer directed causal relationships in fluid-structure systems. They will also be used to identify the best methods for applying TE to this type of data. The results will be validated against expected trends or known values where possible. These experiments are designed to be similar to the experiments of Zhang *et al.* [74] to ensure a high chance of success but are modified to expand our understanding of the application of information theoretic tools to fluid-structure interaction problems. In the present experiments a NACA 0018 airfoil of chord length $l = 0.060$ m is placed in the wake of a circular cylinder of diameter $D = 0.044$ m, both in a free stream of water at velocity $U = 0.04$ m/s, as detailed in Chapter 3. The rotational position of the cylinder is stored in time series C and the rotational position of the airfoil is stored in A . The cylinder is driven by a random decision process with regular decision periods, T_D , similar to a first order Markov Decision Process. When the decision to move is made the cylinder performs a rotational impulse. This rotation disrupts the fundamental vortex shedding that occurs in the wake of the cylinder while it is stationary, discussed in Section 2.4.1. In terms relevant to this topic, this action adds information to the system at a random interval. Information is advected along with the vortex in the form of distortions to the flow field. These distortions act to create a pitching moment on the airfoil. The passive pivot at the quarter chord of the airfoil allows it to rotate in response. This system only allows for communication in one direction; from

the cylinder to the airfoil. The cylinder is coupled to the motor so it cannot be moved by the airfoil. This is the first known condition to which the results will be compared; $TE_{A \rightarrow C}$ must be statistically zero.

The disturbances created by the cylinder are expected to diffuse as they travel downstream which is likely to cause a decrease in airfoil response as the separation between the two structures increases. This was seen to be true and detectable in the experiments of Zhang *et al.* [74], so that condition is used here to validate the method, i.e. the final results should show a decrease in $TE_{C \rightarrow A}$ as separation increases. The other effect of increasing separation is the time taken for the information to pass between structures should increase. As discussed in Section 2.3.1, this should result in an increase in the value of the lag parameters at which peaks in TE are observed. Assuming the information is advected by vortices, it is unclear what the exact communication time is because there are a few complicating factors. Firstly, it is not a good assumption that the vortex will travel at the free stream speed when it is so close to solid structures which alter the local fluid velocity and may interact with the vortex, causing it to accelerate or decelerate. It is also unclear how long after the rotation the vortex sheds and where it will shed from. Assuming that these factors are uniform between cases, the communication delay should be linearly proportional to the separation of the structures, s . By trialling five separations it is possible to evaluate the accuracy of the delay parameter, $\hat{\delta}$, from the trend without having to know the true value of communication lag.

To investigate the effect of the sampling period (Δ) on the results, three different decision periods are trialled. The expectation set by Barnett & Seth [5] is that the sampling period should be close to the event period and Schreiber [50] has shown that the event period is the decision period for systems driven by a Markov process. Therefore, the sampling period should be close to the decision period. The decision periods are chosen to not be harmonic with the fundamental shedding period, which is estimated at 5.6 s. This period is estimated using the Strouhal number for an equivalent stationary, isolated cylinder, so it is not expected to be exactly correct. It is, however, the best estimate available and serves as a reasonable starting point. The chaotic temporal relationship between the fundamental shedding and the rotation is critical because it will decrease the significance of the fundamental shedding interaction in the TE calculation.

The separation distances and decision periods for all 15 cases are given in Table 5.1. Separation is non-dimensionalized using the cylinder diameter, $s^* = s/d$. The decision period is non-dimensionalized by the convection time of the airfoil, $T_D^* = T_D U/l$.

Table 5.1: Case numbers and parameters

| s^*/T_D^* | 1.33 | 2.00 | 2.50 |
|-------------|------|------|------|
| 1.18 | 1 | 2 | 3 |
| 1.77 | 4 | 5 | 6 |
| 2.36 | 7 | 8 | 9 |
| 2.95 | 10 | 11 | 12 |
| 3.55 | 13 | 14 | 15 |

Based on this preliminary analysis of the system, the following hypotheses are defined to validate the use of the proposed methodology for this type of fluid-structure system.

1. In all cases $TE_{A \rightarrow C}$ should be statistically zero.
2. $TE_{C \rightarrow A}$ should decrease as separation between the cylinder and airfoil increases.
3. TE should peak at lags (δ) matching actual communication times. The actual communication lag is unknown but the trend with respect to separation should be linear; $\hat{\delta} \propto s$.
4. The sampling period (Δ) should be approximately the decision period.

5.2 Methodology

5.2.1 Captured Data

Optical tracking is used to measure the angular displacement of the two structures. A camera placed below the test section acquires images at 50 Hz. These images are processed in Matlab to track the position of two green markers attached to each of the structures. From these positions, the structure angles are measured to create the time-series shown in Fig. 5.1. This angle is relative to the camera frame, so it is not exactly the structure angle. No viewing angle correction is made because the data will be symbolized in such a way that the small error will disappear. The example time series shows the cylinder is mostly stationary with short impulses of movement. The markings on the cylinder are approximately aligned to be 45 degrees to the free-stream and the rotation moves it through 270 degrees creating two distinct cylinder positions. This is seen as the plateaus of the angle at approximately ± 0.75 rad and was done to simplify the symbolization process. The spacing between impulses demonstrates how the driving process can be used to approximate a random signal despite the regularity of the decision period. The airfoil is seen in the first plot of Fig. 5.1 to be constantly oscillating but demonstrates little periodicity. The

second plot is for a case with greater separation between the structures which results in lower amplitude responses.

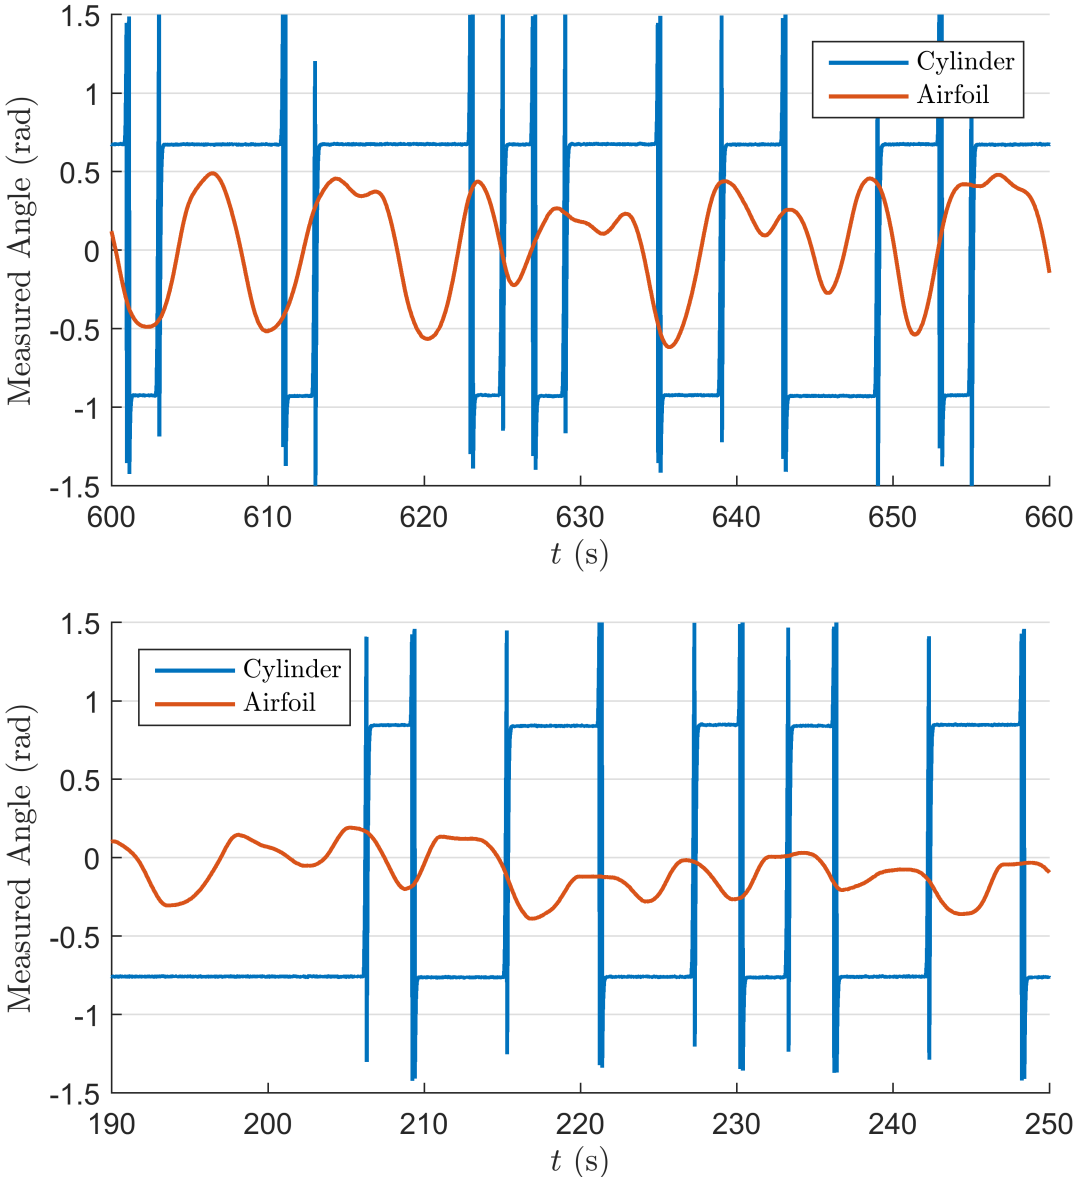


Figure 5.1: Samples of raw data for single airfoil experiments. Top: case 1. Bottom: case 5. The angle is measured relative to the camera frame which is approximately aligned with the free-stream.

5.2.2 Symbolization

Symbolization reduces the spatial information of the data so the rules must be carefully constructed to avoid removing important behaviour information. It was demonstrated by Zhang *et al.* [74] that two symbols was sufficient for their experiment, which is similar to this one. They did demonstrate that increasing the number of symbols to six increased the accuracy of the results but the computational intensity increases by a factor of 32. For minimal computational intensity the decision was made to limit the number of symbols to two. This will be shown in the following sections to be sufficient.

The measured property of the structures is the angular position, but this is not the behaviour of interest. Here the experiment is exploring if movement of one structure causes a movement response of the other structure. For the cylinder it is sufficient to symbolize on the condition “moving” or “not moving” because the details of the movement are constant and unimportant. As such, the symbols for the cylinder data are determined by the rule:

$$c_t = \begin{cases} \beta_t - \beta_{t-1} \geq 0.1, & 1 \\ \beta_t - \beta_{t-1} < 0.1, & -1 \end{cases} \quad (5.1)$$

where β is the measured angle and t refers to a specific time instance. The threshold of 0.1 radians is included to filter out measurement noise. Figure 5.2 shows that the symbolized data faithfully represents the measured data.




Figure 5.2: Sample of cylinder symbolization.

The same rules cannot be used to symbolize the airfoil data because this would not accurately portray the behaviour of the data. Looking at the raw data it is seen that it would almost never be in the “not moving state”. Once again following the framework set out by Zhang *et al.* [74], the symbols are constructed based on the direction of rotation as below:

$$a_t = \begin{cases} \alpha_t > \alpha_{t-1}, & 1 \\ \alpha_t < \alpha_{t-1}, & -1 \\ \alpha_t = \alpha_{t-1}, & a_{t-1} \end{cases} \quad (5.2)$$

where α is the measured airfoil angle. The slow motion of the airfoil makes these data highly sensitive to noise because the change in position per time step is often of the same order of magnitude as the noise. When the data is down-sampled before it is symbolized this is not an issue because the movement per time step is much greater, but when the data is symbolized first it is necessary to use a simple smoothing function. Figure 5.3 shows the symbols created from a sample of airfoil angle data with a 20 point moving average smoothing function applied.




Figure 5.3: Sample of airfoil data symbolization.

Using these symbols strips the data of all amplitude information. Thought experiments are proposed here to show the significance of this. In a first experiment, the airfoil responds to every cylinder rotation by rotating 15 degrees, then rotating back to 0 angle of attack. Then, in a second experiment the airfoil responds to every cylinder rotation by rotating 3

degrees and then back. Intuitively the communication in the second experiment is not as strong as in the first but using these symbols, both will have the same TE. The information that is kept is whether the airfoil responds to the cylinder rotation, the degree of that response is considered less important for this experiment. In other cases, the amplitude data may be necessary. This highlights the importance of understanding the research question before setting the symbolization rules.

Alternate Symbolization Rules

There are many other ways to symbolize these data but testing has identified the above method as the one that gives the best agreement to the known conditions listed in Section 5.1 and it is consistent with the method published by Zhang *et al.* [74]. It will be demonstrated in the next section that this method creates the appearance of two responses of the airfoil, one for the initial rotation and another for the counter-rotation. This can be removed by setting the rules as

$$a_t = \begin{cases} \alpha_t > 0, & 1 \\ \alpha_t < 0, & -1 \\ \alpha_t = 0, & a_{t-1} \end{cases} \quad (5.3)$$

This would still not contain amplitude information but would more faithfully represent the response. In practice this worked well for some cases but generally failed because there is no mechanism to center the airfoil at 0 degrees angle of attack. The pivot is placed at the quarter chord so there is no restoring force at shallow angles and the airfoil often 'settles' at a non-zero angle of attack. This is demonstrated in the bottom plot of Fig. 5.1 where the airfoil pauses at 0.13 rad ($t \approx 212$ s), then next time it settles at -0.12 rad ($t \approx 220$ s). As a result the oscillations in the range $t = (220, 250)$ only become a positive angle once. The symbols using this system then fail to capture most of the response.

5.2.3 Determine Test Length

To test for convergence of the results a twenty minute test was broken into intervals from two to twenty minutes. The on-board memory of the camera set the limit on test length. This could have been extended by a hardware change but the analysis shows that this is not necessary. The expectation is that the convergence time will be related to the number of events that occur during the test, i.e. the number of decisions made. To account for this,

the test was performed using the longest decision period, 3.75 s. The results are shown in Fig. 5.4 to have a spread of less than 10% at ten minutes. A test length of twenty minutes is selected as a balance of accuracy and convenience.




Figure 5.4: Scatter of results for different test lengths.

5.3 Results

5.3.1 Down-Sample First Method

The literature suggested that symbols should be created based on down-sampled data [74], so that method was originally employed. The TE_{net} results for the 15 cases are given in Fig. 5.6. One case is given here in Fig. 5.5 as an example.

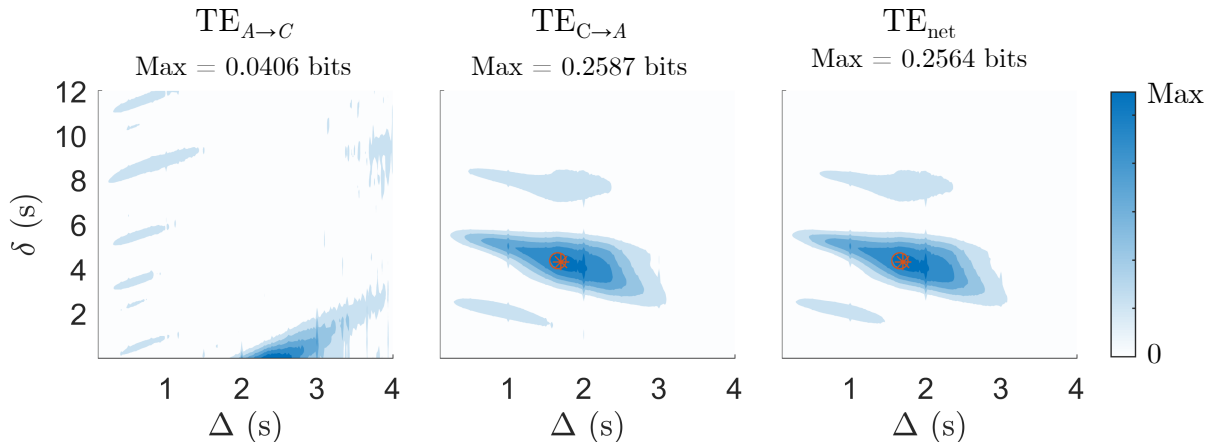


Figure 5.5: TE results for case 1. The plots are self-normalized so the patterns can be compared but the relative magnitudes cannot. The orange asterisk indicates the location of the peak TE_{net} and the orange circle indicates the peak of $TE_{C \rightarrow A}$.

These plots are normalised by the maximum TE for that case allowing the patterns with respect to Δ and δ to be analysed, but mostly obscuring the trends with respect to separation and decision period. Those results will be discussed later. Note that TE_{net} and $TE_{C \rightarrow A}$ are almost identical because $TE_{A \rightarrow C}$ is several orders of magnitude smaller than $TE_{C \rightarrow A}$; as such, only TE_{net} needs to be considered for this analysis. The $|TE_{net}|$ results for all 15 cases are given in Fig. 5.6. It is immediately evident from the steep gradients of the plots that the magnitude of TE is highly sensitive to both the sample period and lag parameter, as it was in Chapter 4. Using this method, the selection of the two parameters is strongly linked. An incorrect selection of Δ will cause a vastly different identification of the critical lag. The results demonstrate the predicted periodicity with respect to both parameters but it is not entirely consistent between cases. For example, cases 7 to 15 exhibit clear diagonal bands where cases 2, 3 and 6 exhibit less structure. In addition to the large scale variations, sharp distortions are seen at distinct Δ 's. One such anomaly is seen at $\Delta = 2$ s in Fig. 5.5, TE_{net} and $TE_{C \rightarrow A}$. It will later be demonstrated that these spikes are erroneous and should be filtered out. Two methods for selecting parameters have been identified in the literature, selecting the parameters which maximise $|TE_{net}|$ or selecting the parameters which maximise the individual components. The parameter selections for these two methods are shown in Fig. 5.5 and Fig. 5.6 as an orange asterisks and circle respectively. The selection method has very little impact for this type of data but the method using TE_{net} is slightly more predictable. The root mean square error (RMSE) of all cases is 1.11 s for the net method and 1.14 s for the component method.

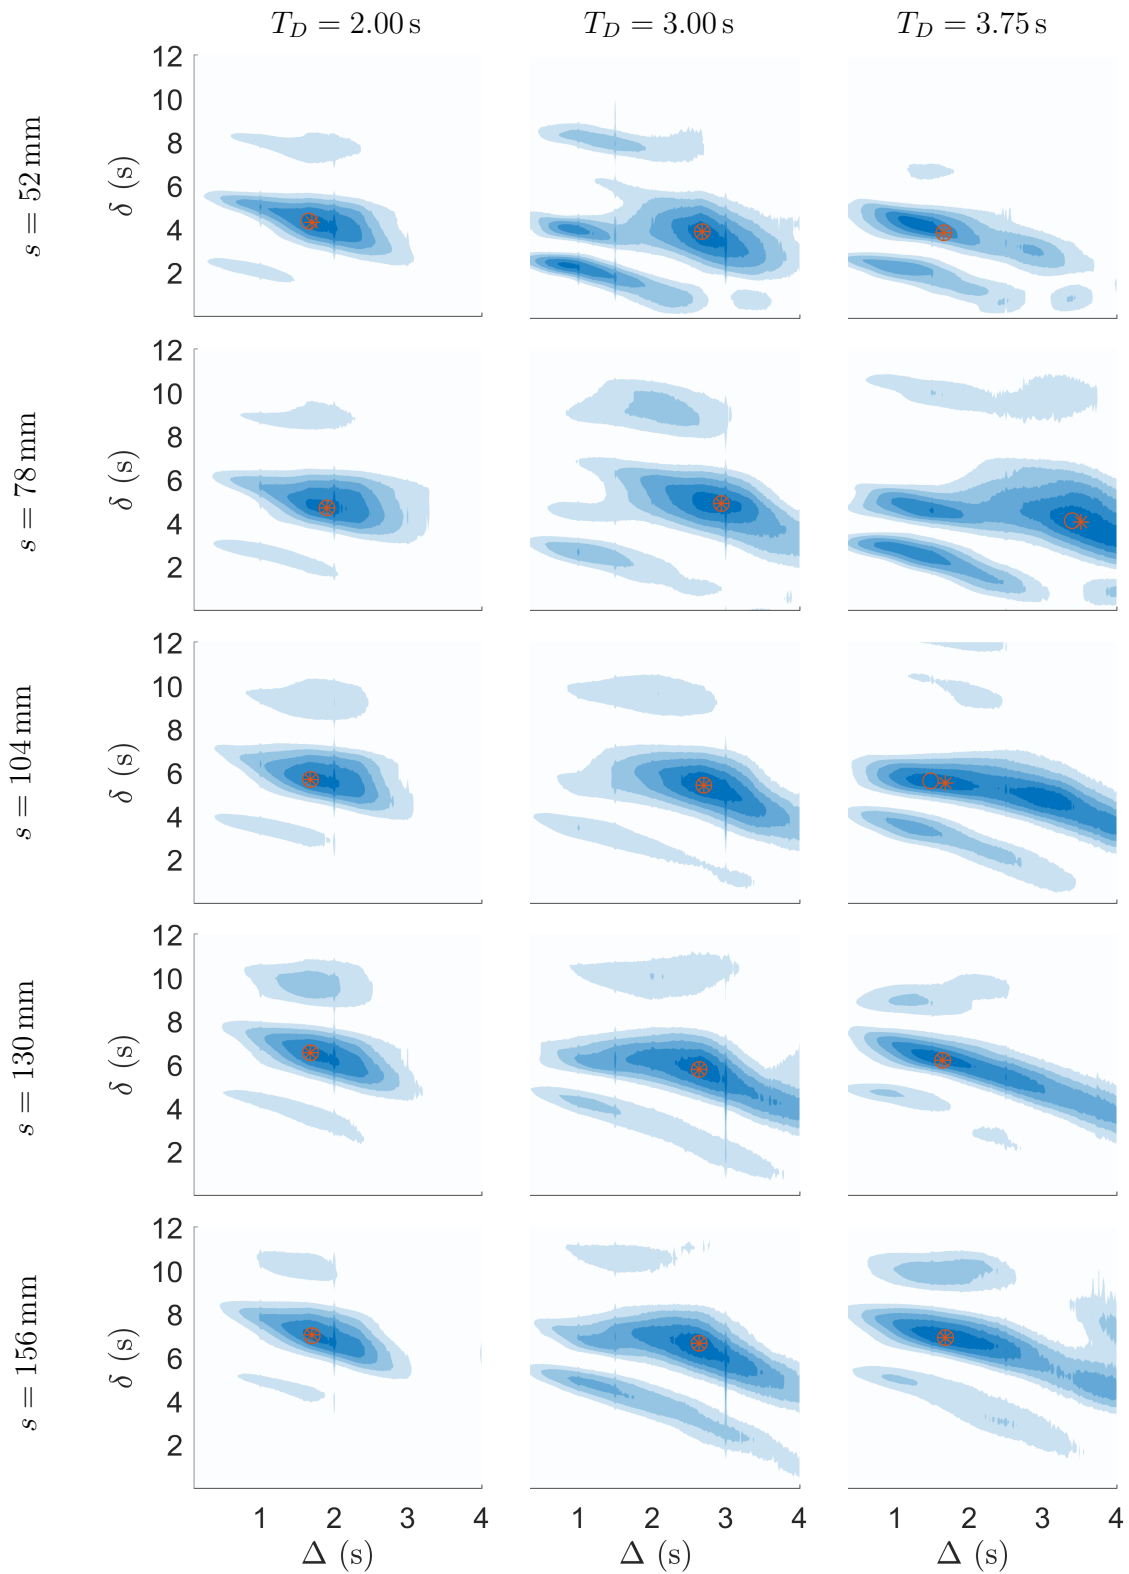


Figure 5.6: $|TE_{\text{net}}|$ for all 15 cases for a range of Δ and δ parameters, calculated using the down-sample first method. The orange asterisk indicates the location of the peak TE_{net} and the orange circle indicates the peak of $TE_{C \rightarrow A}$.




Figure 5.7: Critical lag values selected using TE_{net} (left) and $TE_{C \rightarrow A}$ (right). Black lines are fitted using linear regression. The inverse of the gradient of this line is used to estimate the advection speed.

From the preliminary analysis the peak is expected to be close to $\Delta = T_D$. The selection of Δ is relatively accurate, using either method, for cases with $T_D = 2.00, 3.00$ s; however, for $T_D = 3.75$ s they tend to select $T_D/2$. The reason for occasionally selecting $T_D/2$ is not clear but it being the first harmonic is an encouraging sign that it is not a random error. The true value of $\hat{\delta}$ is not known but it should trend as $\hat{\delta} \propto s$. The location of the peak is seen to increase with s in Fig. 5.6 but the linearity of this increase is better shown in Fig. 5.7. Using linear regression, the advection speed is estimated to be 0.036 m/s (RMSE(δ) 0.26 s) using the parameters selected from TE_{net} and 0.037 m/s (RMSE(δ) 0.31 s) using $TE_{C \rightarrow A}$. These are both reasonable values considering a free-stream velocity of 0.04 m/s. Note that the RMSE is on the estimate of δ which is then used to calculate the advection speed estimate, not the error of the advection speed itself.

To investigate the distortions of the TE results with respect to Δ , the entropy of the two variables are plotted in Fig. D.2 with an example given here in Fig. 5.8. Given the highly oscillatory behaviour of the airfoil, it is expected that the entropy will be close to one because it should spend about the same amount of time in either the clockwise rotating state and the counter-clockwise rotating state. The expected value of the entropy of the cylinder is derived from the expected number of impulses, $T_t \text{Pr}(d = 1)/T_D$ where T_t is the test length in seconds and $\text{Pr}(d = 1)$ is the probability of the driving process returning a decision to move. The amount of time the cylinder spends in the moving state ($C = 1$) is the product of the number of impulses and the amount of time each impulse takes (T_P). The probability of the cylinder being in the moving state is then defined by Eq. 5.4.

$$\text{Pr}(C = 1) = \frac{T_P \text{Pr}(d = 1)}{T_D} \quad (5.4)$$

The expected value of the entropy is then calculated by substituting Eq. 5.4 into Eq. 5.5

which is adapted from Eq. 2.6.

$$H_C = -\Pr(C = 1) \log_2(\Pr(C = 1)) - (1 - \Pr(C = 1)) \log_2(1 - \Pr(C = 1)) \quad (5.5)$$

The expected entropy of C for the three decision periods $T_D = (2.00, 3.00, 3.75)$ s are $H_C = (0.38, 0.29, 0.24)$ bits respectively. Comparing the analytical result to the measured (Fig. D.2) shows a significant difference. It is interesting that the TE results still tend to be accurate and usable. This is possibly due to the formulation of TE which is set up as a difference of entropies (Eq. 2.11). It does however cast doubt on this method so other methods are explored. This analysis was kept because it is an established, published method which gives erroneous results for this type of data. That is not to say that the published results are incorrect because it is seen that the entropy of the airfoil data is correct. This means the algorithm must be tailored to the data and introduces the entropy as a low cost sanity check of the results.




Figure 5.8: Entropy of variables for test case 1 using the down-sample first method. Black dots indicate the sampling periods which will be filtered out by the entropy filter.

5.3.2 Symbolize First Method

In an attempt to remedy the erroneous entropy issue, it was found that the error arose from the symbolization process. The logic of the symbolization rules is sound, but creating them from the down-sampled data caused the issue. It is not clear why this happens, but the

following section provides sufficient evidence to suggest that for this data it is more correct to symbolize the full data then down-sample the symbols. The entropy of the cylinder data for all 15 cases are shown in Fig. D.5 to closely match the analytical result from Eq. 5.5. Using Fig. 5.9 as an example, it can be seen that with the exception of the sharp spikes, the measured entropy is in close agreement with the analytical solution. The spikes are caused by sampling periods which are harmonic to different features of the raw data. This was demonstrated in Section 2.3.2 for a simple square wave. The data here has greater harmonic content because of the random nature of the system. An in-depth analysis of this point adds very little to the understanding of the system, but it is sufficient to say that the spikes are erroneous so they should be filtered out. To that effect, any sampling period which caused a significant spike in the calculated entropy of either variable was not used in the following analysis.




Figure 5.9: Entropy of case 1 using the symbolize first method. The dashed line indicates the analytical result. Black dots indicate the sampling periods which will be filtered out by the entropy filter

As before, the full results are shown in the appendix (Fig. 5.11) and the example of case 1 is given here in Fig. 5.10. The plots show low sensitivity to Δ compared to that of δ . More importantly, the peaks of δ are independent of Δ , meaning that an incorrect selection of one will not cause the incorrect selection of the other as it did in the previous method.

Although it is difficult to see in the plots, TE trends toward zero with the sampling period,¹ as expected. The TE results also have sharp spikes (with respect to Δ) caused by erroneous sampling as discussed in Chapter 4. As before, the sharp spikes can be filtered out by removing Δ 's which correspond to peaks in erroneous entropy; shown as black dots in Fig. D.5. It also has long time scale variations caused by the harmonic relationship with the decision period, similar to that in Chapter 4. Unlike the simulation, the driver of this system (C) is not periodic, meaning the period between events at C can take assume any integer multiple of T_D . This is evidenced by Fig. 5.11, in which the peak TE_{net} ² is often at harmonic values of Δ/T_D . It will be shown in Section 5.3.5 that the mean of the null distribution also increases, so the higher magnitude does not necessarily mean a result is more statistically significant. The reasoning is somewhat convoluted but the consequence is that selecting Δ from the maximum of TE_{net} is not as reliable with this data as it was in the simulations of Chapter 4.

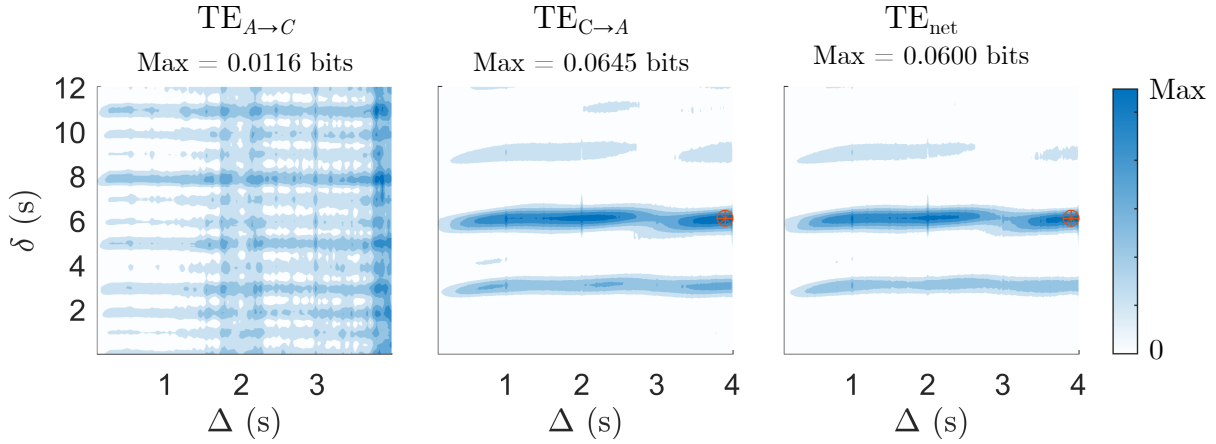


Figure 5.10: TE results for case 1. The plots are self-normalized so the patterns can be compared but the relative magnitudes cannot. The orange asterisk indicates the location of the peak TE_{net} and the orange circle indicates the peak of $\text{TE}_{C \rightarrow A}$.

The behaviour of TE with changing δ appears different with this method. In Fig. 5.11, all cases have at least two identifiable peaks,³ but the distance between them is not very consistent. At first glance, the existence of two peaks suggests that either there are two communication paths from the cylinder to the airfoil or one impulse of the cylinder results in two responses from the airfoil at different delays. However, the actual cause is the

¹ $\text{TE} \rightarrow 0$ as $\Delta \rightarrow 0$ due to detectability issues, discussed in Section 2.3.2.

²Indicated by the asterisk.

³Shown as dark bands at certain δ values.

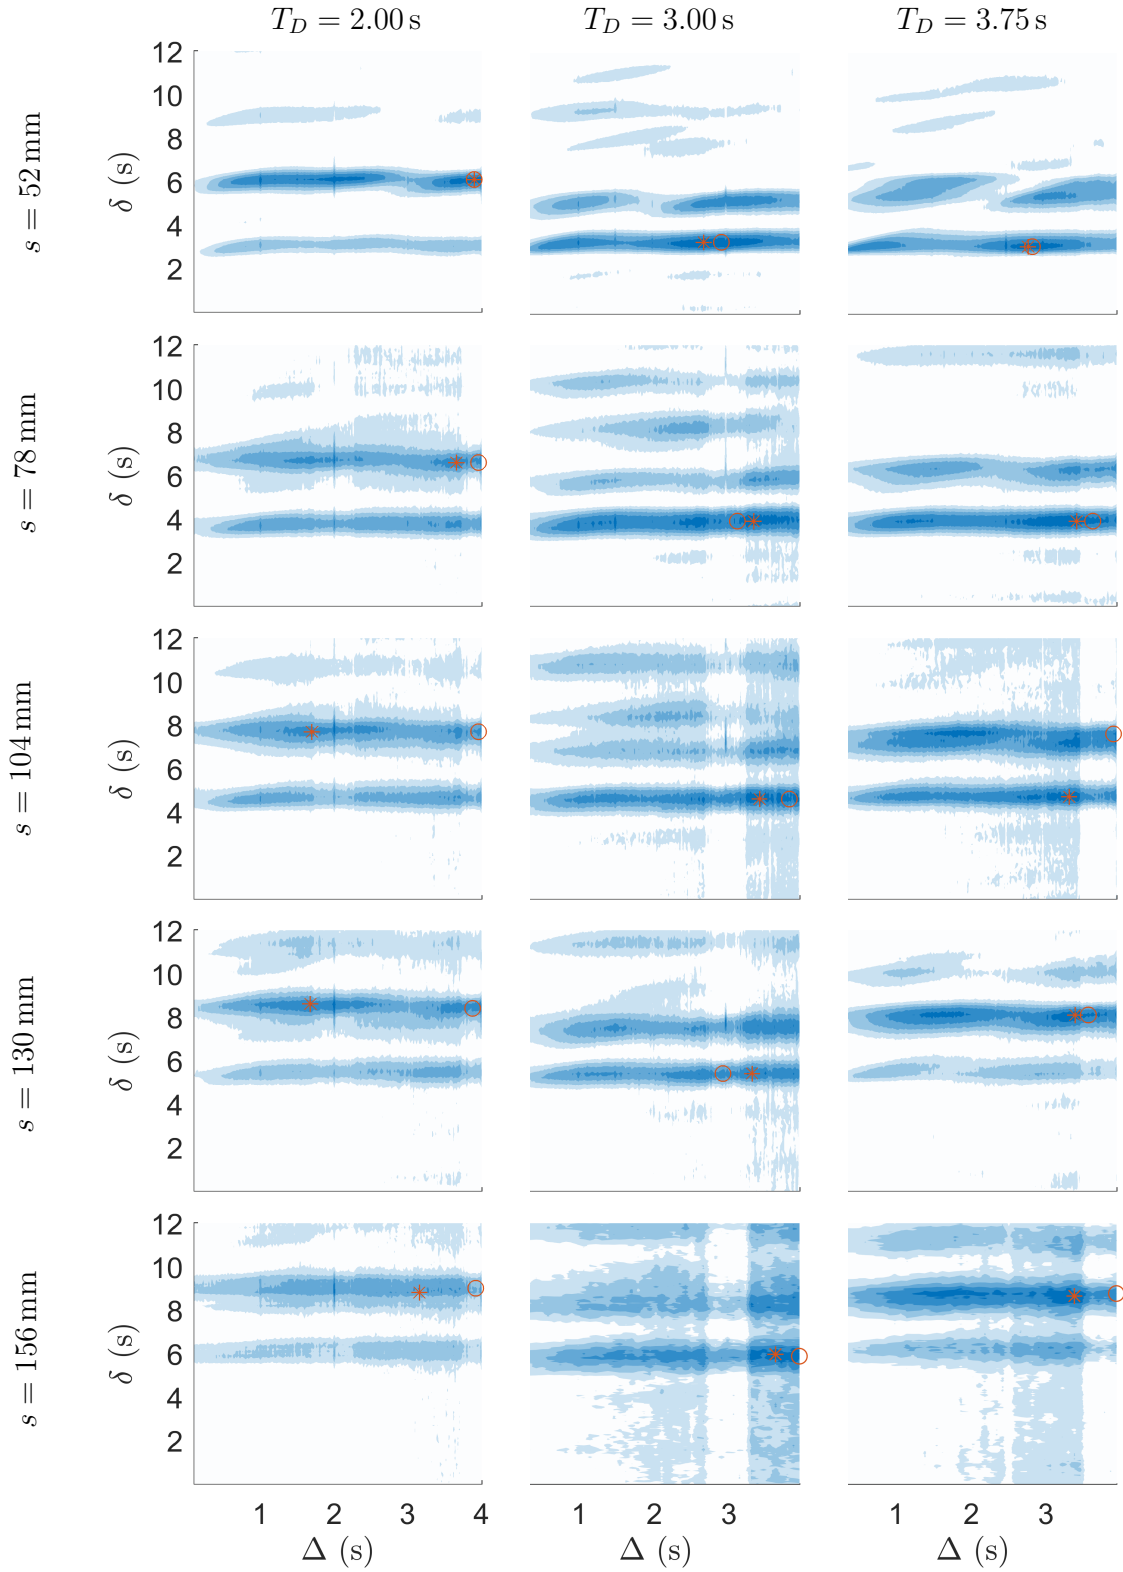


Figure 5.11: $|TE_{\text{net}}|$ for all 15 cases for a range of Δ and δ parameters, calculated using the symbolise first method. The orange asterisk indicates the location of the peak TE_{net} and the orange circle indicates the peak of $TE_{C \rightarrow A}$.

symbolization of the airfoil data. It was stated that the airfoil is expected to rotate, then counter-rotate in response to the passing vortex. With the current symbolisation rules (Eq. 5.2) this creates two symbols for that one response. The separation of the peaks corresponds to the time between the initial rotation and the counter-rotation which is affected by the vortex coherence and the distance between the airfoil and the vortex. In shake-down testing using dye injection, it was noted that in some cases the shed vortex would collide with the airfoil and breakdown. In other cases the vortex would pass as a coherent structure.⁴ This case dependent behaviour may explain the noted variation in the separation of the peaks. In some cases, case 5 for example, third and fourth peaks are seen. These may be numerical anomalies where a response is incorrectly attributed to a cylinder movement or they may be due to the regularity of the driving process. These later peaks are always less significant so they are not considered to be part of the primary behaviour of the system and, for the most part, are ignored in this analysis.

For clarity and because data driven selection of Δ has been shown to be unreliable for these data, the data will be truncated by preselecting Δ . It has been established that Δ should be about T_D , but it cannot be exactly equal, as this results in erroneous sampling; as evidenced by the entropy measurements (Fig. D.5). Therefore the value closest to T_D which is not filtered by the entropy filter is used. This truncated data is presented in the easier to read line plots Fig. D.8 and Fig. 5.12. The Null Threshold line is generated by bootstrapping the data as discussed in Appendix A. Anything below this line is considered not statistically different from 0 with 95% confidence. This accounts for both the inherent bias in the TE calculation and the distribution of the results. It is seen in Fig. D.8 that the component $TE_{A \rightarrow C}$ is never significant, which meets the first known condition from Section 5.1. It is apparent that using either TE_{net} or $TE_{C \rightarrow A}$ to select the dominant communication lags will have no effect on the results for this type of data because $TE_{A \rightarrow C}$ has no spurious results.

⁴A thorough flow visualisation analysis was not feasible for this experiment because of the large number of modes of the system.




Figure 5.12: TE of case 1 as a function of lag for a set sample period, $\Delta = T_D$

Following the work of Wibral *et al.* [66], the dominant communication lags $\hat{\delta}$ are identified by the local peaks of $TE_{C \rightarrow A}$, excluding those below the null threshold. Linear regression is used to fit the trend of the delay as a function of the separation between the structures in Fig. 5.13. The advection speed is estimated as the inverse of the gradient of this trend. To account for the bimodal distribution, the two most dominant delays are treated separately, in order of which occurs first. Due to the low confidence in cases 13 - 15, these cases are excluded from the regression. Both trend lines have a gradient of 28.7 s/m giving an advection speed estimate of 0.035 m/s, which is in good agreement with the free stream velocity of 0.04 m/s. The RMSE of the two regressions are 0.070 s and 0.368 s, confirming the initial observation that the timing of the second peak is inconsistent between cases. Since the actual communication delay is unknown, there is no ground truth with which to validate this method of selecting $\hat{\delta}$. However, the high linearity displayed in Fig. 5.13 is a good indication.




Figure 5.13: Dominant communication delays of cases 1 to 12 with linear fit. The different colours indicate either the first or second peak.

Figure 5.14 presents the measured magnitudes of $TE_{C \rightarrow A}$ for all 15 cases as a function of the separation between the structures. The magnitudes are bias corrected by subtracting the mean of the null distribution from the measured value. Values that are not significantly different from zero are indicated by the + symbol. These results confirm there is a measurable decline in the interaction strength as the separation increases as is expected. A reduction is also seen as the decision period increases in Fig. 5.15, although this appears less prominent.




Figure 5.14: Bias corrected $TE_{C \rightarrow A}$ as a function of separation between structures. Three tests are performed at each separation, corresponding to the different decision periods. The designations first and second when referring to the delays denote the order of occurrence, i.e the lowest magnitude of δ , not the relative magnitude of TE.




Figure 5.15: Bias corrected $TE_{C \rightarrow A}$ as a function of decision period. Five tests are performed at each decision period, corresponding to the different separations. The designations first and second when referring to the delays denote the order of occurrence, not the relative magnitude of TE.




Figure 5.16: Comparison of symbolization methods.

5.3.3 Alternate Symbols

It was stated that the two peaks in $TE_{C \rightarrow A}$ with respect to δ were caused by the symbolization method. To prove this, the method is changed to that defined by Eq. 5.3 and the methods are compared in Fig. 5.16. Using the alternate symbolization method, only one peak is seen in the middle of the two peaks detected by the actual symbolization method (defined by Eq. 5.2). This would be a more intuitive symbol to use because it makes sense to lump the rotation and counter-rotation into a single event. However, this method fails to detect causality in all cases except 1, 2 and 3 where the oscillations are very large as discussed in Section 5.2.2.

5.3.4 Check for Self-Feedback

Measuring TE along the pathways $A \rightarrow A$ and $C \rightarrow C$ (Fig. 5.17 and Fig. D.9) reveals significant feedback in the system. Both components show high feedback at $\delta \rightarrow 0$ due to the time it takes for the movement to be complete; more correctly, the time it takes for the state to change. $TE_{C \rightarrow C}$ has easily explainable spikes at integer multiples of the decision period not caused by real feedback but by the regularity of the movement. $TE_{A \rightarrow A}$ has a significant spike at approximately $\delta = 3$ s, which is independent of either separation or decision period suggesting it may be real. The feedback delay ($\hat{\delta}_{AA}$) is shown in Figure 5.18 and Fig. D.10 to match the separation of the two peaks of $TE_{C \rightarrow A}$. There is no precedent in

the literature to advise whether two peaks in $\text{TE}_{C \rightarrow A}$ would show up as a peak in $\text{TE}_{A \rightarrow A}$. However, Wibral *et al.* [66] found that feedback appears in the pairwise results as peaks at $\hat{\delta}_{CA} \pm \hat{\delta}_{\text{feedback}}$ as discussed in Section 2.3.1. There is sufficient evidence that the twin peaks are caused by the symbolization method so it is concluded that the feedback is spurious. It is interesting to note that there is no data driven method for distinguishing between a system with significant self feedback and one with two communications paths between the structures but this is outside the scope of this experiment.




Figure 5.17: TE results along self feedback pathways.




Figure 5.18: $TE_{C \rightarrow A}$ for case 1 with indicators of communication delays predicted from the self feedback results.

5.3.5 Null Distribution

In the contour plots Fig. 5.11 and Fig. 5.10, it appears that the magnitude of TE does not change significantly with Δ . This is true but it does not tell the whole story because the null distribution is not shown there. In all cases the significance of the measured TE value must be compared to the null distribution shown in Fig. 5.19 to have a large variation for sampling periods close to the decision period. The abstract nature of the TE equation makes it very difficult to identify the cause of the shape of this graph but the deviation from the general trend always occurs about the decision period for all 15 cases.




Figure 5.19: Mean of the null distribution for case 7, $T_D = 2$ s. Note that some of the expected spikes are missing due to the relatively low Δ resolution this data was calculated at. For example, a spike is expected at $\Delta = 2$ s but the null distribution was not calculated at this value.

Marschinski & Kantz [32] suggest that TE can be bias corrected by subtracting the null distribution from the measured value. However, for this experiment, this causes a decrease in TE around $\Delta = T_D$ which ensures the correct Δ is not selected. The selected Δ always corresponds to a value on the lower $A \rightarrow C$ trend-line, meaning it is compared to a lower statistical significance threshold. As a result, there is an increased tendency to falsely infer communication along the path $A \rightarrow C$.

5.4 Single Pathway Conclusions

It has been shown in the preceding analysis that it is possible to infer directed causal relationships in fluid-structure systems using transfer entropy. Additionally, this method has been shown to be useful for identifying dominant communication delays between structures which can provide additional information about the mechanisms acting in the system. By comparing the results obtained by different methodologies, some of the nuances of this technique have been highlighted. The majority of the issues arose from the symbolization of data due to the vast decrease in information. It was demonstrated that changing the symbolization rules and order of operations had a significant impact on the results. The

symbolization process was successful but did require careful design with consideration for the important qualities of the data. At this stage it was necessary in order to maintain an acceptable level of computational intensity. It is impressive that such a complex, chaotic system could be analysed in such depth using such a small amount of information.

Once a reliable symbolization method has been identified, the calculation on TE is straight forward but contains two unknown parameters. No reliable data driven method for selecting Δ was identified here. Selecting it based on the maximum TE value was shown to have a moderate aggregate error but could be off by as much as 100%. With proper symbolization this had no effect on estimates of the communication delays but caused a significant change in the measured TE. Preselecting sample periods using knowledge of the system, as is common in the literature, was shown to give accurate results. For a system driven by a Markov process, which is similar to the driving process used here, it is recommended to use $\Delta \approx T_D$; this proved to be correct. Entropy of the individual variables was introduced as a cheap method for filtering out sampling periods which lead to erroneous results and as a method for checking the accuracy of the symbolization process. Neither of these points have been addressed in any of the reviewed literature.

Interpretation of the results with respect to the communication delay were shown to be more complex than simply selecting one correct TE value. Communication was shown to occur over a range of delays making the case to case comparisons difficult. For this experiment, where it is known that only one communication mechanism exists, it was shown that two delays were detected. Both delays are true because the response of the target takes a finite amount of time. This was shown to be caused by the symbolization method but the question of what is the true delay is still ambiguous and, for the most part, is up to the researcher to decide. For these experiments the logical delay to choose is from the cylinder rotation to the start of the response of the airfoil but this is open to debate. Ultimately it was decided that consistency was most important so the trend of the results were used. By selecting either the first peaks or second peaks, the information advection speed was estimated by linear regression to be 0.875 U.

Comparison of the magnitude of TE is highly interpretive but some definitive statements can be made. By calculating the null distribution, it was shown that no statistically significant communication occurred from the airfoil to the cylinder. Additionally, it was shown that the strength of the communication decreased as the separation between the structures increased and as the decision period increased. It is unclear what the magnitude of this decrease means in real terms.




Figure 5.20: Maximum measured TE for various cross-stream offsets. Magnitude is bias corrected.

5.5 Cross-Stream Offset Experiment

This short experiment forms a bridge between the two body and three body experiments. For the later, it is desired to know the influence of the cylinder on airfoils not directly in its wake. To get this information, an airfoil is placed $s = 1.5 D$ downstream of the airfoil with cross-stream offsets $p = -1.5, -1.0, -0.5, 0.0, 0.5, 1.0, 1.5 D$. Only one decision period is used, $T_D = 2 s$ and the cylinder always rotates counter-clockwise. Entropy of the variables is calculated for a range of sample periods surrounding T_D to aid in selecting that parameter. From this the selected sampling period is $\Delta = 2.02 s$ as it had been for the earlier experiments. TE is calculated over the range $\delta = (0, 12) s$ to ensure all communication is captured.

For this experiment, the only property of interest is the maximum TE. Figure 5.20 shows the asymmetry of the communication path caused by the uniform rotation direction of the cylinder. For a down-stream separation of $1.5 D$, the influence of the cylinder is shown to be significant in the range $p/D = (-1, 0)$.

Chapter 6

Results of Three Body Experiments

These experiments are intended to advance the functionality of transfer entropy in the field of fluid-structure interactions by expanding the number of structures being analysed. Transfer entropy is always a pairwise calculation, measuring the strength of communication from one structure to another. When more than two structures are present, the calculation can be performed for each pair to generate a map of the communication network. Section 2.2.3 discusses that the TE calculation must be conditioned on all variables of the system to avoid spurious detection of communication which is actually caused by mutual influence or cascading information pathways. The physical experiments in this chapter build on the previous experiments by adding a second airfoil, identical to the first, into the system. As before, the cylinder is driven in a manor similar to a first order Markov process. This time the decision period is limited to 2.0s as the effect of changing this parameter has already been covered in Chapter 5. Simulations of highly idealized systems are provided at the beginning of this chapter to help set expectations for the results and to identify the effects of the various indirect communication pathways.

6.1 Simulations

The following simulations represent highly simplified systems to investigate the various effects caused by the presence of a third structure. Conditional TE was introduced in Section 2.3.4 as a method for accounting for the effect of additional variables on the primary pair (the target and source variables) being analysed. There are three proposed methods for analysing these systems, depending on the desired result.

- Plot conditional TE over a range of δ and ζ to highlight patterns or trends.

- Non-uniform results with respect to the ζ parameter indicate the communication from source to target is dependent upon the conditioning variable.
- Analyse the dominant lags, attempting to determine the true cause of that interaction.
 - In Section 2.3.4 it was stated that if Z was a mutual driver of X and Y , $\text{TE}_{X \rightarrow Y}$ would give a false positive detection of communication with a delay $\hat{\delta}_{ZY} - \hat{\delta}_{ZX}$.
 - Additionally, relayed communication $Z \rightarrow X \rightarrow Y$ will give a false positive detection of communication $\text{TE}_{Z \rightarrow Y}$ at delay $\hat{\delta}_{ZX} + \hat{\delta}_{XY}$.
- Hypothesis testing can be used to evaluate specific ideas. For example, ‘is this path significant?’

The first two methods are employed for the simulations below. All three will be applied to the analysis of the experimental data.

6.1.1 Mutual Driver

In this system, both X and Y are driven by Z at delays $\delta_{ZX} = 8$, $\delta_{ZY} = 21$ but no communication occurs between them as shown in Fig. 6.1. Z is driven by a first order Markov process with a decision period of 100 time-steps and a probability of a positive decision of 0.7. When a positive decision is returned, Z switches state for 5 time-steps. The probability of X or Y responding to Z is 0.7. As discussed, it is expected that calculating TE from X to Y will give a spurious, significant measurement of TE at $\delta = \delta_{ZY} - \delta_{ZX}$. The correct peaks $\text{TE}_{Z \rightarrow Y}(\delta = 21)$ and $\text{TE}_{Z \rightarrow X}(\delta = 8)$ as well as the spurious result can be seen in Fig. 6.2a. Conditional TE of the three significant components ($Z \rightarrow Y$, $Z \rightarrow X$, $X \rightarrow Y$) are plotted for a range of δ and ζ in Figs. 6.2b to 6.2d. It can be seen in Fig. 6.2c that the dominant communication delay $\hat{\delta}_{ZX} = 8$, identified by the peak of $\text{TE}_{Z \rightarrow X}$ in Fig. 6.2a, appears in the conditional plot at the same δ . The result is uniform across the full range of ζ confirming that the interaction $Z \rightarrow X$ is independent of Y . In the other two components, the dominant delays, $\hat{\delta}$, again appear as peaks in conditional TE but now the results are not uniform with respect to ζ suggesting some dependence on the conditioning variable. The conditional result $\text{TE}_{X \rightarrow Y|Z}$ (Fig. 6.2d) demonstrates that TE is insignificant if it is conditioned on Z at $\zeta = 21$. Not coincidentally, the conditioning lag which gives the correct result is the real delay from the conditioning variable to the target; i.e, the correct result is obtained at $\zeta = \hat{\delta}_{ZY}$.

If the delays of the system are unknown, as they will be in the experiments, the communication delay $\hat{\delta}_{ZY}$ can be obtained from the conditional TE, $\text{TE}_{X \rightarrow Y|Z}$. Using the reverse

of the logic above, the critical ζ is that which causes a significant reduction in TE.

Care must be taken when using this analysis method. Consider Fig. 6.2b, which has a significant change in TE at $\zeta = 13$. This appears to support the hypothesis that $\text{TE}_{X \rightarrow Y}(\delta = 13)$ is real. However, it has already been shown that $\text{TE}_{X \rightarrow Y}$ is insignificant if conditioned appropriately. The cause of the reduction in $\text{TE}_{Z \rightarrow Y|X}$ is the similarity of the time-series that result from mutual driver and cascading systems as was briefly discussed in Section 2.3.4. This result is a spurious reduction in conditional TE¹ which is related to the relative probabilities of X following Z ($\Pr(x_t = z_{t-\delta_{ZX}})$) and Y following Z ($\Pr(y_t = z_{t-\delta_{ZY}})$). When both probabilities are high, it is impossible to distinguish between a mutual driver effect and a cascading information effect. It is up to the researcher to determine whether one or both of these are possible using some other knowledge of the system. The measurement of TE should be taken from the pairwise calculation because the reduction in conditional TE is spurious. Proof of this comes from knowledge of the system; the coupling strengths ZX and ZY are equal, so $\text{TE}_{Z \rightarrow X}$ should equal $\text{TE}_{Z \rightarrow Y}$, which they are in the pairwise calculation, but not in the conditional.

In summary,

- The components $Z \rightarrow X$ and $Z \rightarrow Y$ are determined to be significant with delays 8 and 21 respectively.
 - The TE measurement for $Z \rightarrow Y$ is taken from the pairwise calculation, not the conditional because the reduction seen there is spurious.
- $X \rightarrow Y$ is determined to be insignificant because it can be completely conditioned out.
- The remaining components are determined to be insignificant because significant TE is never measured.

¹This is why it is not possible to blindly select significant ζ values as those which cause a significant change in TE.


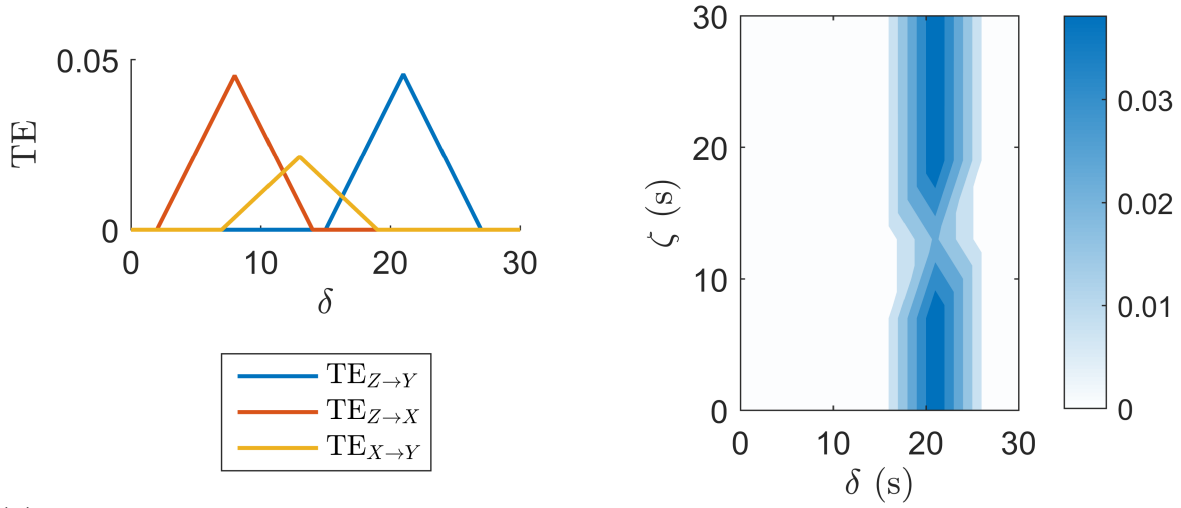



Figure 6.1: Communication pathways for mutually driven X and Y .



(a) Pairwise TE results. For clarity, components with no measurable TE are not shown.

(b) Conditional $TE_{Z \rightarrow Y|X}$ results



(c) $TE_{Z \rightarrow X|Y}$

(d) $TE_{X \rightarrow Y|Z}$

Figure 6.2: TE results for simulation of three variables with Z mutually driving X and Y .

6.1.2 Relayed Information

The behaviour of variables X and Z in this system is the same as that of the previous simulation. However, this time Y is driven only by X with a probability of 0.7. Fig. 6.3 shows that the information transfer from X to Y is wholly relayed from Z . Fig. 6.4 shows the expected false detection of $\text{TE}_{Z \rightarrow Y}$ at $\delta = \delta_{ZX} + \delta_{XY} = 34$. In this case the definition of a false result is less clear than the previous simulation. It can be said that the significant measurement of $\text{TE}_{Z \rightarrow Y}$ is false because there is no direct communication along that pathway. However, regardless of the pathway, the information is generated at Z and received by Y so it is a legitimate pathway. Conversely, a significant measurement of $\text{TE}_{X \rightarrow Y}$ is an accurate indication of the information transferred along that path, but no information is being generated by X . Once again the appropriate definition is up to the researcher. Here it is proposed that a separate definition is used, in which the relay path $Z \rightarrow X \rightarrow Y$ is a separate entity. It will be shown here that the TE measurement along the relay path matches the pairwise measurement $\text{TE}_{Z \rightarrow Y}$ at the specific delay related to the relay path.

The conditional results for $\text{TE}_{Z \rightarrow Y|X}$ in Fig. 6.5b again show that the false path can be conditioned out at a conditioning lag equal to the real delay $\zeta = \hat{\delta}_{XY}$. From this it is determined that there is no direct communication along the path $Z \rightarrow Y$. Looking at the component $\text{TE}_{X \rightarrow Y|Z}$ in Fig. 6.5a, it can be seen that TE is reduced, but still significant, at $\zeta = \hat{\delta}_{ZY}$, where $\hat{\delta}_{ZY}$ comes from the pairwise calculation. This is where the analysis becomes subjective; the researcher must decide what constitutes real communication. It is known that no information is generated at X so it could be argued that this is not real. However, in the next experiment it will be demonstrated that it is not usually possible to distinguish between the system in Fig. 6.3 and one where information is generated at X as in Fig. 6.6. Also, as in the previous simulation, the amount of reduction seen at $\zeta = \hat{\delta}_{ZY}$ is related to the probability of Y responding to X . When $\Pr(y_t = x_{t-13}) = 1$, $\text{TE}_{X \rightarrow Y|Z}(\delta = 13, \zeta = 34) = 0$ and it is impossible to distinguish between mutual driver and cascade effects.

If one was to test the hypothesis, ‘there is direct communication $X \rightarrow Y$,’ the conditional results alone would not be sufficient to disprove it; there is significant TE measured for all conditioning lags. However, the hypothesis ‘there is significant communication along the path $Z \rightarrow X \rightarrow Y$,’ can be tested by evaluating $\text{TE}_{Z \rightarrow Y}$ at $\delta = \hat{\delta}_{ZX} + \hat{\delta}_{XY}$. This results in a significant measurement of TE, thereby inferring the existence of a significant relay. The researcher could then conclude that the communication $X \rightarrow Y$ is caused by the relay.




Figure 6.3: Communication pathways for a system where information generated at Z is relayed to Y by X . The path $X \rightarrow Y$ is drawn starting at the end of the $Z \rightarrow Y$ to symbolize the idea that no information is generated at X .

The next simulation will demonstrate that this is only reasonable if it is known that X cannot generate information.




Figure 6.4: Pairwise TE results. For clarity, components with no measurable TE are not shown.




Figure 6.5: Conditional TE results for simulated system with three variables with a relayed pathway.

6.1.3 Relayed Information With Active Relay

To better reflect the physical experiments in the following sections, the system is modified so X also generates information. This is analogous to the airfoil responding to the fundamentally shed vortices (those not cause by the rotation of the cylinder). This is achieved by driving X with a separate Markov process with a different decision period to that of the process which drives Z . The delay from X to Y is unchanged by the pathway and is drawn separately in Fig. 6.6 to indicate the two different pathways.

The pairwise TE for this system (Fig. 6.7) has peaks at the same δ locations as the previous simulation but the relative magnitudes are significantly different. TE measured from X to Y is by far the most significant due to its involvement in two pathways; directly communicating its own information and relaying information from Z . The mean of the null distribution is increased, resulting in insignificant but non-zero results across all δ . This is likely caused by the multiple event periods which now exist in the system, allowing for more possible apparent delays. For example, an event at Z at $t = 100$ causes a response at Y at $t = 134$. It is possible that another event occurs at X at $t = 110$ which gets a response from Y at $t = 123$. This results in apparent delays $Z \rightarrow Y$ of 23 and 34 as well as delays $X \rightarrow Y$ of 3 and 13. Formally, this is described by the following relations; $\hat{\delta}_{ZY} = \delta_{ZX} + \delta_{XY}$, $\delta_{ZX} - \tau + \delta_{XY}$ and $\hat{\delta}_{XY} = \delta_{XY}$, $\delta_{XY} - \tau$ where τ is the time between independent events at Z and X . τ can be any value because the events are not synchronized, thereby accounting for the increased number of non-zero TE values.

The conditional results show that $\text{TE}_{X \rightarrow Y|Z}$ is slightly reduced by conditioning on Z at $\zeta = 34$ (Fig. 6.8a), matching the relayed pathway. Comparing this system to the previous one, the results are qualitatively similar making the determination of whether X generates information almost impossible unless something is known about the system. Consequently, it is not often possible to rule out direct communication in the latter stages of a relayed path, preventing the researcher from arguing that the existence of the relay $Z \rightarrow X \rightarrow Y$, precludes the existence of a separate, direct pathway $X \rightarrow Y$. This is why it is proposed to define relays as separate entities as shown in Fig. 6.9.

The two relay systems Fig. 6.3 and Fig. 6.6 are summarised together here:

- The component $Z \rightarrow X$ is determined to be significant and independent of Y . The delay of that path is 21.
- The component $X \rightarrow Y$ is determined to be significant with a delay of 13, identified by a peak in $\text{TE}_{X \rightarrow Y}$ and confirmed by significant interaction in $\text{TE}_{Z \rightarrow Y|X}$ at $\zeta = 13$. This could be nullified at the discretion of the researcher using knowledge of the system and contingent on their definition of an interaction. However, here that is not possible.
- The component $Z \rightarrow Y$ is determined to be insignificant² because there are values of ζ which cause $\text{TE}_{Z \rightarrow Y|X}$ to be insignificant. The reason significant TE is measured by the calculation of $\text{TE}_{Z \rightarrow Y}$ is the existence of a relay path. The determining factor is that the dominant delay in $\text{TE}_{Z \rightarrow Y}$ is equal to $\hat{\delta}_{ZX} + \hat{\delta}_{XY}$.

²A caveat to this is included in the next simulation.




Figure 6.6: Communication pathways for a system where information generated at Z is relayed to Y by X .




Figure 6.7: Pairwise TE results. For clarity, components with no measurable TE are not shown.




Figure 6.8: Conditional TE results for simulated system with three variables with information generation at Z and X .




Figure 6.9: Identified communication pathways for the system described in Fig. 6.6

6.1.4 Multiple Pathways

This simulation combines all of the pathways from the previous two simulations. X and Y are mutually driven by Z and information is relayed from Z to Y via X as shown by Fig. 6.10. The pairwise results (Fig. 6.11) are a combination of the mutual driver and relay results. Two peaks are seen in $\text{TE}_{Z \rightarrow Y}$, one at $\delta = 34$ caused by the relay and one at $\delta = 52$ caused by the direct interaction. Similarly, two peaks are seen in $\text{TE}_{X \rightarrow Y}$, one at $\delta = 13$ caused by the direct interaction and the other at $\delta = \hat{\delta}_{ZY} - \hat{\delta}_{ZX} = 31$ caused by the mutual driver effect. Using the results from the previous simulations, a structure for beginning the analysis is worked through here. The goal is to identify the underlying cause of each of the significant delays.

1. Identify significant delays from pairwise TE

As before the significant delays are the delays which cause pairwise TE to peak. The results are give in Table 6.1. The components with no significant delays are declared insignificant.




Figure 6.10: Communication pathways for a system where information generated at Z is relayed to Y by X and also directly.

Table 6.1: Significant delays identified from pairwise TE

| Component | Delays ($\hat{\delta}$) |
|-------------------|---------------------------|
| $Z \rightarrow X$ | 21 |
| $X \rightarrow Z$ | - |
| $Z \rightarrow Y$ | 34,52 |
| $Y \rightarrow Z$ | - |
| $X \rightarrow Y$ | 13,31 |
| $Y \rightarrow X$ | - |




Figure 6.11: Pairwise TE results. For clarity, components with no measurable TE are not shown.

2. Identify components which are dependent upon the conditioning variable

Components which are dependent upon the conditioning variable are identified by variations in conditional TE with respect to the conditioning lag ζ . $TE_{Z \rightarrow X|Y}$ is not shown here but it is not dependent upon the conditioning variable, so that interaction is considered resolved and significant with a delay of 21. Fig. 6.12a and Fig. 6.12b show that the remaining two components are influenced by the conditioning variable at conditioning lags $\zeta_{ZY} = 34, 52$ and $\zeta_{XY} = 13, 31$. The interactions are made insignificant at $\zeta_{ZY} = 52$ and $\zeta_{XY} = 13$, respectively.

At this stage it is possible to construct a simple map of the system. Starting with a map of the delays identified in Table 6.1, as shown in Fig. 6.13. Then remove those pathways which can be made insignificant by conditioning at any ζ value, shown as red arrows in Fig. 6.13. This is seen to match the simulated system but offers minimal information about the dynamics. Additional information is obtained by the following two steps.




Figure 6.12: Conditional TE results for simulation of three variables with both mutual driver and cascade effects.




Figure 6.13: Identified communication pathways for the system described by Fig. 6.10. Red pathways are determined to be insignificant.

3. Identify delays related to mutual driver and cascading effects for unresolved components

The delays associated with the mutual influence of the conditioning variable and those associated with information relayed through the conditioning variable, for the two unresolved components are given in Table 6.2. Note that negative delays are not allowed. More complex pathways could be added but those listed are sufficient for this discussion.

Table 6.2: Delays related to indirect communication

| Component | Mutual Drive Delays | | Relay Delays | |
|-------------------|---|--------|---|--------|
| $Z \rightarrow Y$ | $\hat{\delta}_{XY} - \hat{\delta}_{XZ}$ | - | $\hat{\delta}_{ZX} + \hat{\delta}_{XY}$ | 34, 52 |
| $X \rightarrow Y$ | $\hat{\delta}_{ZY} - \hat{\delta}_{ZX}$ | 13, 31 | $\hat{\delta}_{XZ} + \hat{\delta}_{ZY}$ | - |

4. Relate significant delays to their cause

From this point forward the analysis resembles a puzzle. The goal is to identify the cause of each of the dominant delays. This simulation is relatively simple and has no noise which makes this solvable, but real systems can quickly become unclear. This can be offset by some knowledge of the system, that is an approach which will be avoided here.

It has already been determined that $\hat{\delta}_{ZX} = 21$ is caused by a real interaction from Z to X . Comparing Table 6.1 and Table 6.2, the following can be determined:

- X is not mutually driving Y and Z .
- There is no significant relay $X \rightarrow Z \rightarrow Y$.

The analysis so far has narrowed the possibilities down to two systems; $\delta_{ZX} = 21, \delta_{ZY} = 52, \delta_{XY} = 13$ or $\delta_{ZX} = 21, \delta_{ZY} = 34, \delta_{XY} = 31$.

Focusing on $X \rightarrow Y$, it is seen in Fig. 6.12a this is insignificant at $\delta = 31, \zeta = 52$, suggesting that this is a false result. If $\hat{\delta}_{XY} = 31$ is false, then it must be caused by something else. From Table 6.2, that must mean it is cause by the mutual influence of Z . It also implicates $\hat{\delta}_{ZY} = 52$ as a real delay because the relay $Z \rightarrow X \rightarrow Y$ cannot have a delay of 52 unless $\hat{\delta}_{XY} = 31$.

A similar argument can be made from $TE_{Z \rightarrow Y|X}$ (Fig. 6.12b) which is insignificant at $\delta = 34, \zeta = 13$, suggesting $\hat{\delta}_{ZY} = 34$ is false. This implies that $\hat{\delta}_{ZY} = 34$ is most likely caused by the relay $Z \rightarrow X \rightarrow Y$, leaving $\hat{\delta}_{XY} = 13$ as the real delay.

From this analysis, a map is drawn in Fig. 6.14




Figure 6.14: Identified communication pathways for the system described by Fig. 6.10 using lag analysis.

Admittedly, for this system, not a lot of additional knowledge is gained in the last two steps of this analysis. Ideally, the direct path $X \rightarrow Y$ would have been reasoned away but this would require prior knowledge of the system. It may appear that the relay path exists whenever there is communication along both $Z \rightarrow X$ and $X \rightarrow Y$ given that the sum of the delays is the delay of the relay path, but that is not the case. The significance of the path is determined by a measurement of TE from Z to Y so it is information generated at Z that eventually reaches Y . This becomes an important distinction if X significantly modifies the information before relaying it or it selectively relays information; i.e, if X passes information it creates to Y but does not relay information it receives from Z .

The final section of the analysis should not be employed without considerable thought. There are a number of complications which can be hidden. For example, if the system delays are changed so the real delay from $Z \rightarrow Y$ matches the relay delay, $\delta_{ZY} = \delta_{ZXY}$, conditional TE still goes to zero at $\zeta = \delta_{XY}$. This hides the real interaction making it look like the system in Fig. 6.3. For this analysis to be useful, some knowledge of the system is usually required.

6.1.5 Simulation Conclusions

Mutual driver and cascading pathways were shown to lead to spurious measurements of TE. By conditioning on all other variables in the system, it was possible to identify these spurious results. If an interaction could be conditioned out, at any value of ζ , that interaction is not caused by direct communication from the source to the target. Using that technique, a basic causal map of the system can be drawn. Additional knowledge can be gained by analysing the peak delays. It has been demonstrated that if Z is a mutual driver of X and Y , the following conditions are true;

1. $\text{TE}_{X \rightarrow Y}(\delta = \hat{\delta}_{ZY} - \hat{\delta}_{ZX}) \neq 0$
2. $\text{TE}_{X \rightarrow Y|Z}(\delta = \hat{\delta}_{ZY} - \hat{\delta}_{ZX}, \zeta = \hat{\delta}_{ZY}) = 0$

where $\hat{\delta}_{ZY} > \hat{\delta}_{ZX}$. However, if the communication $X \rightarrow Y$ is relayed through Z , the following are true;

1. $\text{TE}_{X \rightarrow Y}(\delta = \hat{\delta}_{XZ} + \hat{\delta}_{ZY}) \neq 0$
2. $\text{TE}_{X \rightarrow Y|Z}(\delta = \hat{\delta}_{XZ} - \hat{\delta}_{ZY}, \zeta = \hat{\delta}_{ZY}) = 0$

There are some exclusions to this discussed in the previous section.

6.2 Experiments

The simulations in the preceding section set out expectations for the various communication mechanisms which are likely to occur in the physical experiments containing two passive airfoils. Many configurations of the airfoils were trialled to investigate the effectiveness of conditional TE in identifying the correct communication pathways. Obtaining statistically meaningful results from this experiment proved to be difficult so the following example is included to demonstrate the concepts.

6.2.1 Preliminary Analysis

The locations of the two airfoils in this experiment are given in Fig. 6.15. The expectation is that the communication from C to A will be very strong as it is similar to case 1 in Section 5.3.2. With airfoil B being so close behind A, it is expected that that communication link will also be very strong. The cross-stream offset of B makes it possible for direct communication from C to B but it is possible that A will intercept all of the information from C . It is known that it is not possible for either airfoil to affect the cylinder. The decision

period for this experiment is $T_D = 2.0$ s. Based on the results in Chapter 5, the sampling period is selected as $\Delta = 2.02$ s, which is the value closest to T_D which does not cause an erroneous measurement of entropy and symbolization is done before down-sampling, using Eq. 5.1 and Eq. 5.2.




Figure 6.15: Relative placements of structures. Drawn to scale 1:2.

Pairwise TE

The results of the pairwise calculations of TE are given in Fig. 6.16 and the feedback results are shown in Fig. 6.17. They confirm that the components $A \rightarrow C$ and $B \rightarrow C$ are never significant. The dominant delays, $\hat{\delta}$, can be selected from the local peaks of TE and transcribed onto the schematic Fig. 6.18. It is known that the feedback $C \rightarrow C$ is caused by the regularity of the decision process that drives the cylinder, so this can be ignored. Similarly, high TE for $\delta \rightarrow 0$ are not counted as they are caused by the finite time of the response, not communicated information.




Figure 6.16: Pairwise TE results for Test T14




Figure 6.17: Feedback results for Test T14




Figure 6.18: Dominant communication lags identified by pairwise TE, for Test T14.

6.2.2 Conditional TE

Plotting conditional TE over a range of δ and ζ lag values can aid in identifying the true communication paths of the system by highlighting trends or patterns in the results. When the conditional result is significantly different from the pairwise result, it can be said that the communication from the source to the target has some dependence on the conditioning variable. The exact nature of that interaction is not always clear so it takes discretion by the researcher to interpret the results. To create the most basic causality map, the only information needed from conditional TE, Figs. 6.20a to 6.21d, is if there are any $\hat{\delta}$ which are made insignificant at any ζ value. This is the case for both $\hat{\delta}_{CB} = 6.6, 9.6$ s, evidenced by insignificant values at $\delta_{CB} = 6.6, 9.6, \zeta = 3.4$ s.³ Therefore, that path is determined to be insignificant at those delays; they must have a cause other than direct communication from C to B . Similarly, the interactions $\hat{\delta}_{AB} = 0.8, 6.4$ s are shown in Fig. 6.21d to be artefacts of another mechanism related to the self feedback of A .⁴ Also, the interaction described by $\hat{\delta}_{CA} = 11$ s is never significant when conditioned upon A , indicating that interaction is part of a longer pathway.⁵ The results from this analysis, presented in Fig. 6.19, give a much clearer picture of the communication pathways and their associated delays, compared to the pairwise results in Fig. 6.18.

A deeper understanding of the system can be gained by identifying the cause of the differences between the pairwise and conditional results as was done in Section 6.1. The analysis is convoluted for this system, so the important results are summarised in Table 6.3;

³This also implies 3.4 s is the real delay from A to B .

⁴Again the conditioning delay suggests the real delay from $A \rightarrow B$ is 3.4 s.

⁵One must be careful when conditioning on a variable that is also either the source or target. $TE_{A \rightarrow B|A}$ has exactly zero measured TE when $\delta = \zeta$ because it is conditioning out all of the source information. Similarly, $TE_{C \rightarrow A|A}$ has zero measured TE at $\zeta = 0$ because it is conditioning out the target's response.

Table 6.3: Summary of the pathway analysis in Appendix E.

| Component | Delay (s) | Cause |
|-------------------|-----------|--|
| $C \rightarrow A$ | 3.2 | Real interaction. Consistent with results from Chapter 5 |
| $C \rightarrow A$ | 6.2 | The interaction $C \rightarrow A$ plus the self feedback of A |
| $C \rightarrow A$ | 9.6 | A long relay $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$, with the $B \rightarrow A$ delay 0.0s |
| $C \rightarrow A$ | 11.0 | As above but with the longer $B \rightarrow A$ delay |
| $C \rightarrow B$ | 6.6 | Simple relay from C to B via A |
| $C \rightarrow B$ | 9.6 | As above plus the self feedback of A |
| $A \rightarrow B$ | 0.8 | Unknown. Significance is very low so potentially spurious. |
| $A \rightarrow B$ | 3.4 | Real interaction. Consistent with previously identified advection speed |
| $A \rightarrow B$ | 6.4 | The self feedback of A plus the real interaction $A \rightarrow B$ |
| $B \rightarrow A$ | 0.0 | Real interaction. Consistent with pressure communication |
| $B \rightarrow A$ | 4.3 | Unknown. Probably a real interaction but the mechanism is unknown. |

the full analysis is in Appendix E. The summary explains the cause of all of the delays identified by the pairwise TE calculations. One of the most interesting results is that the interaction $C \rightarrow B$ was caused by a relay through A , evidenced by the matching of delays $\hat{\delta}_{CB} = \hat{\delta}_{CA} + \hat{\delta}_{AB}$. Another interesting result was the identification of a long path, $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$. The existence of this path was unexpected but it does explain the longer measured delays from C to A . It also has much lower significance than the direct $C \rightarrow A$ interaction which is expected.

Depending upon the desired information, the conditional TE results (Fig. 6.19) may be a sufficient. However, there are several ways to interpret the results which the researcher should be mindful of. For example, the interaction $C \rightarrow B$ was determined to be entirely dependent upon A , thereby labelling it a false path. However, information generated at C was being communicated to B , even if it was relayed by A . If the research question is, ‘is information from C reaching B ,’ the answer is yes and the pairwise calculation should be employed. If the question is, ‘is there a direct interaction $C \rightarrow B$,’ then the answer is no and the conditional calculation is correct.




Figure 6.19: Causality map determined by detecting statistically significant values of conditional TE at some δ and all ζ values.




Figure 6.20: Conditional TE results for three structure experiment. a; $TE_{C \rightarrow A|B}$, b; $TE_{C \rightarrow B|A}$.

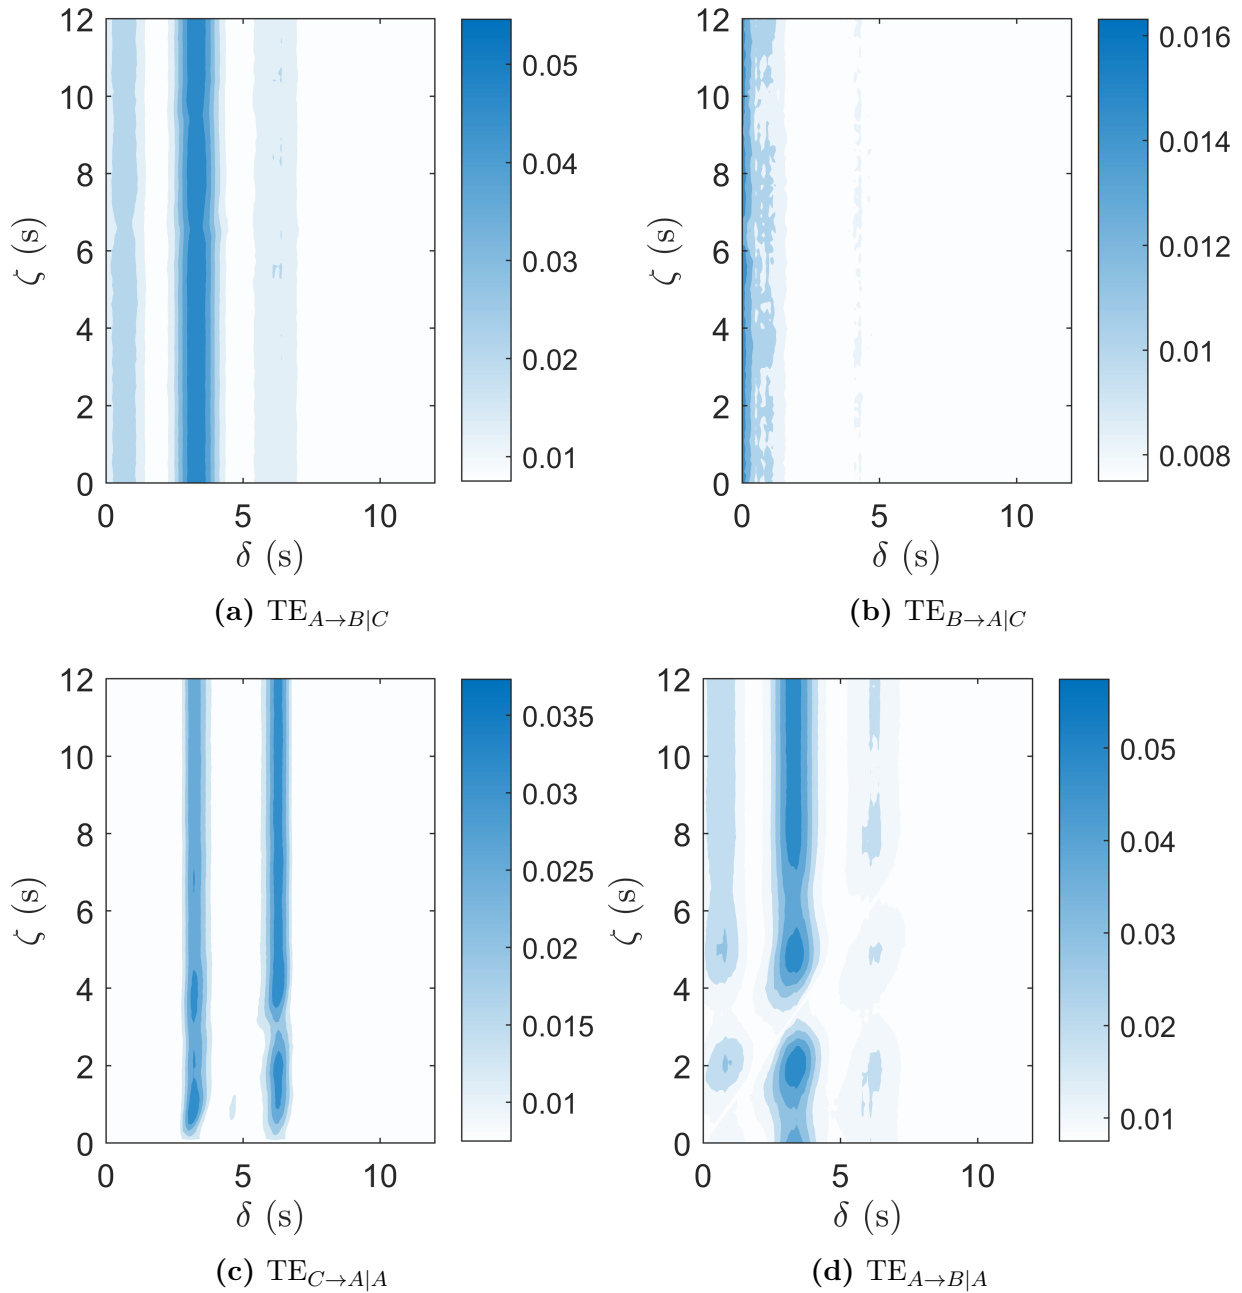


Figure 6.21: Conditional TE results for three structure experiment. a; $TE_{A \rightarrow B|C}$, b; $TE_{B \rightarrow A|C}$, c; $TE_{C \rightarrow A|A}$, d; $TE_{A \rightarrow B|A}$.

6.2.3 Path Significance

In the previous section, the various pathways were determined. A path was considered significant if all of its components were significant. For example, the path $C \rightarrow A \rightarrow B$

is considered significant at a delay of 6.6 s because $C \rightarrow A$ is significant at a delay of 3.2 s and $A \rightarrow B$ is significant at 3.4 s. To compare the relative significance of the different paths it must be possible to measure the TE along that path. The method proposed here is to calculate TE from the first variable (the source) to the last variable (the target) at the delay of the full path. The TE of the path $C \rightarrow A \rightarrow B$ is then $\text{TE}_{C \rightarrow B}(\delta = \hat{\delta}_{CAB} = 6.6 \text{ s}) = 0.0075 \text{ bits}$.⁶ All of the significant pathways, summarized in Fig. E.1, are measured in this way and the results are tabulated in Table 6.4. Interestingly, the result is not always the magnitude of the least significant component as may be expected. The component $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$ is more significant than the interaction $B \rightarrow A$.

Table 6.4: Summary of the delays and TE for the paths in Fig. E.1. Null threshold for all paths is 0.0051 bits.

| Path | Delay | TE |
|---|-------|--------|
| $A \rightarrow A$ | 3.0 | 0.0597 |
| $C \rightarrow A$ | 3.2 | 0.0279 |
| $C \rightarrow A \rightarrow A$ | 6.2 | 0.0313 |
| $A \rightarrow B$ | 3.4 | 0.0501 |
| $A \rightarrow A \rightarrow B$ | 6.4 | 0.0176 |
| $C \rightarrow A \rightarrow B$ | 6.6 | 0.0075 |
| $C \rightarrow A \rightarrow A \rightarrow B$ | 9.6 | 0.0075 |
| $C \rightarrow A \rightarrow B \rightarrow A$ | 11.0 | 0.0073 |
| $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$ | 9.6 | 0.0073 |

6.2.4 Hypothesis Testing Approach

The detection of pressure communication $B \rightarrow A$ raises the question of whether it should also be detected in the other direction. Using the method of selecting the local peaks in TE to get the dominant delays, outlined by Wibral *et al.* [66], does not detect pressure communication $A \rightarrow B$. However, the hypothesis ‘there is pressure communication $A \rightarrow B$,’ cannot be disproved because significant $\text{TE}_{A \rightarrow B}$ is measured at $\delta_{AB} = 0 \text{ s}$. Similarly, the logic used to identify the pathways in Fig. E.1 found that there should be a path $C \rightarrow A \rightarrow B \rightarrow A$ with a delay of 6.6 s. Measuring $\text{TE}_{C \rightarrow A}(\delta = 6.6 \text{ s})$ gives a result of 0.0206 bits which is very significant. It is hypothesised that in both of these cases, the

⁶The null threshold is also taken from the pairwise calculation

proximity of the expected result to another, more dominant delay⁷ has hidden these effects. The question then is ‘what else is being hidden by the wide peaks in TE?’ It is recommended that the approach is guided by the research question, as has been the common theme throughout this work. A lot was learned by analysing the delays as suggested by Wibral *et al.* [66] but this method neglects much of the information in the results. Sometimes it may be more appropriate to do specific hypothesis testing.

6.2.5 Conclusions

The applicability of transfer entropy to fluid-structure problems with more than two bodies has been explored. It has been shown that using pairwise transfer entropy can give a false interpretation of the causal relationships that exist in such systems. By conditioning the calculation on the third structure it was possible to identify interdependent relationships such as mutual driver effects and relayed pathways. It was determined that a path identified by significant pairwise TE was not a direct link if it could be made insignificant by conditioning at any conditioning delay, ζ . In the simplest sense, if conditional TE was significant at any δ and all ζ , a causal relationship can be said to exist from the source to the target; assuming all dependent variables are conditioned out.

Identifying the true nature of an interaction could, in part, be achieved by comparing all of the critical delays, $\hat{\delta}$. Through simulations, it was shown that mutual influence caused an apparent interaction at $\hat{\delta}_{XY} = \hat{\delta}_{ZY} - \hat{\delta}_{ZX}$. Similarly, relay paths caused an apparent interaction at $\hat{\delta}_{XY} = \hat{\delta}_{XZ} + \hat{\delta}_{ZY}$. In the physical experiment, where many pathways were possible, this proved difficult to employ and required significant *a priori* knowledge of the system.

Many nuances specific to applying TE to systems with more than three structures were discovered. The common theme was the importance of having a clearly defined research question and following an appropriate path. To some, the question of where information is generated may be important, while others may want to know if an communication occurs between two structures regardless of the origin of the information.

Wibral *et al.* [66] suggests using the maximum of TE to identify the critical communication delays. This was demonstrated to be a useful tool but it also likely masks other interactions if the delays are too similar. In some cases hypothesis testing may yield very

⁷The delay related to pressure communication ($\delta_{AB} = 0$) is dominated by the unexplained interaction with delay 0.8s. The path $C \rightarrow A \rightarrow B \rightarrow A$ with delay 6.6s is dominated by the interaction $C \rightarrow A \rightarrow A$ at delay 6.2s.

different results. Again, the researcher is advised to be rigorous in defining the research question. Further work on this masking effect is likely to be useful.

Overall, conditional transfer entropy has proven a useful tool for fluid-structure problems with more than two bodies. More experiments of different fluid-structure arrangements are likely to yield interesting results. The increased computational intensity that comes with more degrees of freedom may prove prohibitive for large arrays. To this end, methods more efficient than brute force should be investigated in the near future; perhaps before increasing the number of structures.

Chapter 7

Conclusions and Recommendations

The applicability of the information theoretic measure, transfer entropy, to fluid-structure systems has been investigated through a combination of experiments and simulations. It has been shown that treating structure displacement as information events sufficiently captures the dynamics of the particular systems in this study. The intended purpose of TE is to identify directed causal relationships. Not only was it possible to detect these relationships, but changes in the relative interaction strength were identified and significant knowledge of the system dynamics was gained. However, there are many nuances to the application of TE and to the interpretation of results.

Analysis of analytical solutions highlighted the difference between true causality and Granger causality, to which TE is analogous. This demonstrated the need for some degree of randomness in the system for TE to be applicable. This could either be in the behaviour of the source variable or by having imperfect coupling from the source to the target. Comparing simulations to analytical results confirmed that correct selection of the sampling period and lag parameter is critical. It was demonstrated that the sampling period is related to the time between events in the system and the two should be approximately equal. However, at some sampling periods the data could be misrepresented due to a harmonic relationship between the sampling and the event timing. This error presented as spurious measurements of Shannon entropy, providing a cheap method for filtering out erroneous sampling periods. The lag parameter was shown to be equivalent with the communication delay from the source to the target variables. Two methods for selecting the correct lag were discussed throughout the simulations and experiments; the correct method was dependent upon the type of data.

The first set of experiments was performed with two bodies; one randomly, rotation-

ally impulsed cylinder and one passively pivoted airfoil. The separation of the structures and the period between potential cylinder rotations (decision period) were varied to give 15 unique cases. Various methods were compared due to inconsistencies or gaps in the published literature. These comparisons demonstrated the importance of matching the methodology to the type of data. For example, data from similar experiments conducted by Zhang *et al.* [74] could be calculated on by down-sampling the time-series before symbolizing. However, this caused errors in results from the cylinder data in this experiment, so the order had to be reversed.

The two body experiment showed that the sampling period should approximately equal the decision period, similar to systems driven by Markov processes. Again, the down-sampling process was shown to have the potential to misrepresent the data due to the harmonic relationship with the event period, but this was also able to be filtered out using the measured entropy. Once the appropriate sampling period was selected, the existence of a directed causal relationship could be inferred if any measured $TE(\delta)$ value was greater than the null threshold. The null threshold was identified by bootstrapping the data, destroying the temporal relationship between the variables. Additional knowledge of the system could be gained by analysing the significant communication delays. Communication delays were identified as the lag parameters which maximised $|TE_{net}|$. The validity of the results obtained through TE were confirmed through comparison with *a priori* knowledge of the system. For example, it was known that the airfoil could not affect the cylinder, the cylinder to airfoil interaction should weaken as the separation increases and the communication delay should increase.

For the first time, TE was applied to fluid-structure systems with more than two bodies. This increased the complexity of the possible pathways. Not only are there $n(n - 1)$ direct pathways but there is also the potential for mutual driver effect and relayed pathways. This experiment proved to be less clear than anticipated due to the difficulty in isolating the structures. As a result, the interpretation of the results was very manual and time consuming; only one of more than twenty experiments was fully presented. Still, significant knowledge of the system was obtained, including some unexpected but feasible results. Ultimately, while more research is needed, TE appears to be a highly useful tool for the analysis of fluid-structure systems and, with care, is relatively easy to apply.

Throughout this work, a common theme has been the importance of clearly defining the research question as this affects the method of interpreting results. Identification of directed causal links is, perhaps, the simplest result from this tool. If any measured

value is statistically significant, and cannot be conditioned out, then such a link exists. It is however, possible to obtain a deeper understanding of the system. An exploratory analysis can be performed by calculating TE over a range of parameters and analysing the trends in the results. By comparing the various significant communication delays, it is possible to gain an understanding of the communication mechanisms that drive the system. Alternately, hypothesis testing at specific delays can confirm or rule out the existence of some phenomenon of interest. As is common in statistical methods, using TE it is easier to disprove a hypothesis than to prove it. In all cases, understanding what behaviour of the system is of interest is essential for creating an appropriate symbolization algorithm.

This is a relatively new field of study so there are many opportunities for development. Further experimentation on different fluid-structure systems will help identify which systems are most suitable for this type of analysis. From a theoretical understanding of TE, the most likely candidates are chaotic systems where traditional fluid-structure analysis techniques break down. The usability of this tool would be greatly expanded if direct fluid measurements could be used as a variable.

Chapter 8

Key Ideas

The following is a summary of the key points of this thesis. They are listed in order of occurrence in the main text.

1. Fluid-structure interactions are common in many engineering and science fields. However, they tend to be non-linear, multi-physics problems so they are not easily resolved. **Page 1**
2. Information content is a contextual measure based on the probability of an event occurring. **Page 9**
3. Entropy is a measure of the uncertainty about the state of a variable. **Page 11**
4. Transfer entropy is the difference between the uncertainty about the state of the target variable given knowledge about its past and the uncertainty about the state of the target variable given knowledge about its past and about the state of the source variable. **Page 14**
5. The communication time of the system is directly related to the lag parameter δ and can be determined from the local maxima of $TE(\delta)$. A great deal of care needs to be taken when interpreting the results when the mode of communication is unknown because different systems can have similar spectra. **Page 16**
6. The sampling period should be scaled by the event period of the system to minimise detectability errors and sample fidelity errors. **Page 20**
7. Symbolisation reduces the cardinality of a variable by recasting the data as predefined behaviours. This takes care by the researcher to ensure the critical information is retained. **Page 22**
8. Conditional transfer entropy can be used to identify dependence of an interaction on additional variables and can condition out mutual driver or cascading effects. **Page**

9. Transfer entropy is equivalent with Granger causality. **Page 27**
10. The TE equation can be recast as a collection of joint probabilities. These probabilities are calculated by counting the number of times the described event occurs. **Page 27**
11. A simulation is used to validate the probabilistic form of the TE equation by comparing it to an analytical result. **Page 59**
12. Granger causality, and therefore TE, is undetectable if there is no randomness in the input data. **Page 60**
13. The sampling period, Δ , has a harmonic relationship with the event period which can cause the data to be poorly represented. These error causing sampling periods can be filtered out because they cause sharp changes in the calculated entropy. **Page 61**
14. Experiments place a randomly impulsed cylinder in a uniform fluid flow, upstream of a passively pivoted airfoil. The impulses of the cylinder generate information which is expected to be transmitted to the airfoil in the form of advecting vortices; this is expected to incite a rotational response in the airfoil. The structural displacements will be used to generate the information data. **Page 71**
15. The trend of TE with respect to the Δ and δ parameters demonstrate their relationship to the event time and the communication delay between structures, respectively. Analysis of these trends can yield valuable knowledge of the system dynamics. **Page 82**
16. Simulations and experiments with three bodies show that many more communication pathways exist, some of which can cause spurious TE results if the pairwise form of the equation is used. Additional variables must be accounted for using conditional TE. **Page 95**
17. Mutual driving of X and Y by Z results in a false positive detection of $TE_{X \rightarrow Y}$ at $\delta = \hat{\delta}_{ZY} - \hat{\delta}_{ZX}$. This can be conditioned out by evaluating $TE_{X \rightarrow Y|Z}$ at $\zeta = \hat{\delta}_{ZY}$. **Page 96**
18. Relayed communication along the path $X \rightarrow Z \rightarrow Y$ results in a false positive detection of $TE_{X \rightarrow Y}$ at $\delta = \hat{\delta}_{XZ} + \hat{\delta}_{ZY}$. This can be conditioned out by evaluating $TE_{X \rightarrow Y|Z}$ at $\zeta = \hat{\delta}_{ZY}$ **Page 100**
19. A causality network can be drawn by using the condition, ‘a directed causal relationship exists if conditional TE is significant at any δ , which cannot be made insignificant by conditioning.’ Assuming all variables are included in the conditioning. **Page 114**

20. Significant knowledge of the dynamics of systems with three structures can be obtained by analysing the trends of conditional TE with respect to δ and ζ . This process is very manual and requires some *a priori* knowledge of the system. **Page 115**

References

- [1] ANDERSON JR, J. D. 2001 *Fundamentals of aerodynamics*. 3rd ed. Tata McGraw-Hill Education. ↩
- [2] ANDERSON, R. P., JIMENEZ, G., BAE, J. Y., SILVER, D., MACINKO, J., & PORFIRI, M. 2016 Understanding policy diffusion in the us: an information-theoretical approach to unveil connectivity structures in slowly evolving complex systems. *SIAM journal on applied dynamical systems* **15** (3), 1384–1409. ↩
- [3] ASTE, T. & DI MATTEO, T. 2017 Sparse causality network retrieval from short time series. *Complexity* **2017**. ↩
- [4] BARNETT, L., BARRETT, A. B., & SETH, A. K. 2009 Granger causality and transfer entropy are equivalent for gaussian variables. *Physical review letters* **103** (23), 238701. ↩
- [5] BARNETT, L. & SETH, A. K. 2017 Detectability of granger causality for subsampled continuous-time neurophysiological processes. *Journal of neuroscience methods* **275**, 93–121. ↩
- [6] BETZ, A. 1933 Behavior of vortex systems. *TECHNICAL MEMORANDUM No. 713* ↩
- [7] BLOOR, M. S. 1964 The transition to turbulence in the wake of a circular cylinder. *Journal of Fluid Mechanics* **19** (2), 290–304. ↩
- [8] BOSSOMAIER, T., BARNETT, L., HARRÉ, M., & LIZIER, J. T. 2016 *An introduction to transfer entropy*. Springer. ↩
- [9] BUTAIL, S., MWAFFO, V., & PORFIRI, M. 2016 Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Physical Review E* **93** (4), 042411. ↩

- [10] CHEW, Y., CHENG, M., & LUO, S. 1995 A numerical study of flow past a rotating circular cylinder using a hybrid vortex scheme. *Journal of fluid mechanics* **299**, 35–71. ↩
- [11] CHICHARRO, D. & LEDBERG, A. 2012 When two become one: the limits of causality analysis of brain dynamics. *PLoS One* **7** (3), e32466. ↩
- [12] COUTANCEAU, M. & MENARD, C. 1985 Influence of rotation on the near-wake development behind an impulsively started circular cylinder. *Journal of Fluid Mechanics* **158**, 399–446. ↩
- [13] COVER, T. M. & THOMAS, J. A. 2012 *Elements of information theory*. John Wiley & Sons. ↩
- [14] CRAMER, D. 1997 *Basic Statistics for Social Research*. Taylor & Fransis. ↩
- [15] DANAHER. 2007 CT series step motors. ↩
- [16] DEKKING, F. M., KRAAIKAMP, C., LOPUHAÄ, H. P., & MEESTER, L. E. 2005 *A Modern Introduction to Probability and Statistics: Understanding why and how*. Springer Science & Business Media. ↩
- [17] DIXON, P. M. 2006 Bootstrap resampling. *Encyclopedia of environmetrics* **1**. ↩
- [18] FALNES, J. 2002 *Ocean waves and oscillating systems: linear interactions including wave-energy extraction*. Cambridge university press. ↩
- [19] FEY, U., KÖNIG, M., & ECKELMANN, H. 1998 A new strouhal–reynolds-number relationship for the circular cylinder in the range $47 < re < 200000$. *Physics of Fluids* **10** (7), 1547–1549. ↩
- [20] FILLER, J., MARSTON, P., & MIH, W. 1991 Response of the shear layers separating from a circular cylinder to small-amplitude rotational oscillations. *Journal of Fluid Mechanics* **231**, 481–499. ↩
- [21] GLAUERT, M. 1957 The flow past a rapidly rotating circular cylinder. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **242** (1228), 108–115. ↩
- [22] GRANGER, C. W. J. 2003 Prize lecture: time series analysis, cointegration, and applications. <https://www.nobelprize.org/prizes/economic-sciences/2003/granger/lecture/>. [Accessed: March 2019]. ↩

- [23] GRANGER, C. W. 1969 Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society* 424–438. ↩
- [24] HIPEL, K. W. & MCLEOD, A. I. 1994 *Time series modelling of water resources and environmental systems*, vol. 45. Elsevier. ↩
- [25] HOU, G., WANG, J., & LAYTON, A. 2012 Numerical methods for fluid-structure interaction — a review. *Communications in Computational Physics* **12** (2), 337–377. ↩
- [26] HOWE, M. 1996 Energy conservation and the damping of flexural waves by vorticity production. *Journal of sound and vibration* **190** (1), 1–19. ↩
- [27] JEON, D. & GHARIB, M. 2004 On the relationship between the vortex formation process and cylinder wake vortex patterns. *Journal of Fluid Mechanics* **519**, 161–181. ↩
- [28] JOHNSON JR, P. D., HARRIS, G. A., & HANKERSON, D. 2003 *Introduction to information theory and data compression*. Chapman and Hall/CRC. ↩
- [29] KAISER, A. & SCHREIBER, T. 2002 Information transfer in continuous processes. *Physica D: Nonlinear Phenomena* **166** (1-2), 43–62. ↩
- [30] LIZIER, J. T., PROKOPENKO, M., & ZOMAYA, A. Y. 2008 Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E* **77** (2), 026110. ↩
- [31] LJUNG, G. M. & BOX, G. E. 1978 On a measure of lack of fit in time series models. *Biometrika* **65** (2), 297–303. ↩
- [32] MARSCHINSKI, R. & KANTZ, H. 2002 Analysing the information flow between financial time series. *The European Physical Journal B-Condensed Matter and Complex Systems* **30** (2), 275–281. ↩
- [33] MEYERS, J. & MENEVEAU, C. 2012 Optimal turbine spacing in fully developed wind farm boundary layers. *Wind Energy* **15** (2), 305–317. ↩
- [34] MITTAL, S. & KUMAR, B. 2003 Flow past a rotating cylinder. *Journal of fluid mechanics* **476**, 303–334. ↩
- [35] MOORE, D. 1957 The flow past a rapidly rotating circular cylinder in a uniform stream. *Journal of Fluid Mechanics* **2** (6), 541–550. ↩
- [36] MORTON, B. 1984 The generation and decay of vorticity. *Geophysical & Astrophysical Fluid Dynamics* **28** (3-4), 277–308. ↩

- [37] NORBERG, C. 1994 An experimental investigation of the flow around a circular cylinder: influence of aspect ratio. *Journal of Fluid Mechanics* **258**, 287–316. ↩
- [38] NORBERG, C. 1987 Effects of reynolds number and a low-intensity freestream turbulence on the flow around a circular cylinder. *Chalmers University, Goteborg, Sweden, Technological Publications* **87** (2), 1–55. ↩
- [39] NORMAN, G. R. & STREINER, D. L. 2003 *PDQ statistics*. BC Decker Inc. ↩
- [40] OKAJIMA, A., TAKATA, H., & ASANUMA, T. 1975 Viscous flow around a rotationally oscillating circular cylinder. *NASA STI/Recon Technical Report N* **76**. ↩
- [41] ORANGE, N. & ABAID, N. 2015 A transfer entropy analysis of leader-follower interactions in flying bats. *The European Physical Journal Special Topics* **224** (17-18), 3279–3293. ↩
- [42] PALUŠ, M., KOMÁREK, V., HRNČÍŘ, Z., & ŠTĚRBOVÁ, K. 2001 Synchronization as adjustment of information rates: detection from bivariate time series. *Physical Review E* **63** (4), 046211. ↩
- [43] PALUŠ, M. & VEJMEĽKA, M. 2007 Directionality of coupling from bivariate time series: how to avoid false causalities and missed connections. *Physical Review E* **75** (5), 056211. ↩
- [44] PEARL, J. 2009 *Causality*. Cambridge university press. ↩
- [45] PORFIRI, M. 2018 Inferring causal relationships in zebrafish-robot interactions through transfer entropy: a small lure to catch a big fish. *Animal Behavior and Cognition* **5** (4), 341–367. ↩
- [46] ROCKWELL, D. 1998 Vortex-body interactions. *Annual review of fluid mechanics* **30** (1), 199–229. ↩
- [47] ROSHKO, A. 1993 Perspectives on bluff body aerodynamics. *Journal of Wind Engineering and Industrial Aerodynamics* **49** (1-3), 79–100. ↩
- [48] RUBINOV, J. T. L. ; M. 2012 Multivariate construction of effective computational networks from observational data. *Preprint*, Max Planck Institute for Mathematics in the Sciences. [URL](#). ↩
- [49] SCHAY, G. 2016 *Introduction to probability with statistical applications*. Birkhäuser. ↩
- [50] SCHREIBER, T. 2000 Measuring information transfer. *Physical review letters* **85** (2), 461. ↩

- [51] SHANNON, C. E. 1948 A mathematical theory of communication. *Bell System technical journal* **27**, 379–423, 623–656. [↔](#)
- [52] SOMMER, D. 2014 *Development of a coupled numerical-experimental facility to model the fluid-structure interactions of the human vocal folds.* [↔](#)
- [53] SPINNEY, R. E., PROKOPENKO, M., & LIZIER, J. T. 2017 Transfer entropy in continuous time, with applications to jump and neural spiking processes. *Physical Review E* **95** (3), 032319. [↔](#)
- [54] STANIEK, M. & LEHNERTZ, K. 2008 Symbolic transfer entropy. *Physical Review Letters* **100** (15), 158101. [↔](#)
- [55] STMICROELECTRONICS. 2011 dSPIN fully integrated microstepping motor driver with motion engine and SPI. Datasheet 16737 Rev 3. [↔](#)
- [56] STREITLIEN, K., TRIANTAFYLLOU, G. S., & TRIANTAFYLLOU, M. S. 1996 Efficient foil propulsion through vortex control. *Aiaa journal* **34** (11), 2315–2319. [↔](#)
- [57] SWANSON, W. 1961 The magnus effect: a summary of investigations to date. *Journal of Basic Engineering* **83** (3), 461–470. [↔](#)
- [58] TOKUMARU, P. & DIMOTAKIS, P. 1991 Rotary oscillation control of a cylinder wake. *Journal of Fluid Mechanics* **224**, 77–90. [↔](#)
- [59] TOMARU, T., MURAKAMI, H., NIIZATO, T., NISHIYAMA, Y., SONODA, K., MORIYAMA, T., & GUNJI, Y.-P. 2016 Information transfer in a swarm of soldier crabs. *Artificial Life and Robotics* **21** (2), 177–180. [↔](#)
- [60] TRIBUS, M. & MCIRVINE, E. C. 1971 Energy and information. *Scientific American* **225** (3), 179–190. [↔](#)
- [61] TSAY, R. S. 2005 *Analysis of financial time series*, vol. 543. John wiley & sons. [↔](#)
- [62] VAN DYKE, M. 1982 *An album of fluid motion.* Parabolic Press Stanford. [↔](#)
- [63] VERDES, P. 2005 Assessing causality from multivariate time series. *Physical Review E* **72** (2), 026222. [↔](#)
- [64] VICENTE, R., WIBRAL, M., LINDNER, M., & PIPA, G. 2011 Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of computational neuroscience* **30** (1), 45–67. [↔](#)
- [65] WALKER, J. 1978 The boundary layer due to rectilinear vortex. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **359** (1697), 167–188. [↔](#)

- [66] WIBRAL, M., PAMPU, N., PRIESEMAN, V., SIEBENHÜHNER, F., SEIWERT, H., LINDNER, M., LIZIER, J. T., & VICENTE, R. 2013 Measuring information-transfer delays. *PloS one* **8** (2), e55809. [↔](#)
- [67] WIBRAL, M., VICENTE, R., & LIZIER, J. T. 2014 *Directed information measures in neuroscience*. Springer. [↔](#)
- [68] WIDNALL, S. E. & WOLF, T. 1980 Effect of tip vortex structure on helicopter noise due to blade-vortex interaction. *Journal of Aircraft* **17** (10), 705–711. [↔](#)
- [69] WILLIAMSON, C. H. 1996 Vortex dynamics in the cylinder wake. *Annual review of fluid mechanics* **28** (1), 477–539. [↔](#)
- [70] WOOD, W. 1957 Boundary layers whose streamlines are closed. *Journal of Fluid Mechanics* **2** (1), 77–87. [↔](#)
- [71] WU, J., MO, J., & VAKILI, A. 1989 On the wake of a cylinder with rotational oscillations. In *2nd Shear Flow Conference*, p. 1024. [↔](#)
- [72] YU, Y., AMANDOLESE, X., FAN, C., & LIU, Y. 2018 Experimental study and modelling of unsteady aerodynamic forces and moment on flat plate in high amplitude pitch ramp motion. *Journal of Fluid Mechanics* **846**, 82–120. [↔](#)
- [73] ZDRAVKOVICH, M. 1981 Review and classification of various aerodynamic and hydrodynamic means for suppressing vortex shedding. *Journal of Wind Engineering and Industrial Aerodynamics* **7** (2), 145–189. [↔](#)
- [74] ZHANG, P., ROSEN, M., PETERSON, S. D., & PORFIRI, M. 2018 An information-theoretic approach to study fluid–structure interactions. *Journal of Fluid Mechanics* **848**, 968–986. [DOI](#). [↔](#)

Appendix A

Statistical Significance

The determination of whether two samples are significantly different is performed by a t-test. t-tests estimate the likelihood that the means of two samples is zero. For the cases presented here, the paired t-test is used because the systems are not independent; this is the same approach as [74]. A specific example is worked through here for clarity. For the case $s = 104$ mm, $T_D = 2$ s the experiment is repeated four times as tabulated in table A.1. To gain statistical relevance the number of repeated experiments should be higher but the time each experiment takes is prohibitive. One method for increasing the power of the test will be discussed below.

The null hypothesis (H_0) is that the difference of the mean result of $TE_{C \rightarrow A}$ is not significantly different to the mean result of $TE_{A \rightarrow C}$, the measured values are listed in table A.1. As explained in [39] the t-statistic is simply the ratio of the mean difference to the standard error (SE), the probability that the mean difference is zero is found by converting the t-statistic to a p-statistic using a standard table (discussed further below). The mean difference is the average of the difference between $TE_{C \rightarrow A}$ and $TE_{A \rightarrow C}$ in each test ($\mu_{\text{difference}} = 0.1463$). The SE is calculated from eqn. A.1 as the standard deviation of the differences ($SD_{\text{differences}} = 0.0327$) over the square root of the number of observations ($N=4$), giving a $SE = 0.0164$.

$$SE = \frac{SD}{\sqrt{N}} \quad (\text{A.1})$$

The t-statistic is calculated from eqn. A.2 giving a value of 8.9.

$$t = \frac{\mu_{\text{difference}}}{SE} \quad (\text{A.2})$$

The p-statistic, or probability of the null hypothesis being true, is found using a standard

t-table although it is more accurate to use a software package to interpolate the table. For this work, MATLAB is used for this step. Note that to find a p-statistic on the t-table two additional pieces of information are needed (as well as the t-statistic); the degree of freedom, which for all relevant cases is $N-1$, and the type of test (one-tail or two-tailed) [39]. A one-tailed distribution is used when the hypothesis is set up to include the direction of the difference in the mean (is x greater than y) but for these tests the hypothesis will be phrased to ask if the two values are different so the two-tailed test is employed. For this case the p-statistic is 0.0030, so there is a 0.3% chance of the two values being the same. As is common the threshold used here is 5% [39] [74] so the null hypothesis is rejected; there is a reasonable level of confidence that $TE_{C \rightarrow A}$ is greater than $TE_{A \rightarrow C}$.

Table A.1: Results from repeated experiments of TE. $s = 104$ mm, $T_D = 2.0$ s.

| | 1 | 2 | 3 | 4 | Mean |
|------------------------|--------|--------|--------|--------|--------|
| $TE_{C \rightarrow A}$ | 0.1421 | 0.1414 | 0.1137 | 0.1948 | 0.1480 |
| $TE_{A \rightarrow C}$ | 0.0010 | 0.0006 | 0.0013 | 0.0038 | 0.0017 |
| difference | 0.1411 | 0.1408 | 0.1124 | 0.1910 | 0.1463 |

A.0.1 Bootstrapping

The t-test relies on the data taking on a known distribution. The distribution of TE results cannot be known a priori so multiple measurements must be taken. The more measurements taken, the better the estimate of the distribution will be. However, with each test taking 20 minutes, the number of trials is limited. Bossomaier *et al.* [8] suggest a number of methods for estimating the TE distribution from a small number of experiments. The method chosen here is to repeat the test a few times then do simple bootstrapping to re-sample the data. Verdes [63] demonstrates that the source and conditioning variables should be bootstrapped, leaving the target variable intact. This generates a large number of sampled sets from the original set. The distribution of the means of these sets can be determined empirically. There are a large number of bootstrap methods [17] but the simple case resampling method is selected. The more complex methods are useful for increasing the accuracy of the distribution estimations but the simple resampling method is shown by Dixon [17] to be suitable for estimating the SE when high accuracy is not a priority. The overestimation of SE is also helpful in this case to avoid spurious results. Simple case resampling is the process of randomly sampling the original data many time

to generate new sets called resamples (fig A.1).

$$\begin{array}{ccc}
 & X_1^* & \mu_1^* \\
 & X_2^* & \mu_2^* \\
 & \vdots & \vdots \\
 X \Rightarrow & X_N^* & \Rightarrow \frac{\mu_N^*}{\mu_X} \\
 & & SD_X
 \end{array}$$

Figure A.1: An overview of the case resampling bootstrapping method

The resamples can contain repeated values from the original dataset. For example, the variable $X = (x_1, x_2, x_3, x_4, x_5)$ may become $X_1^* = (x_3, x_2, x_5, x_2, x_1)$. Resamples are indicated by the asterisk, and the subscript differentiates it among the resample variables. The mean of each resample (μ_i^*) is approximately normally distributed even if the original data is not. A thorough discussion of the distribution of means is provided by Norman & Streiner [39]. Once the resampled means are obtained for both series, the t-test is performed as before except the resampled means are the inputs. Returning to the example above; the two variables $TE_{C \rightarrow A}$ and $TE_{A \rightarrow C}$ are resampled $N = 100$ times.

For brevity, a notation change will be employed, $TE_{C \rightarrow A} = C$ and $TE_{A \rightarrow C} = A$, the resamples are C_i^* are A_i^* respectively and the means are $\mu_{C,i}^*$ and $\mu_{A,i}^*$, where i equals all integers 1 to N . The differences of the means is calculated as

$$d = \mu_{C,i}^* - \mu_{A,i}^* \tag{A.3}$$

As before the mean and standard deviation of the differences are taken to calculate the SE (eqn. A.1) and t (eqn. A.2). This time the result is far more confident, the t-statistic is 105.9 with 99 degrees of freedom so the two-tailed p-statistic is much less than 0.001; the results are almost definitely different. The difference between the methods is highlighted in table A.2. It can be seen that increasing the number of samples reduces the SD and SE which has a large effect on the t-statistic.

Table A.2: A comparison of the statistical properties when bootstrapping is applied compared to when it is not.

| Quantity | No Bootstrap | Bootstrap |
|---------------------------|--------------|-------------|
| $\mu_{\text{difference}}$ | 0.1463 | 0.1471 |
| $SD_{\text{difference}}$ | 0.0327 | 0.0139 |
| SE | 0.0164 | 0.0014 |
| t-statistic | 8.921 | 105.9 |
| p-statistic | 0.0030 | $\ll 0.001$ |

A.0.2 Statistical Comparison to Zero

The same method can be used to determine whether a measured value is statistically different from zero. This is done by finding the null distribution, that is, the distribution of the TE calculation when the two variables are independent. This is done by bootstrapping the source variable multiple times and calculating TE of all of the resamples; e.g., $TE_{X_i^* \rightarrow Y}^*$. The mean of these measurements is equivalently the mean of the null distribution and the bias error which exists due to noise in the data. The null threshold is defined by the mean, standard deviation and desired confidence interval;

$$\text{NullThreshold} = \overline{TE_{X_i^* \rightarrow Y}^*} + 1.96SD \quad (\text{A.4})$$

for 95% confidence. If the measured value is greater than the null threshold, it is statistically different from zero.

Appendix B

Derivation of Probabilistic TE

B.1 Pairwise TE

This section derives a probabilistic form of the standard TE equation, Eq. 2.11, which is easier applied to the data in this work. First, a general probabilistic form of conditional entropy is derived using A and B to denote general variables. Local conditional entropy, is defined as [13],

$$h(a|b) = -\log(\Pr(A|B)) \quad (\text{B.1})$$

The expected value of this is conditional entropy,

$$H(A|B) = -\sum_{a,b \in \Omega} \Pr(A, B) \log(\Pr(A|B)) \quad (\text{B.2})$$

where $a, b \in \Omega$ means all possible pairs of a and b . The first term of the Eq. 2.11 is obtained by replacing A and B with Y_t and Y_{t-1} , respectively.

$$H(Y_t|Y_{t-1}) = -\sum_{y, \dot{y} \in \Omega} \Pr(Y_t, Y_{t-1}) \log(\Pr(Y_t|Y_{t-1})) \quad (\text{B.3})$$

where the over dot indicates that event belongs to to the lagged variable, $y \in Y_t$ and $\dot{y} \in Y_{t-1}$.

The second term of Eq. 2.11 is obtained by substituting $A = Y_t$ and $B = (Y_{t-1}, X_{t-\delta})$ into Eq. B.2.

$$H(Y_t|Y_{t-1}, X_{t-\delta}) = -\sum_{y, \dot{y}, x \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}) \log(\Pr(Y_t|Y_{t-1}, X_{t-\delta})) \quad (\text{B.4})$$

To combine the two terms, Eq. B.3 and Eq. B.4, the summations must be made to match. This is achieved by applying the definition $\Pr(A) = \sum_b \Pr(A, B)$ to Eq. B.3, so it becomes,

$$H(Y_t|Y_{t-1}) = - \sum_{y, \dot{y}, x \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}) \log(\Pr(Y_t|Y_{t-1})) \quad (\text{B.5})$$

Therefore, Eq. 2.11 becomes,

$$\text{TE}_{X \rightarrow Y} = \sum_{y, \dot{y}, x \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}) [\log(\Pr(Y_t|Y_{t-1}, X_{t-\delta})) - \log(\Pr(Y_t|Y_{t-1}))] \quad (\text{B.6})$$

which is simplified using the identity $\log(A) - \log(B) = \log\left(\frac{A}{B}\right)$

$$\text{TE}_{X \rightarrow Y} = \sum_{y, \dot{y}, x \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}) \log\left(\frac{\Pr(Y_t|Y_{t-1}, X_{t-\delta})}{\Pr(Y_t|Y_{t-1})}\right) \quad (\text{B.7})$$

Although this is an applicable probabilistic form of the TE equation, it is desirable to rearrange so all of the probabilities are joint. This simplifies the code written to calculate the probabilities but it is not an essential step. Using the identity $\Pr(A|B) = \frac{\Pr(A, B)}{\Pr(B)}$,

$$\Pr(Y_t|Y_{t-1}, X_{t-\delta}) = \frac{\Pr(Y_t, Y_{t-1}, X_{t-\delta})}{\Pr(Y_{t-1}, X_{t-\delta})} \quad (\text{B.8})$$

and

$$\Pr(Y_t|Y_{t-1}) = \frac{\Pr(Y_t, Y_{t-1})}{\Pr(Y_{t-1})} \quad (\text{B.9})$$

Therefore, the final form of Eq. 2.11 is,

$$\text{TE}_{X \rightarrow Y} = \sum_{y, \dot{y}, x \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}) \log\left(\frac{\Pr(Y_t, Y_{t-1}, X_{t-\delta}) \Pr(Y_{t-1})}{\Pr(Y_{t-1}, X_{t-\delta}) \Pr(Y_t, Y_{t-1})}\right) \quad (\text{B.10})$$

To obtain the formula for the calculation in the opposite direction, the variables are simply switched.

$$\text{TE}_{Y \rightarrow X} = \sum_{x, \dot{x}, y \in \Omega} \Pr(X_t, X_{t-1}, Y_{t-\delta}) \log\left(\frac{\Pr(X_t, X_{t-1}, Y_{t-\delta}) \Pr(X_{t-1})}{\Pr(X_{t-1}, Y_{t-\delta}) \Pr(X_t, X_{t-1})}\right) \quad (\text{B.11})$$

B.2 Conditional TE

The process for deriving the probabilistic form of the conditional TE equation, Eq. 2.13, is similar to that of pairwise TE. The first term is obtained by substituting $A = Y_t$, $B = Y_{t-1}, Z_{t-\zeta}$ into Eq. B.2.

$$H(Y_t|Y_{t-1}, Z_{t-\zeta}) = - \sum_{y, \dot{y}, z \in \Omega} \Pr(Y_t, Y_{t-1}, Z_{t-\zeta}) \log(\Pr(Y_t|Y_{t-1}, Z_{t-\zeta})) \quad (\text{B.12})$$

where Z is the conditioning variable, lagged by ζ . As before, the identity $\Pr(A) = \sum_b \Pr(A, B)$ is used to make the summations of the two terms uniform.

$$H(Y_t|Y_{t-1}, Z_{t-\zeta}) = - \sum_{y, \dot{y}, x, z \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}) \log(\Pr(Y_t|Y_{t-1}, Z_{t-\zeta})) \quad (\text{B.13})$$

Similarly, the second term is obtained by substituting $A = Y_t$, $B = Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}$ into Eq. B.2.

$$H(Y_t|Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}) = - \sum_{y, \dot{y}, x, z \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}) \log(\Pr(Y_t|Y_{t-1}, X_{t-\delta}, Z_{t-\zeta})) \quad (\text{B.14})$$

Combining the two terms gives,

$$\text{TE}_{X \rightarrow Y|Z} = \sum_{y, \dot{y}, x, z \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}) \log \left(\frac{\Pr(Y_t|Y_{t-1}, X_{t-\delta}, Z_{t-\zeta})}{\Pr(Y_t|Y_{t-1}, Z_{t-\zeta})} \right) \quad (\text{B.15})$$

This can be recast as purely joint probabilities using the identity $\Pr(A|B) = \frac{\Pr(A, B)}{\Pr(B)}$, such that Eq. 2.13 becomes,

$$\text{TE}_{X \rightarrow Y|Z} = \sum_{y, \dot{y}, x, z \in \Omega} \Pr(Y_t, Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}) \log \left(\frac{\Pr(Y_t, Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}) \Pr(Y_{t-1}, Z_{t-\zeta})}{\Pr(Y_{t-1}, X_{t-\delta}, Z_{t-\zeta}) \Pr(Y_t, Y_{t-1}, Z_{t-\zeta})} \right) \quad (\text{B.16})$$

which can be reversed by substitution to get,

$$\text{TE}_{Y \rightarrow X|Z} = \sum_{x, \dot{x}, y, z \in \Omega} \Pr(X_t, X_{t-1}, Y_{t-\delta}, Z_{t-\zeta}) \log \left(\frac{\Pr(X_t, X_{t-1}, Y_{t-\delta}, Z_{t-\zeta}) \Pr(X_{t-1}, Z_{t-\zeta})}{\Pr(X_{t-1}, Y_{t-\delta}, Z_{t-\zeta}) \Pr(X_t, X_{t-1}, Z_{t-\zeta})} \right) \quad (\text{B.17})$$

Appendix C

Comparison to ARIMA Residuals

Method

The following section contains a report submitted to Prof. Keith Hipel in December 2017, toward the completion of SYDE 631, Time-Series Analysis. It is included to demonstrate the application of a parametric model to FSI data. It shows that, while possible, these models are time consuming to apply and rely on the data being normally distributed. Minor formatting changes have been made but the content is unchanged.

C.1 Introduction

Identification of causal relationships in arrays of structures is incredibly difficult and not well developed. The main issue is that in most cases when one structure acts on another, that structure usually acts on the first creating a feedback loop. This makes it difficult to differentiate interactions between structures from the base excitation or noise. In the field of fluid dynamics, these types of interacting arrays are very common. For example wind farms are arrays of windmills which all modify the air flow and thereby interact with one another, as do tall buildings in cities, structural supports of off-shore oil platforms, moving cars on a highway and any other collection of structures that are acted upon by a fluid. Traditional methods for resolving the fluid interactions can be prohibitively difficult and rarely yield results regarding the degree of feedback between structures. This report details a comparison of two techniques for determining causal interaction between two structures in a fluid flow using the time series of their displacement.

The experiment used to generate the time series used here, places two structures in series

in a flow and measures their displacements with respect to time as detailed in Section C. The time series of the displacement angles is used to calculate the Transfer Entropy in the system. For comparison Auto Regressive Moving Average (ARMA) models are fitted to the individual series (Sections C, C) and the cross correlation of the residuals is used to identify causal interactions (Section C). Exploratory analysis (Section C) shows the observation data to be non-seasonal and stationary thereby indicating that the basic form of ARMA models is appropriate, it does however indicate that Box-Cox transformations may be necessary. The order of the model and any transformation coefficients are determined through a combination of exhaustive and targeted techniques (Section C). Diagnostic tests are used to confirm that the models are appropriate before the Akaike Information Criterion (AIC) is used to discriminate between models. All calculations were performed using MATLAB, as this is a learning exercise everything was coded from scratch where feasible but native functions were used sparingly as detailed in Section C.

The comparison between the two methods of identifying causal interactions is important because, while Transfer Entropy has been widely applied to other networks, it has almost never been applied to fluid-structure interactions. In fact, the only study of this known to the author is the one that generated the data used in this report. This work is directly related to the Master's research of the author who is studying the applicability of Transfer Entropy to fluid-structure interactions.

C.1.1 Description of Experiments

The observation series used throughout this report were obtained at New York University Tandon School of Engineering, Brooklyn, NY, USA. A complete description of the experimental method and results are in the process of being published by Zhang *et al.* [74] as listed in the reference section. A brief outline of the experiment is included here to aid understanding of the results. All calculations and numerical results in this paper were independently calculated from data shared by Maxwell Rosen of New York University.

The experiment placed two airfoils in a channel of flowing water, one directly upstream of the other (Fig. C.1). The airfoils are both placed on pivots, one is driven and the other is allowed to passively change its angle. The driven airfoil is driven to either ± 15 deg, which position it takes is decided randomly every 0.5 seconds to emulate a stochastic process. Six test cases are utilised, three different distances between the airfoils and two driving cases, upstream active and downstream active as listed in Table C.1. This experiment is

used because the fluid dynamics are relatively well understood so the qualitative results are predictable. The expected behaviour is that vortices shed periodically from the upstream airfoil will act to rotate the downstream airfoil. When the active airfoil is driven to a new position the shedding process is disrupted a pressure disturbance is introduced. Vortices are carried by the fluid and as such cannot advect upstream, so the downstream airfoil cannot communicate to the upstream airfoil with vortices. The pressure disturbance caused by the rotation of the airfoil can act upstream but its range of communication is biased downstream. Therefore the expected results are:

- The driven airfoil will have more effect on the passive than the passive has on the active.
- The upstream airfoil will have more effect on the downstream than the downstream has on the upstream airfoil.
- Causal interactions will decrease as distance between the airfoils increases.

The time series used throughout this report are of the airfoil angles with respect to time. Figure C.9 illustrates the relationship between the two airfoils over a small sample period in the case where the upstream airfoil is driven and the downstream airfoil is passive.




Figure C.1: Schematic of airfoil placement in experimental setup.

Each case contains a data set for the angular position of the upstream airfoil, from here called U and a set for the angular position of the downstream airfoil, henceforth called D . This means there are 12 data sets which require individual models to be fit to them.

C.1.2 Auto-regressive Moving Average Models

Auto-regressive moving average (ARMA) models are a highly versatile class of model

| Case | Active | L/c |
|------|--------|-----|
| 1 | U | 0.2 |
| 2 | U | 0.5 |
| 3 | U | 1.0 |
| 4 | D | 0.2 |
| 5 | D | 0.5 |
| 6 | D | 1.0 |

Table C.1: Test Cases

which can be fit to time series data, primarily for the purposes of forecasting or simulation. As described by Hipel & McLeod [24], the model is made up of observations (z_t) and random innovations (a_t), auto-regressive parameters (ϕ) determine the importance of past observations while moving average parameters (θ) act on past innovations. The form is defined by eqn. C.1

$$\phi(B)(z_t - \mu) = \theta(B)a_t \quad (\text{C.1})$$

where B is the back step operator which operates on t and μ is the mean of the series. Note if any transformation is performed on the series, μ is the mean of the transformed series. $\phi(B)$ and $\theta(B)$ can be expanded to

$$\phi(B) = 1 - \phi_1 - \phi_2 - \dots - \phi_p \quad (\text{C.2})$$

$$\theta(B) = 1 - \theta_1 - \theta_2 - \dots - \theta_q \quad (\text{C.3})$$

where p and q are the orders of the AR and MA parameters respectively. ARMA models of this form require observation series to be non-seasonal and stationary with innovations that are identically independently distributed (IID). It is also common to impose the additional constraint of normality on the innovations. The power of ARMA models comes from the ability to add on additional processes to help meet those requirements. For example non-stationarity can be removed by including a differencing operator, seasonality can be accounted for in various ways and residuals can be made normal through data transformations, see Section C. In this study the models are used to whiten the data so the residuals can be compared in an effort to identify causal relationships between the series.

C.1.3 Exploratory Analysis

In this section the data is used to inform decisions regarding appropriate models to use. The sample auto correlation function (ACF) and partial auto correlation function (PACF) are used to guide the selection of parameter ranges for the exhaustive study and the data

distribution histogram helps identify the need for a data transformation.

Case 1, U

This series is of the upstream airfoil which, in this case, is driven. This is evident by the almost binary distribution of the observations. The raw time series (Fig. C.2) clearly shows the process is stationary, has constant variance and is not seasonal. This is confirmed by knowledge of the physical process which is mechanically driven to either 15 or -15 deg.

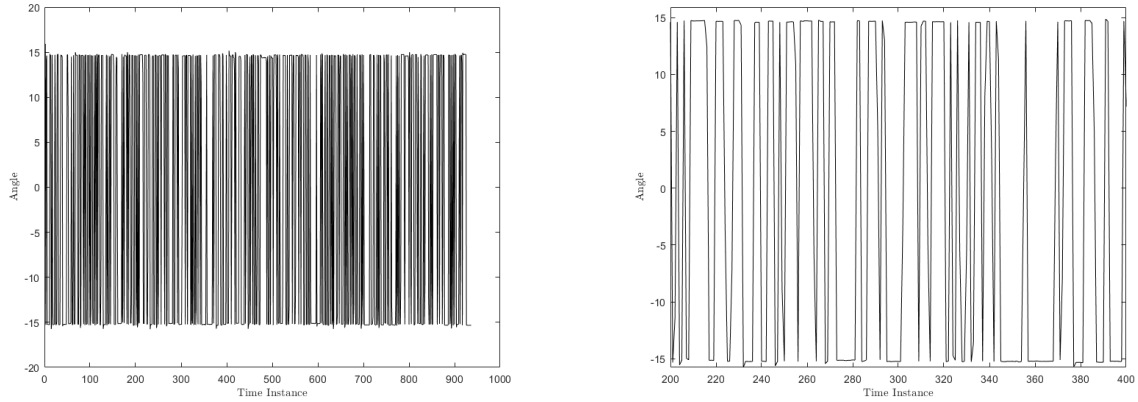


Figure C.2: The raw observation series for case 1, U

Taking the auto-correlation function (Fig. C.3a) and partial auto-correlation function (Fig. C.3b) of this series shows indicates an MA(1) model is likely as demonstrated by the truncation of the ACF after lag 1 and the attenuation of the PACF. However the distribution of the series (Fig. C.3c) shows that it is very likely that a transformation will be required because it is far from a normal distribution.

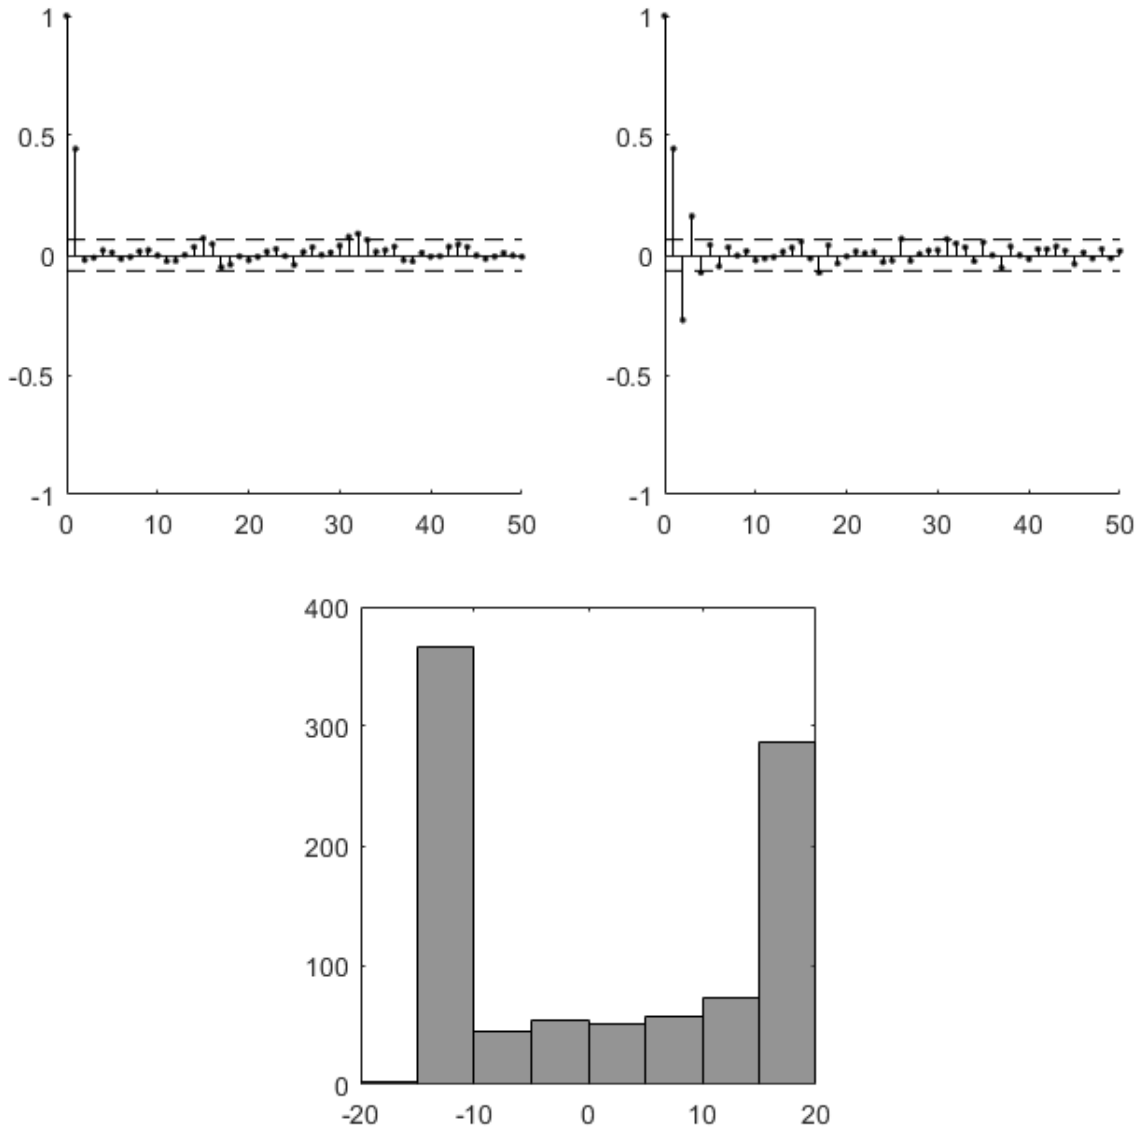


Figure C.3: Sample ACF, PACF and data distribution for case 1, U

From this analysis it appears that a moving average model is likely, however the near binary distribution of the data makes it very likely that a Box-Cox transformation will be needed and this may introduce the need for auto regressive parameters. Therefore ARMA models with a wide range of Box-Cox coefficients will be trialed.

Case 1, D

This is the series of the downstream airfoil which is passive. A plot of the time series shows that the observations are somewhat correlated but are overall stationary. The large

scale fluctuations do not appear to be periodic or seasonal, especially considering the Sample ACF which demonstrates non-periodic cycling. No change in variance is obviously visible.




Figure C.4: The raw observation series for case 1, D

Both the sample ACF (Fig. C.5a) and sample PACF (Fig. C.5b) attenuate, indicating the need for an ARMA model. The data distribution looks to be approximately normally distributed but not exactly so it is unclear if a transformation will be required.

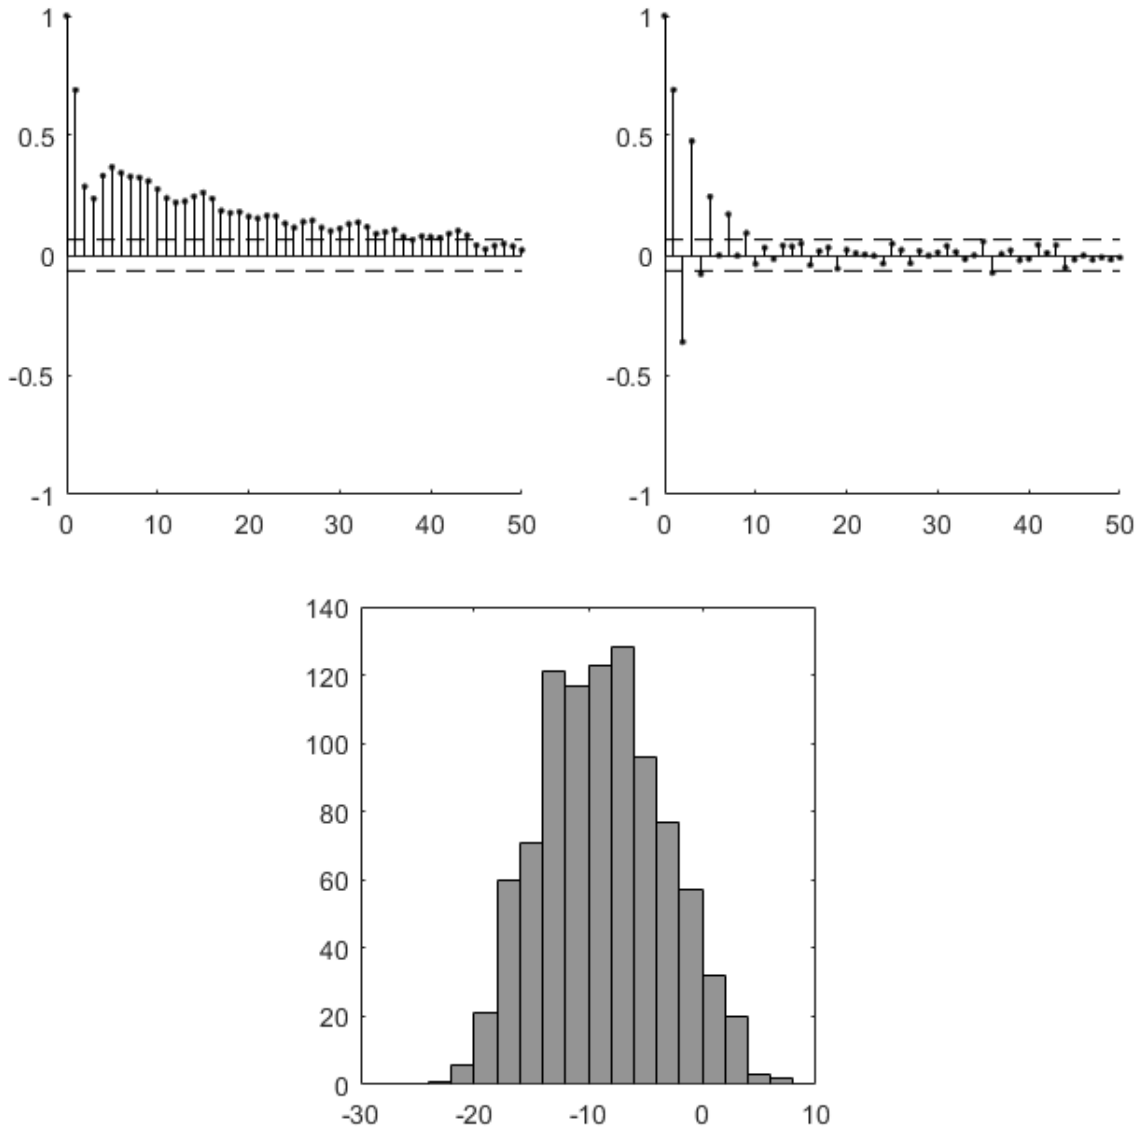


Figure C.5: Sample ACF, PACF and data distribution for case 1, D

Unlike the model for U, it is not immediately clear what order of model is needed so an exhaustive study is carried out in the next section. From this analysis it is concluded that an ARMA model should be an appropriate model class, a Box-Cox transformation may or may not be required.

C.1.4 Confirmatory Analysis

The confirmatory analysis phase of model construction involves fitting models with estimated parameters to the data, performing diagnostic checks to confirm the model is

appropriate and measuring the relative suitability of the model compared to other models. Note that this stage only applies to fitting the ARMA models to the data, Transfer Entropy calculations are done independently.

The approach employed in this study uses a combination of exhaustive and targeted methods to identify the best possible models to fit to the time series data. Exploratory analysis in Section C was used to set parameter limits and identify the appropriate class of model to use. A coarse exhaustive study runs through all possible combinations of p , q and λ and selects the best model based on the Akaike Information Criterion (AIC). Initially the values of p , q and λ are set to the following ranges; $p = 1, 2, \dots, 10$; $q = 1, 2, \dots, 10$; $\lambda = -2, 0, -1.75, \dots, 2.00$, a total of 1600 combinations per data set. Of course $\lambda = 1$ is equivalent to no transformation.

This study automatically excludes any models which do not pass the diagnostic checks. The diagnostics are based on the core assumptions required by ARMA models, those being; the model must be invertible and stationary, the residuals must be white, normally distributed and heteroscedastic. Invertibility and stationarity are confirmed automatically by the MATLAB native function which estimates the parameters, see Section C. Residual whiteness is checked by a version of the Portmanteau test, normality is measured from the $g1$ and $g2$ statistics and heteroscedasticity is checked using the Engle test.

The results are then used to inform refinement decisions, such as model constraints, on a case by case basis.

Exhaustive Study

The code used to perform the exhaustive study follows the procedure:

- For all cases
 - Read in data for case
 - For $\lambda =$ all λ
 - Transform data
 - For $p =$ all p
 - for $q =$ all q
 - Create a model template in the form ARMA(p,q)
 - Estimate model parameters using the Maximum Likelihood Estimation method.
 - Perform diagnostics checks
 - If diagnostic checks pass
 - Calculate AIC

- Find the best (lowest) AIC

Data transformations are performed by the Box-Cox transformation in eqn. C.4 [24].

$$z_t^{(\lambda)} = w_t = \begin{cases} \lambda^{-1} [(z_t + c)^\lambda - 1], & \lambda \neq 0 \\ \ln(z_t + c), & \lambda = 0 \end{cases} \quad (\text{C.4})$$

where the coefficient c ensures the series is all positive values. In the code this is done by finding the minimum value of z_t and rounding it up to the nearest integer. The rounding step is not necessary but it helps keep things a bit more clean and has no residual effects. MATLAB has native tools for setting up ARIMA models and estimating the parameters based on the maximum likelihood of the input data. The model is setup by the command `model = arima(p,d,q)`, this can then be fed into the function `estimate(model,data)` which outputs MLEs for the parameters, the maximized log likelihood value and the variance-covariance matrix of the estimates. Standard errors are calculated from the variance-covariance matrix by squaring the diagonal.

The AIC is calculated from eqn. C.5.

$$AIC = -2\log(\text{ML}) + 2k \quad (\text{C.5})$$

where $\log(\text{ML})$ is output by the estimate function and k is the number of degrees of freedom based on the number of parameters plus one for variance.

Diagnostics

As previously mentioned the invertibility and stationarity of the estimated model are confirmed automatically by the estimate function.

Whiteness of the residuals is checked by graphically from the ACF of the residuals. All values for lags > 0 should be less than the 95% CI where the confidence interval is

$$CI = \frac{1.96}{\sqrt{n}} \quad (\text{C.6})$$

where n is the number of samples. It is confirmed via a Ljung-Box Q-test which is a modified version of the Portmanteau test. The null hypothesis of this test is the data is independently distributed. The test statistic is calculated by the method of Ljung & Box

[31].

$$Q = n(n+2) \sum_{k=0}^h \frac{\hat{\rho}_k^2}{n-k} \quad (\text{C.7})$$

where h is the number of lags, in this case $h = \ln(n)$ as recommended by Tsay [61]. Q is χ^2 distributed with $h - (p + q)$ degrees of freedom.

Normality of the residuals is measured by the skewness and kurtosis parameters as outlined in Cramer [14]. They are calculated as $g1$ (eqn. C.8) and $g2$ (eqn. C.9) receptively and compared to their standard errors given by eqns. C.10 and C.11.

$$g1 = \frac{\frac{1}{n} \sum_{t=1}^n \hat{a}_t^3}{\left(\frac{1}{n} \sum_{t=1}^n \hat{a}_t^2\right)^{3/2}} \quad (\text{C.8})$$

$$g2 = \frac{\frac{1}{n} \sum_{t=1}^n \hat{a}_t^4}{\left(\frac{1}{n} \sum_{t=1}^n \hat{a}_t^2\right)^2} - 3 \quad (\text{C.9})$$

$$SES = \sqrt{\frac{6n(n-1)}{(n-2)(n+1)(n+3)}} \quad (\text{C.10})$$

$$SEK = 2(SES) \sqrt{\frac{n^2 - 1}{(n-3)(n+5)}} \quad (\text{C.11})$$

The skewness coefficient ($g1$) should be within the bounds $\pm 1.96(SES)$ and the kurtosis coefficient should be within $\pm 1.96(SEK)$ to prove that the residuals are normally distributed with a confidence on 95%. If the residuals are not normally distributed, then a Box-Cox transformation is required to be performed on the observations.

C.1.5 Case 1, U

The exploratory analysis suggested that a MA(1) model may be appropriate, however this model failed diagnostic checks as the residuals were not normally distributed. This confirmed that a Box-Cox Transformation would be required so the exhaustive study was performed. Unfortunately no model was found which would result in normally distributed residuals, the values of the Box-Cox parameter that were tried ranged from +5 to -5. It is noted in Hipel & McLeod [24] (pg. 247) that the kurtosis of the residuals is a less important assumption than residual whiteness so this condition was removed from the auto-selection criteria of the exhaustive study. The results of the coarse exhaustive study with no kurtosis requirement are given in Table C.2. Only the best model for each Box-Cox transformation is given to avoid a table of 2000 elements being included.

| λ | p | q | AIC | g1/SES | g2/SEK |
|-----------|----|----|--------|---------|---------|
| 1.00 | 0 | 1 | 7469.5 | 0.0832 | -7.6122 |
| 1.00 | 3 | 4 | 7457.4 | 0.1920 | -7.3450 |
| 0.75 | 3 | 4 | 6257.4 | 0.0819 | -7.2247 |
| 0.50 | 7 | 5 | 5127.1 | -0.1521 | -7.6295 |
| 0.25 | 10 | 10 | 4071.8 | -0.2771 | -6.7509 |
| 0.00 | 9 | 6 | 3102.7 | -1.0466 | -7.3167 |
| 0.00 | 0 | 1 | 3110.5 | -1.0469 | -7.3796 |

Table C.2: Results of an exhaustive study for case 1, U with parameter ranges $p = 1, 2, \dots, 10$; $q = 1, 2, \dots, 10$; $\lambda = -2, 00, -1.75, \dots, 2.00$

The model identified in the exploratory phase is included for comparison, it is clearly not the best fit. This is likely due to the data transformation. Interestingly, looking at the ACF and PACF of the transformed data (Fig. C.6) no significant change is noted, it still appears that an MA(1) model would be appropriate. The significantly lower AIC score shown in Table C.2 seems to be unfounded. It is suspected that the highly non-normal distribution of the data has made the AIC less accurate. With this in mind the model MA(1) with $\lambda = 0$ was tested. The AIC score is slightly higher than that of the ARMA(9,6) with a value of 3110.5 but there is an obvious gain in model parsimony and the sample ACF and sample PACF clearly implicate this model.




Figure C.6: Sample ACF, PACF and histogram of transformed data with $\lambda = 0$ for case 1, U

In an attempt to select between the models MA(1) and ARMA(9,6), the parameter estimates and their SE are shown in Table C.3. This seems to indicate that all of the parameters of the ARMA(9,6) model are significant, with the possible exception of ϕ_8 .

| Parameter | MLE | SE |
|------------|----------|---------|
| ϕ_1 | -0.5531 | 0.08586 |
| ϕ_2 | -0.2962 | 0.07064 |
| ϕ_3 | -0.4451 | 0.07307 |
| ϕ_4 | -0.1397 | 0.06868 |
| ϕ_5 | -0.7968 | 0.04806 |
| ϕ_6 | -0.417 | 0.06441 |
| ϕ_7 | 0.1813 | 0.04715 |
| ϕ_8 | -0.04418 | 0.03675 |
| ϕ_9 | 0.08226 | 0.03598 |
| θ_1 | 1.008 | 0.05865 |
| θ_2 | 0.5003 | 0.07292 |
| θ_3 | 0.551 | 0.08583 |
| θ_4 | 0.3109 | 0.07825 |
| θ_5 | 0.8927 | 0.06157 |
| θ_6 | 0.8222 | 0.04935 |

| Parameter | MLE | SE |
|------------|--------|---------|
| θ_1 | 0.4752 | 0.05121 |

Table C.3: Parameter estimates and standard errors for ARMA(9,6) and MA(1) models with $\lambda = 0$

Knowing that the residuals for all of these models significantly deviate from the normality assumption, the standard errors and AIC are likely incorrect. Also, because the ACF does not rely on the normality assumption and it strongly indicates an MA(1) as being appropriate (see section C), that is the chosen model. For due diligence both models were used to calculate RCCF, see Section C, and no significant difference was noted. This result is shown in Appendix C.

C.1.6 Case 1, D

As discussed in the exploratory analysis the process driving the data is known to be stationary, to confirm that the series is statistically stationary, an Augmented Dickey-Fuller test is employed. The null hypothesis is that the data is non-stationary and is rejected with a value < 0.0001 compared to a 95% confidence threshold 0.05.

The results of the coarse exhaustive study are given in Table C.4, it found that the best model to fit to the data is an ARMA(7,8) model with a Box-Cox Transformation of

$\lambda = 0.50$. Transformations with $\lambda < 0.50$ resulted in non-normal residuals.

| λ | p | q | AIC |
|-----------|----|---|--------|
| 1.00 | 0 | 3 | 5109.0 |
| 0.75 | 10 | 7 | 3783.6 |
| 0.50 | 7 | 8 | 2626.3 |
| 0.25 | - | - | - |
| 0.00 | - | - | - |

Table C.4: Coarse exhaustive study for case 1, D

To reduce model fit errors, the exhaustive study is now refined for this data set. Based on the high order of the indicated models the new parameter ranges are $p = 1, 2, \dots, 15$; $q = 1, 2, \dots, 15$. The resolution of the transformation parameter is increased and its range reduced so the new range is $\lambda = 0.30, 0.35, 0.40, \dots, 0.95$. The refined results are presented in [Table C.5](#)

| λ | p | q | AIC |
|-----------|----|----|--------|
| 0.95 | 0 | 5 | 4848.4 |
| 0.90 | 9 | 9 | 4507.9 |
| 0.85 | 10 | 9 | 4262.5 |
| 0.80 | 10 | 8 | 4029.2 |
| 0.75 | 10 | 7 | 3783.6 |
| 0.70 | 10 | 10 | 3543.7 |
| 0.65 | 9 | 8 | 3316.5 |
| 0.60 | 13 | 10 | 3082.7 |
| 0.55 | 7 | 10 | 2853.1 |
| 0.50 | 7 | 8 | 2626.3 |
| 0.45 | - | - | - |
| 0.40 | - | - | - |
| 0.35 | - | - | - |
| 0.30 | - | - | - |

Table C.5: Refined exhaustive study for case 1, D

The best AIC value is for the model ARMA(7,8) with a Box-Cox transformation with $\lambda = 0.5$. The MLEs and SEs in Table C.6 shows that ϕ_4 , θ_6 and θ_7 are possibly insignificant as the MLE is less than 1.96 SE.

| Parameter | MLE | SE |
|------------|---------|--------|
| ϕ_1 | 0.8988 | 0.1119 |
| ϕ_2 | 0.0309 | 0.1203 |
| ϕ_3 | -0.5820 | 0.1074 |
| ϕ_4 | -0.0030 | 0.1123 |
| ϕ_5 | 0.5993 | 0.1073 |
| ϕ_6 | -0.6004 | 0.1173 |
| ϕ_7 | 0.5555 | 0.0843 |
| θ_1 | 0.0478 | 0.1186 |
| θ_2 | -0.7424 | 0.1096 |
| θ_3 | 0.4107 | 0.1520 |
| θ_4 | 0.7557 | 0.1390 |
| θ_5 | -0.3782 | 0.1542 |
| θ_6 | -0.0357 | 0.1494 |
| θ_7 | -0.0517 | 0.0971 |
| θ_8 | -0.2430 | 0.0877 |

Table C.6: MLEs and SEs for an ARMA(7,8) model with $\lambda 0.5$ for case 1, D

The insignificant values indicate that it may be appropriate to employ a constrained model. An exhaustive study that tests constrained models would require an unfeasible amount of computation time. Therefore targeted constrained models are chosen and fitted for all of the possible combinations of ϕ_4 , θ_6 and θ_7 . The AIC results of this study are given in Table C.7, the parameter estimates for each of these models are listed in Appendix C. Where no AIC is listed, the model failed diagnostic checks.

| λ | p | q | constrained | AIC |
|-----------|---|---|--|--------|
| 0.50 | 7 | 0 | ϕ_4, ϕ_5, ϕ_6 | 2644.2 |
| 0.50 | 7 | 8 | $\phi_2, \phi_4, \theta_1, \theta_6, \theta_7$ | 2618.3 |

Table C.8: Refined constrained model AICs

| λ | p | q | Constrained | AIC |
|-----------|---|---|------------------------------|--------|
| 0.50 | 7 | 8 | $\phi_4, \theta_6, \theta_7$ | 2637.7 |
| 0.50 | 7 | 8 | ϕ_4, θ_6 | - |
| 0.50 | 7 | 8 | ϕ_4, θ_7 | - |
| 0.50 | 7 | 8 | ϕ_4 | - |
| 0.50 | 7 | 8 | θ_6 | 2624.3 |
| 0.50 | 7 | 8 | θ_7 | - |
| 0.50 | 7 | 8 | θ_6, θ_7 | 2622.5 |

Table C.7: AICs for select constrained models fitted to case 1, D

Where no AIC is given, the model failed at least one of the diagnostic checks. In an effort to achieve model parsimony, other constraints were applied based on the parameter estimates of the above combinations. This iterative process led to the constraint combinations in Table C.8.

The chosen model is now a constrained ARMA(7,8) with no $\phi_2, \phi_4, \theta_1, \theta_6, \theta_7$ and with a Box-Cox transformation with $\lambda = 0.5$. The MLEs for this model are:

| Parameter | MLE | SE |
|------------|---------|---------|
| ϕ_1 | 0.9297 | 0.01946 |
| ϕ_3 | -0.5952 | 0.02423 |
| ϕ_5 | 0.6286 | 0.03150 |
| ϕ_6 | -0.6307 | 0.02657 |
| ϕ_7 | 0.5454 | 0.01983 |
| θ_2 | -0.7306 | 0.02133 |
| θ_3 | 0.4866 | 0.01232 |
| θ_4 | 0.7487 | 0.02145 |
| θ_5 | -0.4402 | 0.02231 |
| θ_8 | -0.2416 | 0.03167 |

Table C.9: MLEs and SEs for the chosen model for C1, D

After two exhaustive studies of 4700 possible models plus nine constrained models, the chosen model is a constrained ARMA(7,8) with no $\phi_2, \phi_4, \theta_1, \theta_6$ or θ_7 and a Box-Cox transformation of $\lambda 0.5$. Diagnostic checks were automatically performed during the fitting process, the details of this for the chosen model are discussed here.

The Residual ACF (Fig. C.7a) shows that the residuals of this model are white. This is confirmed by a Portmanteau test value of 13.7, less than the critical value 31.4 given by the χ^2 distribution for 95% confidence. Normality of the residuals is confirmed by the g1 and g2 statistics. g1 has a magnitude 0.1486 which is less than $1.96 \text{ SES} = 0.1566$ and g2 has a magnitude 0.1606 which is less than $1.96 \text{ SEK} = 0.3130$, the residuals are therefore normal. This is supported by the distribution in Fig. C.7b. Heteroscedasticity is confirmed by the Engle test for heteroscedasticity.

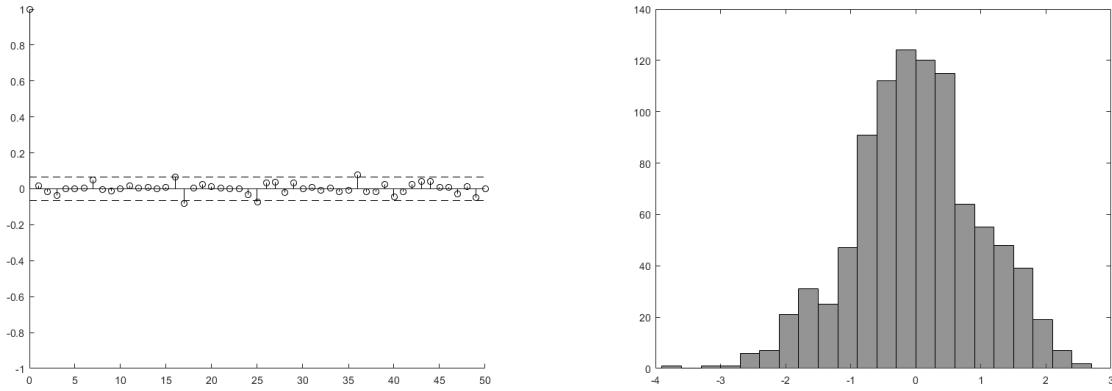


Figure C.7: Residual ACF and histogram for case 1, D after applying a Box-cox transformation $\lambda = 0.50$ and constrained ARMA(7,8) model.

C.1.7 Cases 2 to 6

Due to the time constraints of this project, models are selected based on the results of the exhaustive study; further fitting will be applied at a later date. The selected models are listed in Table C.10, parameter estimates and standard errors are listed in Appendix C.

| Case | Series | p | q | λ |
|------|--------|---|----|-----------|
| 2 | U | 9 | 7 | 0.00 |
| 2 | D | 9 | 7 | 0.25 |
| 3 | U | 6 | 7 | 0.00 |
| 3 | D | 0 | 2 | 1.00 |
| 4 | U | 2 | 7 | 0.00 |
| 4 | D | 5 | 8 | 0.00 |
| 5 | U | 7 | 10 | 0.75 |
| 5 | D | 9 | 10 | 0.00 |
| 6 | U | 8 | 8 | 0.75 |
| 6 | D | 7 | 10 | 0.00 |

Table C.10: Models selected by the coarse exhaustive study for cases 2 to 6. Note that the active airfoil cases do not meet the normal residuals condition

C.1.8 Causality

The models developed in Sections C and C are fitted to the data to generate the residual series. As discussed below the cross correlation of the residual series is used to identify causal interaction between the airfoils. These results are compared to the Transfer Entropy of the series pair which is calculated independent of this report.

C.1.9 Case 1

By overlaying the raw time series of U and D, it is fairly apparent that U affects D.




Figure C.8: Overlaying the time series of case 1, U and case 1, D.

A method for determining causality between two time series is outlined in Hipel & McLeod [24]. The series are pre-whitened by fitting ARMA models then the cross-correlation of the residuals (RCCF) is calculated. Causality is indicated by significant cross correlation of the residuals of two series. A significant value at lag 0 indicates contemporaneous causality, significant values at positive lags indicates U affects D while negative lags indicates D affects U. Significant values at both positive and negative lags indicates feedback.

The RCCF is calculated by eqn. C.12 [24].

$$r_{\hat{U}\hat{D}}(k) = \frac{c_{\hat{U}\hat{D}}(k)}{\sqrt{c_{\hat{U}}(0)c_{\hat{D}}(0)}} \quad (\text{C.12})$$

where $c_{\hat{U}\hat{D}}$ is the estimated cross-covariance given by

$$c_{\hat{U}\hat{D}} = \begin{cases} \frac{1}{n} \sum_{t=1}^{n-k} \hat{U}_t \hat{D}_{t+k}, & k \geq 0 \\ \frac{1}{n} \sum_{t=1-k}^n \hat{U}_t \hat{D}_{t+k}, & k < 0 \end{cases} \quad (\text{C.13})$$

and $c_{\hat{U}}(0), c_{\hat{D}}(0)$ are the sample variances given by

$$c_{\hat{U}}(0) = \frac{1}{n} \sum_{t=1}^n \hat{U}_t^2 \quad (\text{C.14})$$

$$c_{\hat{D}}(0) = \frac{1}{n} \sum_{t=1}^n \hat{D}_t^2 \quad (\text{C.15})$$

Figure C.9a shows that there is some instant negative correlation and U affects D at time lags 2, 3 and 4, whilst feedback (D also affects U) occurs at lag 1. This matches the expected results which predicts that U will act on D through vortex advection at approximately time lag 2, surface wave communication and mechanical connection communication happen in both directions at sufficient speed to happen in < 1 lag period; due to discretisation it is expected that these affects will bleed into the surrounding bins.




Figure C.9: a. Cross correlation of the residuals of case 1, U and case 1, D. b. Transfer Entropy of case 1

The Transfer Entropy plot is read as separate plots for U affecting D ($U \rightarrow D$) and D affecting U ($D \rightarrow U$) where significant values of TE indicate a causal interaction. TE values cannot be negative so it is only a measure of affect and cannot be used to determine between positive and negative correlation. It is also not possible to compare the magnitudes of the results between the two methods because they are measuring different properties. The only appropriate comparison is of the lags at which significant values exist. If the two methods agree, the comparison of RCCF and TE plots should give the following results:

- Significant values at a positive lag on the RCCF should be matched by a significant $TE_{U \rightarrow D}$ value at the same lag on the TE plot.
- Significant values at a negative lag on the RCCF should be matched by a significant $TE_{D \rightarrow U}$ value at the same lag magnitude on the TE plot.

Figure C.9b shows that significant values of Transfer Entropy exist at lags 0 and 2 for U acting on D and at lag 0 for D acting on U. If the significance threshold of the ARMA method is expanded slightly, the two results would match perfectly. It would be reasonable to do so because the confidence interval is based on the assumption that the residuals are normally distributed, which we have confirmed they are not. This is a promising indication that TE is a viable method for detecting causality.

C.1.10 Cases 2 to 6




Figure C.10: a. Cross correlation of the residuals of case 2. b. Transfer Entropy of case 2




Figure C.11: a. Cross correlation of the residuals of case 3. b. Transfer Entropy of case 3




Figure C.12: a. Cross correlation of the residuals of case 4. b. Transfer Entropy of case 4




Figure C.13: a. Cross correlation of the residuals of case 5. b. Transfer Entropy of case 5




Figure C.14: a. Cross correlation of the residuals of case 6. b. Transfer Entropy of case 6

C.1.11 Discussion

Comparing the lags at which significant interaction is identified by the two methods shows relatively good agreement. It is easier to compare using the summary table C.11 which orders the lags by significance for each case and method. Remember that $TE_{U \rightarrow D}$ is comparable to RCCF + and $TE_{D \rightarrow U}$ is comparable to RCCF -. The most obvious property of this comparison is that the RCCF method always identifies more significant lags than the TE method. It is not immediately obvious which is more accurate since discussions of causality rarely have a ground truth. The lack of ground truth coupled with the idea that falsely indicating causality is worse than failing to identify causality leads to the conclusion that the significance level of the RCCF method should be increased. The agreement between the most significant values is excellent. Cases 1, 4, 5 and 6 all perfectly agree on the two most significant lags while case 2 agrees on the two most significant values but in the opposite order for $U \rightarrow D$. Case 3 is very noisy with a wide range of lags indicated and very small values of TE and correlation.

Both methods confirm the hypothesis that for the cases with a driven upstream airfoil (cases 1-3) the communication time increases as separation distance increases. Case 1 is the closest case with a primary lag of 2, case 2 has a primary lag of 2 or 3 and case 3 is 3 or 4. Both methods also confirm that the driven airfoil has a greater effect on the passive

airfoil than the passive has on the driven. This cannot be seen in Table C.11 but is evident in Figures C.9 - C.14. For cases 1-3 where U is driven, the RRCCF is more significant at positive lags than at negative lags and $TE_{U \rightarrow D}$ is more significant than $TE_{D \rightarrow U}$. The inverse is true for cases 4-6 where D is driven.

| Case | $TE_{U \rightarrow D}$ | $TE_{D \rightarrow U}$ | RCCF + | RCCF - |
|------|------------------------|------------------------|-------------|----------|
| 1 | 2,0,1 | 0 | 2,0,4,3 | 0,1 |
| 2 | 3,2 | 0 | 2,3,1,5 | 0 |
| 3 | 4,2,1,3,6 | - | 3,9,1,2,8,7 | - |
| 4 | 0,1 | 0,2 | 0,1 | 0,2,4,11 |
| 5 | 0 | 0,2 | 0,1 | 0,2,1 |
| 6 | 0 | 0,2 | 0,1 | 0,2,1 |

Table C.11: Summary of significant lags predicted by the two methods ordered by significance.

Now that it is proven that the models are significantly dependent it is theoretically possible to fit multivariate ARMA models to these cases. Contemporaneous ARMA (CARMA) models are not appropriate given the significant values at lags other than 0. Transfer function noise models could be fitted to cases 2 and 3 because the interaction in uni-directional however the feedback in cases 1, 4, 5 and 6 mean that the general form of the multivariate ARMA models must be employed. That is not ideal because it leads to a high number of parameters.

C.1.12 Conclusion

Transfer Entropy is a relatively new method for determining causal relationships between nodes of a network. Despite having been applied successfully to a range of industries, Zhang *et al.* [74] claim for the first time that this method can be applied to arrays of structures interacting via a fluid. To test this claim, their method for applying the principles of Transfer Entropy is compared to the residual cross-correlation function. The residual series were obtained by fitting ARMA models to each series individually. This class of model was proven to be acceptable for the passive airfoils because the observation series are stationary and non-seasonal. Box-Cox transformations were required in almost all cases to obtain normally distributed residuals. For the driven airfoils it was never possible to obtain normally distributed residuals therefore it was determined that the ARMA models were not ideal. However, the models were still able to produce uncorrelated residual series

and the results of the causality analysis strongly agreed with both the expected results and the TE method.

ARMA models were selected based on the Akaike Information Criterion after a coarse exhaustive study which was informed by an exploratory analysis, then for case 1 model fit errors were reduced through targeted model fitting experiments based on those results. For cases 2-6, the results of the coarse exhaustive study were used, further model fitting for these series is recommended at a later date. Diagnostics were performed to ensure the models were stationary and invertible, and the residuals were uncorrelated, heteroscedastic and normally distributed. For case 1, the chosen models are

$$(u_t^{(0)} - \mu) = a_t - 0.4752a_{t-1} \quad (\text{C.16})$$

$$(1 - 0.9297B + 0.5952B^3 - 0.6286B^5 + 0.6307B^6 - 0.5454B^7)(d_t^{(0.5)} - \mu) = (1 + 0.7306B^2 - 0.4866B^3 - 0.7487B^4 + 0.4402B^5 + 0.2416B^8)a_t \quad (\text{C.17})$$

Models for cases 2-6 are given in Appendix C.

Figures C.9 to C.14 and Table C.11 are used to show a high level of agreement between the two methods, especially when the significance threshold of the RCCF method is increased. Both models confirmed the working hypotheses that causal interactions are stronger when the distance is less and that the driven airfoil has a greater affect on the passive airfoil than the passive has on the driven. In general the RCCF method appears to be more likely to falsely identify causality where none exists. There is also a significant difference in the complexity of the methods. Fitting the ARMA models to the data took considerable time, the coarse exhaustive study took 8 hours to compute plus several hours for the targeted studies and some time for the exploratory analysis; over 20,000 models were trialed. A certain amount of artfulness was also required to select an appropriate model. By comparison the TE calculation to 0.6 seconds to calculate all 12 series. It also required some knowledge of the physical system to appropriately define the symbols but the efficiency gain is self evident. The RCCF model did prove to be robust given that the model was not a perfect fit for the driven series but the results worked none the less. The result was also barely affected by fitting of different models as proven by the two models fitted to case 1, U (Appendix C).

Due to the high level of non-normality in the ARMA residuals of the driven series, it is unlikely that these models would be useful for forecasting or simulation experiments.

Therefore, while it is potentially possible to fit multivariate ARMA models to the data, there is no practical application for such a model. This study has shown that they are useful for detecting causality and in the future they are likely to be used to strengthen the argument that Transfer Entropy is also a viable method.

C.2 Additional Information

C.2.1 Case 1, U RCCF comparison

Knowing that the AIC could not be relied upon to select an appropriate model for this case, the two most likely models were compared by their affect on the AIC.




Figure C.15: A comparison of the affect of the two most likely U models on the RCCF for case 1

C.2.2 Models 2 to 6

| Parameter | MLE | SE | Parameter | MLE | SE |
|--------------------|----------|---------|--------------------|----------|---------|
| phi ₁ | -0.7946 | 0.1066 | phi ₁ | -0.2526 | 0.1666 |
| phi ₂ | -0.6542 | 0.1779 | phi ₂ | -0.4313 | 0.1199 |
| phi ₃ | -0.5655 | 0.1979 | phi ₃ | -0.09315 | 0.1156 |
| phi ₄ | -0.1258 | 0.1981 | phi ₄ | 0.3985 | 0.09544 |
| phi ₅ | 0.1953 | 0.1844 | phi ₅ | 0.03784 | 0.1286 |
| phi ₆ | 0.562 | 0.1313 | phi ₆ | 0.09646 | 0.09146 |
| phi ₇ | 0.4912 | 0.06014 | phi ₇ | 0.4766 | 0.08161 |
| phi ₈ | -0.06989 | 0.04693 | phi ₈ | 0.03224 | 0.05882 |
| phi ₉ | 0.1901 | 0.04452 | phi ₉ | 0.2761 | 0.05614 |
| theta ₁ | 1.275 | 0.1171 | theta ₁ | 1.664 | 0.1666 |
| theta ₂ | 1.044 | 0.2401 | theta ₂ | 1.344 | 0.2753 |
| theta ₃ | 0.8709 | 0.3045 | theta ₃ | 0.8448 | 0.2447 |
| theta ₄ | 0.4189 | 0.3237 | theta ₄ | 0.1424 | 0.1854 |
| theta ₅ | -0.09381 | 0.3048 | theta ₅ | -0.1134 | 0.1876 |
| theta ₆ | -0.5617 | 0.2318 | theta ₆ | 0.2902 | 0.1844 |
| theta ₇ | -0.6291 | 0.1059 | theta ₇ | 0.1672 | 0.07918 |

Table C.12: Parameter estimates and SEs for Case 2

| Parameter | MLE | SE |
|--------------------|---------|---------|
| phi ₁ | 0.8743 | 0.1471 |
| phi ₂ | -0.5689 | 0.1575 |
| phi ₃ | 0.6136 | 0.152 |
| phi ₄ | -0.851 | 0.1364 |
| phi ₅ | 0.3391 | 0.1749 |
| phi ₆ | 0.3884 | 0.1032 |
| theta ₁ | -0.4206 | 0.1776 |
| theta ₂ | 0.1483 | 0.1229 |
| theta ₃ | -0.3466 | 0.09029 |
| theta ₄ | 0.5773 | 0.07858 |
| theta ₅ | 0.08057 | 0.1357 |
| theta ₆ | -0.5734 | 0.0556 |
| theta ₇ | -0.1859 | 0.08398 |

| Parameter | MLE | SE |
|--------------------|--------|---------|
| theta ₁ | 1.414 | 0.02337 |
| theta ₂ | 0.7317 | 0.02287 |

Table C.13: Parameter estimates and SEs for Case 3

| Parameter | MLE | SE |
|--------------------|----------|---------|
| phi ₁ | 0.344 | 5.407 |
| phi ₂ | 0.6228 | 5.306 |
| theta ₁ | 0.4362 | 5.409 |
| theta ₂ | -0.6694 | 1.083 |
| theta ₃ | -0.5043 | 2.938 |
| theta ₄ | -0.1521 | 0.8647 |
| theta ₅ | -0.0185 | 0.279 |
| theta ₆ | -0.02547 | 0.08212 |
| theta ₇ | -0.02216 | 0.1878 |

| Parameter | MLE | SE |
|--------------------|----------|---------|
| phi ₁ | 0.8766 | 0.5149 |
| phi ₂ | -0.7215 | 0.6576 |
| phi ₃ | 0.7951 | 0.5961 |
| phi ₄ | -0.1455 | 0.6563 |
| phi ₅ | -0.2165 | 0.3911 |
| theta ₁ | -0.3975 | 0.4961 |
| theta ₂ | 0.3027 | 0.4595 |
| theta ₃ | -0.4686 | 0.2901 |
| theta ₄ | -0.2543 | 0.3677 |
| theta ₅ | 0.3178 | 0.2134 |
| theta ₆ | 0.04857 | 0.2016 |
| theta ₇ | -0.01026 | 0.04604 |
| theta ₈ | -0.05732 | 0.03936 |

Table C.14: Parameter estimates and SEs for Case 4

| Parameter | MLE | SE |
|---------------------|----------|---------|
| phi ₁ | 1.134 | 0.3352 |
| phi ₂ | 0.02237 | 0.5549 |
| phi ₃ | -0.6249 | 0.4655 |
| phi ₄ | 0.2837 | 0.4554 |
| phi ₅ | -0.2681 | 0.4852 |
| phi ₆ | 0.4683 | 0.4083 |
| phi ₇ | -0.03938 | 0.2426 |
| theta ₁ | -0.1933 | 0.3322 |
| theta ₂ | -0.5982 | 0.2926 |
| theta ₃ | 0.4309 | 0.2318 |
| theta ₄ | 0.1958 | 0.314 |
| theta ₅ | 0.2154 | 0.2043 |
| theta ₆ | -0.3358 | 0.1759 |
| theta ₇ | -0.2121 | 0.09941 |
| theta ₈ | 0.157 | 0.1057 |
| theta ₉ | -0.08876 | 0.04774 |
| theta ₁₀ | -0.138 | 0.06273 |

| Parameter | MLE | SE |
|---------------------|----------|---------|
| phi ₁ | 0.5199 | 0.1802 |
| phi ₂ | 0.1585 | 0.1849 |
| phi ₃ | -0.5481 | 0.1543 |
| phi ₄ | 0.1913 | 0.1827 |
| phi ₅ | -0.1974 | 0.1701 |
| phi ₆ | 0.1484 | 0.1519 |
| phi ₇ | 0.1524 | 0.1373 |
| phi ₈ | -0.3695 | 0.1513 |
| phi ₉ | 0.5928 | 0.1345 |
| theta ₁ | -0.02937 | 0.1757 |
| theta ₂ | -0.4356 | 0.1509 |
| theta ₃ | 0.4681 | 0.178 |
| theta ₄ | 0.1265 | 0.1604 |
| theta ₅ | 0.1401 | 0.1546 |
| theta ₆ | -0.1449 | 0.1444 |
| theta ₇ | -0.2694 | 0.1385 |
| theta ₈ | 0.4014 | 0.155 |
| theta ₉ | -0.4685 | 0.1164 |
| theta ₁₀ | -0.345 | 0.07369 |

Table C.15: Parameter estimates and SEs for Case 5

| Parameter | MLE | SE | Parameter | MLE | SE |
|--------------------|------------|---------|---------------------|----------|---------|
| phi ₁ | 0.3268 | 0.06346 | phi ₁ | 0.1263 | 0.1898 |
| phi ₂ | 0.2772 | 0.06377 | phi ₂ | -0.3765 | 0.1394 |
| phi ₃ | 0.4055 | 0.06249 | phi ₃ | 1.0 | 0.1228 |
| phi ₄ | 0.1113 | 0.07345 | phi ₄ | 0.07859 | 0.258 |
| phi ₅ | -0.0004029 | 0.0665 | phi ₅ | 0.3345 | 0.1308 |
| phi ₆ | -0.09919 | 0.06126 | phi ₆ | -0.543 | 0.1131 |
| phi ₇ | -0.5925 | 0.05946 | phi ₇ | -0.449 | 0.1825 |
| phi ₈ | 0.5583 | 0.03745 | theta ₁ | 0.3754 | 0.2167 |
| theta ₁ | 0.6235 | 0.07283 | theta ₂ | 0.3452 | 0.06708 |
| theta ₂ | -0.1598 | 0.0868 | theta ₃ | -0.83 | 0.08792 |
| theta ₃ | -0.6754 | 0.08917 | theta ₄ | -0.5778 | 0.2312 |
| theta ₄ | -0.6313 | 0.09926 | theta ₅ | -0.3896 | 0.07067 |
| theta ₅ | -0.3591 | 0.09986 | theta ₆ | 0.3918 | 0.08386 |
| theta ₆ | -0.1682 | 0.08146 | theta ₇ | 0.7334 | 0.1537 |
| theta ₇ | 0.574 | 0.07999 | theta ₈ | 0.2276 | 0.1166 |
| theta ₈ | 0.185 | 0.05432 | theta ₉ | -0.05092 | 0.03805 |
| | | | theta ₁₀ | -0.03673 | 0.03753 |

Table C.16: Parameter estimates and SEs for Case 6

Appendix D

Additional Figures

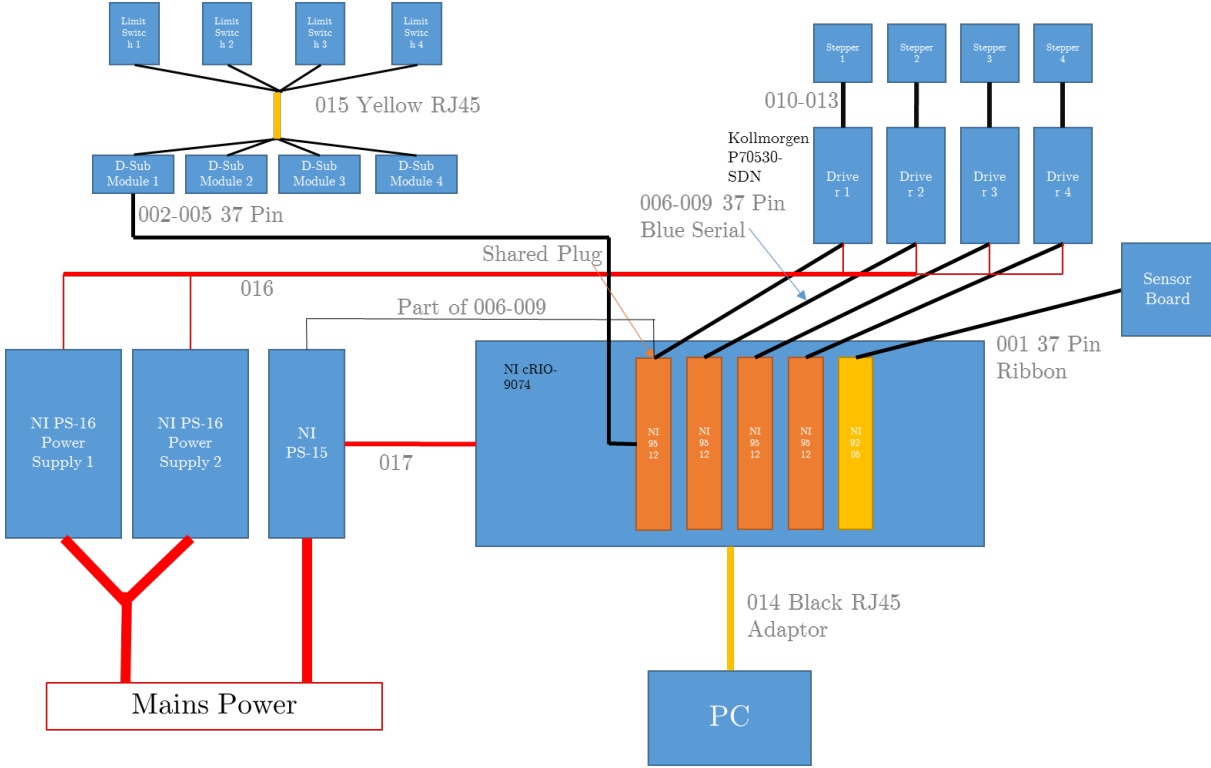


Figure D.1: Schematic of electronics used to drive the cylinder from the cRio

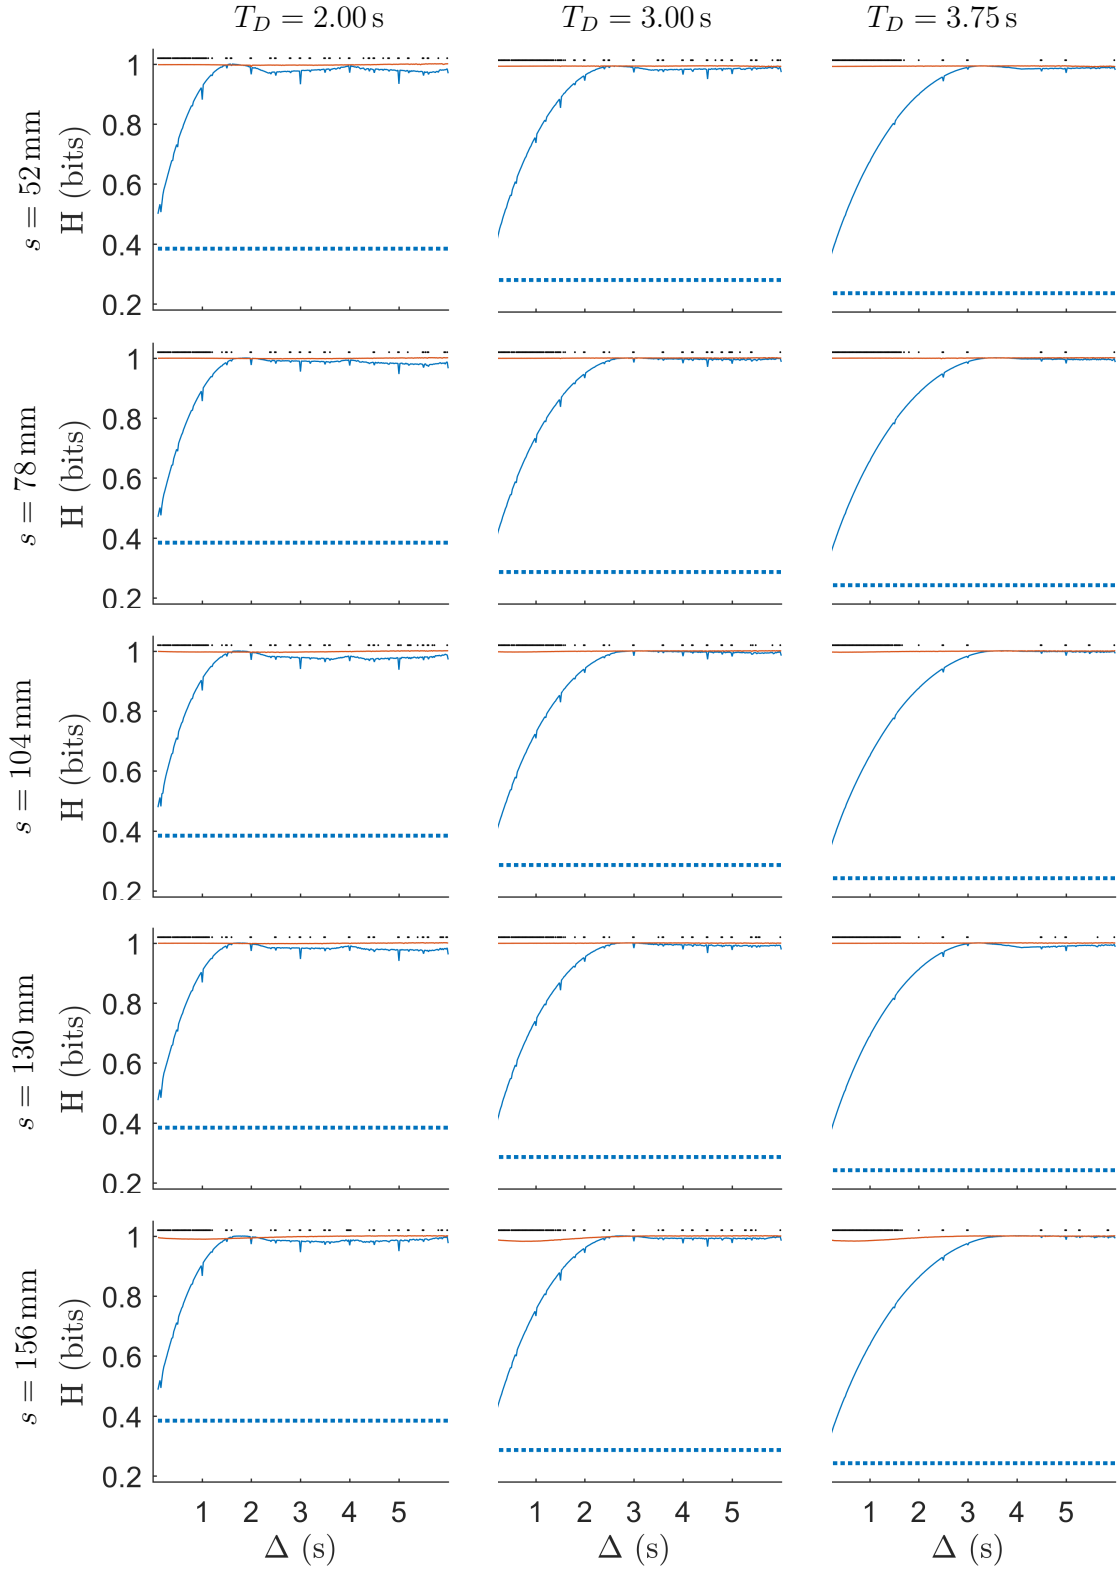


Figure D.2: Entropy as a function of sampling period, calculated using the down-sample first method. Orange line: $H(A)$, blue solid line: $H(C)$ measured, blue dashed line: $H(C)$ analytical, black dots indicate Δ will be filtered out by the entropy filter.

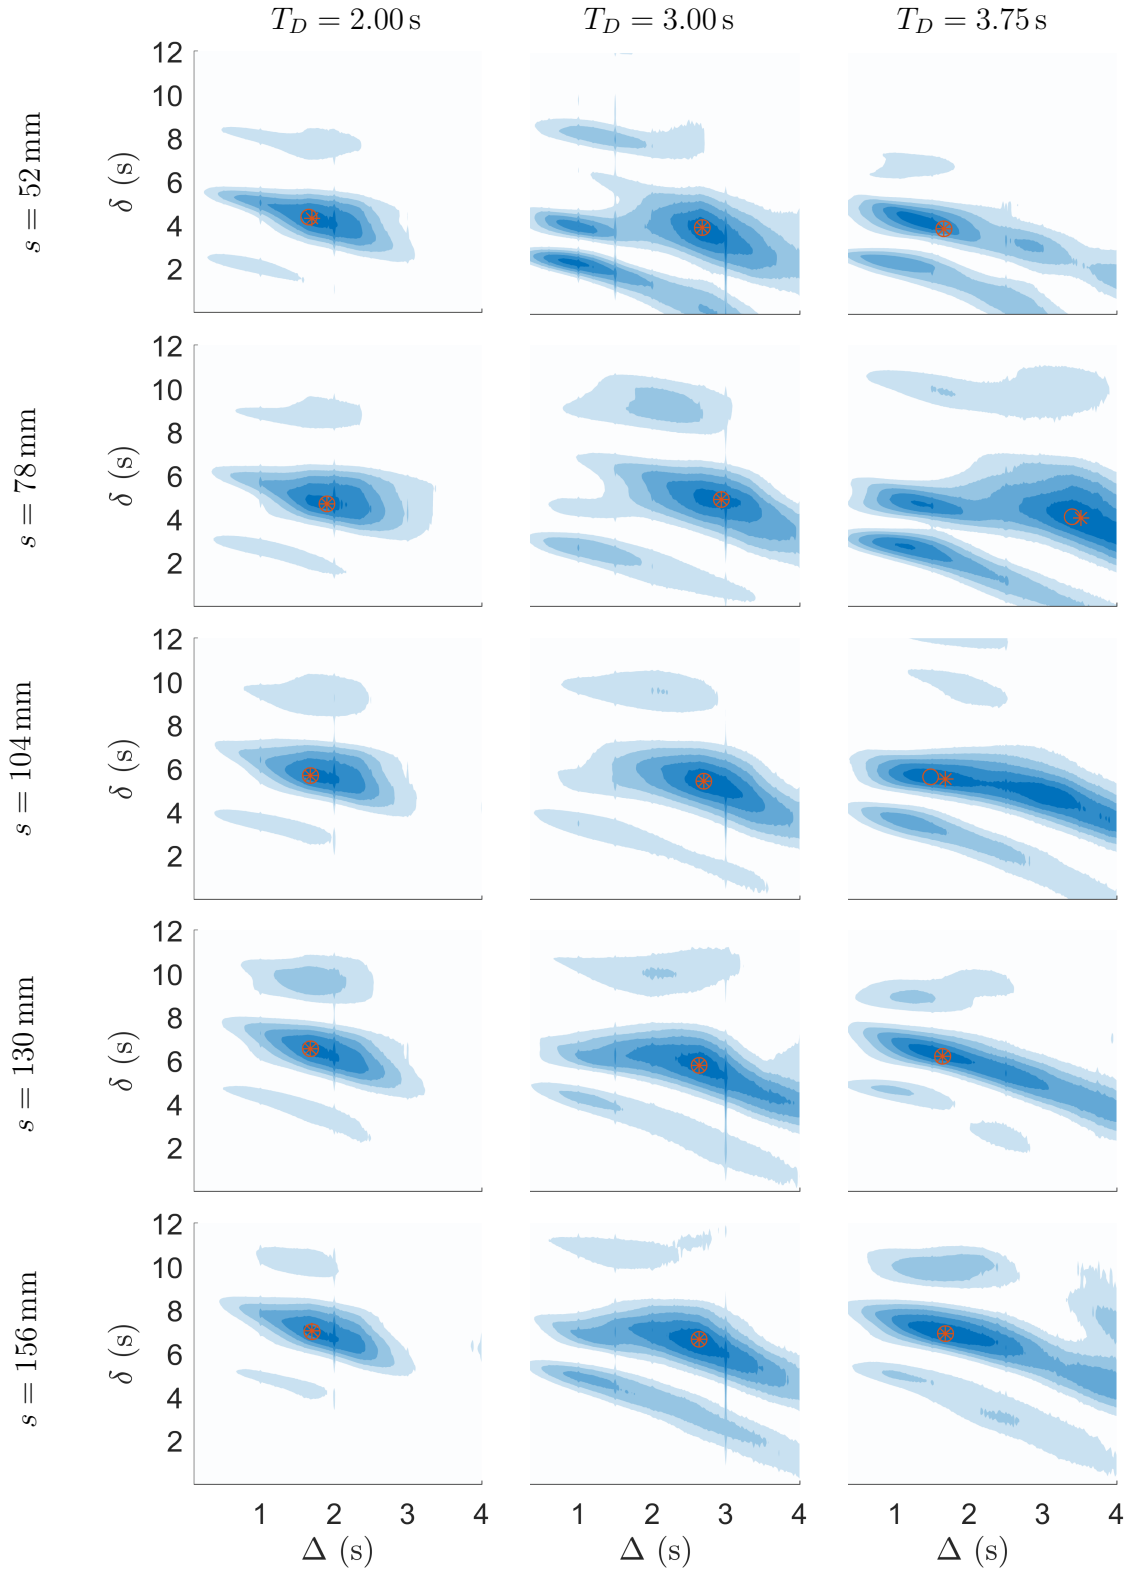


Figure D.3: $TE_{C \rightarrow A}$ for all 15 cases of the single airfoil experiments, calculated using the down-sample first method.

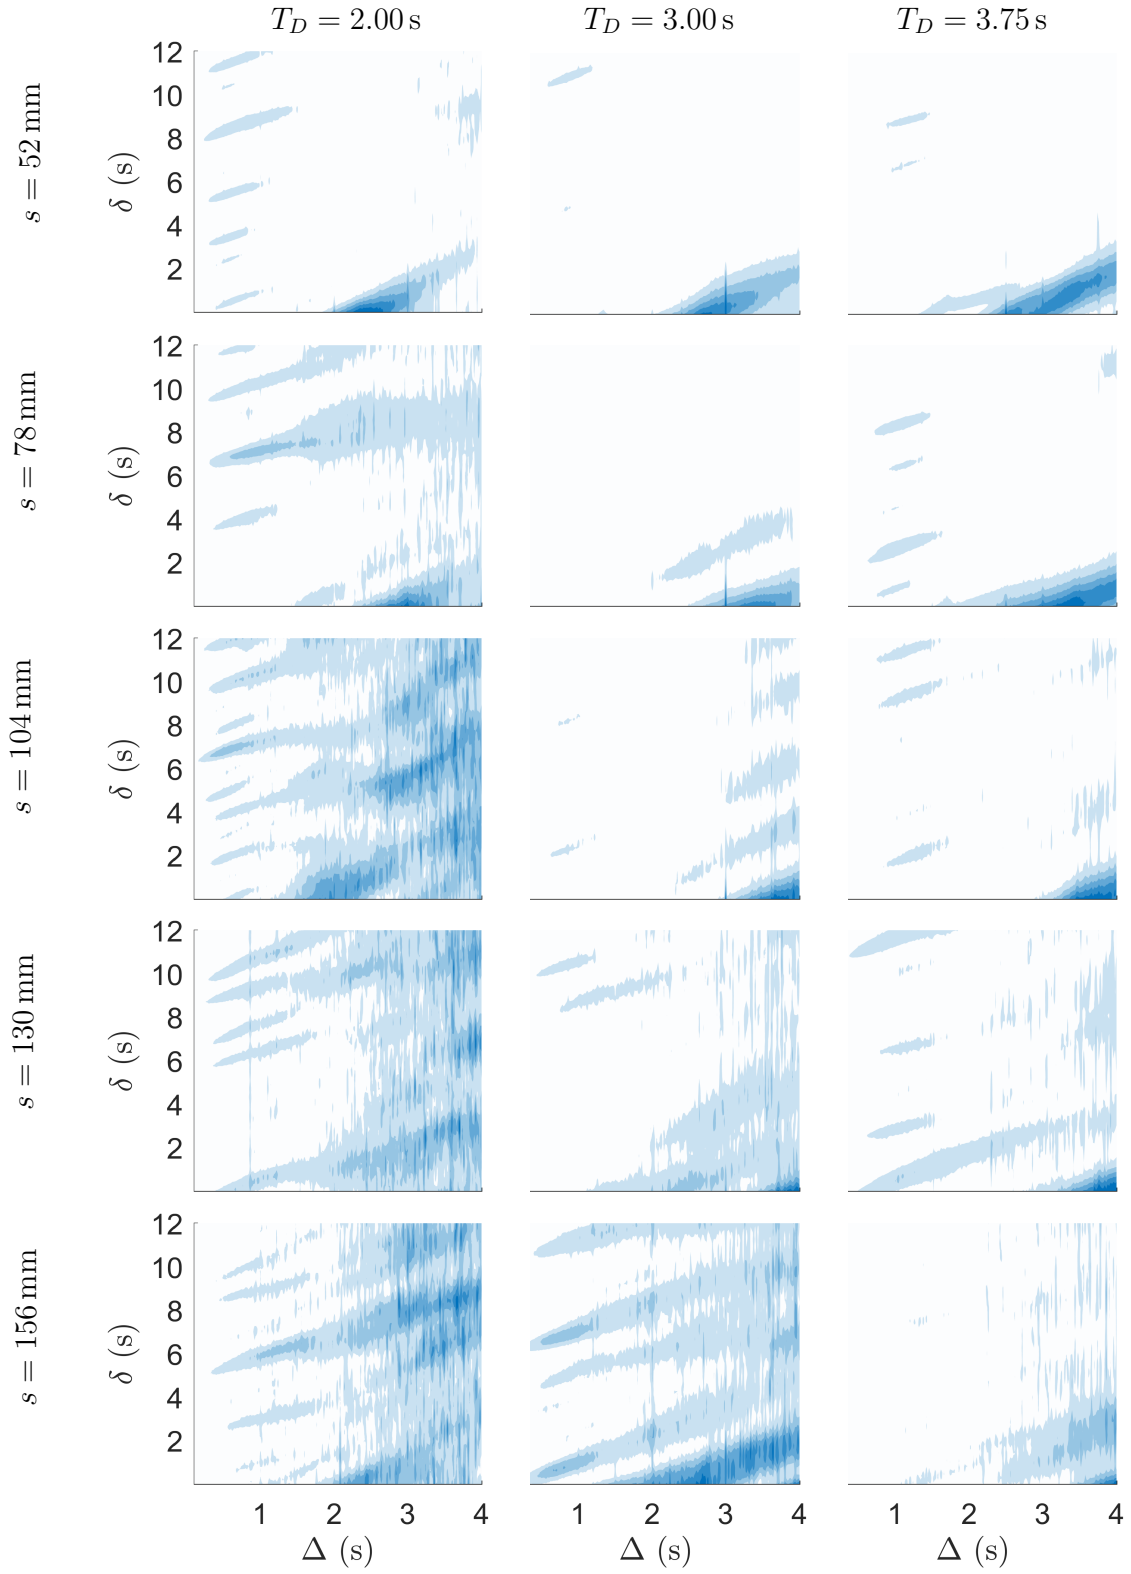


Figure D.4: $TE_{A \rightarrow C}$ for all 15 cases of the single airfoil experiments, calculated using the down-sample first method.

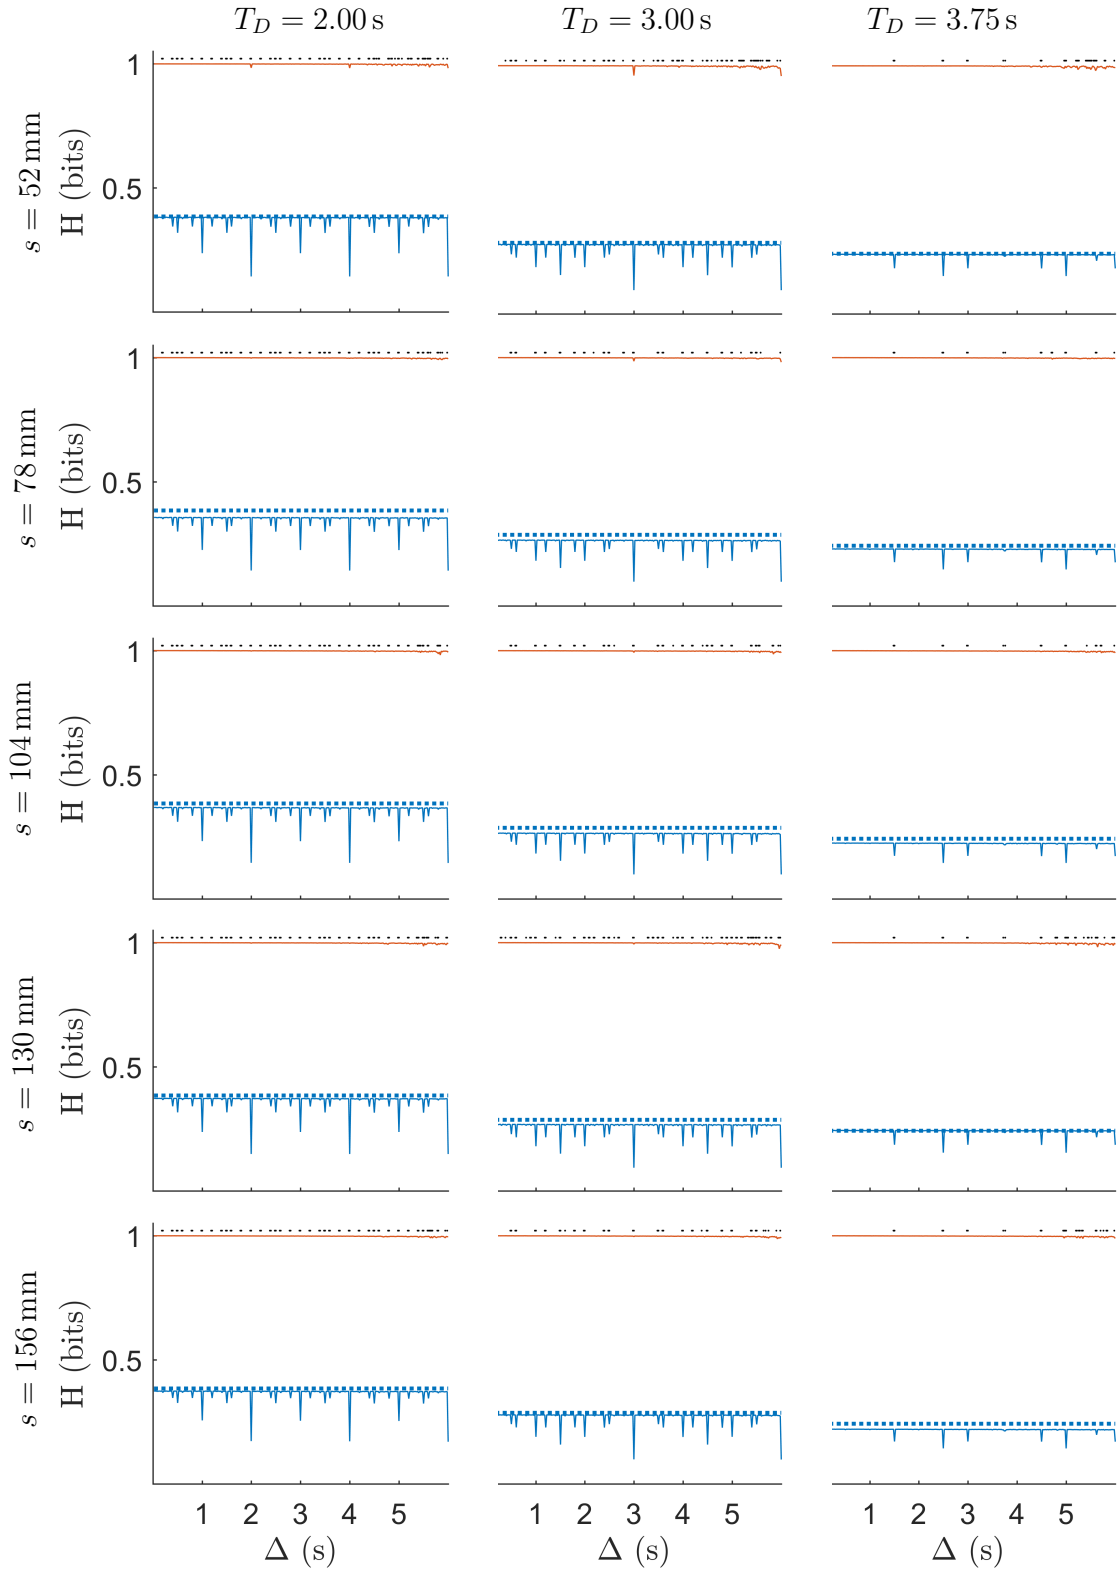


Figure D.5: Entropy as a function of sampling period, calculated using the symbolise first method. Orange line: $H(A)$, blue solid line: $H(C)$ measured, blue dashed line: $H(C)$ analytical, black dots indicate Δ will be filtered out by the entropy filter.

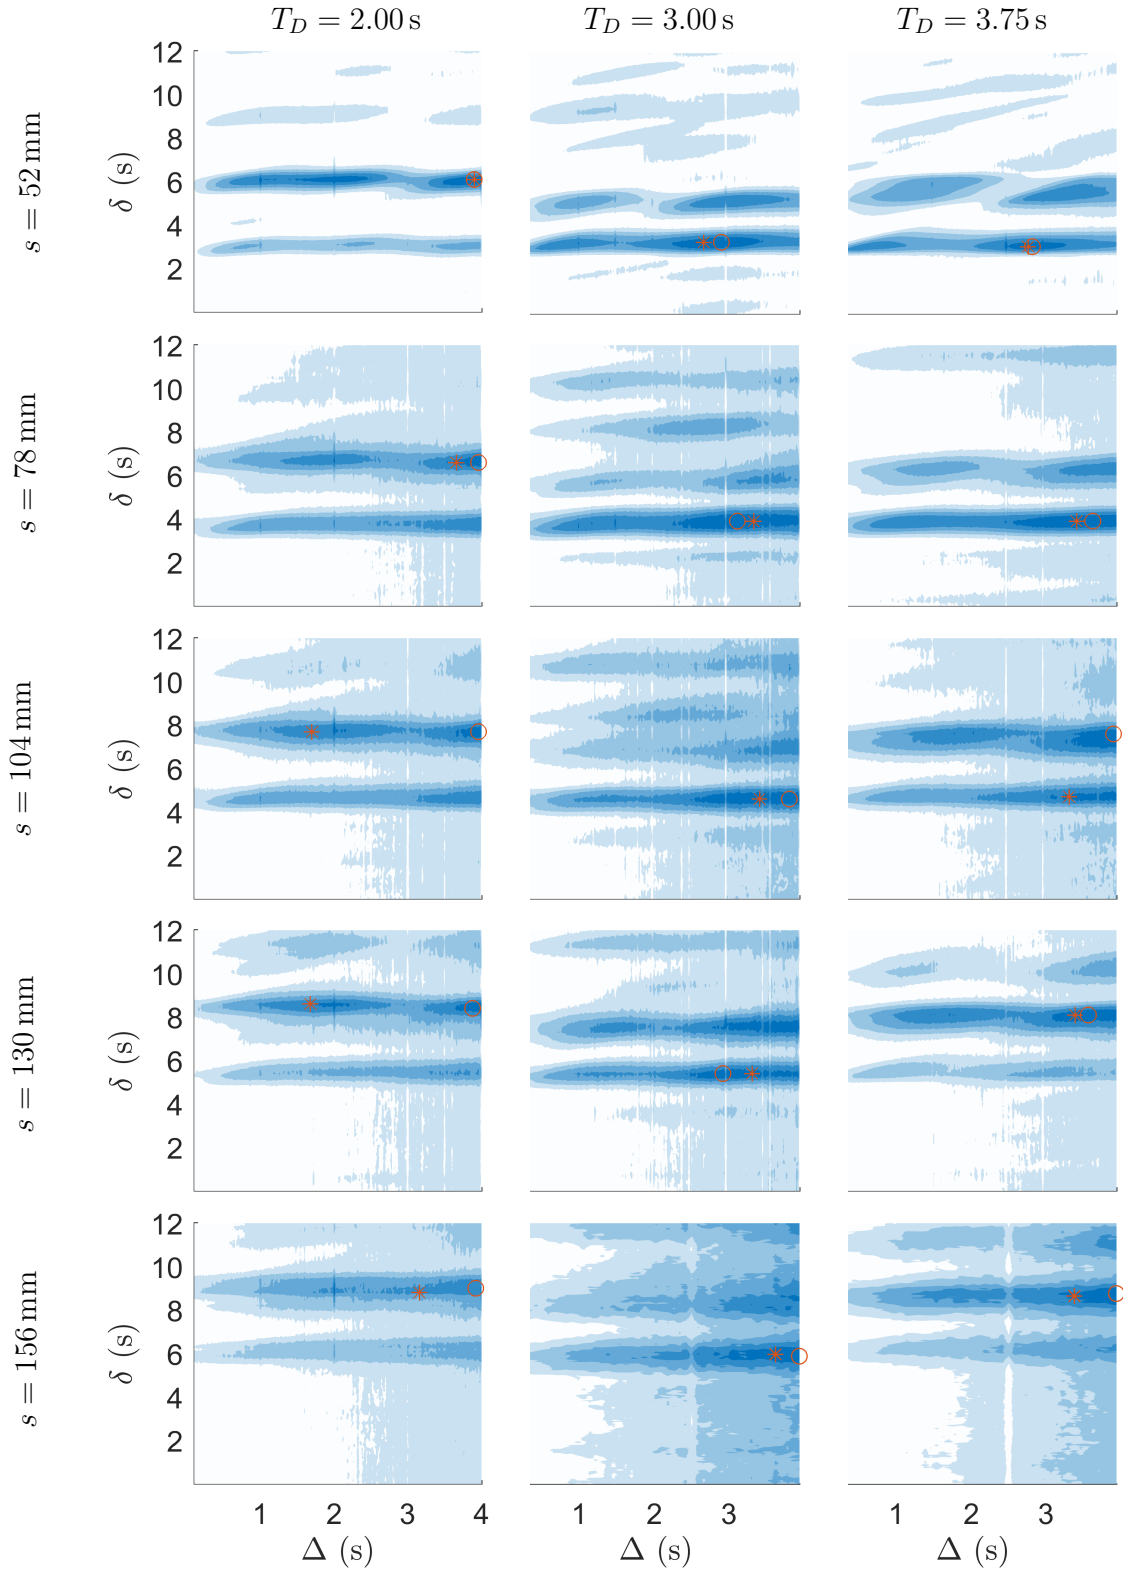


Figure D.6: $TE_{C \rightarrow A}$ for all 15 cases of the single airfoil experiments, calculated using the symbolise first method.

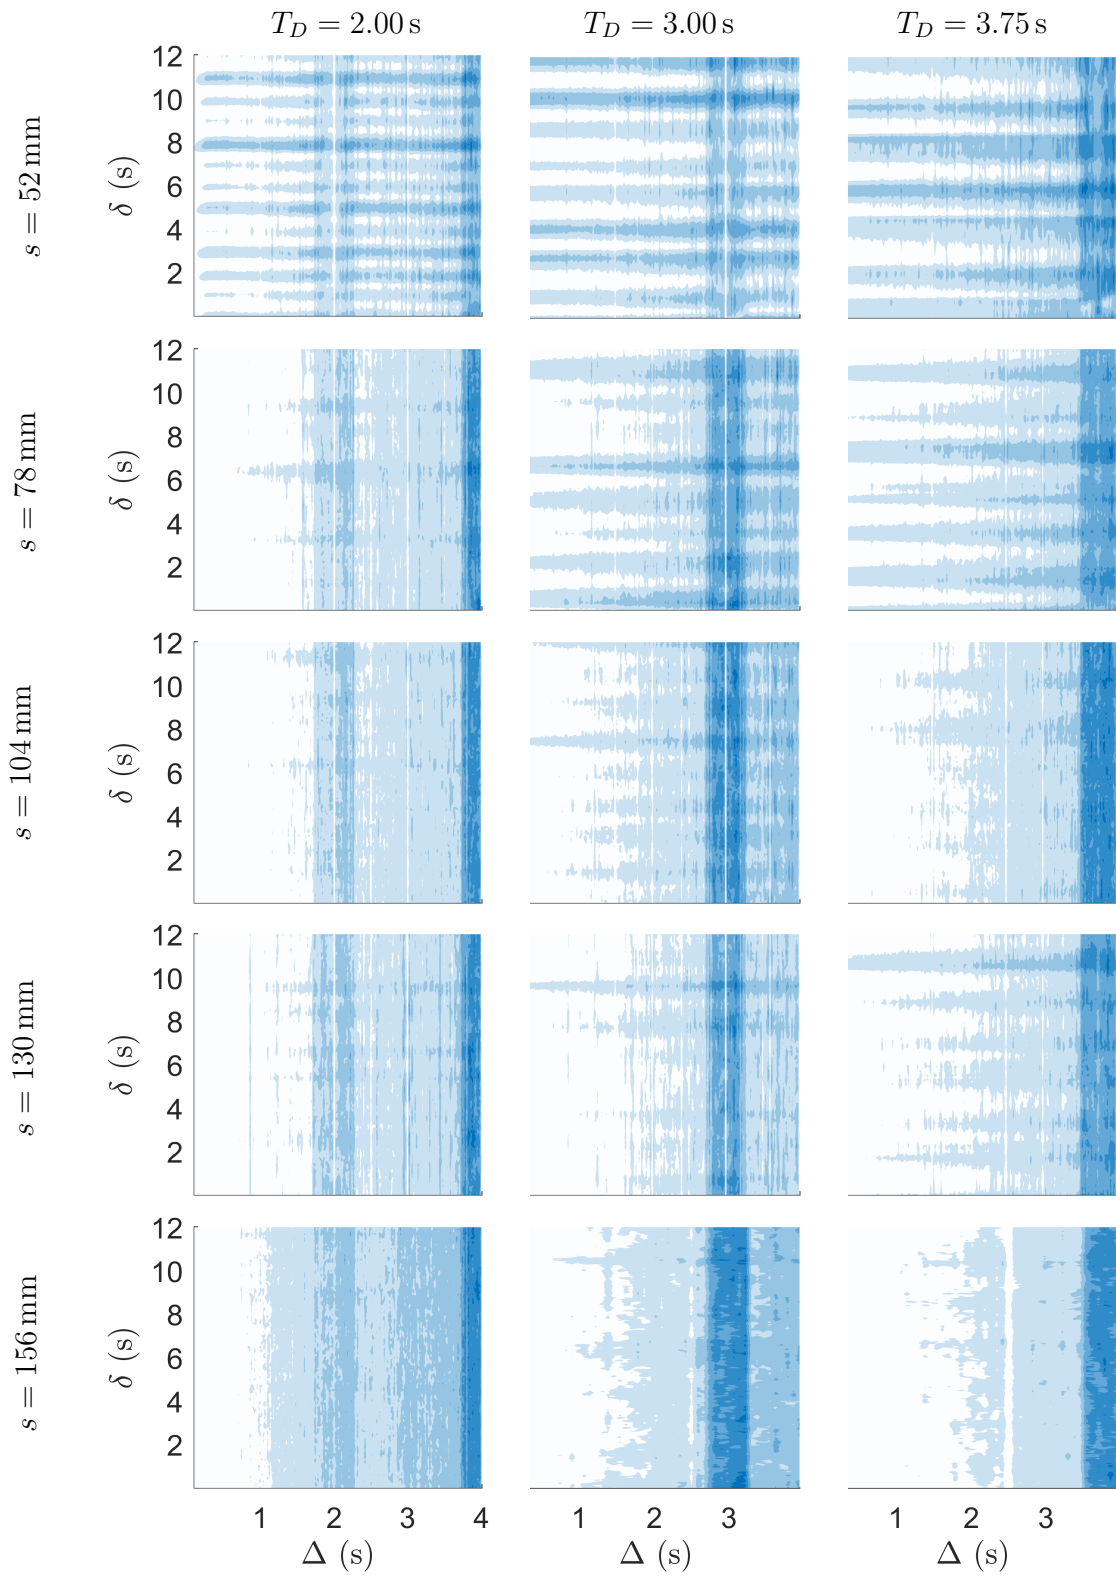


Figure D.7: $TE_{A \rightarrow C}$ for all 15 cases of the single airfoil experiments, calculated using the symbolise first method.

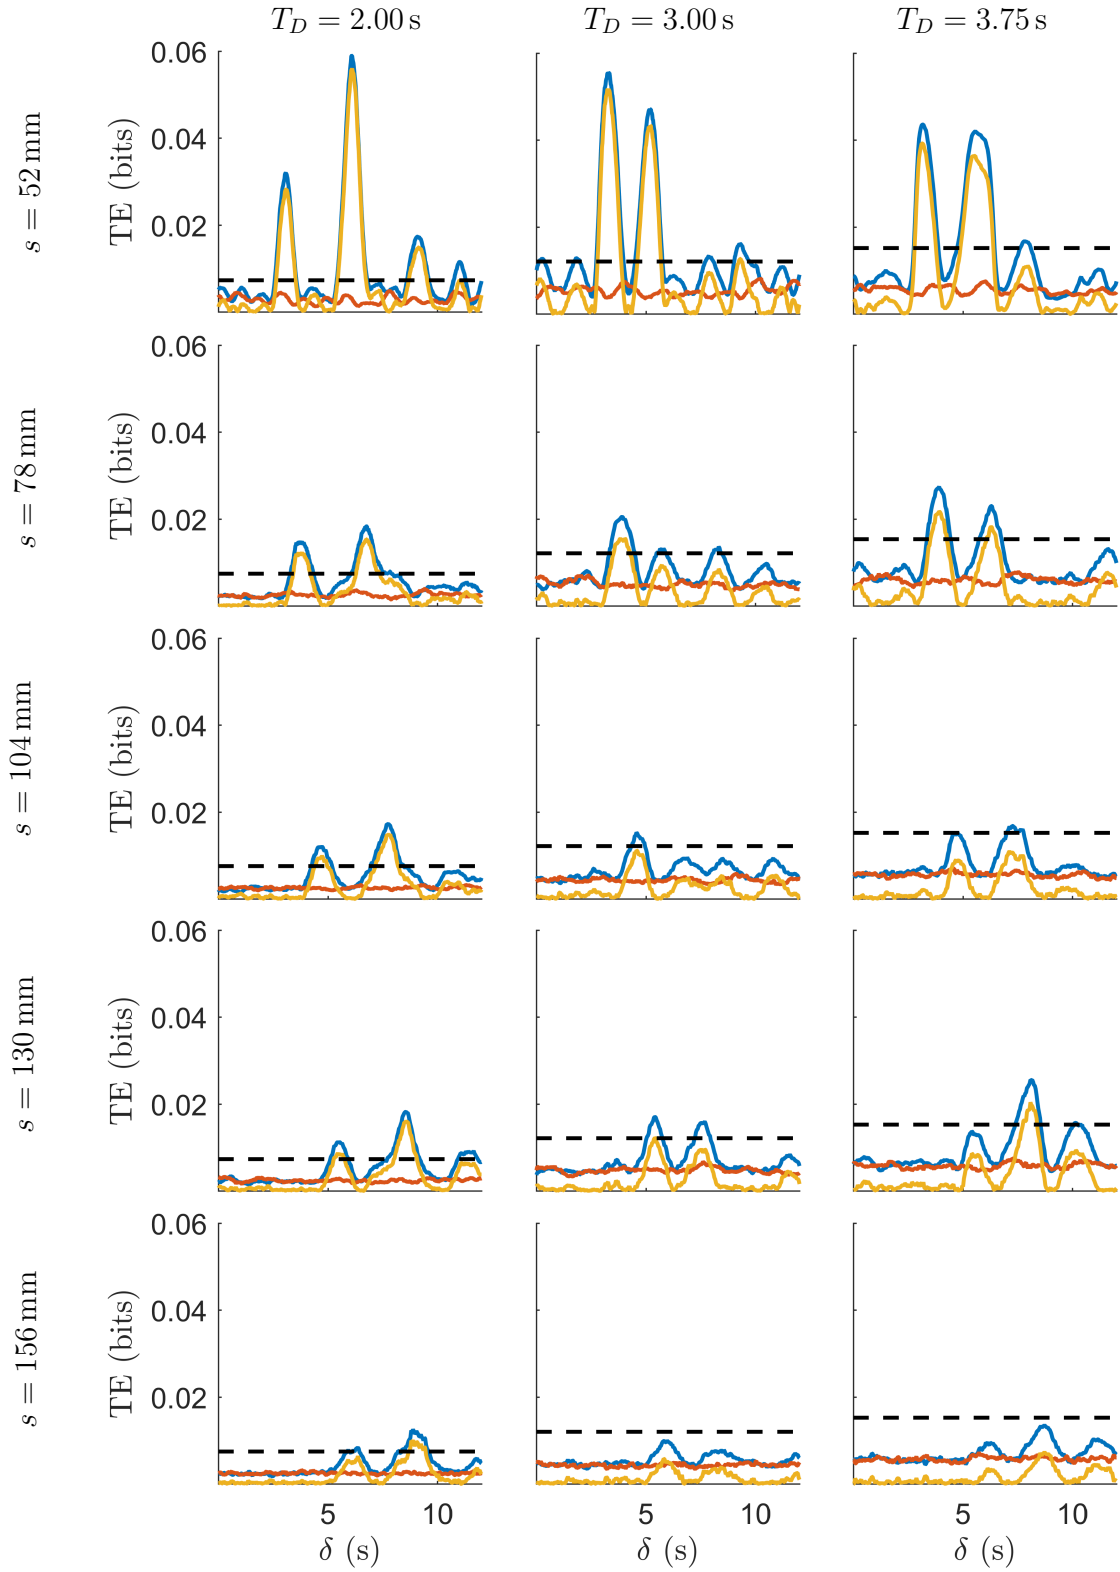


Figure D.8: TE as a function of δ , for a single Δ , for all 15 cases of the single airfoil experiments, calculated using the symbolise first method. Blue line; $TE_{C \rightarrow A}$, orange line; $TE_{A \rightarrow C}$, yellow line; $|TE_{\text{net}}|$, black dashed line; null threshold.




Figure D.9: Self feedback of A (orange line) and C (blue line) compared to the null threshold (black dashed line). Data is presented for a single Δ and a range of δ for all 15 cases.




Figure D.10: $TE_{C \rightarrow A}$ for all cases with indicators of communication delays predicted from the self feedback results.

Appendix E

Analysis of Pathways: Three Body Experiments

In Section 6.1.4, the first step in the delay analysis was to identify the delays which could possibly be related to mutual driver or cascading effects. That is done here again but with a slight modification to suit this data. In the simulation the pathways were simple and obvious. For example, there was only one possible relay path $Z \rightarrow Y$. Here there are many paths that have the same start and end points; these are listed in Table E.1. The current version of the software written to calculate conditional TE cannot process more than one conditioning variable, so any path with more than three nodes cannot be assessed with that tool.¹ The significance of a pathway is then evaluated by assessing the pairwise TE of the component at the delay of the path. For example, the significance of the path $C \rightarrow A \rightarrow A$ is measured by $TE_{C \rightarrow A}$ at delays 6.2 and 9.2 s; that path is significant.

It is acknowledged that this method is incomplete, being entirely new, and that this experiment ended up not being a great example so the following analysis contains a few assumptions to allow for the demonstration to be completed. Here it will be assumed that if the delay of a relayed path matches one of the delays previously attributed to a direct interaction ($\hat{\delta}$) and it is feasible given what is known about the system, then that is the most likely cause of the detected interaction.² The goal is to attribute each of the direct

¹Time restrictions prevent the software being rewritten.

²Ideally, each of the identified pathways would be evaluated by conditional entropy. Just as the path $C \rightarrow A \rightarrow B$ was shown to have insignificant TE when evaluated by $TE_{C \rightarrow B|A}(\delta = \hat{\delta}_{CA} + \hat{\delta}_{AB}, \zeta = \hat{\delta}_{AB})$, $TE_{C \rightarrow A|A}(\delta = \hat{\delta}_{CA} + \hat{\delta}_{AA}, \zeta = \hat{\delta}_{AA})$ should also be insignificant. This is shown not to be case in Fig. 6.21c suggesting that either, $\hat{\delta}_{CA} = 6.2$ s is real or another path is matching the delay. The most likely candidate for that is $C \rightarrow A \rightarrow B \rightarrow A$ with delay 6.6 s.

delays to either one of the relay paths or a feasible physical interaction.

$\hat{\delta}_{CA} = 3.2\text{ s}$ does not match any of the other path delays, but it is a reasonable physical delay as it closely matches similar experiments in Chapter 5. $\hat{\delta}_{CA} = 6.2\text{ s}$ only appears in $C \rightarrow A \rightarrow A$, making that a likely cause. The true cause of the 3.0 s delay $A \rightarrow A$ was identified in Chapter 5 to be caused by the combination of the symbolization algorithm and the dynamics of the airfoil's response,³ so this pathway is feasible. $\hat{\delta}_{CA} = 9.6\text{ s}$ occurs in both $C \rightarrow A \rightarrow B \rightarrow A$ and $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$. This may indicate that both are true but there is no better tool for assessing the validity of this, at this stage. $\hat{\delta}_{CA} = 11.0\text{ s}$ does not exactly match any of the paths, but the low magnitude and wide spread of $\text{TE}_{B \rightarrow A}$ around $\delta = 4.3\text{ s}$ makes the estimation of the delay along that path susceptible to noise. Accepting that this method is imperfect and noting the low significance this determination has on the final result, it is decided that $\hat{\delta}_{CA} = 11.0\text{ s}$ is most likely caused by the path $C \rightarrow A \rightarrow B \rightarrow A$.

$\hat{\delta}_{CB} = 6.6\text{ s}$ matches the relay path $C \rightarrow A \rightarrow B$, that coupled with the conditional TE being insignificant suggests that that is the real path. $\hat{\delta}_{CB} = 9.6\text{ s}$ occurs in both alternate paths but logic dictates that if $C \rightarrow A \rightarrow B$ and $A \rightarrow A$ are both significant, the cause is most likely the path $C \rightarrow A \rightarrow A \rightarrow B$.

The significant TE values at $\hat{\delta}_{AB} = 0.8, 3.4\text{ s}$ appear to be real as they do not match the alternate path. $\hat{\delta}_{AB} = 3.4\text{ s}$ is a reasonable value given the advection speed measured in Chapter 5 and the separation of the airfoils. $\hat{\delta}_{AB} = 0.8\text{ s}$ is also possible because the airfoils are separated by 16 mm at their closest position. It is possible that the communication delay is affected by the amplitude and/or the relative angles of the airfoils. This is not accounted for in any way so there is no way to validate this. $\hat{\delta}_{AB} = 6.4\text{ s}$ is most likely caused by the path $A \rightarrow A \rightarrow B$ as the delays match and conditional TE $\text{TE}_{A \rightarrow B|A}(\delta = 6.4, \zeta = 3.4)$ is insignificant.

³This also implicates $A \rightarrow A$ as the cause of $\hat{\delta}_{AA} = 3.0\text{ s}$. That hypothesis is supported by the lack of matching values for other paths.

Table E.1: Possible relay pathways. A limit is imposed to prevent recurring loops. For example, the path $A \rightarrow B \rightarrow A \rightarrow B$ is excluded. Delays greater than 12s are also excluded. Delays which result in insignificant TE for the corresponding path are indicated by *

| Component | Path | Delays (s) |
|-------------------|---|---|
| $C \rightarrow A$ | $C \rightarrow A$ | 3.2, 6.2, 9.6, 11.0 |
| | $C \rightarrow A \rightarrow A$ | 6.2, 9.2 |
| | $C \rightarrow A \rightarrow B \rightarrow A$ | 4.0, 6.6, 7.0*, 8.3*, 9.6, 10.4*, 10.9, 11.3, 11.8* |
| | $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$ | 7.0*, 9.6, 10.0, 11.3 |
| $C \rightarrow B$ | $C \rightarrow B$ | 6.6, 9.6 |
| | $C \rightarrow A \rightarrow B$ | 4.0*, 6.6, 7.0*, 9.6, 10.4*, 11.8* |
| | $C \rightarrow A \rightarrow A \rightarrow B$ | 7.0*, 9.6, 10.0* |
| $A \rightarrow A$ | $A \rightarrow A$ | 3.0 |
| | $A \rightarrow B \rightarrow A$ | 0.8, 3.4, 5.1*, 6.4, 7.7, 10.7* |
| | $A \rightarrow A \rightarrow B \rightarrow A$ | 3.8, 6.4, 8.1, 9.4*, 10.7* |
| $A \rightarrow B$ | $A \rightarrow B$ | 0.8, 3.4, 6.4 |
| | $A \rightarrow A \rightarrow B$ | 3.8, 6.4, 9.4* |

A reasonable conclusion of this analysis is to summarise the cause of each of the dominant delays identified by the pairwise calculation, and group them to suit the intended purpose of the research. Here the compound communication paths are listed

- $C \rightarrow A$ has real delay of 3.2 seconds. The most likely communication mechanism is advection because those results are similar to those in Chapter 5.
- $C \rightarrow A \rightarrow A$ feeds the information from C back to A with a delay of 3.0s, resulting in a delay of 6.2 seconds for this path. This was shown in Chapter 5 to be caused by the symbol selection so it could be removed although it does provide some insight into the time scale of the airfoil's response.
- $A \rightarrow B$ has a real delay of 3.4 seconds which is consistent with the advection speed previously estimated. It is unclear if this delay exists solely due to its role in the relay paths or if information generated at A is also being communicated at this delay. The 0.8 second delay is not satisfactorily explained. Its most suspicious property is its absence from the relay paths, although it is conceded that the time of the delay is feasible. More testing is required to evaluate this delay.
- $A \rightarrow A \rightarrow B$ has a delay of $\hat{\delta}_{CA} + \hat{\delta}_{AA} = 6.4$ seconds.

- $C \rightarrow A \rightarrow B$ is a simple relay from C to B , through A with a delay of $\hat{\delta}_{CA} + \hat{\delta}_{AB} = 3.2 + 3.4 = 6.6$ s.
- $C \rightarrow A \rightarrow A \rightarrow B$ is similar to the above relay but with the $A \rightarrow A$ feedback included, resulting in a delay of 9.6 seconds.
- $C \rightarrow B$ likely has no direct communication although this statement cannot be made definitively. Both of the detected communication delays can be explained by the alternate paths above and independent testing showed that if airfoil A is fixed, no significant communication is seen $C \rightarrow B$.
- $B \rightarrow A$ has delays of 0 and 4.3 seconds. The 0s delay is consistent with pressure communication but the physical mechanism that leads to the 4.3 second delay has not been identified but its effect is seen in the path $C \rightarrow A \rightarrow B \rightarrow A$ so it is suspected to be real. It does not meet the criteria for being an artefact of mutual influence of C on A and B but it may be part of a higher order communication path. From knowledge of the system, it is expected that no information is generated at B so, at the discretion of the researcher, this direct path could be removed.⁴
- $C \rightarrow A \rightarrow B \rightarrow A$ combines the relay $C \rightarrow A \rightarrow B$ and the direct interaction $B \rightarrow A$ resulting in delays of 6.6 and 11 seconds. The 11 s matches one of the delays previously identified by the pairwise calculation as being directly $C \rightarrow A$. The 6.6 s delay does not match a peak but could still be significant. This point is discussed in a later section.
- $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$ adds the $A \rightarrow A$ feedback loop to the path above resulting in delays of 9.6 and 14 seconds. The 14s delay is ignored because the calculation stops at $\delta = 12$ s. The 9.6s delay is the last of the delays previously identified by the pairwise calculation as being directly $C \rightarrow A$.

From the information presented here, a map of the communication pathways is created (Fig. E.1) based on the philosophy, any marked path would still exist if any variable not included in that path were removed. For example, the path $C \rightarrow B$ only exists because of A , so it is not included. Conversely, A likely generates information from the fundamental shedding of the cylinder⁵, so the path $A \rightarrow B$ is kept, both as a direct link and in the relay paths.

⁴The effect of this delay is still accounted for in the relay path, so no information is discarded.

⁵It is important to remember that the cylinder structure is a separate concept from the variable C which is a representation of the rotation events of the cylinder. The information generated by A is in response to the cylinder structure, the information relayed through A is in response to the rotation events C .




Figure E.1: One interpretation of the communication pathways of the three structure system.

This is just one interpretation of the results, other simplifications can be made by considering which information is important to the research question. Here the feedback loop $A \rightarrow A$ can be reinterpreted as two responses of A to C , or it could be removed all together by acknowledging that it is a construct of the dynamics of the response, thus, not part of the information transfer in the system. Fig. E.3 uses this logic argument to create a more digestible map of the communication pathways⁶. That figure appears to be simpler but each path has multiple delays associated with it which may make it difficult to draw conclusions.




Figure E.2: Another interpretation of the communication pathways of the three structure system, assuming that the feedback $A \rightarrow A$ is part of the system dynamics rather than information transfer.

⁶Both methods result in the same map but the associated delays are different.

For this system, the most useful presentation of the information is the longest pathway, $C \rightarrow A \rightarrow A \rightarrow B \rightarrow A$. The delays of each section of the pathway represent the real (physically explainable) communication delays between those structures. This removes the complication of having multiple paths between the same structures and, as long as it is noted that direct communication can occur concurrently with any section of the relay path, all of the information is preserved.

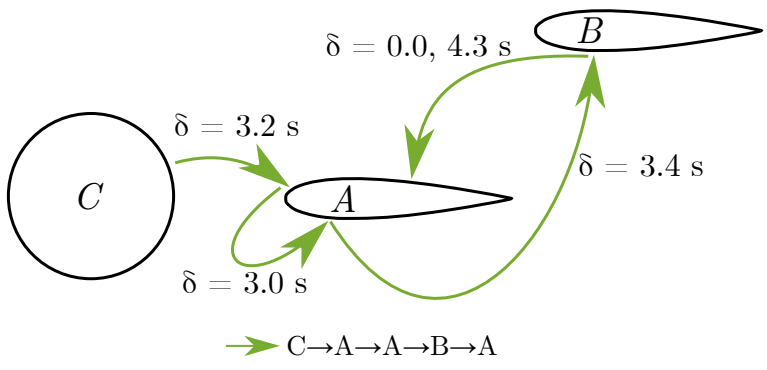


Figure E.3: Another interpretation of the communication pathways of the three structure system, assuming that the feedback $A \rightarrow A$ is part of the system dynamics rather than information transfer.