Transfer Dynamics Cancellation in Real-Time Dynamic Substructuring

J. L. du Bois, B. Titurus, N. A. J. Lieven
University of Bristol, Department of Aerospace Engineering
Queens Building, University Walk, Bristol BS8 1TR, U.K.

Abstract

Real-time dynamic substructuring, also referred to as mechanical hardware in the loop, model in the loop, or hybrid numerical-physical testing, is gaining increasing interest from industrial sectors as a practical means of testing components of large structures under realistic operating conditions. Such tests commonly employ hydraulic actuation with force or displacement feedback to couple the physical substructure to the numerical model of the system being tested. A significant challenge in this type of configuration lies in the cancellation of the transfer system dynamics. The control electronics and hydraulics as well as the measurement equipment and feedback mechanisms all introduce their own dynamics to the system which need to be compensated for in order to provide meaningful results. This paper details efforts that have been made to model the transfer dynamics of a hydraulic actuator and its controllers and to employ these models in the cancellation of the transfer dynamics. Several models are proposed, ranging from first order transfer functions to more complete process models, and their performance is assessed in qualification tests as well as simple substructuring exercises.

1 Introduction

Real-time dynamic substructuring is an experimental technique which is rapidly attracting interest in a number of fields. Improvements in the availability of actuation hardware, computing power and control techniques have recently allowed significant advances in the fidelity of these hybrid dynamic simulations.

The technique variously goes by the names “experimental dynamic substructuring”, “mechanical hardware-in-the-loop testing”, “model-in-the-loop testing”, “real-time hardware-in-the-loop testing”, “real-time pseudo-dynamic testing”, “hybrid dynamic testing”, and assorted permutations of these terms. Drawing from two established experimental testing platforms, hardware-in-the-loop (HiL) and substructured pseudo-dynamic testing (PDT), the method combines the advantages of numerical modelling and simulation with those of experimental testing to allow large and complex structures to be tested in the laboratory under realistic operating conditions.

The principle behind the method is to split a structure into two or more substructures. At least one of these substructures will be a physical piece of hardware while the remainder of the substructures will be simulated numerically. The physical substructure(s) will either be a critical part, where its exact performance is to be studied, or will be a complex system which is difficult to model numerically. Alternatively it may simply have unknown characteristics which are most easily determined experimentally. In contrast, the pertinent features of the numerical substructure will generally be well understood and easily modelled. Reasons for omitting these parts from a physical study are usually based around the logistics of laboratory testing: the parts could be too large to fit in a laboratory, too expensive, or they could have demanding load and displacement requirements for which the test equipment is unavailable. The latter may include distributed forces, coupled aero- or hydro-dynamic forces, or simply forces and displacements which are too large or too numerous for the available test equipment.

Real-time dynamic substructuring differs from HiL and substructured PDT in the critical nature of the experimental coupling between the numerical and physical parts. For example, HiL is used to evaluate real electronic hardware with simulated operating conditions. Here, the coupling between the two systems is comprised of electrical connections, which will closely match those of the in-service equipment. There are no significant elements in the interface which will introduce spurious behaviour. Like real-time dynamic substructuring, PDT is concerned with structural dynamics, where instead of electrical signals, forces and displacements are passed between the substructures. Here the interface is comprised of force and displacement transducers, with actuation devices (typically hydraulic or electromagnetic) providing displacement on the physical side. The control and response of the actuators, combined to some extent with the signal processing from the transducers, introduce significant dynamics into the coupled physical-numerical system. In PDT the tests are not conducted in real time, and these dynamics can effectively be eliminated. In real-time dynamic substructuring, however, the real-time response is critical and the extra dynamics will have a detrimental effect on the accuracy of the test results. Compounding this effect, the extent to which the results are affected is very difficult to determine without a reference system for comparison (and the existence of such a system would obviate the need for the real time dynamic substructuring).

Early attempts to counter the problems of the transfer dynamics centred around compensating for a perceived delay in the system, as in the experiments of Horiuchi *et al.* [1] and Darby *et al.* [2]. Wallace *et al.* [3] extended this idea in the development of an adaptive delay and gain compensation method, using polynomial based forward prediction to counter the delay, and Bonnet *et al.* [4] performed similar tests for a range of delay compensation strategies. Other studies have used used lag compensation instead, inverting a first order transfer function to form a feedforward controller [5]. These efforts produce measurable performance improvements, but they are ultimately limited due to the nature of the dynamics of the actuators, valves and control hardware. The behaviour of these components can not be fully represented by a simple delay or first order lag. Plummer [6] develops a more appropriate model for a hydraulic actuation system, resulting in a fifth order transfer function plus delay. He also comments on sources of nonlinearity including the valve flow gain and the displacement sensitivity of the main cylinder stiffness. While comprehensive, such a model requires extensive resource to characterise all of the components, and it may prove difficult or practically impossible to measure some of the necessary quantities in the course of a real time substructuring exercise.

Standard feedback control approaches may also be taken, minimising the error between the demand signal and the measured load or displacement. Proportional, integral and derivative (PID) controllers are commonly employed in this capacity as the proprietary controllers supplied with the actuation equipment. These often form an inner loop controller with further control applied as an outer loop. Other linear state feedback approaches such as LQR are equally applicable here.

The topology of the system lends itself immediately to model reference adaptive control (MRAC), as the physical system must follow the dynamics of the numerical model as closely as possible. Stoten and Hyde [7] use minimal control synthesis (MCS), a variant of the MRAC approach, to create an adaptive linear controller with state error feedback and forward loop gain. This effectively combines the benefits of the feedforward and feedback systems discussed above with an adaptive scheme to tune their parameters. The MCS algorithm has the additional advantage of being asymptotically stable.

The work contained herein starts from the premise that successful real-time dynamic substructuring needs to be supported by a high-performance feedforward controller. While error feedback can be used to take up the slack, reducing errors from disturbances (both random noise-based and systematic), the objective as far as possible is that there are no errors to correct. The feedback of physical measurements into the numerical system means that it is difficult to create a feedback controller whose dynamics will not interact with those of the system being simulated. Thus it is desirable to reduce the influence of such dynamics from the outset. Specifically, it is important that a suitable topology is chosen for the feedforward controller, even where adaptive control is used. This paper investigates several models of the transfer system dynamics, incorporating first and second order transfer functions as well as delay, and uses these to create open loop controllers to minimise the forward path errors. The models span a void between the simple lag or delay models described above and the fifth order plus delay described by Plummer [6]. They are supplemented by

a feedback controller in the form of the proprietary PID controller.

In the next section the equipment is described, highlighting some interesting features of the test piece and offering a detailed investigation of the characteristics of the complete physical rig. In Section 3, parametric models of the transfer system dynamics are fitted to the test data and these are used to create feedforward controllers for transfer dynamics cancellation. Their performance is evaluated using prescribed test signals. A real-time dynamic substructuring framework is then laid out in Section 4, and the controllers' efficacy is assessed.

2 Equipment Characterisation

2.1 Description

The physical substructure used for these studies is a hydraulic orifice damper. It was chosen arbitrarily, but has several features which make it interesting in this context. Firstly, it is relatively 'stiff' with respect to velocities, such that relatively small velocities can require high forces. This means that the limits of the actuator's performance will be tested, and the control problem becomes more critical. Secondly, it has a nonlinear force-velocity relationship due to the turbulent oil flow induced by the small piston orifice. Thirdly, the high forces are accompanied by high fluid pressures which lead to significant compressibility in the damper. Fourthly, the device includes poppet valves which open up additional, parallel flow paths when the pressure difference exceeds a given threshold. Finally, operational wear has lead to the development of a non-trivial amount of backlash, or free play, in the bearings used to mount the damper. All of these factors combine to produce a device with parametric uncertainty and a full state dependent input-output relationship. The damper features are illustrated in the experimental data in Fig. 1, produced in response to a nominally sinusoidal piston displacement. The data are normalised for commercial reasons. As well as the features described in the caption, it can be seen that the response is skewed from a symmetric profile (as would be seen for a purely velocity-sensitive device) by compressibility and elasticity within the damper.

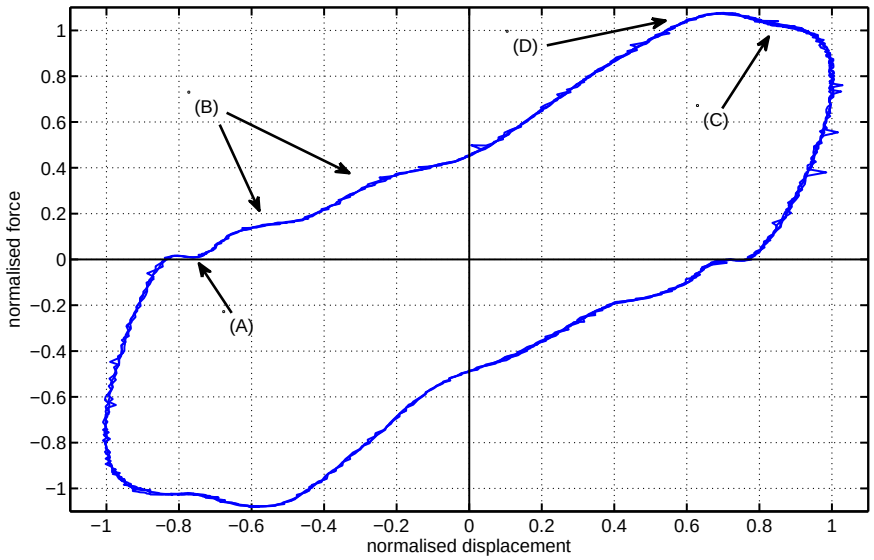

Figure 1: Features of the damper manifested in experimental data: (A) backlash; (B) transient effects due to force/pressure reversal (could be induced by PID controller response to backlash, compressibility effects, or cavitation); (C) poppet valves limiting damping forces; (D) pressure overshoot on opening of poppet valves.

The interface between the damper and the numerical substructure is facilitated by a Hydropuls^{TM} 25kN servohydraulic actuator. The damper is mounted in line with the actuator, in a frame comprised of a rigid steel floor, two heavy duty angle brackets and an I-beam. The displacement of the actuator piston is measured by means of an internal linear variable differential transformer (LVDT), and the force is measured using a 25kN force transducer forming the coupling between the actuator piston and the damper rod. An Instron Labtronic 8800 provides the inner-loop control of the actuator, using a serial PID controller in displacement control mode, shown in Fig. 2. The PID gains used for the results in this section are proportional: K_P =30dB, integral: K_I =0.5s^{-1}, derivative: K_D =0s. The outer-loop control is provided by a dSpace DS1103 board, with analogue connections to the Labtronic 8800. In the first instance this outer-loop control consists purely of displacement demand signals, with no feedback. In Section 3, feedforward control is added, followed in Section 4 by substructuring feedback.

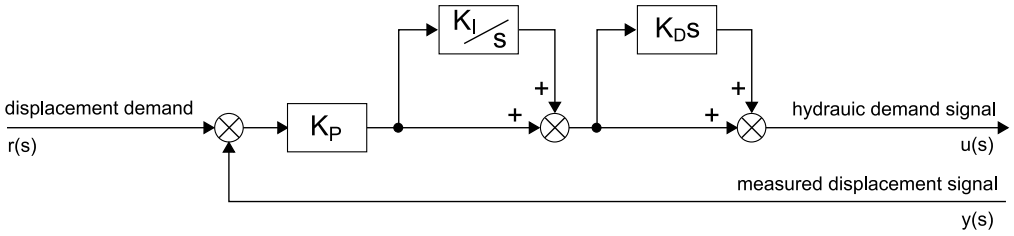

Figure 2: Serial PID controller used for inner loop displacement control.

2.2 Dynamic Response

The displacement response is influenced not only by the tuning of the inner-loop PID controller, but also by the dynamics of the servovalves, pistons, and hydraulic power train of the actuator. To a lesser extent, the electrical systems will also contribute to errors in the displacement tracking in the form of latency and discretisation. To quantify the response of the actuator, a set of single-harmonic displacement demands are applied and the actual displacement is measured for each. The harmonic input signals cover a range of frequencies and amplitudes. In each case the response is given time to settle to a steady periodic state before the measurements commence. A sampling frequency of 1kHz is used throughout.

Initially, a linear response is assumed and a fine mesh of data points is used to determine transfer function surfaces for the system. This identification is performed both with the damper removed and with the damper installed using two different orifice configurations: the standard orifice and a new, larger orifice. The transfer functions can be seen in Fig. 3, with the amplitudes normalised to the test range (suffice to say that displacements are of the order of millimetres). The empty rig exhibits a more or less flat response surface to around 50Hz, beyond which it resembles a heavily damped second order response. On first inspection it does not vary significantly with the amplitude of the harmonic excitation. Similarly, the damper with the large orifice demonstrates second order characteristics with no pronounced variation over the range of amplitudes tested. In this case, somewhat surprisingly, the response is less damped. It appears that the large orifice does not offer significant damping, while the overall characteristics of the damper contribute to the resonance of the system. Using the damper with the standard, small orifice produces a more damped response again at low amplitudes. This time, however, a markedly different response is seen at higher amplitudes, demonstrating pronounced nonlinearity.

Fig. 4 gives a clearer picture of the behaviour of the rig with the standard damper at higher amplitudes. Supporting the indisputable evidence of nonlinearity already seen, the curves in Fig. 4 show a behaviour in keeping with a nonlinear spring hardening characteristic. It is found that the onset of the high amplitude response coincides with the activation velocity for the poppet valves in the damper, which lower the damping constant. This feature would be expected to produce an effect similar to the observed spring hardening

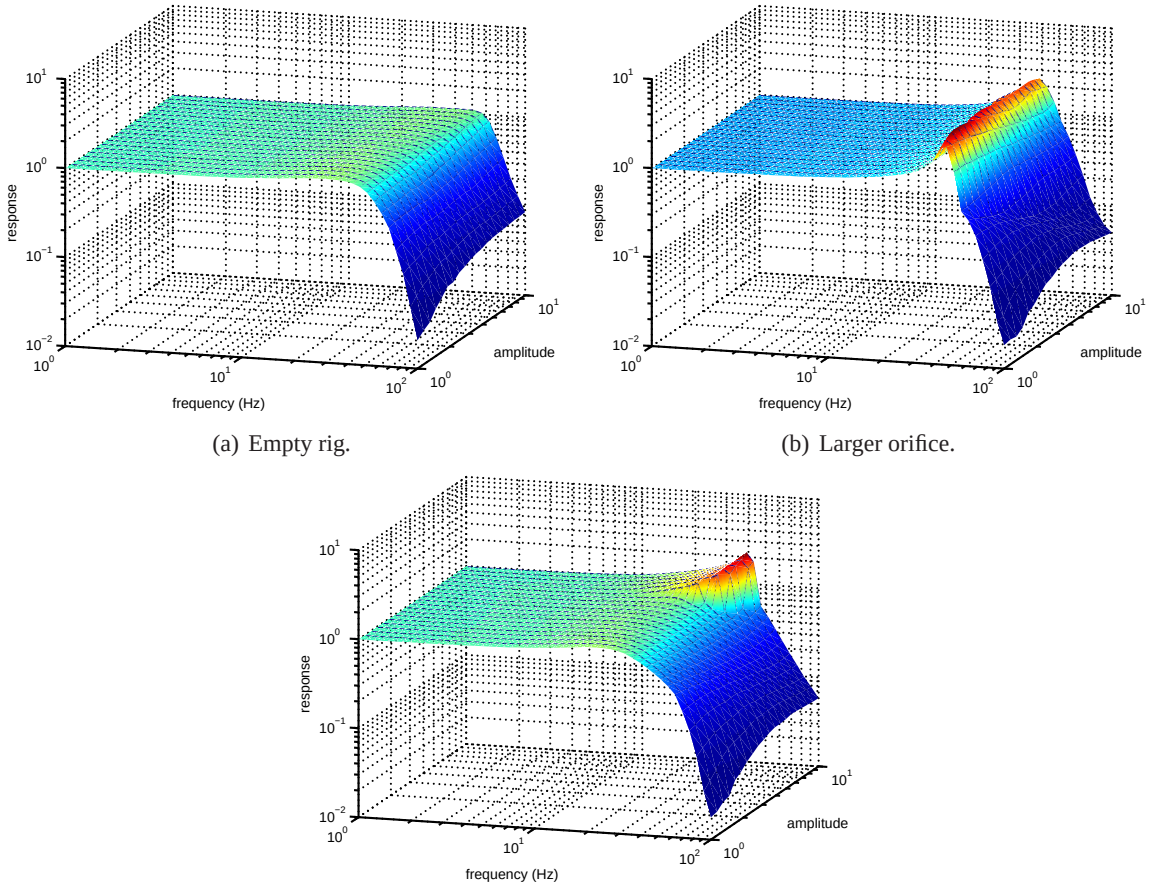

(a) Empty rig.

(b) Larger orifice.

(c) Standard damper.

Figure 3: Transfer function surfaces for three different rig configurations: one with no damper mounted, one with the standard damper mounted, and one with a modified damper with larger orifice.

characteristic. Another interesting feature of Fig. 4 is that while the gain curves appear to show a second order response, the phase curves extend well beyond the -180° that would be expected of a second order system. This is discussed in more depth in Section 3.1.

Further information regarding the nature of the nonlinear response is obtained through examining the power spectral densities (PSDs) of the measured displacement signals for each of the demand signal frequencies. Each of the measured response signals are split into 8 segments, windowed with a Hamming function and a fast Fourier transform (FFT) is applied to produce PSD surfaces showing the response spectrum across the excitation range. Fig. 5 offers some illuminating examples. In each, the dominant response is the excitation frequency, but the harmonics of this frequency also show clearly elevated response levels. In the low amplitude examples, there is little response at any other frequency. This is particularly true for the empty rig, where even the harmonics are hard to discern. The disturbances to these signals are therefore periodic distortions of the waveform, with the empty rig demonstrating a near perfect reproduction of the demand signal. The high amplitude examples show much greater pollution of the frequency spectrum, with strong harmonic components in addition to more general excitation. It seems that the level of structural response in the whole rig is so large as to violate linear assumptions. In particular, the distinct nonlinear curves depicted in Fig. 4 are associated with a strong band at around 40Hz which spans the whole response spectrum.

The important observations from this section are:

- Firstly, the response of the hydraulic actuator to displacement demands is heavily dependent upon the physical substructure being tested.

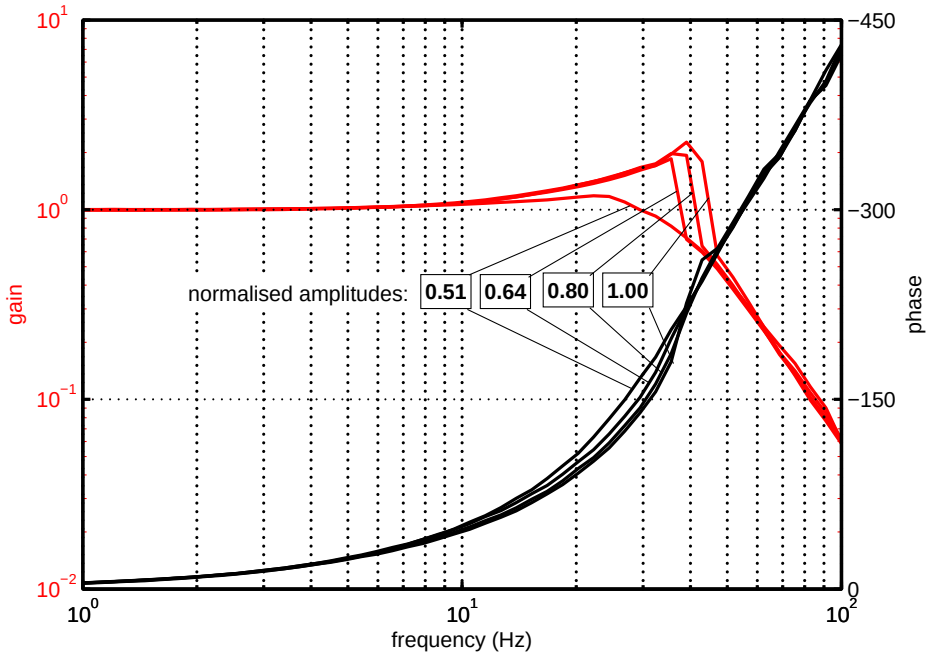

Figure 4: Transfer functions of the rig with the standard damper mounted, showing the onset of nonlinear artefacts.

- In addition to this, nonlinearities in both the rig and the test piece can be significant so linear control strategies may be inadequate.
- Finally, nonlinearity needs to be considered not only in the context of the macroscopic signal amplitude, but also in terms of the signal distortion throughout the piston stroke.

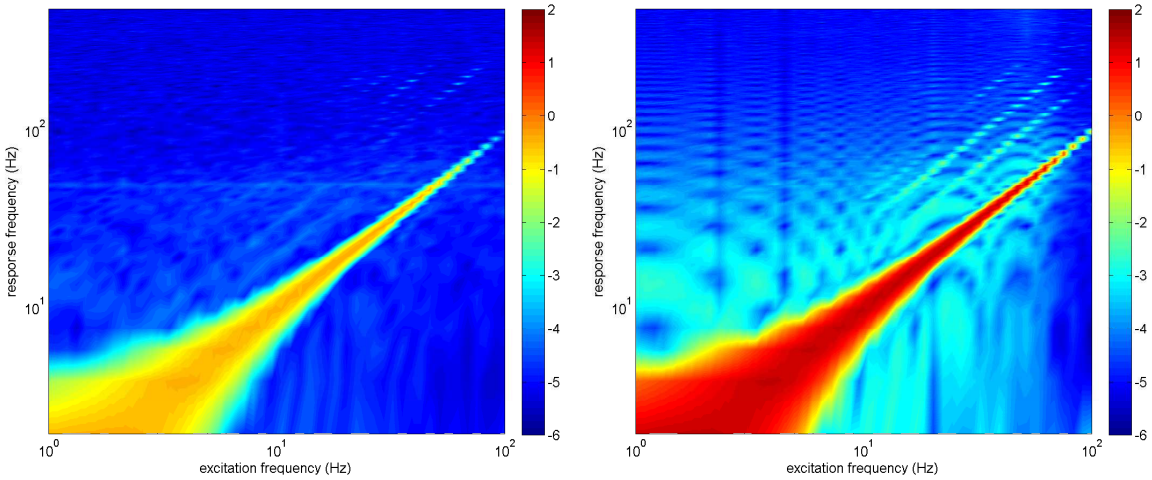

(a) Empty rig, small amplitude.

(b) Empty rig, large amplitude.

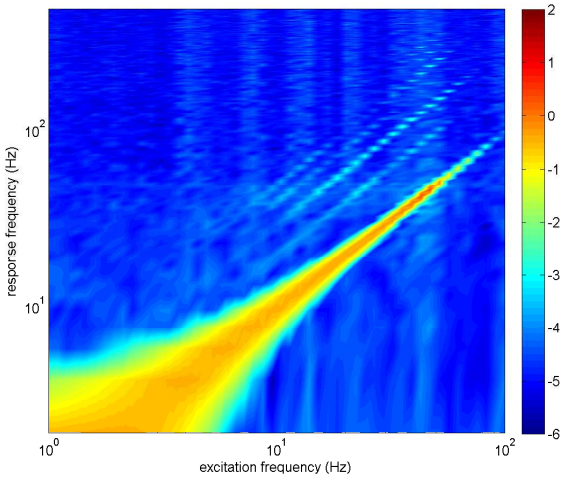

(c) Larger orifice, small amplitude.

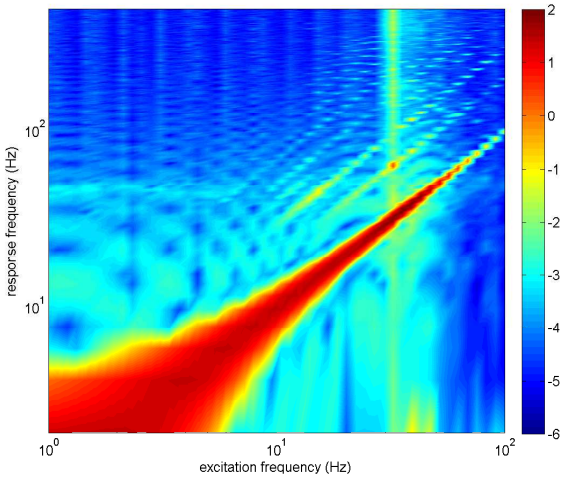

(d) Larger orifice, large amplitude.

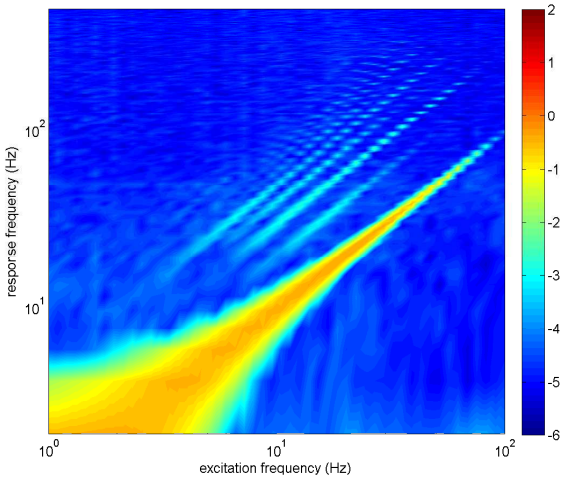

(e) Standard damper, small amplitude.

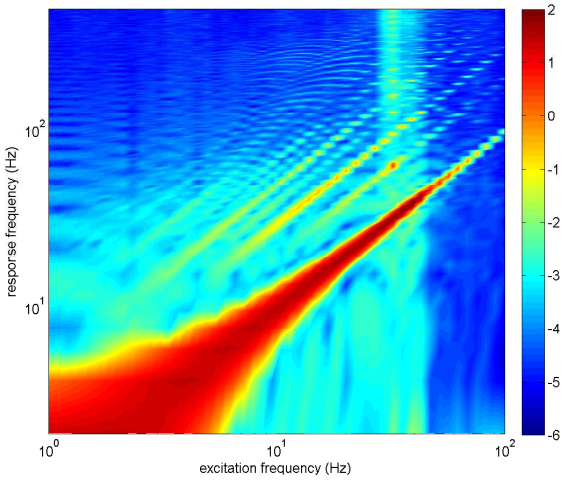

(f) Standard damper, large amplitude.

Figure 5: Power spectral densities for the three different rig configurations, given for the smallest and largest amplitude excitations in the test range. The colourbar scale is logarithmic, base 10.

3 Transfer Dynamics Cancellation

3.1 Process Model Identification

In this section, approximate models are derived for the transfer dynamics: the dynamics of the hydraulic actuator, its controllers, and its interaction with the test piece. Ideally the outer loop controller for the actuator would be capable of operating independently of the test piece dynamics, but this capability is well beyond the scope of the current work.

From the investigations of the foregoing chapter, the transfer dynamics gain can for the most part be represented by a second order transfer function. In some cases, a first order approximation may suffice. It was also noted, however, that the phase shifts through more than 180° over the frequency range, indicating that higher order dynamics are in fact present. One method for incorporating the effect of these higher order dynamics is to augment the transfer functions with a pure delay. Correspondingly, the four process models investigated here are a first order transfer function ($G_1(s)$), a second order transfer function ($G_2(s)$), a first order transfer function with delay ($G_{1\mathrm{D}}(s)$) and a second order transfer function with delay ($G_{2\mathrm{D}}(s)$). These are represented in the frequency domain as:

$$\begin{aligned} G_1(s) &= \frac{K}{\tau_p s + 1} & G_2(s) &= \frac{K}{s^2 + 2\zeta\omega_0 s + \omega_0^2} \\ G_{1\mathrm{D}}(s) &= \frac{K e^{-\tau_d s}}{\tau_p s + 1} & G_{2\mathrm{D}}(s) &= \frac{K e^{-\tau_d s}}{s^2 + 2\zeta\omega_0 s + \omega_0^2} \end{aligned} \tag{1}$$

Each process model has between two and five parameters which need to be identified, from the following selection: K is the DC gain, τ_p is the first order time constant, ζ is the damping ratio, ω_0 is the undamped natural frequency, and τ_d is the time delay. These parameters are determined from the empirical transfer function estimates using a nonlinear optimisation algorithm, where a time domain error-based cost function is minimised with respect to the model parameters.

Before any outer loop control was developed, the PID controller was tuned manually to provide the optimum response with the damper in place. The PID gains used for the remainder of the tests documented here are proportional: K_P =30dB, integral: K_I =0.2s^{-1}, derivative: K_D =3.6ms. New transfer functions are identified, once again using a sampling rate of 1kHz and identifying the Fourier coefficients of the input and response signals over an integer number of periods at each excitation frequency. It is important to realise that these are only transfer function approximations to the full, nonlinear system response. It is because of the system nonlinearity that the stepped sine excitation approach is adopted (in favour of broadband, chirp or sweep excitation, for example). Four amplitudes are tested but this time the high response regime is avoided, with the demand signals covering only half the amplitude range of the previous tests.

The identified process models are compared with the empirical transfer functions in Fig. 6. Each figure displays the process model being used and the parameters used to obtain the best fit curves. The first order transfer function is the least effective at reproducing the measured data, with the phase diverging at 10Hz and the gain starting to drop before that. Even at low frequencies the gain shows discrepancies. The second order transfer function performs better, with the phase in close agreement until around 30Hz and the gain starting to diverge at around 10Hz. Despite these limited correlations, it is clear that neither function accurately reproduces the transfer dynamics. When a delay is added to the process model, however, the correlation is dramatically improved, and both the first and second order process models with delay offer a convincing embodiment of the significant dynamic features.

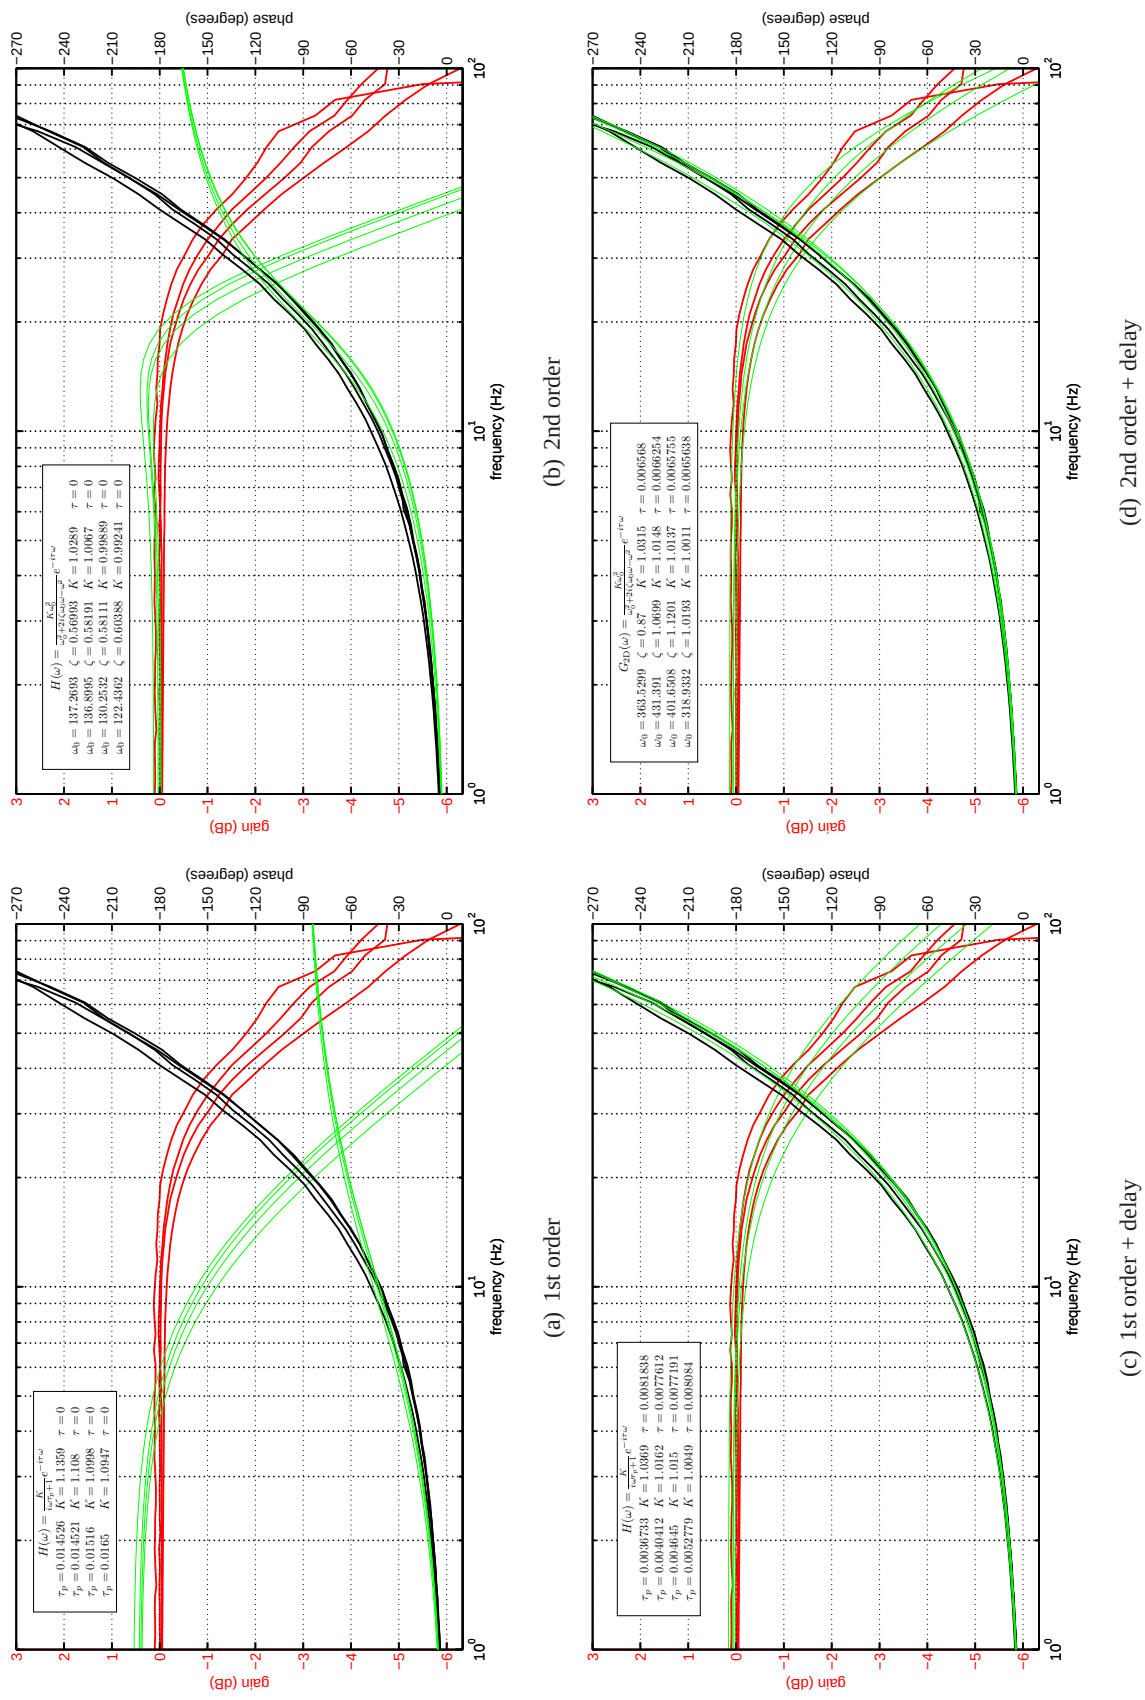

Figure 6: Process model identification for the transfer system dynamics. The parametric transfer functions are shown in green, corresponding to the process models given in the key. The highest amplitude demand signals correspond with the first gains to drop off and the largest phase lags.

3.2 Process Model Inversion

To take advantage of these models in an open-loop controller, they must be inverted. In theory a transfer function can not be inverted for control applications if the denominator is of a higher order than the numerator. In practice, however, approximately numerical evaluation of such functions is relatively easy. The difficulty lies in the inversion of a delay. There is a wealth of control literature on the topic of delay compensation. Chen and Ricles [8] give a brief review of available strategies with emphasis on those currently used in real time dynamic substructuring, and develop a methodology for the analysis of such systems. Here, the polynomial forward predictive algorithm used by Wallace *et al.* [3] is employed, and combined with the inverse first- and second-order transfer functions.

At each time step the forward predictive algorithm takes a sample set of data points from the displacement demand history, and uses these to fit a least squares polynomial curve. The polynomial is evaluated at a forward time specified by the identified delay τ_d, and the resulting signal is used as the input to the inverse transfer function. The parameters which can be used to tune the performance of the predictive stage are the sampling period, T_s, spanned by the data sample set, the number of points, n_s, used for the least-squares fit, and the order of the polynomial. For these tests a sixth order polynomial was chosen arbitrarily. The effects of changes to the sampling period and number of data points are investigated later, although the number of data points was ultimately limited by the minimum time step that could be achieved with the dSpace controller as configured.

The sixth order polynomial identification is obtained from

$$\begin{Bmatrix} a_6 \\ a_5 \\ a_4 \\ a_3 \\ a_2 \\ a_1 \\ a_0 \end{Bmatrix} = \begin{bmatrix} (-(n_s-1)\Delta t)^6 & (-(n_s-1)\Delta t)^5 & \dots & -(n_s-1)\Delta t & 1 \\ (-(n_s-2)\Delta t)^6 & (-(n_s-2)\Delta t)^5 & \dots & -(n_s-2)\Delta t & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix}^+ \begin{Bmatrix} x_{-(n_s-1)\Delta t} \\ x_{-(n_s-2)\Delta t} \\ \vdots \\ x_0 \end{Bmatrix} \tag{2}$$

where x_t is the demand signal at time t relative to the current time, $\Delta t = T_s/(n_s - 1)$ is the time step used for the data sampling, and $+$ denotes a pseudo-inverse. Note that the pseudo-inverse only needs to be performed once and can be reused at each time step. The forward predicted signal is then given by

$$x_{\tau_d} = a_6\tau_d^6 + a_5\tau_d^5 + a_4\tau_d^4 + a_3\tau_d^3 + a_2\tau_d^2 + a_1\tau_d + a_0. \tag{3}$$

Inverting the first and second order transfer functions from Eqn. (1) and converting back to the time domain gives

$$r = \frac{1}{K}(\tau_p\dot{x}_{\tau_d} + 1) \qquad r = \frac{1}{K}(\ddot{x}_{\tau_d} + 2\zeta\omega_0\dot{x}_{\tau_d} + \omega_0^2 x_{\tau_d}) \tag{4}$$

and the derivatives of the demand signal are estimated as

$$\dot{x}_{\tau_d-\delta/2} = x_{\tau_d} - x_{\tau_d-\delta} \tag{5}$$

$$\dot{x}_{\tau_d-3\delta/2} = x_{\tau_d-\delta} - x_{\tau_d-2\delta} \tag{6}$$

$$\ddot{x}_{\tau_d} \approx \ddot{x}_{\tau_d-\delta} = \dot{x}_{\tau_d-\delta/2} - \dot{x}_{\tau_d-3\delta/2} \tag{7}$$

$$\dot{x}_{\tau_d} = \dot{x}_{\tau_d-\delta/2} + \ddot{x}_{\tau_d-\delta}\delta/2 \tag{8}$$

$$\tag{9}$$

where δ is the controller time step. These derivatives could equally be found using the derivatives of Eqn. (3). Measurement noise should not cause difficulties with the derivative calculations as the displacement demand, x, is either a smooth prescribed signal or it comes from a numerical substructure model which serves as a

filter for the physical measurements. To complete the control, the signal r is passed to the inner-loop PID controller.

The inverted process models are tested by applying the same set of harmonic demand signals as used in the identification stages above. For these tests, the forward predictive algorithm is not necessary: the test signals are known *a priori*, and so the delay compensation consists simply of time shifting the known signal by τ_d. The transfer functions produced are seen in Fig. 7. The performance of the first order model is the least impressive, with the gain deviating from unity and the phase increasing from zero at the same points that the parametric model deviates from the measured data in Fig. 6(a). In this regard, however, it has performed as well as can be expected. The second order model shows better performance and corresponds equally well with the behaviour expected from the identification procedure. The two delayed models show exceptional performance, with gains close to unity and phase close to zero throughout most of the test regime. A surprising result is that the first order plus delay model performs better than the second order plus delay model, with the latter's gain deviating more significantly at the high end of the frequency spectrum.

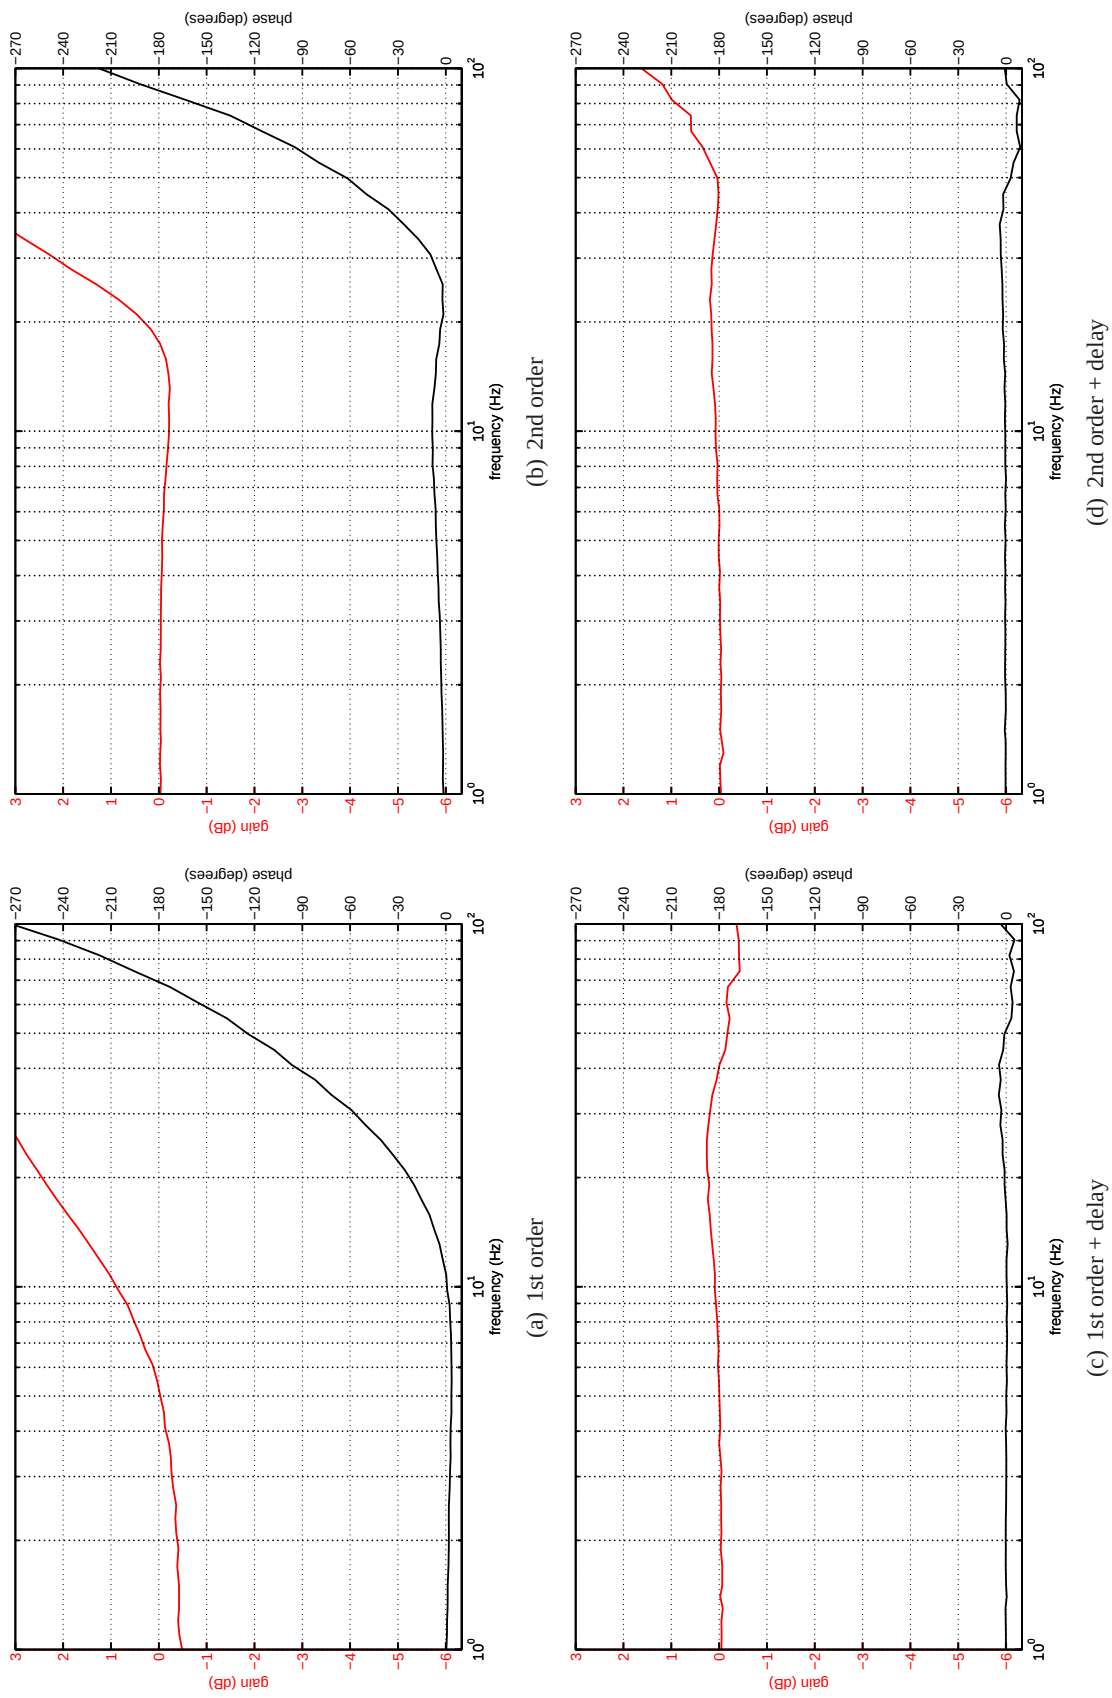

(a) 1st order (b) 2nd order

(c) 1st order + delay (d) 2nd order + delay

Figure 7: Transfer functions relating demand displacement to measured displacement with the inverse process models augmenting the transfer system dynamics.

4 Real-time Dynamic Substructuring

The foregoing sections detail the identification of appropriate process models for the transfer dynamics, and the implementation of these models in open-loop controllers to reduce the influence of the transfer dynamics. In this section, a real-time dynamic substructuring exercise is undertaken to evaluate the performance of the inverse process models in practice. The numerical substructure developed here is simple, being comprised of only two degrees of freedom (DOFs). Nonetheless, this setup introduces challenges beyond those faced in the earlier tests:

- The coupling between the numerical simulation and the physical system is now full and reciprocal. So far the physical displacements have been determined by the numerical controllers but there has been no feedback. Now the measured forces at the interface are fed back into the numerical system. In this situation, significant lag in displacement tracking can produce artificial negative damping in the numerical-physical system, leading to instability and divergence.
- Because the dynamics of the hybrid system are now influenced by the physical forces, the nonlinearities and discontinuities of the real hardware will produce distortions to the sinusoidal numerical excitations. It is unknown to what extent these will affect the linear assumptions made in the transfer dynamics cancellation methods.
- The real-world feedback also means that the displacement demand signal is no longer predetermined. In the foregoing tests the advance signal used to compensate for delay was simply a time-shifted duplicate of the known demand signal. Here, the forward predictive algorithm will be put through its paces.

4.1 Numerical Substructure

The numerical substructure consists of the two-DOF lumped spring-mass-damper system seen in Fig. 8. In addition, a model of the physical substructure is created, so that the two can be interchanged in operation. The damper model is discussed further in the next section. The numerical substructure is represented by two second order transfer functions, one for each mode of the two-DOF system. This approach has been adopted to provide a framework to allow the future simulation of larger systems, which may have many degrees of freedom but only a limited subset of significant modes to be modelled. When the dSpace controller code is compiled, the transfer functions are transformed into a state space representation to be solved using a 4th order explicit Runge-Kutta ordinary differential equation (ODE) solver.

The creation of the transfer functions begins with a finite element (FE) model of the structure, in the form of the mass, stiffness and damping matrices, $\mathbf{M}$, $\mathbf{K}$ and $\mathbf{C}$. The equations of motion are given by

$$\mathbf{M}\ddot{x}(t) + \mathbf{C}\dot{x}(t) + \mathbf{K}x(t) = f(t) \tag{10}$$

where $f(t)$ is the external force vector, $x(t)$ is the displacement vector, and $\dot{x}$ and $\ddot{x}$ are its first and second derivatives respectively. Because of the second order transfer function system adopted, only real mode shapes can be accommodated and the FE model must be proportionally damped, so that

$$\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}. \tag{11}$$

The displacement vector can be expressed in terms of the modal displacements, $q(t)$, such that

$$x(t) = \mathbf{\Phi}q(t) \tag{12}$$

where $\mathbf{\Phi}$ is a matrix whose columns are the system eigenvectors. Substituting Eqn. (12) into Eqn. (10), premultiplying by $\mathbf{\Phi}^T$ and using Eqn. (11) along with the orthogonality conditions gives a diagonal matrix equation comprising the independent modal equations of motion:

$$\mathbf{I}\ddot{q}(t) + (\alpha\mathbf{I} + \beta\mathbf{\Lambda})\dot{q}(t) + \mathbf{\Lambda}q(t) = \mathbf{\Phi}^T f(t) = p(t) \tag{13}$$

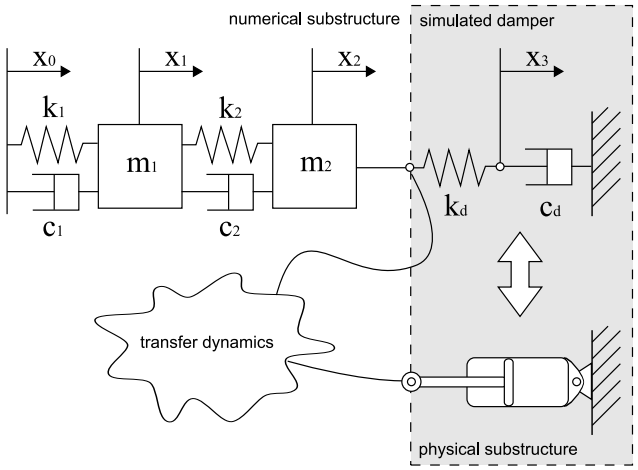

Figure 8: The numerical substructure, pictured with the physical and simulated physical substructures.

where $\mathbf{\Lambda}$ is a diagonal matrix of eigenvalues, $\lambda_{ii} = \omega_{ii}^2$, ω_{ii} is the undamped natural frequency of mode i and $p(t)$ is a vector of modal forces. From Eqn. (13), the relevant modal equations can be extracted and converted to transfer functions of the form

$$\mathrm{H}_i(s) = \frac{q_i(s)}{p_i(s)} = \frac{1}{s^2 + (\alpha + \beta\lambda_i)s + \lambda_i}. \tag{14}$$

The numerical substructure's algorithm is then comprised of three stages: firstly, convert the external forces to modal forces with $p_n(t) = \mathbf{\Phi_n}^T f(t)$ (N.B. $f(t)$ includes both the structural excitation source and the force feedback from the physical substructure). Secondly, determine the modal displacements from the transfer functions $\mathrm{H}_i(s)$. Thirdly, convert the modal forces back to spatial coordinates using $x(t) = \mathbf{\Phi_n} q_n(t)$. In these stages, $\mathbf{\Phi_n}$, $p_n(t)$ and $q_n(t)$ are reduced matrix and vectors corresponding to the modes of interest. Depending on the forcing applied and the displacement outputs needed, only a select few rows of $\mathbf{\Phi_n}$ may be required.

If a displacement excitation is to be used, the corresponding row and column is removed from Eqn. (10) (as the force at this DOF is now irrelevant), and $f(t)$ is augmented by a force due to the motion of the input DOF:

$$f(t) = f_{-a}(t) - M_a\ddot{x}_a - C_a\dot{x}_a - Kx_a \tag{15}$$

where f_{-a} is the external forcing vector with the a^{th} coordinate removed, M_a, K_a and C_a are the a^{th} columns of the mass, stiffness and damping matrices and x_a is the a^{th} displacement coordinate. If necessary, it is then a simple matter to compute the force at the a^{th} DOF from the a^{th} row of Eqn. (10).

4.2 Physical Substructure Model

A model of the physical damper is included to help provide stability to the system and to allow a smooth transition to the fully coupled state once the desired operating conditions are attained. The switch from fully numerical to hybrid physical-numerical operation is explained in Section 4.4 below. The damper is modelled as a symmetric piston with compressible fluid, and orifice flow which can be characterised as a serial laminar-turbulent flow path. The model has a single state variable in the form of the pressure difference between the chambers, Δp, and the fluid properties are treated as lumped parameters for each chamber. From reference [9] the model can be written

$$\Delta\dot{p} = \frac{B_{eff}(V_1 + V_2)}{V_1 V_2}\left[A\dot{y} - \frac{\mathrm{sign}(\Delta p)}{2c_2}\left\{-c_1 + \sqrt{c_1^2 + 4c_2|\Delta p|}\right\}\right] \tag{16}$$

where B_{eff} is the effective fluid compressibility (also accounting for elasticity of the damper body), V_1 and V_2 are the fluid volumes in chambers 1 and 2, respectively, A is the piston area, $\dot{y}$ is the piston velocity, and c_1 and c_2 are the laminar and turbulent flow coefficients. Most of these quantities are measured from the physical device, while the flow coefficients are determined experimentally. Assuming small piston motions about a central position, $V_1 \approx V_2 \approx V_0$. Substituting into Eqn. (16) and integrating with respect to time gives

$$\Delta p = \frac{4B_{eff}}{V_0}\left[Ay - \frac{1}{2c_2}\int \mathrm{sign}(\Delta p)\left\{-c_1 + \sqrt{c_1^2 + 4c_2|\Delta p|}\right\} dt\right]. \tag{17}$$

This equation can be transcribed to the simulation diagram in Fig. 9. The poppet valves are included by means of a second loop, identical to the bottom loop of the main orifice flow, but with larger flow coefficients. The continuous activation of the poppet valves is simulated with a tanh function, governed by a first order transfer function of Δp representing the poppet dynamics. These latter dynamics are not modelled accurately, and are included more to promote numerical stability than to capture the detail of the damper response.

Fig. 10 compares the simulated damper forces to the measured damper forces for two examples, both at 3.5Hz. The displacement demand is nominally the same for the simulated and real dampers, although only the standard inner loop PID control is employed for the real damper, meaning that the actual displacement suffers some distortion.

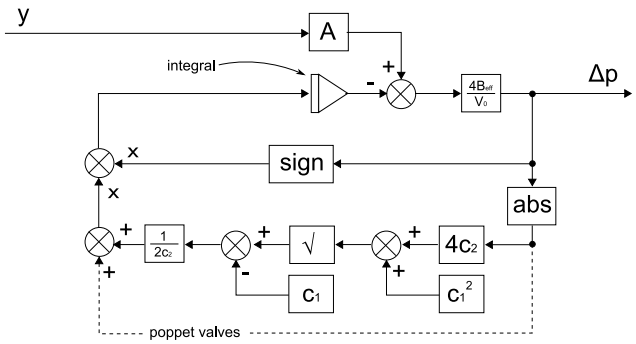

Figure 9: Simulation diagram for the damper model.

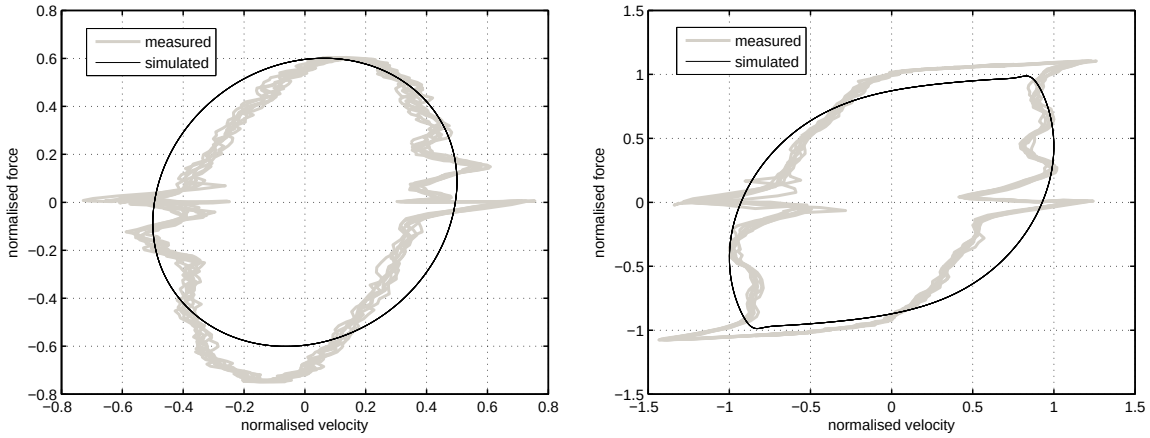

(a) Below poppet valve activation threshold

(b) Poppet valves activated

Figure 10: Simulated damper forces compared to measured forces on a force-velocity plot, subject to the same nominal displacement input.

4.3 Hybrid System Overview

In creating a hybrid numerical-physical system, it is useful to have an estimate of the expected system behaviour before commencing tests on the full rig. To this end, a model of the coupled system is created.

The two-DOF numerical substructure has already been characterised in terms of modal transfer functions. Here, a variation on that approach is taken, casting the transfer functions instead in terms of the mass, damping and stiffness properties of the substructure. Note that while this is a simple process for the two-DOF case, the strategy outlined in Section 4.1 is more versatile and can be applied easily to much larger systems. The equations of motion for the two masses are

$$m_1\ddot{x}_1 + c_1(\dot{x}_1 - \dot{x}_0) + c_2(\dot{x}_1 - \dot{x}_2) + k_1(x_1 - x_0) + k_2(x_1 - x_2) = 0 \tag{18}$$

$$m_2\ddot{x}_2 + c_2(\dot{x}_2 - \dot{x}_1) + k_2(x_2 - x_1) + f(\dot{x}_2) = 0 \tag{19}$$

The function $f(\dot{x}_2)$ is the force generated by the damper. For this analysis, a simple linear approximation of the damper is made. It is taken to be a linear viscous damper in series with a spring (to represent the compressibility), as pictured in Fig. 8. The parameters k_d and c_d are approximated from empirical data under mid-range operating conditions. The equations describing this system are

$$f = k_d x_3 \tag{20}$$

$$f = c_d(\dot{x}_2 - \dot{x}_3). \tag{21}$$

Differentiating Eqn. (20) and substituting into Eqn. (21) produces

$$c_d\dot{f} + k_d f = k_d c_d \dot{x}_2 \tag{22}$$

and transforming to the frequency domain yields

$$f = \frac{c_d k_d s}{c_d s + k_d} x_2. \tag{23}$$

Now transforming Eqns. (18) and (19) to the frequency domain and manipulating the three equations, expressions for the displacements x_1 and x_2 can be found in terms of the displacement excitation at x_0:

$$x_2 = \frac{(k_1 + c_1 s)(k_2 + c_2 s)(c_d s + k_d)}{(m_1 s^2 + (c_2 + c_1)s + k_2 + k_1)((m_2 s^2 + c_2 s + k_2)(c_d s + k_d) + c_d k_d s) - (k_2 + c_2 s)^2(c_d s + k_d)} x_0$$

$$x_1 = \frac{(k_2 + c_2 s)x_2 + (k_1 + c_1 s)x_0}{m_1 s^2 + (c_2 + c_1)s + k_1 + k_2} \tag{24}$$

The roots of the denominator in the expression for x_2 yield the five complex poles of the theoretical hybrid system. Assuming the vibrational modes are underdamped, this will lead to two second order pairs of complex conjugate poles and a single first order pole.

Using the above methods, a system was designed with features appropriate to the frequency range of interest. The transfer functions can be seen in Fig. 11, along with the information about the poles. Both poles are located under 10Hz, where the controllers are all thought to be stable, but this does not diminish the quality of the demonstration; where the controllers perform badly the real substructured system can resonate or diverge even without poles of the nominal system in the vicinity.

4.4 Substructure Coupling and Stability Controls

To complete the substructuring rig, all that is needed is the coupling between the numerical and physical systems. Two switches are employed for this purpose, both of them taking continuously variable values from

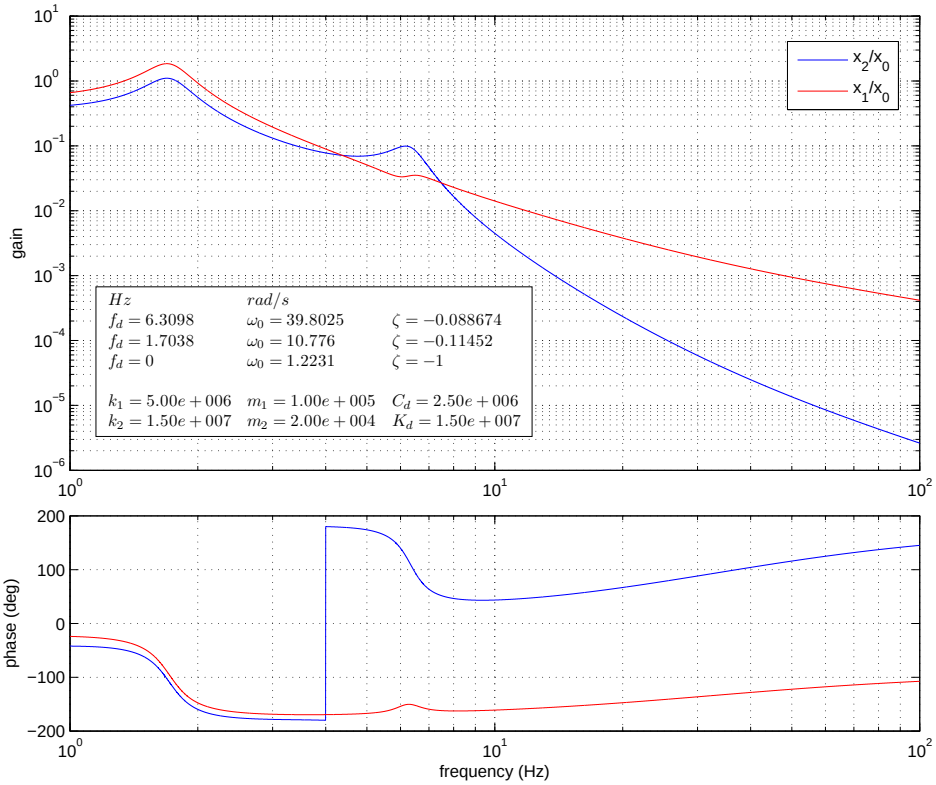

Figure 11: Transfer functions and properties for the theoretical hybrid system. f_d is the damped natural frequency in Hz, ω_0 is the undamped natural frequency in rad/s, and ζ is the modal damping ratio. The mode with $f_d = 0$ is a first order pole with $s = -\zeta\omega_0$.

0 to 1. The first switch governs the displacement coupling between the two substructures, while the second governs the force coupling.

With the displacement switch set to zero, the physical substructure remains at rest, regardless of what the numerical substructure is doing. When it is set to 1, the physical substructure tries to track the numerical displacement, with the aid of the transfer dynamics cancellation. Values between 0 and 1 scale the physical displacement proportionally.

When the force switch is at zero, the forces fed back into the numerical substructure are determined by the simulated damper outlined in Section 4.2. When this switch is set to 1, the measured force is fed back to the numerical substructure. This is the critical setting in the coupling, because when the numerical model is subject to real-world loads, any lag in the real-world displacement tracking can lead to spurious energy injections into the system. This in turn can cause instability and divergence. The exact force feedback is determined by

$$f = \gamma f_p + (1 - \gamma) f_n \tag{25}$$

where f_p is the measured force from the physical damper, and f_n is the force in the simulated damper.

A stability trip switch is also implemented, to catch any instabilities before they risk damaging the hardware. This trip switch detects discrepancies in the energy transfer seen by the physical system and that seen by the numerical system at the interface. Both sides of the interface will be subject to the same measured force, more or less simultaneously, but the velocities of the physical and numerical systems will differ, dependent upon the quality of the displacement tracking. Thus not all the power absorbed by the physical damper will necessarily leave the numerical substructure, and in some cases the numerical substructure may even see a net input of power from the damper (which would ideally be physically impossible). Of course, the reverse is also possible and the most common scenario encountered in the course of these tests is that of a rapid,

artificially induced power flow backwards and forwards between the two substructures. The magnitude of the flow will often grow steadily, even though the net energy transfer remains at zero.

The instability detection facility thus measures the *magnitude* of the power discrepancy between the two systems, regardless of the direction it flows. The power difference is squared then low-pass filtered, so that transient power spikes in either direction are ignored but persistent discrepancies are registered. Once a threshold is passed, the force switch, γ, is changed continuously but quickly to zero, and the simulated damper resumes the task of providing the force feedback.

4.5 Substructuring Tests

To test the efficacy of the proposed inverse process models, prescribed displacement signals are applied to the numerical model, at the input denoted x_0 in Fig. 8. Once more a set of periodic, single-harmonic signals are chosen. While such simple signals are unlikely to be encountered in fully fledged substructuring exercises, they do have a number of advantages in the context of this evaluation.

The most obvious advantage *is* their simplicity. Faced with a steady, periodic input at a single frequency, the response is easy to interpret: primarily, a periodic response should be expected. If this is not attained then either the system dynamics are chaotic, or the test rig and physical hardware has a time-variant component. Assuming a periodic response is achieved, this choice of signal will readily highlight any nonlinearities, as a linear system would respond only at the excitation frequency. Furthermore, where the system *is* found to be more or less linear, the response to any other signal can be constructed from a superposition of the response to sinusoidal inputs. In any case, the transfer dynamics cancellation techniques employed here are all based on the assumption of an approximately linear system, so it is their ability to perform under these conditions that is of primary interest.

A further advantage of harmonic signals is that a systematic test sequence can be planned, thus covering a broad range of operating conditions, and helping to isolate the causes of any unexpected results. Simultaneously, however, this is a disadvantage: to cover all possible operating conditions in this manner would require an exhaustive set of signals. For example, all of the tests conducted here will use zero offset, thus obscuring any sensitivity of the response to the mean displacement. The most convincing argument for the use of harmonic excitation, however, is that it permits an analysis of the phase response. As has been discussed, phase is important because too much phase lag will lead to instability. Put succinctly, stepped sine excitation is the best choice for performing a linear analysis, while simultaneously highlighting the existence of and minimising the influence of nonlinearity.

For this evaluation, several sets of harmonics are used, each set covering a range of frequencies at a single amplitude. From Fig. 11, the response at the interface, x_2, is expected to drop off rapidly after the second natural frequency at 6.3Hz. Thus the excitation amplitudes must increase for the higher frequencies in order to avoid the measurements dropping below the noise threshold. The frequency bands are chosen to overlap, providing results for the two extremes of excitation levels in several frequency bands. In these regions, the amplitudes are limited at the top end of the spectrum by acceptable acoustic noise levels in the lab, and at the bottom end by the signal noise floor. Fig. 12 shows a typical example of the numerical and physical displacements at DOF 2.

The gain and phase results can be seen for each of the four inverse process models in Fig. 13. The dark lines represent the substructuring results, while in the background the thick pale lines indicate the trend of the verification tests from Section 3.1. The results using the first and second order process models match those of the verification tests quite reliably, although the new tests are limited to 40Hz instead of 100Hz. (This is because of the rapidly diminishing response of DOF 2 beyond this frequency.) In these first two examples, the only difference in the setup is the presence of the substructure with force feedback. In contrast, however, the two process models that incorporate delay match the verification tests only up until around 10Hz, after which they diverge rapidly. (For this reason these tests were halted early at 30Hz.) The difference here is that the inverse process models can no longer rely on time-shifted versions of the known demand signals

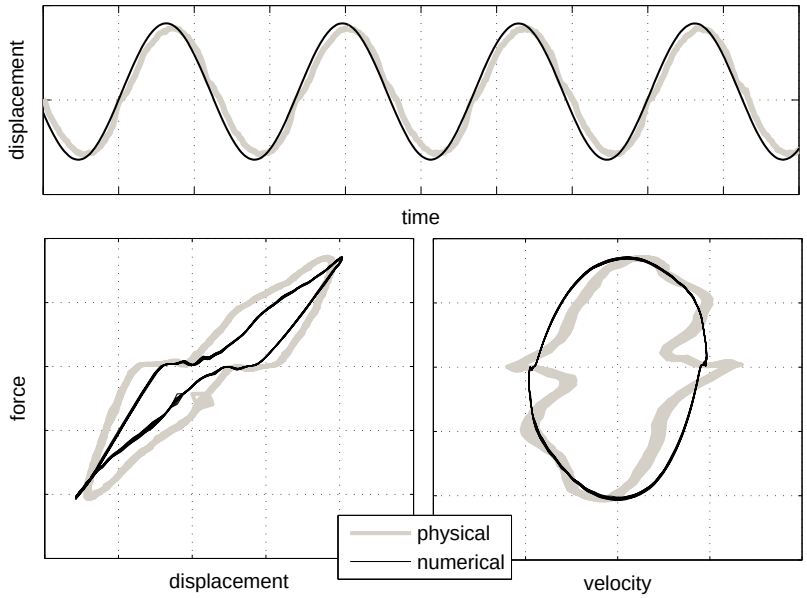

Figure 12: Typical displacement-velocity-force results for the numerical and physical substructures at their interface, using a second order transfer function for the process model and exciting the system at 8.6Hz.

but must now rely on the capabilities of the forward predictive algorithm. Evidently, this algorithm does not perform well beyond 10Hz; this shortcoming will be examined in more depth presently. As a consequence, however, the latter two process models struggle to outperform even the pure first order transfer function, with the displacement tracking from both of them deteriorating beyond about 10Hz. The second order transfer function offers the best results, with acceptable gain and phase agreement beyond 20Hz. It is interesting to note that despite testing the upper and lower limits of the permissible excitation amplitudes, relatively little discrepancy is seen in the overlapping sections of the response curves.

The transfer function-style results presented so far only convey part of the picture. A metric is needed to describe how well the measured response agrees with the numerical model. A straightforward correlation was considered but it was found that the signal noise distorted the results at lower amplitudes. Instead, the root mean square (RMS) value of the error was investigated. This produces higher error values for higher amplitudes, but this is physically representative and not just an artefact of the measurement methods. The results are shown in Fig. 14. The RMS values are normalised to one standard deviation of the signal noise, and it can be seen that as the frequency increases and the substructure response decreases, the RMS error for many of the curves disappears into the noise floor. Comparing the results from the first and second order transfer functions, it is difficult to define one as having better performance than the other. The second order function produces better results at higher frequencies, but has markedly higher responses in the 5-10Hz band.

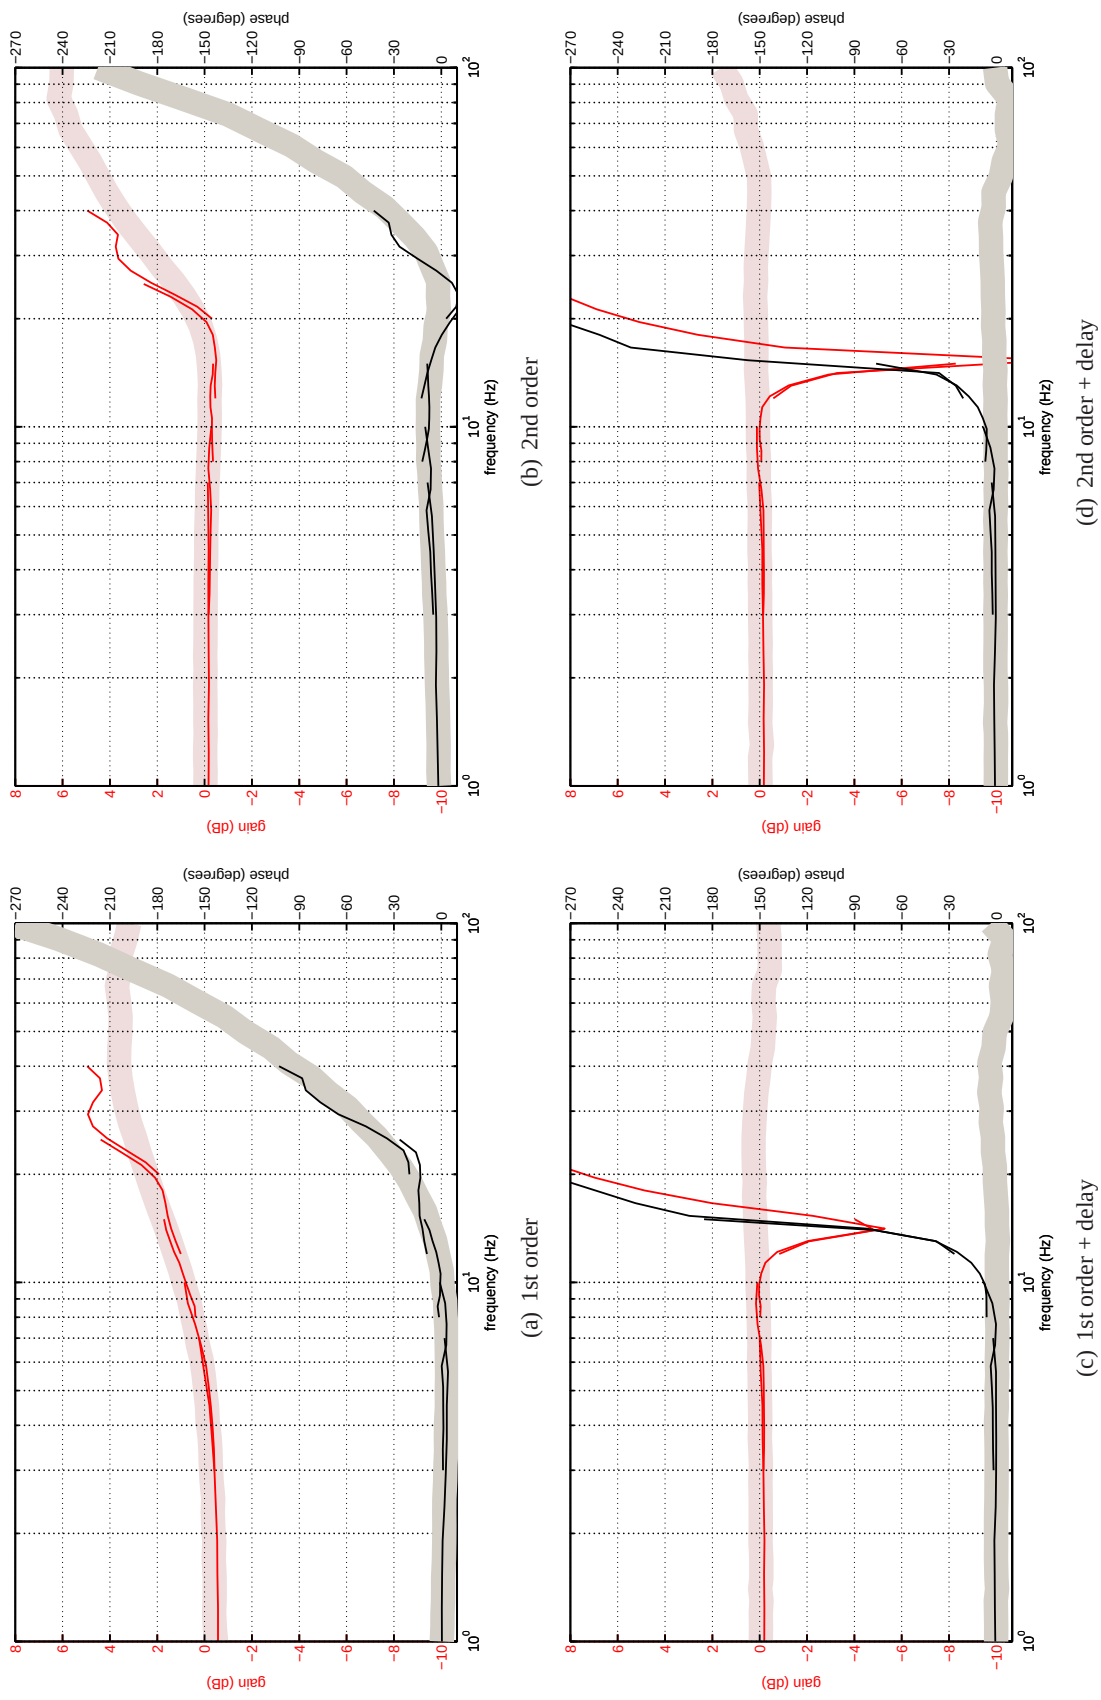

Figure 13: 'Transfer functions' for the output of the substructured system, showing the gain and phase relationship between the nominally identical first harmonic of the numerical and physical displacements. The equivalent transfer functions obtained from prescribed harmonic demand signals are shown in the background as pale, thick lines.

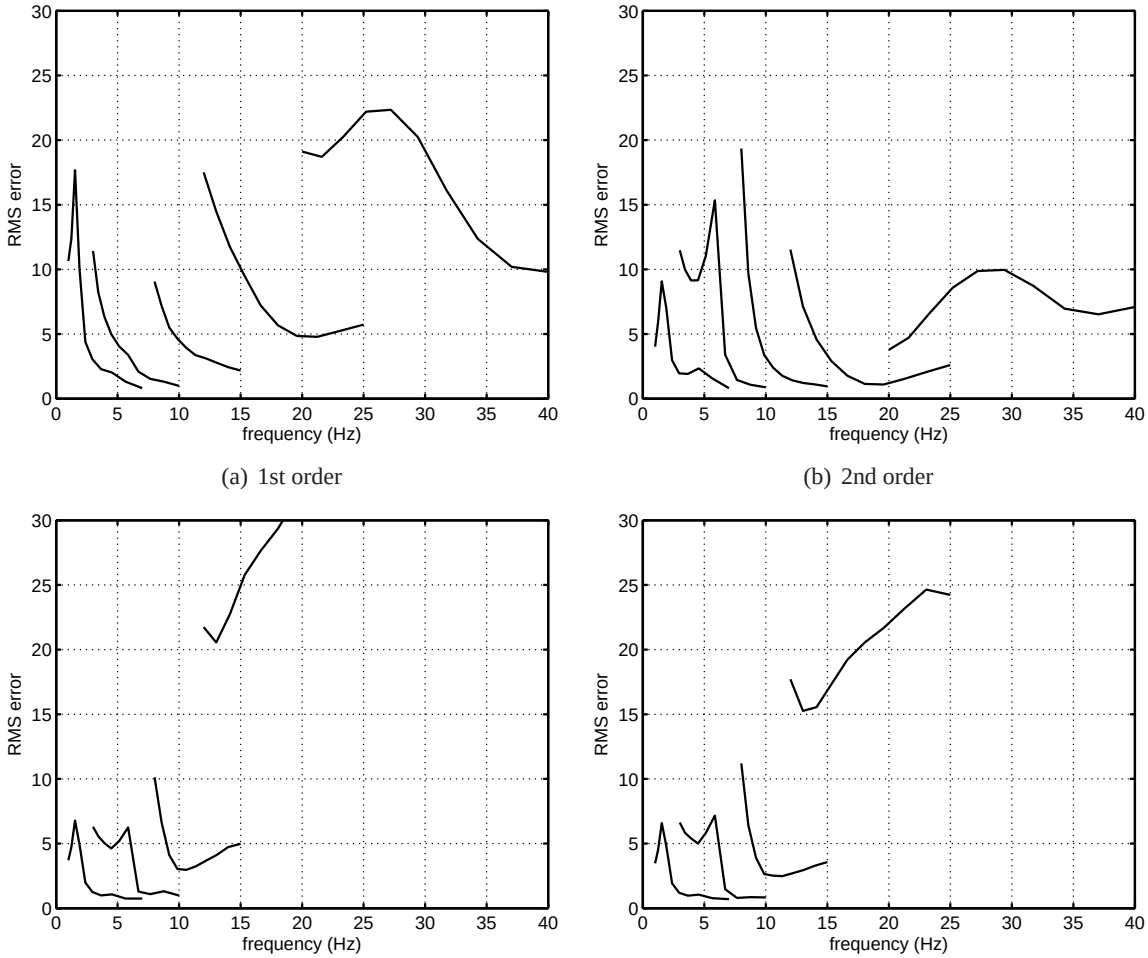

(a) 1st order

(b) 2nd order

(c) 1st order + delay

(d) 2nd order + delay

Figure 14: RMS error in physical substructure displacement for the four process models over a range of single harmonic excitation signals. The RMS values are normalised to the standard deviation of the measurement signal noise.

As would be expected following the gain and phase results, the RMS plots for the two models with delay show large errors for frequencies above 10Hz. In contrast, however, the RMS errors reported below 10Hz are all around half the corresponding errors in the first two models. As before, there is no significant difference in performance to be observed between the two delayed models. The exception to this is in the range above 12Hz, but the performance of both controllers has already been shown to be unreliable at these frequencies. From this analysis it appears that the delay in the latter two models does contribute to better performance, albeit contingent upon the capabilities of the forward predictive algorithm.

To investigate the limitations of the polynomial forward predictive method, its predictions are tested over a range of frequencies using a known demand signal. The predictions are compared with the time shifted demand signal and the gain and phase relationship is documented in Fig. 15. Two sets of curves are shown. The first set is determined using the forward predictive parameters used for the previous tests, in particular a sample period of $T_s = 0.1$s. The second set is created after experimenting with the parameters to produce the best results, and settling on a sample period of $T_s = 0.01$s. The first curves are seen to degrade shortly after 10Hz, corresponding with the drop in performance in the previous tests. In contrast, the second set of curves remain acceptable until around 40Hz. In changing the sample period to 0.01s it was also necessary to drop the number of sample points from 20 to 10 due to the sampling frequency of 1kHz on the dSpace board, but this does not seem to have had an adverse affect on the predictions.

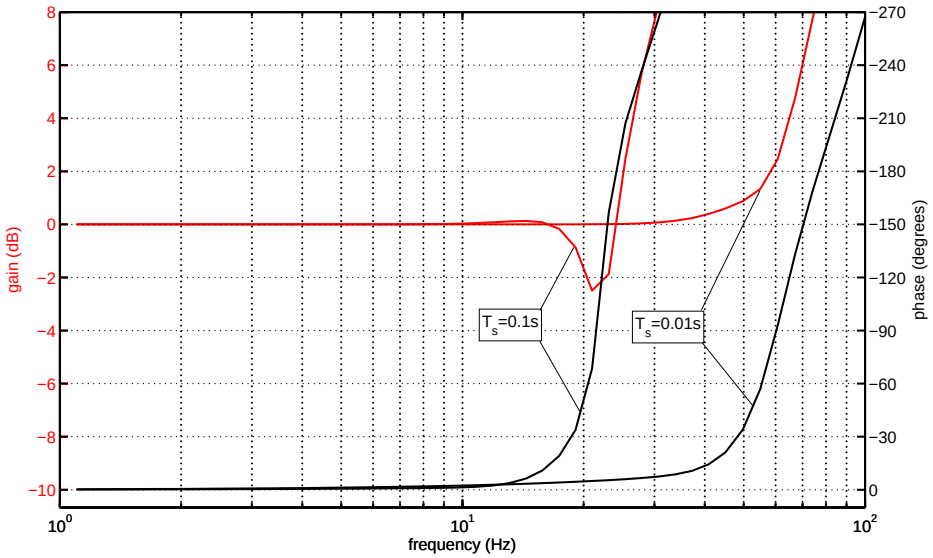

Figure 15: Gain and phase of fundamental harmonic of forward predictive signal with respect to a known sinusoidal input signal with forward predictive time constant optimiseed for the second order transfer function.

With these new forward predictive parameters, it is thought that the transfer function plus delay process models should significantly outperform their pure transfer function counterparts. Future work will first test this hypothesis, before performing a more thorough investigation of the effect of the forward predictive parameters. Other delay compensation techniques will also be investigated, as the forward predictive capability seems to be the limiting factor in improving the performance of these controllers.

5 Conclusions

Four open-loop linear controllers have been proposed to help reduce displacement tracking errors in real-time dynamic substructuring caused by the transfer dynamics in hydraulic actuation plant. Rigorous testing

has shown that these controllers all have the capacity to make significant improvements compared to using the existing inner-loop PID controllers in isolation. Preliminary results suggest further improvements should be readily achievable using better delay compensation methods as part of an inverse process model. This should form the basis of further work on this topic as it is the key factor limiting the performance observed in these tests.

Despite the successes noted here, the investigations have shown that only so much can be achieved using linear controllers. Both the actuation rig and the physical substructure being tested can act as sources of significant nonlinearity. In addition, the work has shown that the macroscopic demand amplitude is not the main concern in this regard, but instead it is nonlinear variations and discontinuities throughout each stroke that lead to tracking errors. Importantly, this is not a problem which can be tackled by adaptive algorithms as they will not react fast enough for variations on this time scale.

One approach would be to produce detailed and more realistic process models for the hydraulic plant, incorporating nonlinearities directly. Such an approach would ideally be augmented with state measurements from within the hydraulic plant, for example pressures and valve spool locations. If state measurements are not directly available, some form of state observer could be employed to this end. A less demanding proposal would be to take advantage of available measurements such as force and displacement at the output and to investigate the benefits of gain scheduling.

An ultimate goal of real-time experimental dynamic substructuring should be to minimise the influence of the physical hardware being tested on the control strategy. After all, one of the motivating factors behind using real-time substructuring in the first place lies in testing pieces of hardware about which little is known to start with. Hopefully some of the results presented here will provide insights to help move towards this goal.

References

[1] T. Horiuchi, M. Inoue, T. Konno, and Y. Namita. Real-time hybrid experimental system with actuator delay compensation and its application to a piping system with energy absorber. *Earthquake Engineering & Structural Dynamics*, 28(10):1121–1141, 1999.

[2] AP Darby, A. Blakeborough, and MS Williams. Improved control algorithm for real-time substructure testing. *Earthquake Engineering & Structural Dynamics*, 30(3):431–448, 2001.

[3] MI Wallace, DJ Wagg, and SA Neild. An adaptive polynomial based forward prediction algorithm for multi-actuator real-time dynamic substructuring. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 461(2064):3807, 2005.

[4] PA Bonnet, MS Williams, and A. Blakeborough. Compensation of actuator dynamics in real-time hybrid tests. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, 221(2):251–264, 2007.

[5] MI Wallace, DJ Wagg, SA Neild, P. Bunniss, NAJ Lieven, and AJ Crewe. Testing coupled rotor blade-lag damper vibration using real-time dynamic substructuring. *Journal of Sound and Vibration*, 307(3-5):737–754, 2007.

[6] AR Plummer. Control techniques for structural testing: a review. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, 221(2):139–169, 2007.

[7] DP Stoten and RA Hyde. Adaptive control of dynamically substructured systems: the single-input single-output case. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, 220(2):63–79, 2006.

[8] C. Chen and J.M. Ricles. Analysis of actuator delay compensation methods for real-time testing. *Engineering Structures*, 31(11):2643–2655, 2009.

[9] B. Titurus, J. du Bois, N. Lieven, and R. Hansford. A method for the identification of hydraulic damper characteristics from steady velocity inputs. *Mechanical Systems and Signal Processing*, 2010.