Development of a Relative Motion Facility for Simulations of Autonomous Air to Air Refuelling

Peter Thomas, Jon duBois, Tom Richardson
Department of Aerospace Engineering
University of Bristol
UK, BS8 1TR
{p.thomas, jon.dubois, thomas.richardson}@bristol.ac.uk

***Abstract*—The initial development and results of a comprehensive simulation and testing environment for autonomous air to air refuelling is presented. A thirteen degree of freedom relative motion facility has been installed at the University of Bristol to support the design, testing, and validation of measurement systems and autonomous control algorithms through hardware in the loop simulations. Initially this facility is used to simulate the 'hook-up space' in air to air refuelling. A real-time platform handles the control of the manipulators in synchronisation with streamed data generated by simulated kinematics. An air-to-air refuelling simulation provides the kinematic data of a refuelling probe and drogue, with the feedback loop made by the provision of position measurements from proximity and vision-based sensors. The synthetic environment is real-time and consists of nonlinear models for the receiver and tanker, and accounts for the additional dynamics of the probe and drogue. Data packaging and delay compensation on the network between the real-time platform and the robot controller is addressed in this paper.**

TABLE OF CONTENTS

1. INTRODUCTION

Air to air refuelling (AAR) is predominately an endeavour to increase endurance. Since the 1920s when the first AAR was demonstrated [1] only a few methods have been developed for carrying out the transfer of fuel between aircraft. The looped-hose and wing-to-wing methods have been resigned to history leaving Flight Refuelling Limited's probe-drogue and Boeing's flying boom as the only two methods used in current AAR operations. Both of these methods have developed within the scope of human-controlled flight, and now present interesting challenges in the areas of control, pose estimation, and machine intelligence for unmanned systems to replicate and ultimately improve performance.

With the increasing use of high-altitude, long endurance (HALE) UAVs in the foreseeable future an autonomous air to air refuelling (AAAR) capability is an inevitable necessity to support an expanded mission range. Other applications where AAR would be an effective force multiplier are in aerial search and rescue and airborne communication relays. And whilst AAR is currently used almost exclusively in military scenarios there are potential cost savings to be made in air freight by taking off with lower fuel stores (hence lower take-off weight) and refuelling en-route; an idea proposed back in 1931 by Glover [2] but never pursued commercially. Recent works [3], [4] have revisited the ideas in the context of modern operations. Looking towards the future, unmanned air freighters could benefit in this manner. Consequently AAAR is a topic of much ongoing research that, amongst other aspects, entails integration of sensors that must meet high bandwidth requirements for fast dynamic operation in close proximity of other aircraft. The motion is relative to the tanker, which ultimately becomes the reference point for all motion in the refuelling operation. The risks in the 'hook-up space' (the region behind the tanker where coupling with refuelling apparatus takes place) are significant and the cost of mistakes can be catastrophic. A proposed system for AAAR will require sufficiently rigorous testing to certify its robustness and performance, for which hardware in the loop (HIL) facilities, that reduce the cost and risk from flight testing, are ideally suited. Cobham Mission Equipment and the University of Bristol are developing a Relative Motion Robotics (RMR) facility, based in the university's Advanced Composites Centre for Innovation and Science. The RMR facility can function as a HIL facility to provide a cost effective research and trials capability for evolving technology. Previous works in [5] and [6] used robot manipulators to simulate aircraft motion and refuelling boom movements to replicate the AAAR scenario, albeit at a reduced scale. The RMR facility employs two manipulators capable of supporting full-size refuelling apparatus with the capability to integrate pose estimation systems into a real-time control loop. In doing so the suitability of vision systems, tracking algorithms, control system designs, and refuelling hardware, that are mature in their development can be tested in a safe and repeatable environment. The RMR is also capable of investigating wider technology exploitation and utility to industry and academia for relative motion work as well as supporting research into robotic composites manufacture.

Generating a simulated refuelling environment that is sufficiently representative for HIL testing is not a trivial task [7], [8]. The hook-up space is a relatively compact environment with complex interactions between aircraft, refuelling equipment, and aerodynamic effects from bow waves, wake vortices, and air turbulence. Over the last year we have been developing and integrating a number of different aircraft and atmospheric models into a nonlinear, real-time simulation to provide the synthetic environment for the RMR. The simulation environment supports the research and development

978-1-4577-0557-1/12/$26.00 ©2012 IEEE.
[1] IEEEAC Paper #1423, Version 1.0, Updated 01/01/2012.

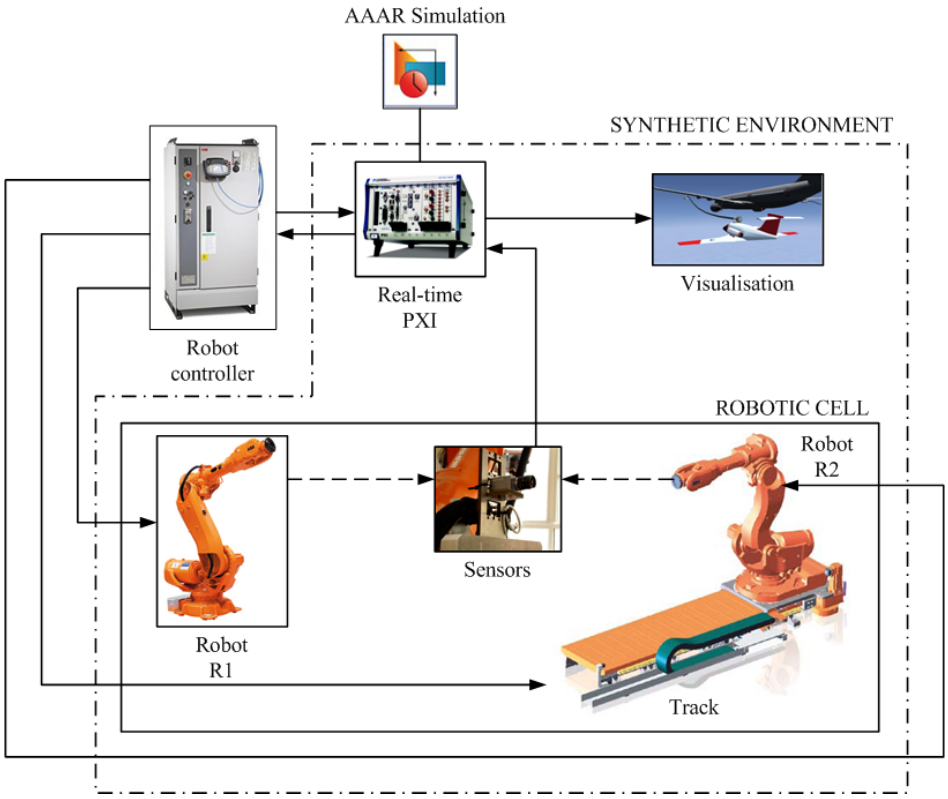

Figure 1. Proposed layout of the relative motion robotics facility as a synthetic environment for air to air refuelling simulations.

part of this investigation as part of the wider ASTRAEA programme, and provides a model framework for testing and evaluating a number of different component models, control laws, and regulatory logic from several parties. It is developed and maintained between Cobham Mission Equipment and the University of Bristol. Third party components, particularly the wake vortex model (which is calculated with a vortex lattice method, originating from the work described in [9]) will not be addressed in this paper.

This paper describes the current work towards the development of the RMR as a HIL facility for AAAR. As such there are two key areas addressed in this paper: the development and integration of the RMR facility, and the simulation environment the RMR aims to replicate. The paper is organised as follows: firstly we discuss the layout of the RMR and its constituent parts, and describe the challenges and solutions used to to interface the parts of the RMR, highlighting issues of time delay compensation and communication flow. Next we summarise the AAR simulation environment and draw attention to the elements critical to the probe-drogue hook-up space. In Section 4 we explain the integration of the physical and simulated systems, in terms of both execution control and positional mappings. Lastly we present some experimental results illustrating the challenges and successes in the development of the integrated RMR facility to date.

2. The Relative Motion Robotics Facility

The RMR is primarily intended for the modelling of relative motion between bodies. The layout of the facility, as shown in Figure 1, is used here to create a synthetic environment for AAR to provide HIL testing capabilities for machine vision and other sensor systems and algorithms towards autonomous refuelling of aircraft. Simulation models of AAAR (along with control systems, sensor models and filters) are executed in real-time on a National Instruments PXIe-8133RT 1.73 GHz control board mounted in a PXIe-1033 chassis. The PXIe system is also runs a TCP/IP client for communicating with the ABB IRC5 robot controller, and a supervisory process responsible for transforming the output of the flight dynamics simulation to position command data suitable for the two robot manipulators, in order to replicate the relative motion between the aircraft or other apparatus used in the refuelling process. In addition, the supervisory process handles the execution control for the flight dynamics simulation, the downsampling of the position data, system performance monitoring and communication timing and coordination, and safety limits for the position demands.

The robotic cell comprises two ABB IRB6640 industrial robots, designated R1 and R2, and having the performance characteristics as detailed in Table 1. R1 is secured to the ground whilst R2 is mounted on a 7.7 m IRBT6004 track to permit translation of the robot base at a rate of 5.2 ft/s (1.6 m/s). The addition of the track provides thirteen degrees of freedom motion. Actual refuelling hardware is used in the rig: a drogue is attached to the end of R1 and a refuelling probe nozzle is mounted to the track-mounted R2 (Figure 2). In this way the RMR can use the combination of joint rotations and movement along the track to place the probe at positions and orientations anywhere in the operational envelope. Similarly R1 can translate and rotate the drogue to simulate turbulence and motion of the drogue in flight.

The absolute operational envelope is a cylindrical working area of length 10 m and 2 m diameter, however practical and safety limits imposed on the arm's movements limits

Table 1. Performance characteristics for the IRB6640 robots and the IRBT6004 track.

		Unit
Maximum acceleration	2	g
Maximum relative velocity	6	m/s
Pose accuracy	0.16	mm
Pose repeatability	0.07	mm
Pose stabilisation time	0.36	s
Track length	7700	mm
Track maximum velocity	1.6	m/s
Track pose repeatability	0.08	mm

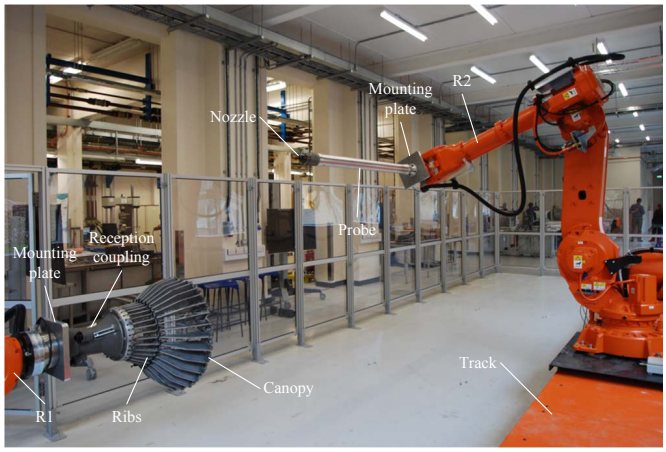

Figure 2. The RMR cell mounted with refuelling drogue and probe apparatus.

the nozzle tip and drogue's position to inside the working area indicated in Figure 3. This provides an approach range of 6.5 metres with 1.5 m lateral and 2.6 m vertical motion. Coordinate frames can be defined in the robot controller and for the work conducted here the position demands are defined relative to a set of inertial axes aligned with the tanker body axes, located at the end of the track nearest the drogue-carrying robot R1.

The flow of position information from the flight dynamics model (FDM) simulations to the actual robot motion is illustrated in Figure 4. Two physical devices are depicted (three or more including the robots and track themselves, although in the discussion that follows they are included as part of the proprietary ABB system). The important elements are the PXIe real-time controller and the IRC5 robot controller. The communication between these elements is by means of

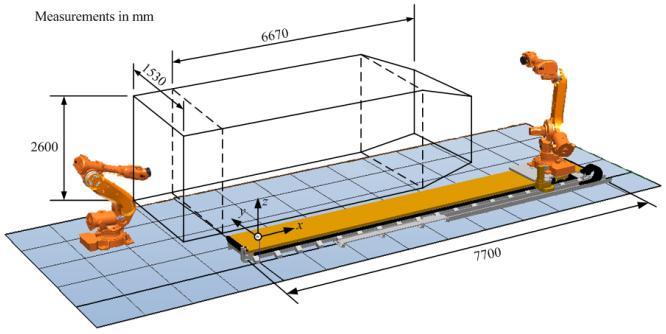

Figure 3. RMR working area.

ethernet TCP/IP streams, carried over 100BASE-TX using a Category 5 crossover cable.

The FDM simulation is shown in the bottom right corner; the complexity of this system is belied by its representation on this diagram but is elaborated in Section 3. The simulation runs at 1 kHz on the real time Veristand Engine operating system of the PXIe box. This operating system is capable of overseeing the deterministic execution of multiple models, or processes at defined rates. The primary control loop executes the FDM model and the supervisory process in turn, both at a rate of 1 kHz. At the start of each time step for a given process the data mappings into that process are read from the buffers, and when it executes, the outputs of the process are written to the buffers. The data mappings between processes are referred to as channels, and this is how information is exchanged between the processes on the PXIe. Critically, no data can be exchanged mid-way through a time step of a particular process. Note that the processes can be configured to run in parallel or consecutively, and in the latter case the outputs of earlier processes will be available for later processes within the same time step. In the current application the FDM is the first process to run each time step, and the position data from the FDM is made available to the supervisor process. This data is passed as 64-bit double precision floating point variables.

The supervisory process performs many tasks, including providing execution control for the FDM, but its most critical task is to control the flow of data to and from the ABB IRC5 controller. The key technical barrier is that while the FDM and supervisory process are both run in real time, and the robot motion can be controlled such that it meets position demands in deterministic time frames, the communication protocols do not mirror this determinism. On the IRC5 side of the communications, the data transmission buffer and process/thread management is handled internally by proprietary firmware and very limited control can be exercised over these processes. On the PXIe side of the communications, the so-called *custom device* process responsible for the TCP/IP transmissions necessarily runs asynchronously with respect to the real time processes to avoid delaying any time steps in the event of a delayed message from the IRC5 controller. The design of the custom device will be described presently. The supervisory process controls the flow of data to and from the IRC5 controller using two sets of counters: the first set is used to synchronise with the cycles of the IRC5 communication loops and will be discussed shortly with reference to the TCP/IP custom device. The second set of counters are the COMM and ACKN indices seen in Figure 4. These are used to orchestrate the motion commands send to the IRC5. The position demands from the FDM are sampled regularly at 20 Hz using a timing pulse trigger to ensure a smooth motion path definition. These are then placed into a FIFO buffer so that no position data will be omitted in the event of communication delays. Each position dataset is sent to the IRC5 with a unique, sequential COMM (command) index. Once the position instruction has been completed on the IRC5 it returns the corresponding ACKN (acknowledge) index. The receipt of this index by the supervisory process provides the ACKN trigger used to send the next buffer entry. In general operation this buffer remains empty, with the position instructions being removed from the buffer in the same time-step that they are placed in it. It is nonetheless a necessary feature to prevent errors in the event of communication speed fluctuations.

The TCP/IP communications process on the PXIe is imple-

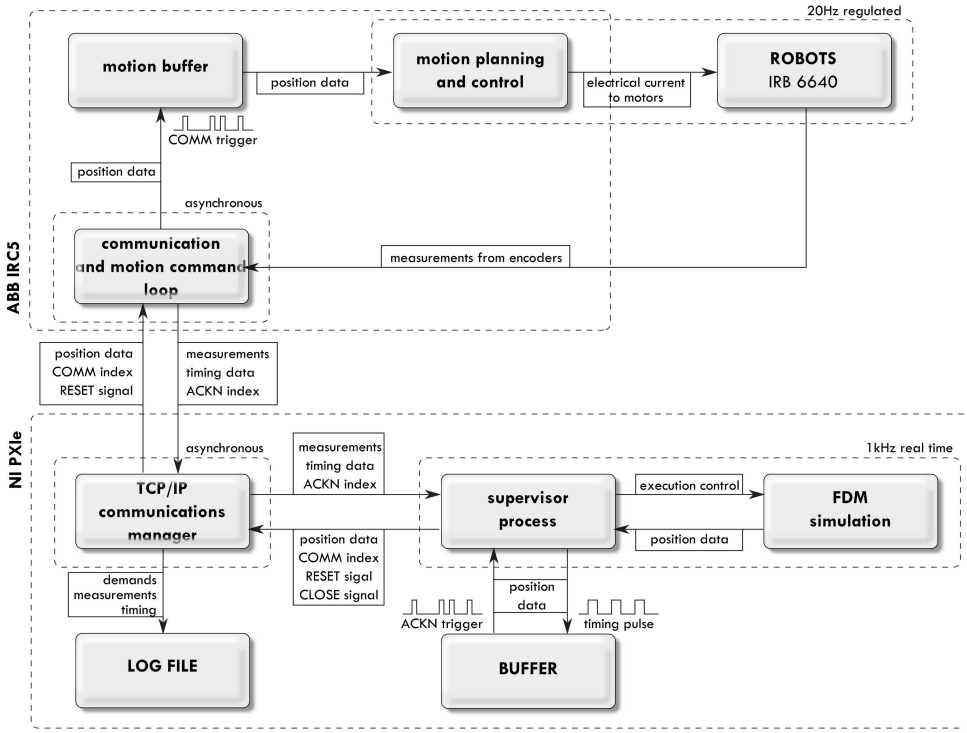

Figure 4. Position and control data flow between processes on the PXIe (real time) and IRC5 (proprietary robot) controllers

mented as a custom device in the Veristand Engine, running asynchronously with respect to the real time processes, interfacing with the IRC5 controller on one side via a TCP/IP socket and with the real-time supervisory process on the other side by means of 64-bit floating point data channels, read from and written to FIFO buffers in the shared memory space at the start and end of each custom device loop. Data received from the IRC5 includes measured positions derived from the joint encoders, timing data used in the supervisor process to optimise position sampling timing, timestamps for measured data, and control variables such as the ACKN index and cycle synchronisation counter. Data sent to the IRC5 includes the position demands from the supervisory process, the COMM index, the cycle synchronisation counter and the sampling time (currently held constant at 20 Hz). The communications run faster than the 20 Hz position demands, to permit measurements to be recorded at a higher frequency. The asynchronous custom device runs as fast as it can, using the COMM/ACKN indices to trigger the motion instruction events. The second set of counters, used to synchronise with the IRC5 cycle and referred to as IRC5iteration and PXIiteration, ensure that the communications to and fro always interleave the processing loops. That is, the supervisory loop will always run once following the receipt of a message from the IRC5 before a message is sent back to the IRC5, and vice versa for the motion control and measurement loop on the IRC5. This involves repeatedly looping following receipt of a TCP message until PXIiteration=IRC5iteration, indicating the supervisory process has processed the received data, and only then sending a message back to the IRC5. Messages from the IRC5 controller are comprised of ASCII string representations of the numerical values, separated by commas and terminated by a CRLF(0D0A) sequence. This is a legacy system and will be replaced in due course, but the performance analysis in Section 5 illustrates that it does not impose a severe performance penalty. It does impact on the precision of the data transmissions, but this does not have a real effect on the accuracy of the position demands and measurements. In contrast, the string parsing functions on the IRC5 are not well suited to processing long strings of numerical values and in this direction the values are now encoded as 32-bit floats, with big-endian bit ordering and little-endian byte ordering, in a fixed-length message with start delimiting header bytes. These can be efficiently reconstructed at the IRC5 end. The secondary responsibility of the custom device is for the logging of all data sent in each direction. This will include all available measured values as well as performance and timing data, and demanded positions.

On the IRC5, the equivalent of the PXIe's supervisor process is the motion planning and control process. This process is handled by proprietary ABB firmware, and currently the only means of influencing the process is to issue move instructions from a high-level scripting code called RAPID. When a move instruction is issued, the motion planning process buffers the position data and once the buffer contains sufficient positions it constructs a smooth path, interpolating at the corners. The move instruction contains position data as well as a time-step indicating the time the robot should take to complete the motion from the previous point to the new point. Provided the buffer is replenished at the same rate the motions are completed, the planned path is iteratively updated to ensure a continuous motion. This buffering process introduces a delay between the FDM simulation and the robot motion (augmented to a small extent by the transmission times, message processing, and position filtering), and this delay must be compensated for as described in Section 4. The IRC5 controller is responsible for control of the electrical

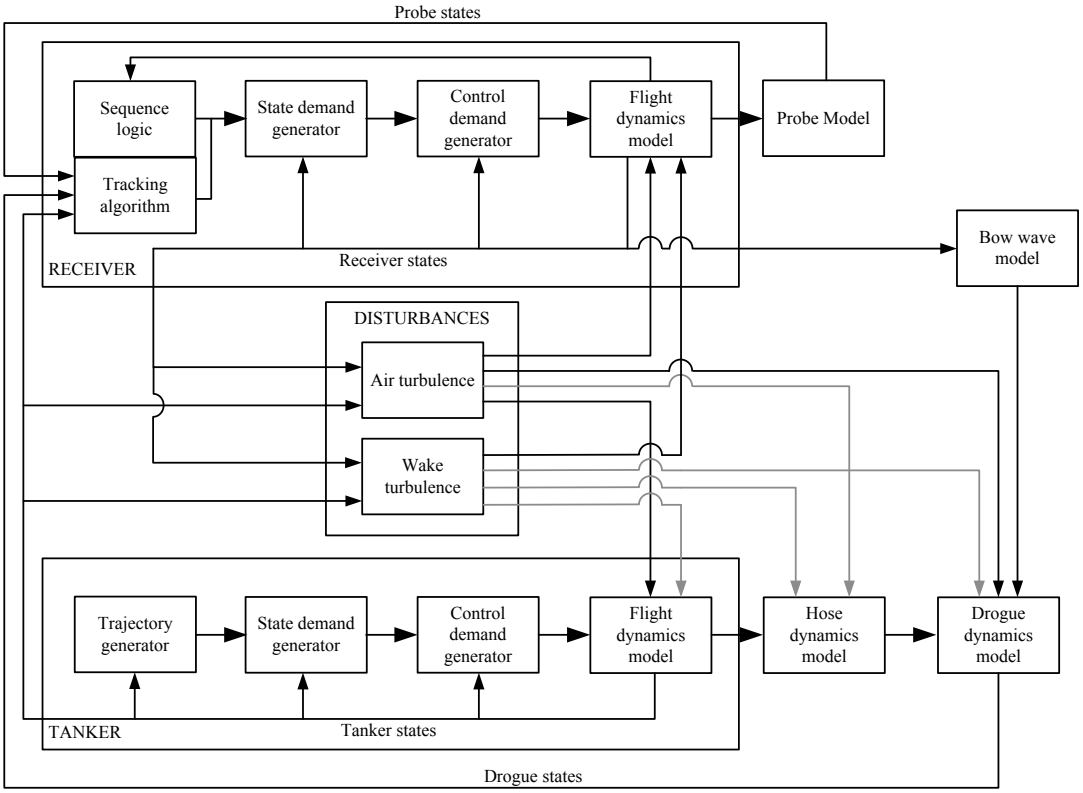

Figure 5. Structure of the AAAR simulation environment. The grey paths indicate links yet to be implemented.

motor supply currents for driving the joints directly, and it is understood that this control is comprised of both feedback and feedforward paths taking into account known properties of the robot limbs and the inertial properties of the mounted hardware to minimise overshoot, rise times and settling times.

The RAPID script loop which serves as the gateway to the IRC5 controller, mirroring the TCP/IP custom device on the PXIe, is independent of the motion planning, and needs only to supply motion instructions as they are made available over the communications link. It performs a simple loop, repeatedly measuring poistions, recording timing information, sending these to the PXIe along with the ACKNindex and PXIiteration counters, and awaiting a response from the PXIe. Once a response is received, if the COMMindex has incremented then a move instruction is executed and the ACKNindex is adjusted. The loop then repeats.

The measurement data from the IRC5 can be received at rates of around 100 Hz under favourable conditions, but the timing of the measurements is not regular and they can be interrupted by the motion planning routines (which are apparently iterative algorithms that take indeterminate time to converge). In addition, the measurements thus obtained make the assumptions of zero backlash, accurate geometry models, and most importantly no structural flexibility. To provide higher fidelity measurements, accelerometers are mounted on the manipulator plates to estimate the high frequency positional data. Each robot uses a single triaxial accelerometer in conjunction with two more collocated transducers and a fourth, non-collocated device, to provide 6DOF acceleration measurements. The accelerometers are piezoelectric low frequency IEPE devices, calibrated to $\pm 5\%$ gain at the extents of the range 0.5–1000 Hz. Signal conditioning and data aquisition is performed on board the robots, with an analogue pre-filter supplying a 50 kS/s analogue-to-digital conversion, followed by digital signal processing and then downsampling. An EtherCAT deterministic distributed measurement system is used to relay the measurements back to the PXIe master device with system-wide jitter rated at less than 1 ms. Data fusion combines this high-frequency data with the low-frequency measurements from the IRC5 using complimentary filtering to provide a full-spectrum estimate of the robot manipulator motion.

3. AIR TO AIR REFUELLING SIMULATION

Simulations are written in Mathworks' Simulink environment and compiled with the Simulink Coder (Real Time Workshop) toolbox for use on the PXI platform using National Instruments' Veristand target language compiler. Simulations cover the wider refuelling scenario in order to develop and investigate control strategies, with the RMR specifically providing the HIL capability for the more complex hook-up space. The simulation environment (Figure 5) takes into account:

1. Tanker trajectory demands and control, FCS and flight dynamics model.
2. Models of the hose and drogue assembly
3. Receiver navigation logic, FCS and flight dynamics model.
4. Atmospheric (gust and wake) disturbance models

The simulation structure is purposely modular such that ongoing improvements to individual components can be made in parallel and swapped in, limiting the changes needed to the simulation environment.

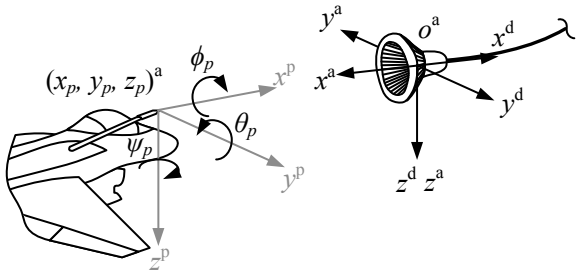

Figure 6. Probe (p), drogue (d), and approach (a) axes definitions.

The purpose of the simulation, in the context of the RMR facility, is to generate position and orientation information for the probe and drogue which can be replicated by the manipulators. To that end we define a set of axes systems in Figure 6 which identifies the refuelling probe (p) and paradrogue (d) objects. The task in probe-drogue configured AAR is to approach and couple the probe with the drogue to close the refuel line. Consequently the probe must track and close the range between it and the drogue, this described in terms of the approach frame (a) which is coincident with the drogue. The probe position is therefore described with the coordinates (x_p, y_p, z_p), relative to the origin o^{a}.

Aircraft models

Both the receiver and tanker are rigid-body, six degrees of freedom objects having nonlinear aerodynamic behaviour in the form of lookup data. The general schematic for the rigid bodies is illustrated in Figure 7. Reference commands from the guidance and navigation systems are used by the flight control system to generate input commands to the actuator models. These in turn, along with the dynamic aircraft states are used to generate the aerodynamic forces and moments on the aircraft at the centre of mass (CM). Clearly the CM will vary throughout the refuelling process, primarily affecting the pitching moment of both receiver and tanker. However up to now we have assumed the variation will have a negligible effect on the performance of the flight control laws and have used a fixed CM at $0.25c$ i.e. 25% from the leading edge of the wing's mean aerodynamic chord. Future improvements to the simulation will determine if this was a valid assumption: it has already been suggested that that mass variation due to fuel transfer compounds the difficulties created by tanker wake turbulence [10]. A generic tanker flight dynamics model is employed but the tanker dynamics are not critical to the simulation - in simpler scenarios the tanker model has been replaced with a reference point moving at constant velocity. Two configurations for the receiver aircraft are used: an F-16 fighter jet and the conceptual Innovative Control Effector aircraft.

A model for an F-16 unmanned jet fighter was derived from the data in [11], which itself is a reduced version from [12]. The simplified model is valid for the aerodynamic range $\alpha \in [-10^\circ, 45^\circ]$, $\beta \in [-30^\circ, 30^\circ]$, which is well within the flight regime for refuelling aircraft. Three first order lags with rate limits and saturations model the actuators similar to those used in [12]. Aerodynamic forces (X,Y,Z) and moment (L, M, N) coefficients about the centre of mass (CM) are calculated in the aerodynamic subsystem using the previous time step aircraft states:

$$C_X(\alpha, q, \delta_e) \quad C_Y(\alpha, \beta, p, r, \delta_a, \delta_r)$$
$$C_Z(\alpha, \beta, q, \delta_e)$$

$$C_L(\alpha, \beta, p, r, \delta_a, \delta_r) \quad C_M(\alpha, q, \delta_e, C_Z)$$
$$C_N(\alpha, \beta, p, r, \delta_a, \delta_r)$$

where α, β are the aerodynamic incidence and sideslip angles and $p,\ q,\ r$ are the rotational rates. The parameters δ_e, δ_a, and δ_r correspond to the elevator, aileron, and rudder deflections. Leading edge flaps and differential tail inputs are not used in the model. The propulsive thrust is calculated through a lag in the power generated by the jet engine simulated with a first order transfer function.

A model for the conceptual Innovative Control Effector (ICE) [13] aircraft is used in addition to the F-16 to investigate control challenges relevant to future aircraft configurations. The ICE is a tailless delta wing fixed-wing vehicle with a 65 degree leading edge sweep and saw-tooth trailing edge. The design for the ICE was driven by the need for a low radar cross section, hence the minimum vertical profile and control surface edges aligned with the external airframe edges. Yaw control is provided through multi-axis thrust vectoring (however structural loads limit its operation to below 200 knots) and clamshells. Consequently a multitude of control effectors are needed to enable the aircraft to operate with sufficient lateral command authority throughout its intended flight envelope. Aerodynamic forces and moments on the CM are tabulated in a similar fashion to the F-16 model, depending on the vehicle's inertial velocity parameters through the air ($\alpha,\ \beta,\ p,\ q,\ r$), and the magnitude of control defections for each of the effectors.

For both F-16 and ICE models the aero-normalised forces and moments are dimensionalised using the aircraft's characteristic dimensions and the current dynamic pressure. The total sum of both aerodynamic and propulsive forces and moments is used to solve the standard equations of motion for a fixed wing aircraft. These equations relate the time derivative of each of the twelve primary states to the current state values and the forces and moments acting on the aircraft. If the sum of the forces and moments on the aircraft are expressed in the form of Newton's second law and subsequently integrated and transformed to the appropriate axes systems, the state equations describing the six velocities (three translational and three rotational) and six positions (again, a translational and a rotational triad) of the aircraft are obtained:

$$\left.\begin{aligned} \dot{u} &= rv - qw - g\sin\theta + \frac{X+T}{m} \\ \dot{v} &= pq - ru + g\sin\phi\cos\theta + \frac{Y}{m} \\ \dot{u} &= qu - pv + g\cos\phi\cos\theta + \frac{Z}{m} \end{aligned}\right\} \tag{1}$$

$$\left.\begin{aligned} \dot{p} &= \frac{pqI_{xz}(I_x - I_y + I_z) + qr[I_z(I_y - I_z) - I_{xz}^2] + I_z + I_{xz}}{I_x I_z - I_{xz}^2} \\ \dot{q} &= \frac{M + pr(I_z - I_x) + I_{xz}(r^2 - p^2)}{I_y} \\ \dot{r} &= \frac{pq[I_x(I_x - I_y) + I_{xz}^2] - qrI_{xz}(I_x - I_y + I_z) + I_{xz} + I_x}{I_x I_z - I_{xz}^2} \end{aligned}\right\} \tag{2}$$

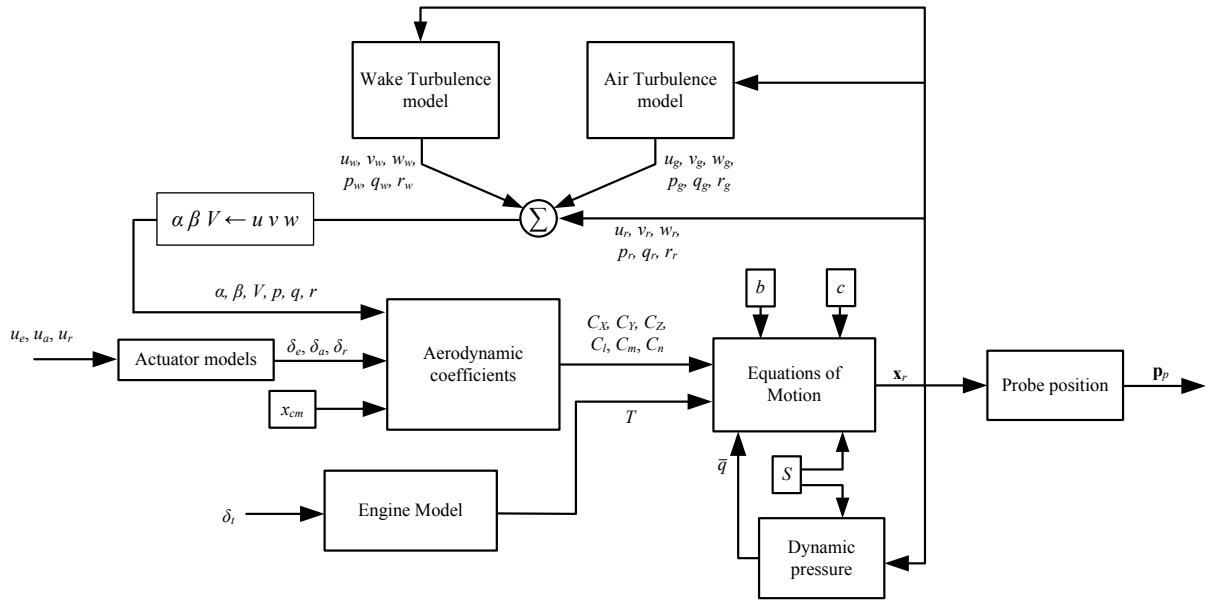

Figure 7. Rigid-body fixed-wing simulation model

$$\left.\begin{aligned}\dot{x} &= u(\cos\theta\cos\psi) + v(\sin\phi\sin\theta\cos\psi - \cos\phi\sin\psi)\\ &\quad + w(\cos\phi\sin\theta\cos\psi + \sin\phi\sin\psi)\\ \dot{y} &= u(\cos\theta\sin\psi) + v(\sin\phi\sin\theta\sin\psi + \cos\phi\cos\psi)\\ &\quad + w(\cos\phi\sin\theta\sin\psi + \sin\phi\cos\psi)\\ \dot{z} &= -u\sin\theta + v\sin\phi\cos\theta + w\cos\phi\cos\theta\end{aligned}\right\} \quad (3)$$

$$\left.\begin{aligned}\dot{\phi} &= p + \tan\theta(q\sin\phi + r\cos\phi)\\ \dot{\theta} &= q\cos\phi - r\sin\phi\\ \dot{\psi} &= \frac{q\sin\phi + r\cos\phi}{\cos\theta}\end{aligned}\right\} \quad (4)$$

The translational velocity equations are transformed into the wind axes to obtain equations of motion for the angle of attack, sideslip angle, and total airspeed:

$$\dot{V} = \frac{\dot{u}V\cos\alpha\cos\beta + \dot{v}V\sin\beta + \dot{w}V\sin\alpha\cos\beta}{V}$$

$$\dot{\beta} = \frac{(V\dot{v} - \dot{V}V\sin\beta)\cos\beta}{V^2\cos^2\alpha\cos^2\beta + V^2\sin^2\alpha\cos^2\beta}$$

$$\dot{\alpha} = \frac{\dot{w}V\cos\alpha\cos\beta - \dot{u}V\sin\alpha\cos\beta}{V^2\cos^2\alpha\cos^2\beta + V^2\sin^2\alpha\cos^2\beta}$$

With the solution to these states the position of the probe nozzle is then calculated taking into account rotations about the CM. Sufficient accuracy is obtained in the simulation model solving these using a third-order Runge-Kutta algorithm with a time step of 10 ms.

Air turbulence

Additional intermittent forces and moments on aero-objects comes from atmospheric instabilities relating to gradients in temperature, pressure, and velocity, resulting in deviations in the air flow from the free stream. Turbulence is observed in individual patches and is characterised by random, homogenous, and isotropic behaviour. It is normally modelled by passing white noise with unity spectral density through a low-pass shaping filter that gives the desired output spectrum.

Mathematical representation—The continuous Dryden form is used, being convenient in that it has rational power spectral densities making modelling far simpler [14]

$$\left.\begin{aligned}\phi_u(\omega) &= \frac{2\sigma_u^2 L_u}{\pi U_0}\frac{1}{1+\left(L_u\frac{\omega}{U_0}\right)^2}\\ \phi_v(\omega) &= \frac{\sigma_v^2 L_v}{\pi U_0}\frac{1+3\left(L_v\frac{\omega}{U_0}\right)^2}{\left[1+\left(L_v\frac{\omega}{U_0}\right)^2\right]^2}\\ \phi_w(\omega) &= \frac{\sigma_w^2 L_w}{\pi U_0}\frac{1+3\left(L_w\frac{\omega}{U_0}\right)^2}{\left[1+\left(L_w\frac{\omega}{U_0}\right)^2\right]^2}\end{aligned}\right\} \quad (5)$$

where

$\sigma_{(\cdot)}$	are the gust intensities,
$L_{(\cdot)}$	are the turbulence scales,
U_0	is the still-air aircraft velocity, and
ω	is the turbulence frequency

By assuming that the turbulence varies linearly over the aircraft's surfaces the aerodynamic effect that is equivalent to an inertial rotation of the aircraft can also be modelled. This leads to spectral densities for the rotational affects of gusts which, for rigid airframes, can be simplified for moderate

angles of attack [15]:

$$\left.\begin{aligned}
\phi_p(\omega) &= \frac{\sigma_w^2}{U_0 L_w}\frac{0.8\left(\frac{\pi L_w}{4b}\right)^{\frac{1}{3}}}{1+\left(\frac{4b\omega}{\pi U_0}\right)^2} \\
\phi_q(\omega) &= \frac{-\left(\frac{\omega}{U_0}\right)^2}{1+\left(\frac{3b\omega}{\pi U_0}\right)^2}\phi_v(\omega) \\
\phi_r(\omega) &= \frac{-\left(\frac{\omega}{U_0}\right)^2}{1+\left(\frac{4b\omega}{\pi U_0}\right)^2}\phi_w(\omega)
\end{aligned}\right\} \qquad (6)$$

where b is the wingspan. Equations (5) and (6) are solved in the time domain by transforming them into canonical state-space form so the turbulent velocity components can be summed to the aircraft's inertial velocity parts prior to solving the equations of motion. For example, in the longitudinal axes the axial and vertical gust perturbations (u_g, w_g) can be written and solved with

$$\left.\begin{aligned}
\begin{bmatrix}\dot{s}_u\\ \dot{s}_{w_1}\\ \dot{s}_{w_2}\end{bmatrix} &= \begin{bmatrix}-\frac{U_0}{L_u} & 0 & 0\\ 0 & 0 & 1\\ 0 & \left(\frac{U_0}{L_w}\right)^2 & -\frac{2U_0}{L_w}\end{bmatrix}\begin{bmatrix}s_u\\ s_{w_1}\\ s_{w_2}\end{bmatrix} + \begin{bmatrix}\delta_u\\ 0\\ \delta_w\end{bmatrix} \\
\begin{bmatrix}u_g\\ w_g\end{bmatrix} &= \begin{bmatrix}\sigma_u\sqrt{\frac{2U_0}{\pi L_u}} & 0 & 0\\ 0 & \frac{\sigma_w}{\sqrt{\pi}}\left(\frac{U_0}{L_w}\right)^{\frac{3}{2}} & \sigma_w\sqrt{\frac{3U_0}{\pi L_w}}\end{bmatrix}\begin{bmatrix}s_u\\ s_{w_1}\\ s_{w_2}\end{bmatrix}
\end{aligned}\right\} \qquad (7)$$

where $s_{(\cdot)}$ are the transfer function states and $\delta_{(\cdot)}$ are the white noise disturbance source. Current airspeed and altitude values throughout the simulation are used to calculate the filters. The values for the turbulence scales are chosen equal (1750 ft), as are the values for each gust intensity in order to satisfy the mathematical requirement for isotropic turbulence [14]. For altitudes above 2000 ft the turbulence intensities, σ, are related to a probability of exceedance: a lower probability represents more severe turbulence, as indicated in Figure 8 .

Implementation—After filtering the white noise the resulting turbulent velocity components are summed with the aircraft's inertial velocity components from the aircraft, prior to the calculation of the aerodynamic forces and moments. This requires a temporary transformation of the aircraft's aerodynamic velocity (α, β, V) from wind to body axes in order to apply the changes. Since there is a correlation between the turbulence in pitch and normal turbulence (w_g, q_g) and in yaw and lateral turbulence (r_g, v_g) [14], the same white noise generator is used for each parameter of the correlation pair.

As far as air turbulence modelling is concerned, the Dryden gust model has become the *de facto* representation for stochastic air turbulence. There are however two limitations with the Dryden model: Firstly the spectrum density decays at ω/V^{-2} at high frequencies [16], greater than the observed rate of $\omega/V^{-5/3}$. This discrepancy would be of concern if high frequency motion, such as structural bending modes, where being considered. The spectral densities for the rotational disturbances are also valid only for low frequencies since the assumption of linear variation of turbulence across the aircraft's surfaces only holds when the wavelength of the disturbance is greater than 8 times the length of the aircraft

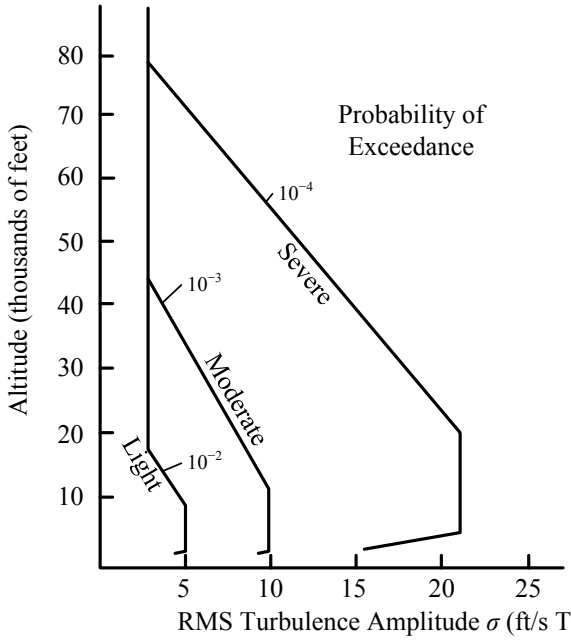

Figure 8. Turbulence severity and exceedance probabilities.

[15]. Secondly, the time-domain-transformed turbulence has, like the white input noise, a Gaussian probability distribution. Atmospheric turbulence is not considered to have a normal distribution; this can be addressed by randomly modulating the filter output to obtain a more realistic probability distribution [17].

Also developed for use in the simulation are a coupled hose-drogue model, with integrated bow wave effects. The inclusion of these elements, the wake vortex model, along with all the elements above serves to provide a high fidelity simulation suited to the proposed technology validation purposes.

4. MANIPULATOR CONTROL

While FDM simulations can be started, stopped, and reset instantaneously, the same is not true of the robot manipulator motions. The supervisory process on the PXIe is responsible for queuing up the robot positions for the simulation being performed, and for determining the end behaviour, including end conditions, halting the robot motion, and returning the robots to the correct positions for the start of the next simulation.

In addition, the important information received from the FDM simulation is simply the *relative* positions of the two pieces of refuelling hardware, represented by 6 degrees of freedom. The supervisory process determines the optimal motion of the two robots and the track to resolve the redundancy offered by these 13 degrees of freedom. In future implementations this motion optimisation will expand the operational envelope of the RMR facility, pushing the limits of the maximum relative accelerations, velocities and positions achievable. For the current work, the individual robot performances are within the required specification, so a datum position is chosen, fixed relative to the tanker aircraft, and the probe and drogue motion relative to this datum are reproduced independently by R2 and R1 respectively. One degree of redundancy remains: the track motion. This is resolved by separating the motion of the drogue into high-

and low-frequency components; the robot axes are used to perform the high-frequency motion and the track moves the robot base to provide the low-frequency, quasi-static response and give the probe its full longitudinal operational range.

As discussed in Section 2, there is an inevitable delay introduced by the motion planning stage in the IRC5 controller. This delay needs to be compensated using predictive control to cast forward the simulation by the equivalent time step. Approximations are inevitably introduced as a result, but previous studies in the context of structural-HIL-style testing have shown that in continuous systems a reasonable approximation can be obtained with a simple polynomial forward predictive capability [18], [19], [20]. This capability is provided in the current implementation of the supervisory process. Future implementations will take into account evaluations of the approaches presented by Chen and Ricles [21] and may also take advantage of more traditional delay compensation methods such as Smith predictors. There is a trade off between temporal and spatial accuracy in the motion reproduced by the robots. Exploratory studies will be made for the final version of this manuscript, quantifying the fidelity of the emulated motion and identifying the optimal point in the trade off.

The start point of the simulations is reached through a smooth transition from the default starting position, 5500 mm directly aft of the datum position (itself 500 mm aft of the drogue canopy starting position). This transition is effected with a triangular velocity profile to accelerate and decelerate uniformly between the default and start positions. The simulation is paused throughout the transition and is commenced from a stationary pose at the start position. The end of the simulations is not currently detected by the supervisory process, but instead safety limits are imposed on the relative positions of the probe and drogue to ensure that the two pieces of hardware do not impact as the simulation progresses through the contact stage.

The robot system has three levels of safety provided as standard: a set of physical joint limit stops and two levels of software limits (for joint positions and tool centre point (TCP) cartesian positions). There is also a motion supervision system which provides emergency stop behaviour in the event of a collision. These safety measures are not comprehensive, however, and high-speed impact of the probe and drogue need to be averted using higher-level control. Furthermore, to optimise motion paths in more advanced testing stages it will be necessary to relax the restrictive joint limits currently in operation. This will require advanced safety control in the supervisory process on the PXIe. In initial tests no contact at all is required between the probe and drogue. As a preliminary safety measure, the probe tip is constrained such that it cannot move forward of the datum position, and the drogue is similarly constrained such that it cannot move aft of the datum. To allow the execution of any simulation without the need for abrupt halts the forward-aft motion of the probe is transformed using the following equation:

$$\hat{x}_2 = (x_2 < x_{safe})\,?\,(x_2)\;:\;\left(\frac{x_{safe}}{2 - x_2/x_{safe}}\right) \qquad (8)$$

where $\hat{x}_2$ is the demanded x-position for R2, x_2 is the respective position from the simulation, and x_{safe} is the distance at which the position modulation begins. The syntax $a?b:c$ denotes the C conditional operator. The output of this function approaches zero smoothly and asymptotically as the demanded position increases past x_{safe}, with a continuous derivative at the point x_{safe} (and elsewhere). Behind x_{safe} the probe motion is mapped directly to the robot motion. An equivalent, negated, expression is derived for the drogue motion. A value of $x_{safe} = -1000$ is chosen arbitrarily for the tests conducted here, where the value represents a distance measured in millimetres.

5. EXPERIMENTAL RESULTS

Some timing data is now presented to illustrate the importance of tuning the communication routes properly, and of optimising the algorithms. Firstly, Figure 9 shows the current case of a 20 Hz motion path update. The stacked bars indicate the times that the respective tasks have taken on the IRC5 controller for each time step. The total height of each bar is proportional to its width, and represents the time for a full cycle to complete. The precision of the measurements is 1 ms, which in some cases is too small to measure a time difference in some of the execution steps. The cycles are divided into six stages: the messgaseCompose stage is where the measurement and timestamp data is aquired and sequenced into a message ready for transmission from the IRC5 controller; the messageSend stage is where this message is sent over the TCP/IP link; the messageReceive stage is where the position demand and control data are revceived over the TCP/IP link; the messageParse stage is where the received message is transcribed into the appropriate variables on the controller; the moveInstruction stage is where the motion command is issued to the motion planning routines in the IRC5 controller, and finally the cycleTime stage is simply the time for the cycle to return to the beginning of the loop. The latter stage takes negligible time but sometimes appears as a millisecond in the figures presented as an artefact of rounding errors.

In the case presented in Figure 9, the communication loop was locked to the main cycle. That is, the only time the PXIe was sending messages to the IRC5 was when a motion command needed to be sent. Accordingly, every cycle takes approximately 50 ms, corresponding to a 20 Hz cycle rate. It is by chance that the bottleneck in the case presented is waiting for the latest sample to be buffered from the FDM simulation by the supervisory process. That is why the IRC5 time data indicates a large proportion of the cycle time is spent waiting to receive a message from the PXIe. If the time was not allocated here, it would be spent waiting to buffer the position data when the move instruction was executed in the IRC5. It can be seen that the composition of the message string to send takes a comparable time to that taken to read the binary message received from the PXIe.

Figure 10 shows an earlier test performed at 10 Hz. In this case, however, the communication cycle is not locked to the position instruction cycle, and intermediate communications relay measurements to the PXIe. In this case a standard communication cycle takes around 10 ms. The cycles where a move instruction is received at the IRC5 can be seen, as the cycle takes longer, with the time taken to process the move instruction accounting for the difference. These cycles are spaced at approximately 0.1 s intervals as expected, and the time spent processing the move command is the time needed to regulate the motion timing at the IRC5 end.

Finally, Figure 11 shows another example, but this time no position data is sent (i.e. the COMM index remains at zero). In this case, however, the messages from the PXIe to the IRC5 are encoded as ASCII strings of numeric values, requiring string parsing on the IRC5. It can be seen here that the

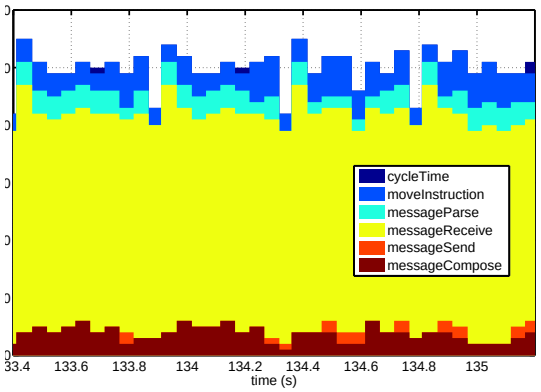

Figure 9. Stacked bar chart depicting breakdown of the times for the tasks in each communication cycle on the IRC5 controller. This data is for a 20 Hz motion rate. The precision of the time measurements is 1 ms. (Case 1)

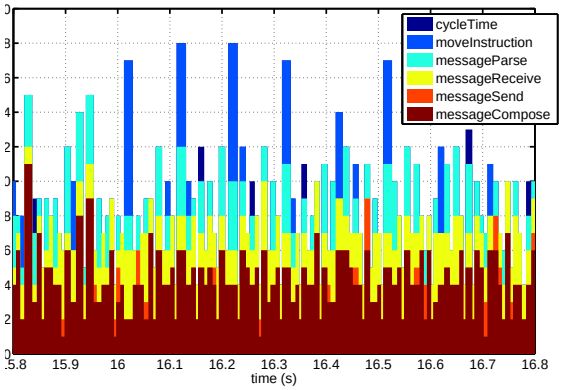

Figure 10. Stacked bar chart depicting breakdown of the times for the tasks in each communication cycle on the IRC5 controller. This data is for a 10 Hz motion rate, after changing to a binary communication protocol. The precision of the time measurements is 1 ms. (Case 2)

message parsing on the IRC5 takes around 30 ms, far greater than the 10 ms or less taken to decode an equivalent binary message in Figure 9. In the final figure the receive stage of the cycle also takes around 30 ms. This was found to be caused by diagnostic console output being written to the screen on the PXIe, slowing down the cycle time. The average times for the six stages over the three cases are given in Table 2

The results presented in this section illustrate the response characteristics of the robot motion. It was expected that a small delay would be observed as a result of the motion path buffering, and that artefacts of the interpolation around position data points would be seen. What was not clear in advance was what the dynamic response of the motors and their proprietary feedback/feedforward controllers would be. To test these effects, a predetermined motion path was implemented. For the tests described herein a script reads position and orientation data from an ASCII file and executes the corresponding motion instructions at a rate of 50 Hz. The target points are provided from an AAAR simulation. Measurements of the actual robot positions along with time stamps provided by the robot controller are streamed over a TCP/IP connection, in this case to a separate PC, at a

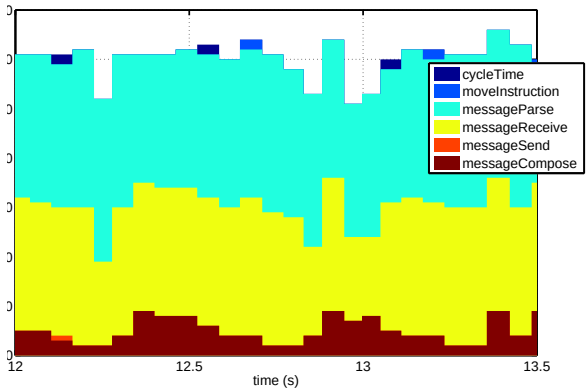

Figure 11. Stacked bar chart depicting breakdown of the times for the tasks in each communication cycle on the IRC5 controller. This data is for a 10 Hz motion rate, using an ASCII communication protocol, prior to changing to a binary communication protocol. The precision of the time measurements is 1 ms. (Case 3)

Table 2. Average times, in milliseconds, for the six stages of the communication and control loop on the IRC5 robot controller, for the three cases described in Figures 9-11.

	Case 1	Case 2	Case 3
messageCompose	3.75	4.09	4.96
messageSend	0.53	0.21	0.04
messageReceive	37.78	2.23	25.14
messageParse	3.15	1.33	29.56
moveInstruction	4.78	0.56	0.2
cycleTime	0.08	0.1	0.24

rate of 50 Hz, and recorded. The resulting data can not be synchronised with the move instructions, but for this analysis has been aligned with the prescribed motion path by minimising an error function. It is thus not possible to identify the static delay component of the response, but all other features of the response should be apparent.

Figure 12 shows the absolute position data, including the 3DOF translational position output from the simulation overlaid with the measured response of the robots. At this scale the lines appear coincident. The simulation shows a position hold approximately 5 m aft of the drogue, followed by an approach to the pre-contact position, which is again held approximately 2 m aft of the drogue, and finishing with an aggressive engagement. The simulated refuelling procedure is conducted in light turbulence.

Figure 13 shows the position error from these plots. It is interesting to note that any effects from corner path artefacts are indiscernible in the presence of other disturbances. The large peak at approximately 30 s corresponds with the rapid approach of the receiver aircraft to the pre-contact position. A similar sharp rise is seen at the end of the plot where the final engagement is made. Although these errors are presented here as positional errors, they are found to be better described as temporal discrepancies; the differences seen are the result of a lag between the demanded motion and the measured robot position when moving at high speeds. What is not apparent in this figure, but can be determined from close examination of Figure 12, is that while the robot motion lags the demand at some points, it leads the demand at others. This

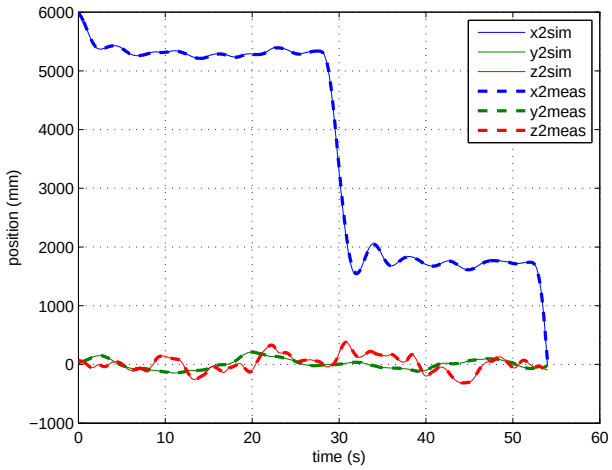

Figure 12. Absolute positions from the simulated data and from the measured robot positions.

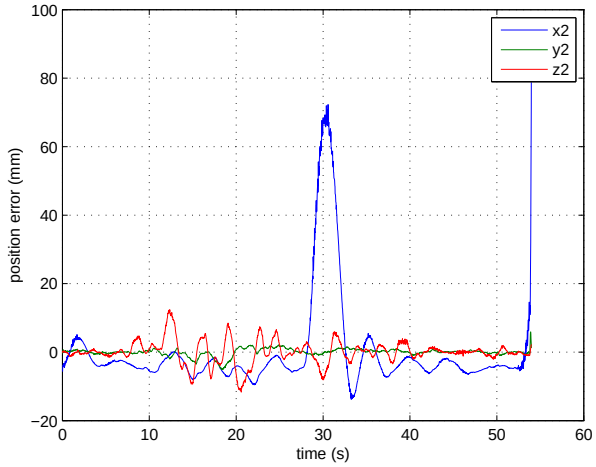

Figure 13. Position error for the robot motion relative to the prescribed simulation data motion path.

lead may not be a true lead, as the alignment of the two signals in this case is not guaranteed, but it nonetheless points to a variable frequency response that may be characterised to the ends of further improving the performance with an inverse approach.

6. Future Work

Development of the RMR facility is ongoing and provides varying levels of functionality even in development. Previous uses have included validation of simulation behaviour in quasi-real time and the evaluation of drogue-tracking algorithms recorded with a mid-range grey-scale camera. Future projects highlight the flexibility of the facility: these include vision-tracking of satellite systems for orbital docking control and various composites-based research projects. To achieve a fully closed loop system the interface between the PXIe and the a suitable vision sensor system will be established, and the performance of the interface between the PXIe and IRC5 robot controller will be fully characterised and managed. This will provide a complete HIL capability for evaluating real sensor systems in a high fidelity refuelling simulation environment. Particular effort will be put into minimising the inevitable delays induced by the kinematic computations and implementing delay compensation strategies.

In terms of the simulation model further areas of research include better modelling the aerodynamic characteristics of the hose and drogue system, including nonlinear effects from hose whip and any from the hose drum unit. The influence on mass variation on the flight mechanics of the receiver should be addressed. At moderate levels of turbulence and in the presence of the wake vortex the aeroelastic effects could be modelled. The bending moments and loading forces on the extended probe could be investigated, which so far has been assumed rigid. Similarly, aeroelastic effects on the drogue could be modelled and captured in the HIL simulation.

7. Conclusions

An overview of the Relative Motion Robotic (RMR) facility at the University of Bristol, developed in collaboration with Cobham Mission Equipment, has been given and the important considerations in implementation and performance optimisation have been discussed. Timing is critical in structural hardware in the loop simulations, and factors affecting performance in this regard have been described. Steps taken to improve the performance and to push the limits of the equipment being used are documented, and preliminary results from a simulated air to air refuelling exercise have been presented. These results demonstrate the suitability of the facility for conducting advanced tests of aerial refuelling hardware and sensors for the purpose of developing automated aerial refuelling capabilities. A high fidelity multi-entity flight dynamics model has been developed and the results to date have highlighted key areas for investigation as being advanced delay compensation strategies and further motion response characterisation of the RMR to facilitate inverse feedforward control to improve on high speed motion reproduction.

Acknowledgments

This work is funded by Cobham Mission Equipment as part of the ASTRAEA Programme. The ASTRAEA programme is co-funded by AOS, BAE Systems, Cobham, EADS Cassidian, QinetiQ, Rolls-Royce, Thales, the Technology Strategy Board, the Welsh Assembly Government and Scottish Enterprise. Website: http://www.astraea.aero/

References

[1] R. K. Smith, *Seventy-Five Years of Inflight Refueling: Highlights, 1923-1998*. United States Government Printing Office Superintendent of Documents, 1998.

[2] W. I. Glover, "Wings across the Atlantic," *Popular Mechanics*, vol. 55, no. 2, pp. 186–191, February 1931.

[3] M. A. Bennington and K. D. Visser, "Aerial refueling implications for commercial aviation," *Journal of Aircraft*, vol. 42, no. 2, pp. 366–375, 2005.

[4] R. Nangia, "Operations and aircraft design towards greener civil aviation using air-to-air refuelling," *The Aeronautical Journal*, vol. 110, pp. 705–721, November 2006.

[5] L. Pollini, M. Innocenti, and R. Mati, "Vision algo-

rithms for formation flight and aerial refueling with optimal marker labeling," in *AIAA Modeling and Simulation Technologies Conference and Exhibit*, August 2005, pp. 1–15.

[6] W. R. Williamson, G. J. Glenn, V. T. Dang, J. L. Speyer, S. M. Stecko, and J. M. Takacs, "Sensor fusion applied to autonomous aerial refueling," *Journal of Guidance, Control, and Dynamics*, vol. 32, no. 1, pp. 262–275, 2009.

[7] A. Dogan, S. Sato, and W. Blake, "Flight Control and Simulation for Aerial Refueling," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, August 2005, pp. 1–15.

[8] J. Wang, V. Patel, Vijay, C. Cao, N. Hovakimyan, and E. Lavretsky, "Novel L_1 adaptive control methodology for aerial refueling with guaranteed transient performance," *Journal of Guidance, Control, and Dynamics*, vol. 31, no. 1, pp. 182–193, 2008.

[9] D. Saban, J. Whidborne, and A. Cooke, "Simulation of wake vortex effects for uavs in close formation flight," *Aeronautical Journal*, vol. 113, pp. 727–738, 2009.

[10] W. Mao and F. Eke, "A survey of the dynamics and control of aircraf during aerial refueling," *Nonlinear Dynamics and Systems Theory*, vol. 8, no. 4, pp. 375–388, 2008.

[11] B. L. Stevens and F. L. Lewis, *Aircraft Control and Simulation*, 2nd ed. John Wiley & Sons, USA, 2003.

[12] L. T. Nguyen, M. E. Ogburn, W. P. Gilbert, K. S. Kibler, P. W. Brown, and P. L. Deal, "Simulator study of stall/post-stall characteristics of a fighter airplane with relaxed longitudinal static stability," NASA TP 1538, Washington, D. C., December 1979.

[13] K. M. Dorsett and D. R. Mehl, "Innovative control effectors (ICE)," Wright-Patterson Airforce Base, Ohio, WL-TR-96-3043, January 1996.

[14] MIL-F-8785C, "Flying qualitites of piloted aircraft," Department of Defense, USA, 1980.

[15] J. Roskam, *Airplane Flight Dynamics and Automatic Flight Controls*. Roskam Aviation and Engineering Corporation, 1979.

[16] S.-T. Wang and W. Frost, "Atmospheric turbulence simulation techniques with application to flight analysis," NASA CR 3305, 1980.

[17] C. G. Justus, C. W. Campbell, D. M. L., and D. L. Johnson, "New atmospheric turbulence model for shuttle applications," NASA TM 4168, 1990.

[18] J. du Bois, B. Titurus, and N. Lieven, "Transfer dynamics cancellation in real-time dynamic substructuring," in *International Conference on Noise and Vibration Engineering, ISMA2010*, Leuven, Belgium, 2010.

[19] M. Wallace, D. Wagg, S. Neild, P. Bunniss, N. Lieven, and A. Crewe, "Testing coupled rotor blade-lag damper vibration using real-time dynamic substructuring," *Journal of Sound and Vibration*, vol. 307, no. 3-5, pp. 737–754, 2007.

[20] M. Wallace, D. Wagg, and S. 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*, vol. 461, no. 2064, p. 3807, 2005.

[21] C. Chen and J. Ricles, "Analysis of actuator delay compensation methods for real-time testing," *Engineering Structures*, vol. 31, no. 11, pp. 2643–2655, 2009.

BIOGRAPHY

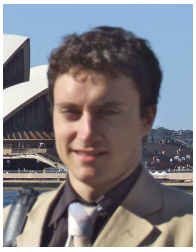

Peter Thomas is a member of the Dynamics and Systems Research Group at the University of Bristol where he is presently working on flight control systems and simulation for autonomous aerial refuelling. Previously he completed his PhD at Cranfield University in 2010, investigating applications of low cost tools and techniques for flight testing and system identification of UAVs. He received a MEng in Mechanical Engineering from Warwick University in 2006 and is an Associate Member of the IMechE. His research interests include flight dynamics simulation, aircraft modelling and system identification, robotic biomimetics, and control theory.

Jonathan DuBois studied as an undergraduate at the University of Bristol and the Kungliga Tekniska Hogskolan in Stockholm. He graduated with an MEng in Aerospace Engineering in 2002 and after two years pursuing competitive sporting interests he returned to study for a PhD in Active Fuselage Response Suppression in Rotorcraft. In 2008 he spent three months as a visiting academic at the University of California, San Diego. Presently he is a member of the Dynamics and Systems research group and AgustaWestland Helicopters University Technology Centre in Vibration, working on the Rotor Embedded Actuator Technology programme. Jonathan's research interests include: Novel methods in parametric dynamic analyses, Experimental parameter identification and Finite Element model updating problems, Eigenvalue curve veering, eigenvector localisation, and modal coupling, Active and adaptive vibration control, and Structural Health Monitoring Substructuring control problems (hybrid numerical-hardware simulations)

Tom Richardson lectures Flight Mechanics and Control in the Aerospace Department at Bristol University. He obtained his first degree (MEng) from the Department at Bristol in 1998 and stayed on to undertake a PhD in Nonlinear Flight Mechanics. Awarded his PhD in 2002 he moved over to the Mechanical Engineering Department in Bristol as a Research Assistant. In June 2003 he moved back to the Aerospace Department when he was appointed to the Lectureship in Flight Mechanics. Dr Richardson is a member of the Dynamics and Systems Research Group. His principal research interest is control, ranging from classical flight control up to high-level decision-making and autonomy. He has been a member of two GARTEUR action groups. In the first, (1999-2002) he was a member of a team which addressed the problem of flight clearance using a combination of Continuation and Optimisation methods. More recently he has been working with members from Engineering Mathematics to address the problem of autonomy within groups of unmanned air vehicles.