Proceedings of the 9th U.S. National and 10th Canadian Conference on Earthquake Engineering
Compte Rendu de la 9ième Conférence Nationale Américaine et
10ième Conférence Canadienne de Génie Parasismique
July 25-29, 2010, Toronto, Ontario, Canada • Paper No 1024

UK-NEES - DISTRIBUTED HYBRID TESTING BETWEEN BRISTOL, CAMBRIDGE AND OXFORD UNIVERSITIES: CONNECTING STRUCTURAL DYNAMICS LABS TO A GEOTECHNICAL CENTRIFUGE

M. Ojaghi[1], I. Lamata Martínez[1], M. Dietz[2], M.S. Williams[1], A. Blakeborough[1],
A. Crewe[2], C. Taylor[2], G. Madabhushi[3], S. Haigh[3], A. Ali[3]

ABSTRACT

Distributed hybrid testing is a natural extension to and builds upon the local hybrid testing technique. Taking advantage of the hybrid nature of the test, it allows a sharing of resources and expertise between researchers from different disciplines by connecting multiple geographically distributed sites for joint testing. As part of the UK-NEES project, a successful series of three-site distributed hybrid tests have been carried out between Bristol, Cambridge and Oxford Universities. The first known multi-site distributed hybrid tests in the UK, they connected via a dedicated fibre network, using custom software, the geotechnical centrifuge at Cambridge to structural components at Bristol and Oxford. These experiments were to prove the connection and useful insights were gained into the issues involved with this distributed environment. A wider aim is towards providing a flexible testing framework to facilitate multi-disciplinary experiments such as the accurate investigation of the influence of foundations on structural systems under seismic and other loading. Time scaling incompatibilities mean true seismic soil structure interaction using a centrifuge at *g* is not possible, though it is clear that distributed centrifuge testing can be valuable in other problems. Development is continuing to overcome the issues encountered, in order to improve future distributed tests in the UK and beyond.

Introduction

The substructuring techniques of hybrid testing (Blakeborough et al. 2001) allow multiple physical and numerical parts of a hybrid experiment to be coupled together however, with distributed testing these parts are no longer necessarily located in the same lab. Instead, through utilising computer networks, multiple geographically distributed testing facilities may collaborate in a joint hybrid experiment, not only much larger than could be independently achieved but, also exploiting the specialist technical expertise and facilities available at other sites. The correct application of this technique is an extremely significant step, especially for

[1] Dept. of Engineering Science, University of Oxford, UK.
[2] Dept. of Civil Engineering, University of Bristol, UK.
[3] Dept. of Engineering, University of Cambridge, UK.

large scale seismic (and similar) testing. There is a real potential to conduct experiments which would not have been otherwise realistically possible either due to sheer size and expense, or due to a lack of expertise and facilities available at just one site. Distributed hybrid testing offers significant advantages but, also introduces new and challenging problems, namely: reliable online data transfer between legacy hardware systems, variable network delay and associated synchronization issues.

UK-NEES

In 2006, in a similar vein to NEES (Network for Earthquake Engineering Simulation), development work began on the UK-NEES, a collaborative grid network setup with the aim of extending seismic and similar testing capabilities within the UK (Ojaghi et al. 2007). UK-NEES aims to facilitate a sharing of expertise in a spirit of joint collaboration to encourage research into earthquake engineering and related topics, while also raising awareness of earthquake engineering issues within the UK. As with NEES, (Spencer et al. 2004), currently one of the main research themes is the development of two site and multi-site distributed hybrid testing techniques at a variety of timescales up to and including real time.

To facilitate the new testing requirements, tele-presence facilities have been installed and work is progressing on a shared data repository and web presence for the network. Attention has also been given to understanding the new user environment in order to cater for the usability challenges of distributed hybrid testing (De la Flor et al. 2009). The UK-NEES network currently connects three of the main earthquake engineering research centres in the UK at Bristol, Cambridge and Oxford universities: each providing complimentary facilities for the network. Bristol, the largest earthquake engineering research centre in the UK hosts, a six degree of freedom shaking table, an arrangement of approximately 200m^2 of reaction walls and a large array of hydraulic actuators. Oxford specialises in real time hybrid testing and hosts a set of 6 high performance hydraulic actuators and a strong floor. Cambridge focuses on research into tunnels and foundations, notably hosting a 10m beam centrifuge capable of achieving a centrifugal acceleration of approximately 125g at 4.125m. While the network has access to a wide range of testing facilities a wider aim is to form part of an international grid connecting with other national NEES type networks, fostering international research collaboration and, taking part in joint distributed testing.

On the 19th February 2009 a series of three site tests were carried out on the UK-NEES network connecting physical substructures at the earthquake engineering lab at Bristol University and the geotechnical centrifuge at Cambridge University with a physical and numerical substructure located at the structural dynamics lab in the University of Oxford. The tests represented a culmination of development work carried out since the start of the project, over a year of two site testing between Oxford and Bristol and, several weeks of testing between Oxford and Cambridge. The aim of the tests were to prove the connection with stable tests and thus, exploring the capabilities for distributed testing, discovering with a view to overcoming any shortcomings with the developed techniques and, in general terms better understanding the issues encountered with distributed multi-site site testing.

This article first discusses the role of the distributed hybrid testing technique within earthquake engineering testing and the applicability of distributed centrifuge testing. Next, a three site distributed hybrid test conducted is introduced together with details of the distributed architecture and communication workflow, before presenting some test results. The differences and challenges of the multi-site distributed environment, both technical and in regards to usability are then briefly discussed, as derived from the experiences gained during the development of the test procedure. In concluding, test learning experiences are outlined together with brief discussions of current work progressing to overcome the shortcomings of the test.

The Role of Distributed Testing

The main impetus behind local hybrid testing is to allow the realistic seismic response of structures to be simulated and, the hybrid nature means that the experiment is split between physical and numerical parts. This is advantageous as numerical modelling of earthquake engineering components alone can be inadequate and the alternative - large scale testing, costly with experiment size severely restricted due to capacity. Even large scale shaking table facilities, the largest of which is E-Defence, while a valuable resource in itself may only test at full scale, medium sized structures - most other shaking tables have significantly lower capacities.

Hybrid testing therefore serves as an alternative technique when the limitations of fully numerical simulation may mean results are not adequate or when fully physical testing is not feasible. Since only certain parts of the structure are physically tested with the remainder modelled numerically, larger structural systems may be tested at full/large scale without meeting the capacity limits of the local testing facility. The technique allows the global seismic response of the structural system to be found with numerical modelling deemed adequate for 'known' parts of the structure and placing emphasis on the local performance of the physical components under development. Thus, hybrid testing develops a synergy between fully physical structural testing and fully numerical simulation. However, with the need for advances in earthquake engineering knowledge and the development of new technologies leading to a growing requirement for testing ever larger and complex structural systems (Nakashima, 2008), the resources at anyone site may become saturated.

Distributed hybrid testing serves to further extend the application of the local technique to not only, allow bigger and large scale substructured experiments, with a single experiment spread over and using the testing facilities of more than one site but, encourages a sharing of expertise and specialist equipment. Within the UK-NEES network it is envisaged that in connecting for distributed testing equipment, including actuator arrays, shaking table and the geotechnical centrifuge, a flexible collaborative testing framework is provided for researchers, maximising the potential of the network to facilitate experiments that may not have been otherwise possible. Testing facilities may be brought online as necessary for conducting complex multidisciplinary experiments for example, large scale seismic soil-structure interaction, attempting to address problems that may be overlooked by structural or geotechnical testing labs alone (Fig. 1).

While large scale multi-site experiments are realistic aims within the UK-NEES network and present exciting possibilities for advancing earthquake engineering and related technologies there are many challenges yet to overcome. The first task is to achieve the connections between the facilities available at each site. In the experiment to be described, existing testing equipment and control units for the actuator arrays at Bristol and Oxford were connected via the Internet, together with the main testing facility at Cambridge, the geotechnical beam centrifuge.

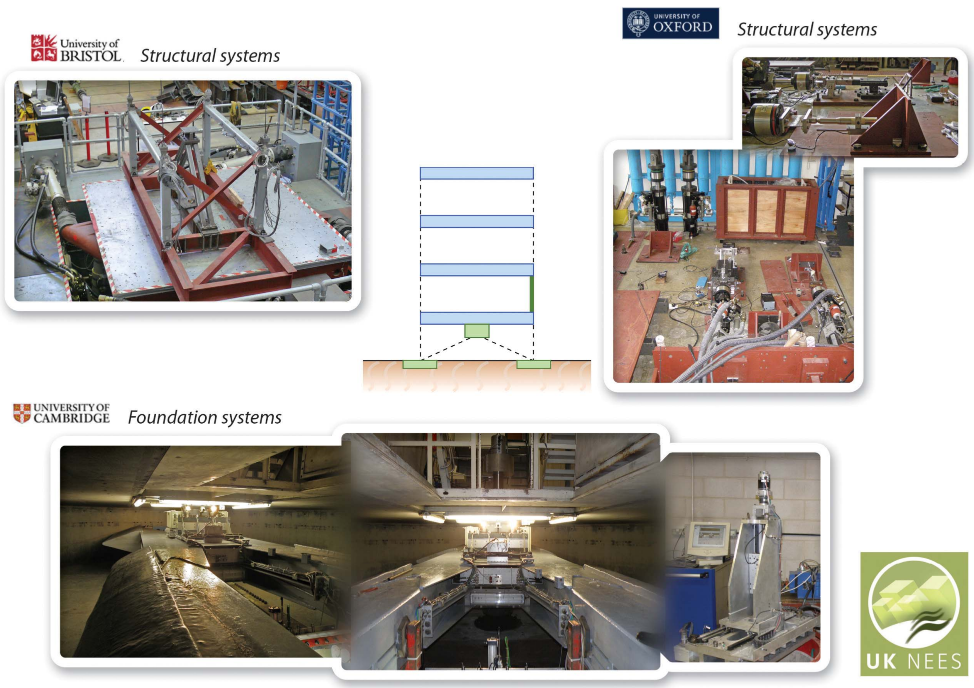

Figure 1. UK-NEES distributed testing framework

Distributed Centrifuge Testing

It is important to highlight the technical incompatibilities between true seismic soil structure interaction as was referred to previously and distributed centrifuge testing. In testing structural systems, experiments may either be carried out in real-time or where rate effects may be neglected, slower than real time. The behaviour of soil (and similar materials) is dependent upon self weight and geotechnical centrifuge testing takes advantage of this to allow typically small scale models of soil when driven at high values of *g* (gravitational acceleration) to represent much larger volumes of soil but, with equivalent stress states. Hence, realistic features including failure models may be simulated without resorting to the otherwise large scale models required. Typical scaling rules mean that a prototype model of a 1m depth of soil at 50g is equivalent to 50m depth of soil. However, as length is scaled in this manner and, as soil densities between full scale and model are essentially the same or cannot be greatly varied, to correctly simulate inertial effects - as induced by an earthquake, time is scaled. Centrifuge time is 1/N of real earthquake time, where N is the number of gravities used. Hence, at 50g, a 50 second earthquake will last only 1 second. To simulate effects such as liquefaction, the pore fluid - water, will often be replaced with a viscous fluid. This is as while pore pressure generation is driven by inertial effects, dissipation is slowed to allow parity between the dynamic time scale and the consolidation time scale ($1/N^2$), (Madabhushi and Schofield 1993). Therefore, experiments with physical parts distributed to structural labs, designed to study seismic soil structure interaction must take place in real time, at 1g and scale models of foundation system tests will represent with reduced accuracy the true soil stress states. These distributed seismic centrifuge experiments are not possible but, if pore pressure generation is correctly applied, the resulting soil stress state is suitable for testing in situations where inertial effects are not significant. For example, harmonic loading due to waves may be applied in an experiment involving a physical substructure of an offshore platform or offshore wind turbine foundation and the experiment may be distributed with another structural unit connected together with a numerically modelled superstructure.

Three Site Test Setup

Towards development of a successful testing framework for accurate distributed hybrid testing, a simplified experiment is devised (Fig 2.). The experiment is designed to test a multi-site distributed testing architecture connecting actuator arrays and physical substructures at Oxford Bristol and Cambridge. It does not aim to model a particular or real engineering problem but, to prove the connections between the sites, further identifying and quantifying issues related to this new distributed environment in order to form a basis on which to develop the testing technique. The test couples a linear five degree of freedom shear structural model experiencing ground motion to three physical substructures. Locally at degrees of freedom 2 and 3, the Oxford 2 storey column rig (Bonnet et al. 2006) is placed. Distributed on the 1st degree of freedom, to allow direct comparison between the distributed sites, the Bristol column rig and the Y direction

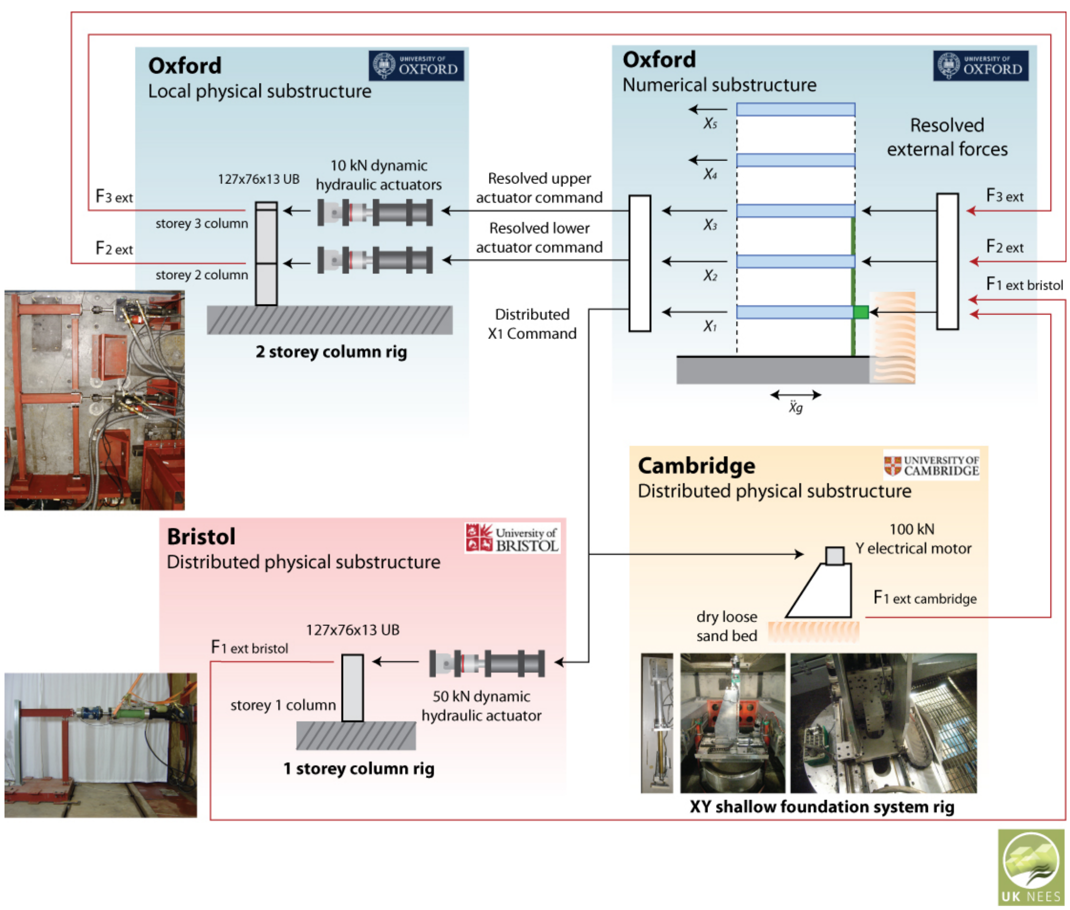

Figure 2. Multi-Site distributed hybrid test setup.

of the Cambridge XY shallow foundation pad rig resting atop a bed of loose dry sand is connected. The foundation pad is controlled using two electrical motors providing linear motion or rotation via a gearing system. While seismic loading is applied, the lifting behaviour of the foundation pad may be similar to that experienced by other types of loading.

UK-NEES Three Site Distributed Hybrid Testing Architecture

The first major challenge to accomplishing this distributed hybrid test, is to connect the different hardware systems at each site and to also ensure that data transfer between them is robust and reliable. This is especially challenging as none of the hardware systems have been designed with such a distributed architecture in mind and remote signals are not treated with any special priority, unlike high priority internal control signals. The UK-NEES distributed architecture as used for the three site test is shown in Fig. 3. This client server architecture has been found to be

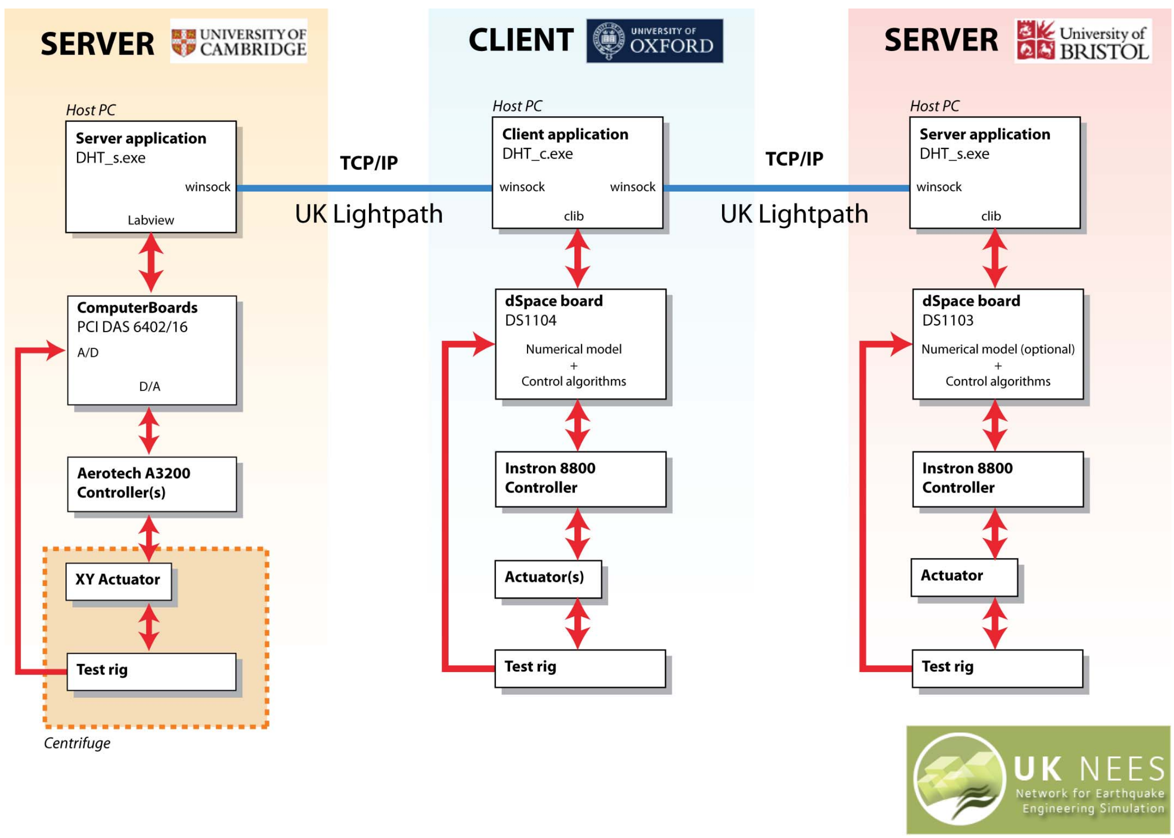

Figure 3. High level representation of the UK-NEES three site distributed architecture used.

the most prudent, following Saleem et al. 2008 and in this view the Oxford node acts as a client hosting a physical substructure as well as the main numerical model connecting all the physical substructures. The Bristol and Cambridge nodes act as servers, each hosting physical parts of the test. The architecture developed uses existing hardware and software and adapts these legacy systems for distributed use by incorporating changes in the control systems and using the UK-NEES Distributed Hybrid Testing communications program, DHT (thus far unreleased) to make the software connections between sites. In this model the UK JANET Lightpath network, a dedicated fibre network connecting the sites is used. Alternatively the Internet may be directly used. Distributed testing poses unusual network requirements: network latency is of greater importance than bandwidth. The Lightpath network was installed to ensure that during a distributed test network usage fluctuations would not interfere with the test.

Local testing environments

Bristol and Oxford share similar testing hardware and software. They both use hard time single tasking processor boards. This allows numerical models and control software to be run onboard with a high resolution hardware clock ensuring accurate and consistent time-steps. The boards, hosted on a Windows machine, directly command the actuator controller and control signals are fed back to them. Network communication may be achieved with the boards by using the dSpace, Clib and Windows, Winsock API's. In testing, both Oxford and Bristol use dynamic hydraulic actuators and have the capability to run real-time hybrid experiments.

As real-time testing is not a priority, Cambridge uses high load capacity electrical motors with gearing, that fulfil power requirements and while there are significant velocity restrictions they are relatively compact - as required for use on the centrifuge basket. The Cambridge systems run LabView on a Windows (multi-tasking) environment to allow communication with a Computer Boards A/D board, regulating time-steps using a software based timer. While LabView is used to interface with the A/D board, direct access to the memory registers of the board is possible via a Computer Boards software library. A LabView based program interfaces via common read and writes files (memory or disk based) with the DHT program, receiving commands to pass to the actuator controller and transferring feedback control signals to the DHT program.

UK-NEES Distributed Hybrid Testing Program, DHT

The distributed environment between Bristol and Oxford is relatively well developed allowing stable communication with 20ms time-steps. Part of the aim of this test was to extend distributed testing capabilities to Cambridge and in this first instance explore the capabilities of the developed distributed environment. In order to achieve communication a custom distributed hybrid testing program, DHT was written and is under development. DHT is specifically written to maximise efficiency. Making use of standard libraries to transmit the required control signals between sites, DHT aims to enable multi-platform connections with little end user customisation. The version of DHT used for this test was adapted for multi-site testing and a high level representation of the workflow is shown in Fig. 4. In this multi-site environment, feedback control signals must be received from both sites before the next command is sent from the client site and, the rate of communication governed by the slowest site - in this case Cambridge. While this synchronous mode of operation is suitable to ensure accurate tests other DHT versions use more robust asynchronous operations and can tolerate some loss of data from individual sites. Communication speed is mainly governed by the PC that is used to host LabView. The PC used outside the centrifuge for 1g testing would allow reliable communication at 200ms time-steps (through use of a read delay loop) however, the centrifuge PC would only permit reliable communication at 1s time-steps. Send operation ordering as depicted does not unduly affect the test performance due to the relative time required to carry out a send operation. Communication ends when the client sends an 'end test' value that both servers interpret and acknowledge.

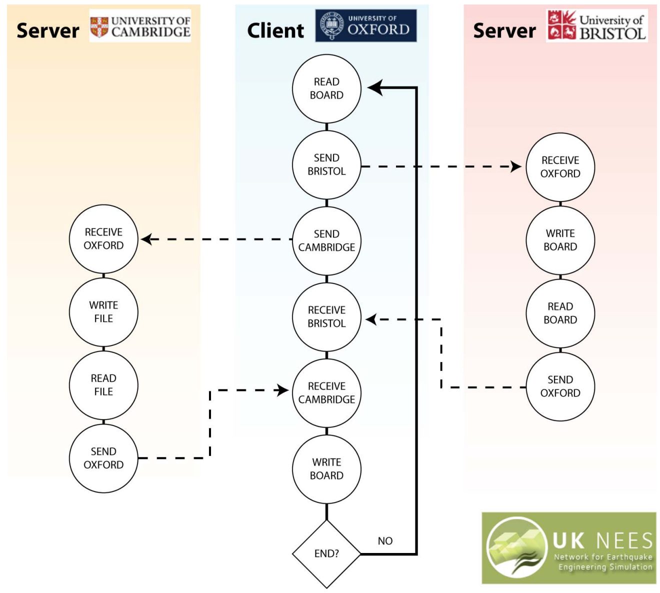

Figure 4. The workflow used by the multi-site DHT program.

The Preliminary Test: Results and Discussions

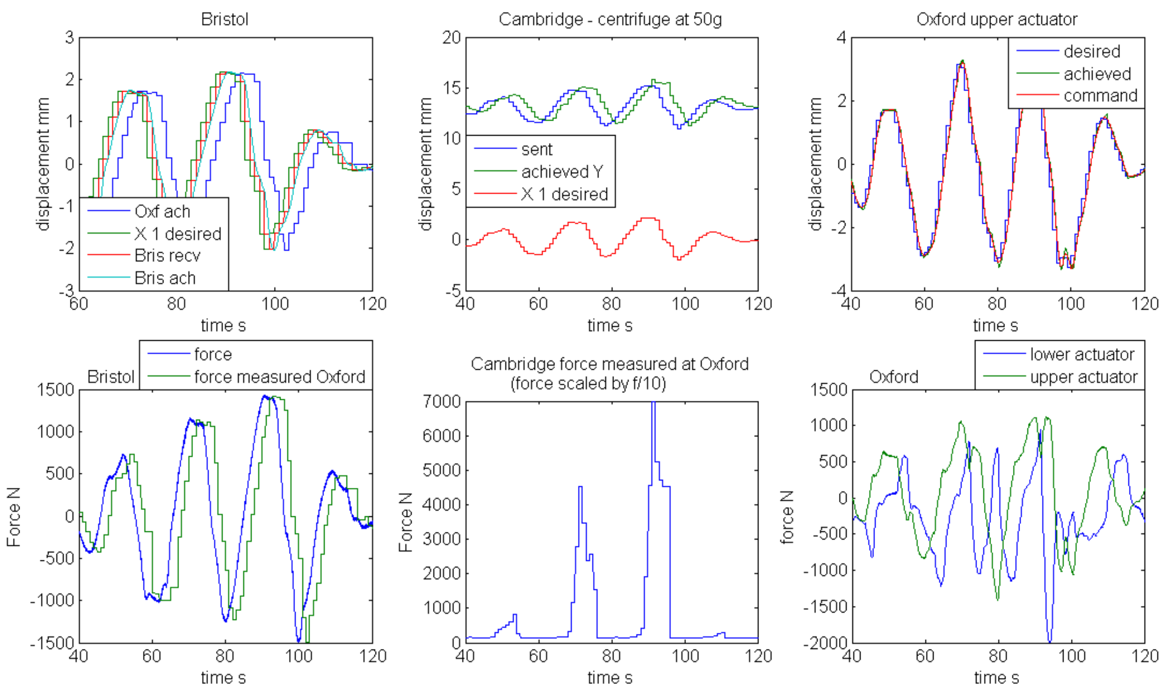

Figure 5. Test results from a distributed multi-site hybrid test using 1s time-steps.

A series of hybrid tests with the centrifuge at 1g and 50g were conducted. Fig. 5 shows results from one seismic test. Here the linear modal model (fixed step, multi-rate with no predictor) is used to provide command signals to the respective actuators as in Fig 2. A third order time-step independent polynomial extrapolator/interpolator is used to smoothly actuate and compensate for local actuator delays. Delay compensation for the remote parts was not attempted in this test to demonstrate the maximum extent of the delays encountered. However, actuator sub-stepping was achieved at Bristol using a 2nd order polynomial interpolator – which has been found to perform better in the distributed environment. At Cambridge on receiving commands, the local actuator controller smoothly actuates towards them and the return force is scaled for stability and compatibility with the rest of the model. While for dynamic hybrid testing the time-step of the numerical model is chosen to adequately represent the dynamic system required, in this case the time-step was chosen as 1 second for technical reasons and the experiment time scaled to slow time by a factor of 20. While the tests are stable, the level of delay is quite large - up to 4/5 time-steps at points and the overall delay dependent on both sites.

Distributed vs Local Hybrid Testing

A typical local hybrid test progresses in time-stepping fashion. A numerical model (and specifically in a multi-time step strategy) which may be modal or implicit (using a predictor/compensator and correction algorithm) or explicit, is used to output the next desired command to the physical parts. These are smoothly and continuously actuated at a control time-step to the desired position (in displacement control) through the numerical model time-step and the achieved force at the end of this time-step is fed back to the numerical model to be used to calculate the next desired displacement. A complication (of more significance in real-time testing) arises due to the response time and 'inner loop' control of hydraulic and electrical actuators subject to input commands. Often regarded as a phase delay between commanding and achieving a value, it depends on the relative performance of the actuator and the properties of the physical part. The delay will vary and an amplitude error may also exist, with compensation techniques used to overcome these (Bonnet, 2006).

In a distributed environment, test complexity increases. There are restrictions on data transfer rates between sites and an additional variable data transport delay is encountered when sending commands and receiving feedback. Additionally there is an issue of time-step synchronization between sites: while local and remote time-steps cannot be controlled to exactly coincide, their relative position must be taken into account and may be a source of an additional delay in the feedback loop. Finally, data transfer in the distributed environment does not share en-route the same high priority that it is locally accustomed to due to the hardware infrastructures in use. However, as large time-steps were used for communication, data loss due to hardware saturation was not significant, though it is clear the soft time, multitasking environment is less robust, especially if smaller time-steps are used. 'Data loss' strategies must be adopted, to mitigate this.

Usability: The New Distributed User Environment

The distributed user environment presents usability challenges. Hybrid testing labs are often noisy and can be stressful working environments. The local test operator has to contend with the control of multiple user interfaces and systems while ensuring that the physical component is tested safely and correctly - especially in regards to samples that break. In moving to a two and multi-site test setting, additional systems are overlaid and merged with existing ones, greatly increasing the chances of operator error. Therefore, there is a need to develop systems and software designed to assist the operator to deal with the challenges of the distributed user environment. This became clear in early two site tests and a test strategy based on a 'traffic light' scheme and user interface to guide the operator has been developed and is successful for two site testing between Bristol and Oxford. In three site testing, two remote sites must be handled by the client operator. As the main test controller, operational demands are greatly increased and further improvements to the developed two site strategy are required to suit multi-site testing.

Conclusions and Current Work

A successful series of three site distributed hybrid tests have been carried out connecting a foundation system inside a geotechnical centrifuge to two distributed structural components linked via a common numerical model. These stable experiments tested the connections, developing a better understanding of the distributed multi-site environment. The next step is to ensure accurate tests. The large delay is the major factor leading to inaccurate test results and problems with stability as system damping is reduced or the influence of the physical portions of the test are increased. Current work has been focused on eliminating this delay through the use of explicit numerical models, developing new large delay compensation strategies and improvements to the architecture of the DHT program to ensure faster data transfer.

References

Blakeborough A., Williams M.S., Darby A.P., Williams. D.M., 2001. The Development of Real-Time Substructure Testing, *Phil. Trans. R. Soc. London*, 359, 1869-1891.

Bonnet P.A., 2006. The Development of Multi-Axis Real-Time Substructure Testing. *DPhil. Thesis* University of Oxford.

Bonnet P.A., Williams M.S, Blakeborough A., Ojaghi M., 2006. Real-Time Hybrid Earthquake Simulation of a Steel Column in a 20-Storey Building, *First European Conference on Earthquake Engineering and Seismology*, paper 634.

De la Flor G., Ojaghi M., Lamata Martínez I., Jirotka M., Blakeborough A., Williams M.S., 2009. Reconfiguring practice: the interdependence of experimental procedure and computing infrastructure in distributed earthquake engineering, *All Hands Meeting*, Oxford.

Madabhushi, S.P.G. and Schofield, A.N., 1993. Centrifuge modelling of tower structures subjected to earthquake perturbations, *Geotechnique*, Vol.43, No.4.

Nakashima, M., 2008. Roles of Large Structural Testing for the Ad*a*vancement of Earthquake Engineering, *The 14th World Conference on Earthquake Engineering, Beijing*, keynote, K009.

Ojaghi M., Parra-Fuente J., Saleem K., Blakeborough A., Williams M.S., 2007. Grid Based Distributed Hybrid Testing, *All Hands Meeting*, Nottingham.

Saleem K., Parra-Fuente J., Ojaghi M., Williams M.S., Blakeborough A., 2008. UK-NEES: Grid Services Architecture for Earthquake Engineering, *Applied Computing Conference (ACC '08)*, Istanbul.

Spencer B. F., Elnashai A., Nakata N., Saliem H., Yang G., Futrelle J., Glick W., Marcusiu D., Ricker K., Finholt T., Horn D., Hubbard P., Keahey K., Liming L., Zaluzec N., Pearlman L., Stauffer E., 2004. The MOST Experiment: Earthquake Engineering on the Grid, *NEES grid* whitepaper 1.0.