Current Capabilities and Future Roles for Internet Distributed Large Scale Real-Time Seismic Testing

M. Ojaghi
KBR, London; University of Oxford, UK.

M. Dietz
University of Bristol, UK

I. Lamata Martínez, M.S. Williams, A. Blakeborough
University of Oxford, UK

SUMMARY:
Real-time distributed hybrid testing (DHT) presents an efficient way to rapidly expand current seismic testing capabilities. In real-time DHT, geographically distributed substructures are coupled together as part of a single hybrid experiment. The combined system is tested under strict real-time loading conditions hence ensuring that rate effects prevalent in many aseismic technologies can be investigated. In this paper multi-axis real-time DHT using a realistic test setup is presented. Conducted across the Internet between the universities of Bristol and Oxford, it is used to demonstrate the effectiveness of the UK-NEES distributed testing system. An overview of the testing system, believed to be the only one currently capable of conducting such tests is given, explaining key aspects of the technology developed. Current capabilities are discussed together with considerations for the next phase of development. Finally, a vision for the future wider application of this method in earthquake engineering is presented.

Keywords: Internet, real-time, distributed, hybrid, testing

1. INTRODUCTION

Synergy occurs when a combined system enables the accomplishment of tasks the sum of its individual parts cannot. This concept is at the heart of grid based technologies and, in the field of earthquake engineering is exemplified by the real-time distributed hybrid testing (DHT) technique.

Real-time DHT is a new technique for enabling very large scale seismic testing for applications where the rate of loading is crucial for test accuracy (Ojaghi et al. 2010a; Ojaghi, 2010). The technique has been developed through UK-NEES (the UK Network for Earthquake Engineering Simulation) as a viable answer to the call for increased testing capacity worldwide, particularly for testing in real-time, i.e. when earthquake loading is not time-scaled (Nakashima, 2008). The technique combines the resources of multiple geographically distributed laboratories including actuator arrays, sensors and computational capacity, using them in a single hybrid experiment. Sharing not just scientific equipment but also laboratory space, time and expertise, the technique fosters a spirit of sharing and collaboration, and enables experiments to be conducted that would be out of the reach of a single laboratory, working alone.

In this paper results are presented from multi-axis real-time distributed hybrid testing. The tests connect two of the UK-NEES nodes at the universities of Bristol and Oxford, a roundtrip network distance of approximately 700km. The results together with an equivalent local test are used to demonstrate that robust (stable and accurate), continuous and repeatable real-time DHT with a realistic test setup is possible using existing laboratory equipment. Prior to presenting the results a brief background to hybrid testing is given. Following this a brief overview of the UK-NEES real-time distributed testing system is given. Believed to be the only system currently capable of connecting coupled geographically distributed scientific equipment in real-time across the Internet, the network topology is presented as well as a high level control view relevant to the presented test. This highlights key details of the testing environment and the technologies developed. The test is then introduced and results presented. Following this, current capabilities are discussed, describing limitations of the setup and presenting directions for the next stage of development. Finally in concluding, a future vision for the wider application of the technique is discussed.

2. BACKGROUND

Real-time DHT builds upon existing hybrid techniques. Hybrid testing was first developed in a local laboratory setting for testing at expanded timescales by the 1980's (e.g. Mahin et al. 1989). In order to make the test suitable for capturing rate effects in structural response, the technique was later extended to allow testing in real-time (e.g. Blakeborough et al. 2001). A hybrid test is usually conducted with substructuring. In a typical test the structure under investigation is split into a main numerical substructure and one or more physical substructures. The physical substructures are usually the focus of the experiment, tested physically as they often exhibit unknown or difficult-to-model response characteristics. These numerical and physical substructures are coupled together under seismic loading via a transfer system of actuators and sensors. A feedback control loop passes data between the various substructures throughout the test to ensure that the coupled system realistically simulates the response of the complete structure.

Hybrid testing is attractive since only part of the structure is represented physically; meaning much larger scale structural systems may be tested than in pure shaking table tests. Although multi-axis (actuator) tests have been conducted, most recently in real-time (Bonnet 2006), as the size of the distributed experiment grows it is not long before the capacity of a single lab is saturated. To overcome traditional laboratory capacity limitations the concept of distributed hybrid testing was introduced. Taking advantage of the widespread deployment of Internet infrastructure to make data communications between geographically distributed sites feasible, individual substructures no longer need be located within the same laboratory. Instead by passing control data across a computer network, individual parts of a hybrid experiment may be located in different cities and even across the world. This concept was originally introduced through the US NEES programme (Spencer et al. 2004; Mosqueda et al. 2004) and by others notably, Tsai et al. (2008). However, to date and outside of UK-NEES there have been limited attempts at distributed hybrid testing. A detailed review of DHT research conducted worldwide may be found in Ojaghi, 2010. In all cases the true quality of results are not clear. Common themes are that distributed tests are slow, primarily due to the communication link, and that there are networking issues related to delay or loss of data. Indeed most of the tests conducted to date have been at greatly expanded time-scales. A few have attempted to run faster tests. Mosqueda et al. (2006) conducted a multi-site test across the USA, the test taking approximately 0.3 hours to represent a 15s earthquake. Notably, Schellenberg et al. (2008) developed a system for fast DHT. The system utilized a predictor-corrector algorithm, and allowed for continuous actuation even if network delays exceeded a specific proportion of the model calculation time-step. Demonstrated for a simple test set-up connecting two sites in California at a distance of around 95km, a test running using approximately 20 ms time-steps could run without initially scaling the earthquake time. While this test approaches the requirements for a real-time test, communication issues meant that it would slow down in response to network delays.

3. OVERVIEW OF THE UK-NEES REAL-TIME DISTRIBUTED TESTING SYSTEM

UK-NEES is a scientific grid network setup in 2006 with the aim of promoting distributed collaboration and extending seismic and similar testing capabilities within the UK and beyond (Ojaghi, 2010b). Inspired by US NEES activities in system integration (Foster & Kesselman, 2009) UK-NEES has aimed to develop the potential of distributed collaboration for the next generation of earthquake engineering knowledge and technology development. Consisting of three sites at Oxford, Bristol and Cambridge universities and with a strong e-Science ethos the network has had the opportunity to conduct extensive testing and development. One of the main areas of research has been the development of real-time DHT, where in 2009 the first series of robust tests were carried out. The development process highlighted two key areas to focus work. The first was to deal with the complex challenges of test control, the second, was to facilitate working in the distributed environment.

Prior to discussing the former it is worth highlighting the importance of the latter. Developing policies to manage the social impact and needs of working in a distributed testing environment and the use and development of auxiliary tools to assist distributed working was critical to the tests. It is also crucial for successful distributed collaboration. A summary of UK-NEES activities in implementing auxiliary systems and developing policies is found in Ojaghi (2011). Issues considered include network security policy and access rights and its impact on testing. When transitioning work to the distributed environment it is necessary to reconfigure practice, providing tools and procedures to support work. The direct impact of this has been to develop a usability strategy implemented within the testing system, the aim to simplify the real-time distributed test operation process. This is summarized in de la Flor et al (2010) & Ojaghi (2010). UK-NEES has experimented with a variety of support tools both to set up experiments in the distributed environment (e.g. video conferencing) and during the running of tests. The latter includes support for audio communication (telephone); audio/visual communication between remote operators and for viewing of remote test rigs and systems (tele-presence/remote desktop etc.); and, purely textual communication (instant messaging - useful for multisite tests). The issue of data sharing and ownership has also been investigated. Shared online documents have been used to allow live editing of test notes by distributed parties. This allows local operators to give first hand input and reduces the work load on the main test operator. Other more recent work at the Oxford node in support of the EU SERIES project distributed database has superseded early UK-NEES work. The database will provide a valuable resource in archiving and allowing sharing of test data. This and the further development and integration of auxiliary tools for distributed collaboration will only improve the ease of running tests, the user experience and reduce operator error.

3.1 UK-NEES Test Topology and Distributed Testing System

In this section a brief summary of the testing architecture is presented, further details are found in Ojaghi (2010). Real-time DHT is particularly challenging. Tests must be conducted across an environment where real-time communication is not native. The Internet is not designed for real-time communication of low latency control signals; there has been no prior application for this. Data which on local real-time controller boards is accustomed to dedicated hardware controlled real-time priority must share access with other Internet users and contend with routing. Although dedicated network connections can be installed, existing communication protocols are inadequate to enable real-time distributed communication and data still has to contend with switches. Furthermore, legacy hardware at local and remote sites must be synchronised, the additional time delays introduced when data is transmitted between sites can be prohibitive and there exists a potential for loss or irregular arrival of data between distributed controllers.

To overcome these challenges three design issues have been addressed. Firstly, the testing system ensures regular arrival of data between hardware controllers at each site at a high rate (around 50-80Hz). Secondly, existing hardware controllers have been adapted to allow network communication. Lastly, the delay introduced by the dynamics of the distributed environment is overcome through the system architecture to minimise the delay and, to accurately compensate for much larger delays than are accounted for locally, new delay compensation algorithms have been developed.

The first step in setting up a test is to choose the test site topology and the control strategy used to run it. Careful application of these serves to minimise the communication delay between distributed controllers, controlling time-step synchronisation between sites. In Fig. 1 a summarized overview of the UK-NEES test site topology is shown. A client server architecture as is typical in sockets communication was applied. However, in the UK-NEES case the implementation is not the usual one where multiple clients connect to a single server. Instead multiple server nodes, hosting distributed physical or numerical substructures wait to provide a service to the single client site. At the client, the main numerical model is hosted along with local numerical/physical substructures. It is found that by having the servers wait for connection from the client the operational burden on the client as the central node is reduced. The onus is placed on the waiting servers to be ready at the agreed time for the main test operator at the client to connect to when ready. In the case of a three site test, if computational capacity allows, it can be advantageous from a test performance point of view to make the client site the site which has the lowest maximum latency to connect to the other sites. This will allow lower model time-steps and higher test frequencies to be achieved.

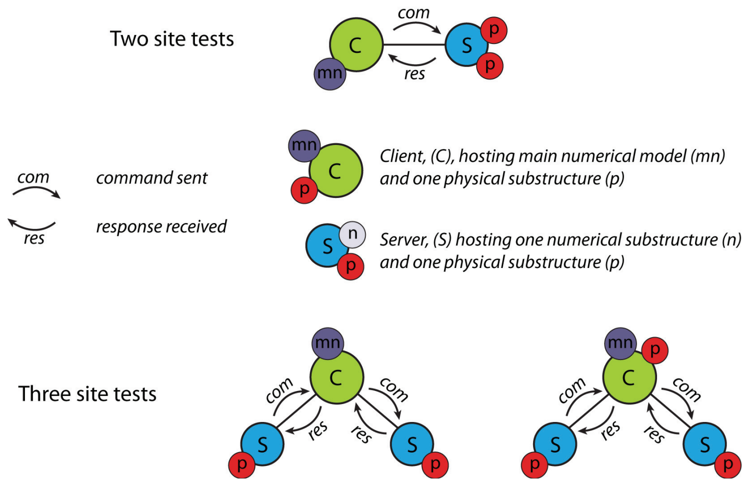

Figure 1. Client-Server distributed test architecture topology (adapted from Ojaghi, 2010)

In the case of the UK-NEES network, testing is primarily conducted on the Ja.net network. This is a high performance network which provides Internet access to research institutions within the UK. A separate dedicated UK-light connection was also installed for this project. A map showing the geographical layout of the network connections is shown in Fig. 2.

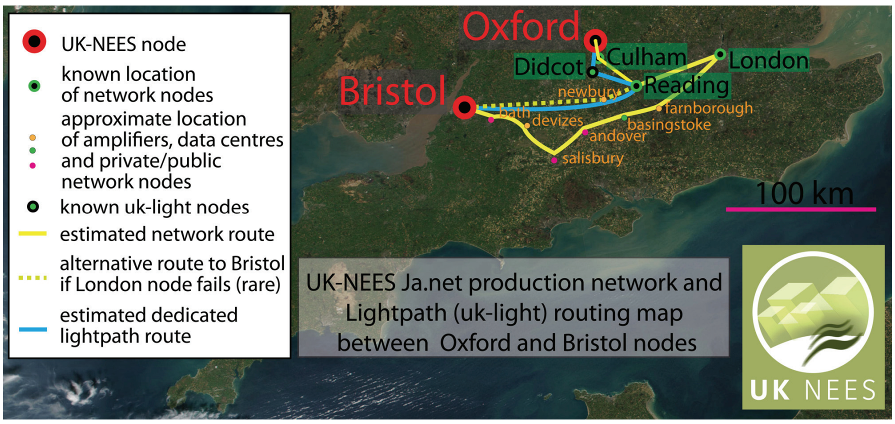

Figure 2. UK-NEES geographical network route between Oxford and Bristol (adapted from Ojaghi, 2010)

It is clear that normal routing is not conducted via the shortest geographical path. It is determined by network administrators based on their judgments of network traffic to maintain ‘best service’ - a reliable high bandwidth connection for the majority of users. This is in direct contrast to the low-bandwidth, low-latency need for real-time DHT. The route between the Oxford test machine to Bristol is largely fixed and has 12 hops. The Internet connection between institutions is generally of a high quality. Although connections are shared, routing is robust and conducted using high performance multiprocessor routers. Although previous research has suggested the network is the chief source of delay, UK-NEES findings are that that once communications are properly set-up it is more likely that saturation of local node hardware is the primary cause of delay and data loss. Interestingly although following a shorter route, the dedicated connection did not offer performance benefits. This is likely due to the software/hardware employed and could be improved, budget allowing (Ojaghi, 2010).

Fig. 3 presents an overview of the control architecture for a typical two site, two-axis test indicating the main control functions (A detailed description together with the corresponding network architecture model is presented by Ojaghi, 2010). Testing is conducted using dSpace programmable real-time controller boards connected to Instron PID-L actuator controllers. dSpace boards are hosted on Windows XP based multi-processor PCs.

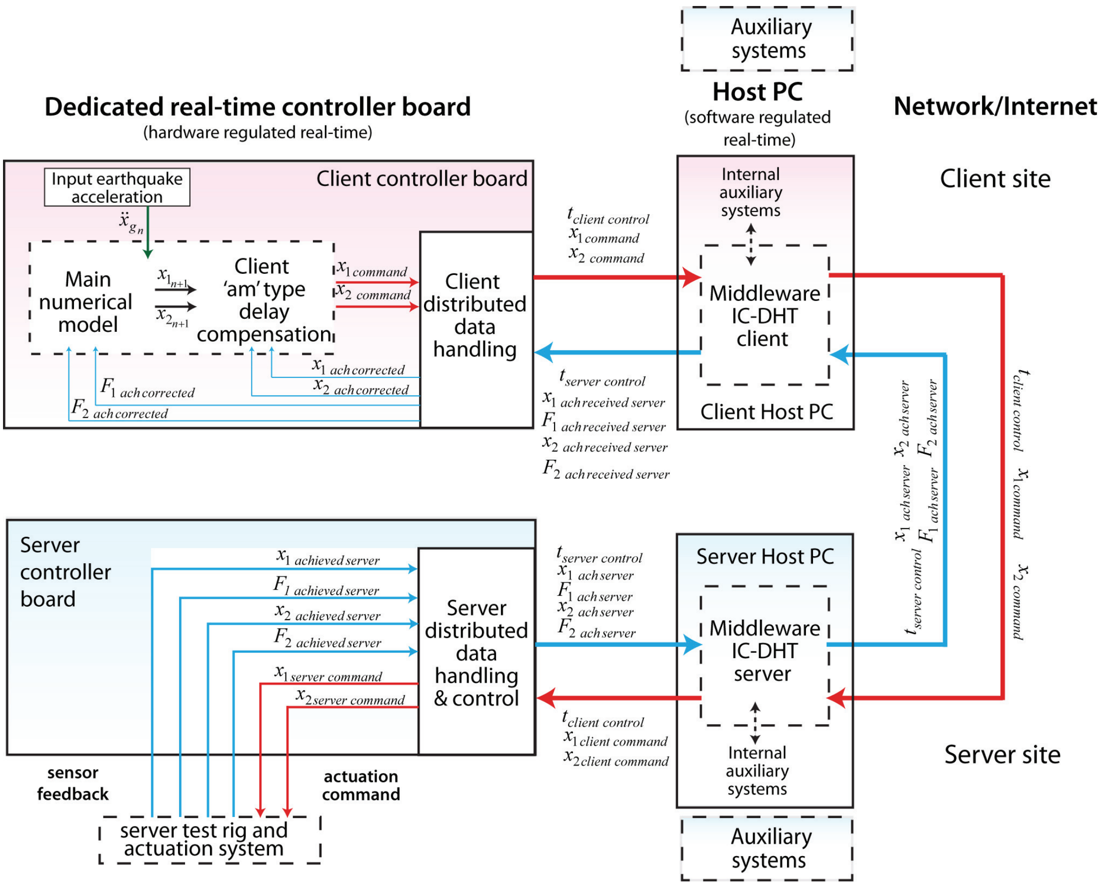

Figure 3. Generic control architecture for a multi-axis two site test (high level overview - test running state)

Real-time control is guaranteed on the hardware controllers at each site. To enable communications between distributed controller boards a communications program, the test middleware named IC-DHT has been developed. Located on the Host PC's at each site it is responsible for setting up the test topology and time-step synchronisation between sites, necessary to determine which control strategy is used to run the test. The program interfaces with internal auxiliary systems to support work in the distributed environment (external auxiliary systems although applied have not been fully integrated into the testing system as yet). To maximise communication performance at the local node, the test middleware function has been optimised and runs on a software controlled real-time operating environment. To achieve communication between distributed controllers the program implements a form of the UDP/IP network protocol. Carefully balanced to avoid saturation it allows high frequency data transfer, minimising delay in its operation. However, on its own this system is unreliable and cannot ensure regular availability of data when required by the real-time controller boards at each site. The data handling controllers located on each controller board effectively act as a higher level protocol acting over UDP/IP to ensure data is available when required. Using prediction algorithms and time-stamps they ensure regular data arrival by managing data arriving early, on-time, late or not at all, in the two latter cases providing accurate corrections. It is found that in a well balanced test correction can be avoided.

Prior to running the control strategy for the test is chosen. Various control strategies have been developed based on the relative synchronisation of time-steps at distributed sites. The results for two are presented here. The CSF T3 (client start first type 3) and the SSF T3 (server start first type 3). The relative synchronisation of time-steps is either to ensure that the servers are running before the client starts or the reverse. Each has its own advantages (Ojaghi, 2010). In the case of the server starting first the server data handler uses algorithms to predict the desired client command, smoothly updating this when it arrives. The CSF T3 algorithm conducts variable *am* type delay compensation at the client. The SSF T3 conducts fixed *am* type delay compensation at the client, although further delay control is inherent in its functioning. Time-step synchronisation at the servers means that SSF tests achieve lower overall test delays but are more complex to set-up. The optimised communication environment discussed enables reliable and regular communication between sites at high frequency, with real-time communication ensured by data handling. While the set-up minimises delay, delays in the distributed environment are much higher than locally encountered. The performance of the popular polynomial delay compensators is inadequate. To overcome this issue new *am* type algorithms have been developed (Ojaghi, 2010) with superior performance characteristics.

4. AN EXAMPLE MULTI-AXIS REALTIME DISTRIBUTED HYBRID TEST

As part of its development programme and in order to prove that robust and repeatable real-time DHT may be achieved, UK-NEES has conducted several hundred real-time distributed hybrid tests using a variety of realistic and challenging test setups (see Ojaghi, 2010). These have allowed the current systems to be evaluated. Here results are presented from multi-axis two site real-time DHT between Oxford and Bristol. A high level overview of the test set-up is shown in Fig. 4. In this test a three degree of freedom shear building model located on the Bristol real-time controller board is coupled to the two-storey single-column test rig (Bonnet, 2006) located in Oxford. The physical part of the test comprises of 50% of the columns (and hence approximate stiffness) on the 24 columned lower two floors. This is modelled by multiplying the measured force from each actuator load cell by 12.

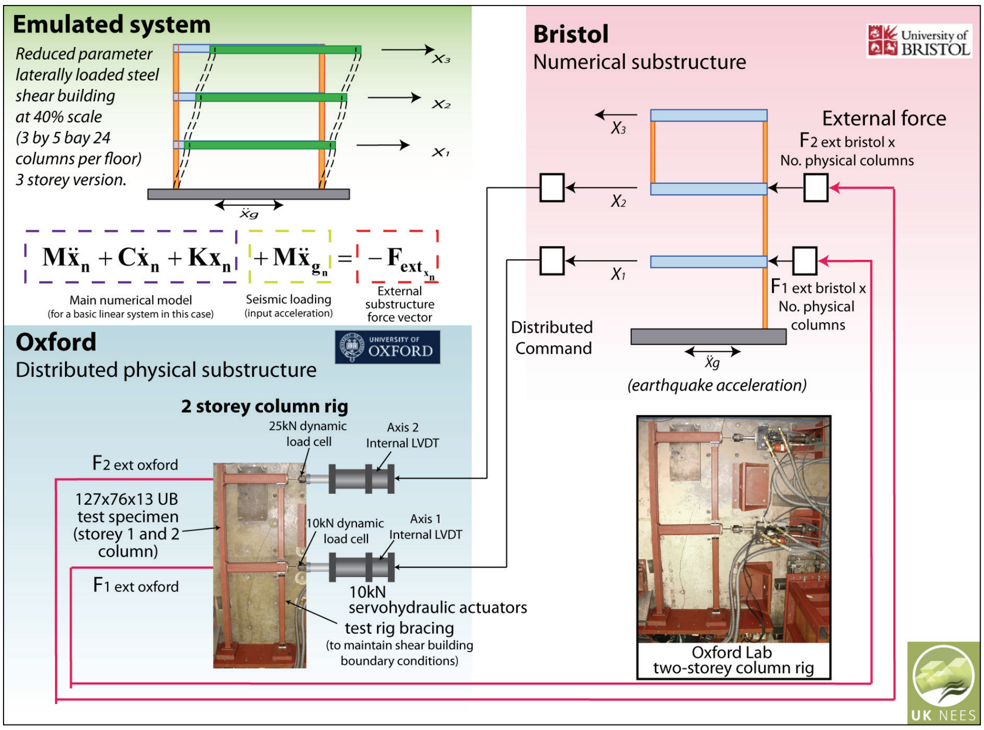

Figure 4. High level overview of the multi-axis two site real-time DHT

The structure is scaled at 40% and subjected to the El-Centro NS earthquake (further scaled to 50% for comparison with other tests). Stiffness proportional damping at 5% of the first mode is applied. The Newmark explicit scheme is used to solve the equations of motion at 20ms time-steps. Hence this is the rate at which client commands are generated. This time-step is chosen since the highest modelled natural frequency is approximately 4.7Hz. This will allow accurate distributed control up to 5Hz (actuator permitting). A series of approximately 20 real-time distributed hybrid tests were conducted using this setup. Results from this test setup have not been previously presented and are selected to indicate the general performance of CSF T3 and SSF T3 real-time DHT controllers. The tests are compared to each other and a local equivalent (where the main numerical model is solved on the Oxford real-time controller board). The local test uses variable second order polynomial delay compensation, a state of the art local controller developed prior to start of the project (Bonnet, 2006). Although the test rig was originally designed to maintain linear response for repeatability, several hundred earthquakes applied just at or near yield have significantly degraded the test specimen. This causes slight nonlinearity. The test is also particularly challenging due to the high level of force coupling on the lower actuator. This can cause local testing issues (Ojaghi, 2010; Bonnet 2006).

A comparison of achieved response for the two distributed actuators versus desired structural response is shown in Fig. 5 demonstrating the performance of the CSF T3 controller. The results match very well and demonstrate the good performance of the testing system. This is selected here since this controller exhibits higher test errors as compared to the SSF T3 and local controller.

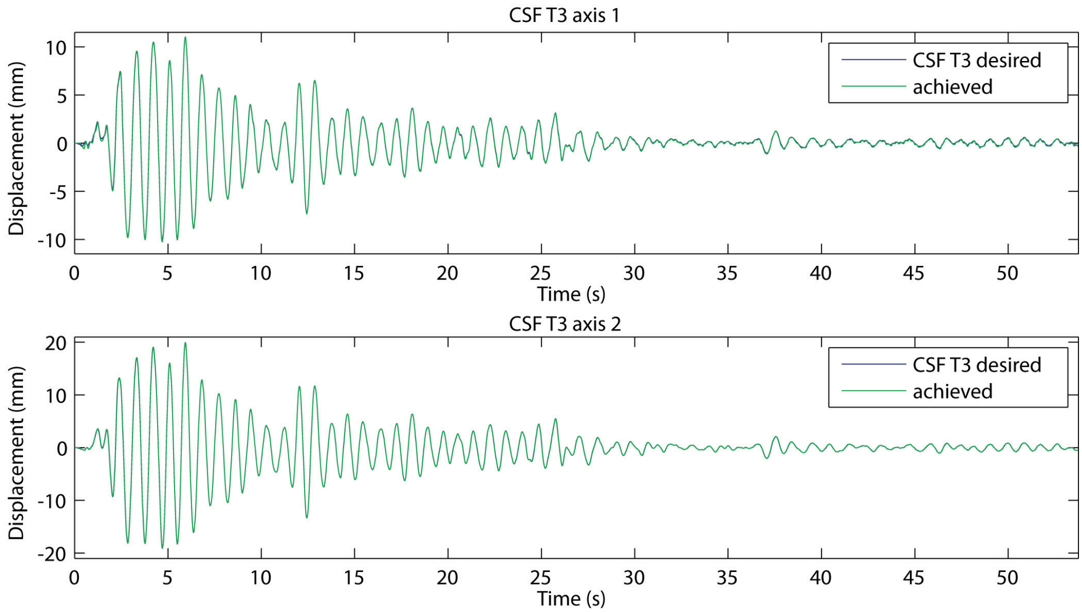

Figure 5. Client desired vs. achieved response for axis 1 and 2, results from a typical CSF T3 real-time DHT

In Fig. 6 zoomed comparisons of response at floors 1 and 2 are shown together with force hysteresis plots. This compares the response of local, CSF T3 and SSF T3 control systems. The results match well. It should be noted that due to specimen nonlinearity (see force hysteresis comparison), tests are not expected to be entirely repeatable. However, a good match between them may be found. In Fig. 7 a zoomed comparison of local/client control signals is shown further demonstrating the quality of the tests and the effectiveness of the control measures used. In every case desired and achieved response match well. One of ten client data handling correction events with the SSF T3 test is shown. The correction is quite accurate. The number of corrections here is higher than found in other tests/test set-ups (Ojaghi, 2010) where many tests have no client correction events. This is due to increased rates of data capture increasing computational load. The level of delay compensation required in the CSF T3 case varies between 49 - 56ms and in the SSF T3 case around 40ms. In both cases *am* type delay compensation (Ojaghi, 2010) works quite effectively. The local delay is of the order of around 32ms. It should be noted that 20ms of the delay is accounted for by the explicit formulation used. An additional 20ms delay is caused by sub-stepping the delay compensated actuation command to 1ms time-steps suitable for actuation. The sub-stepped signal is shown in the local test command. A good measure of the test errors can be found by comparing the hysteresis between desired and achieved results (Fig. 8). Though error is related to the quality of force feedback, the lower the hysteresis in the displacement signals the better the quality of the test. Error hysteresis comparisons with CSF T3 and SSF T3 controllers are shown and are also compared to the local case. It is clear that the error hysteresis between local and distributed tests match well. Slight hysteresis is found is all cases and is chiefly caused by actuator coupling. Test errors between SSF T3 and local tests match very well. In fact SSF T3 tests errors are lower. This is due to the superior performance of the *am* delay compensator used. Test errors with the CSF T3 controller are only slightly higher than the local test. This is due to the higher levels of delay compensation applied in CSF T3 control.

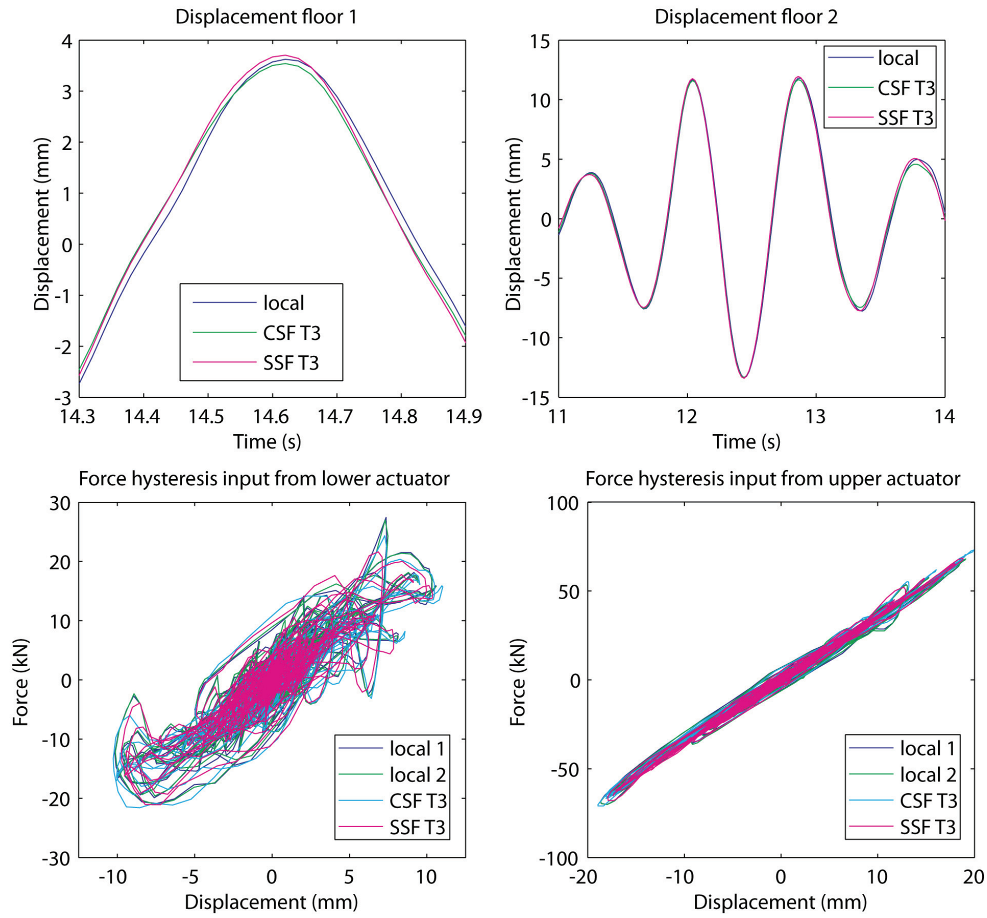

Figure 6. Zoomed views of 1st and 2nd floor displacement comparing local, CSF and SSF tests (above). Force hysteresis input to numerical model from actuator load cells for two local and two real-time DHT tests (below)

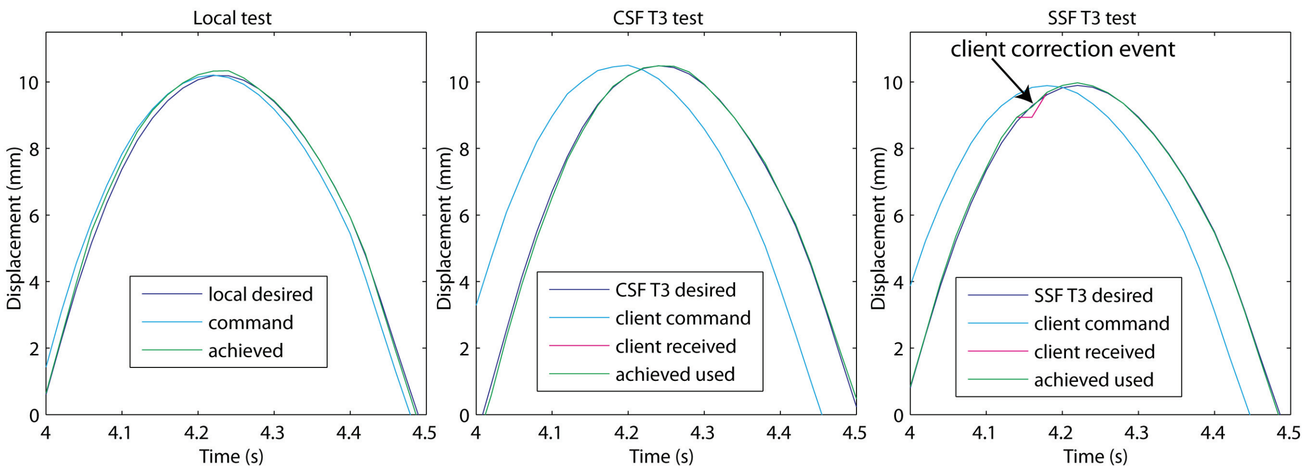

Figure 7. Local/client control lines between desired, command and achieved for various tests

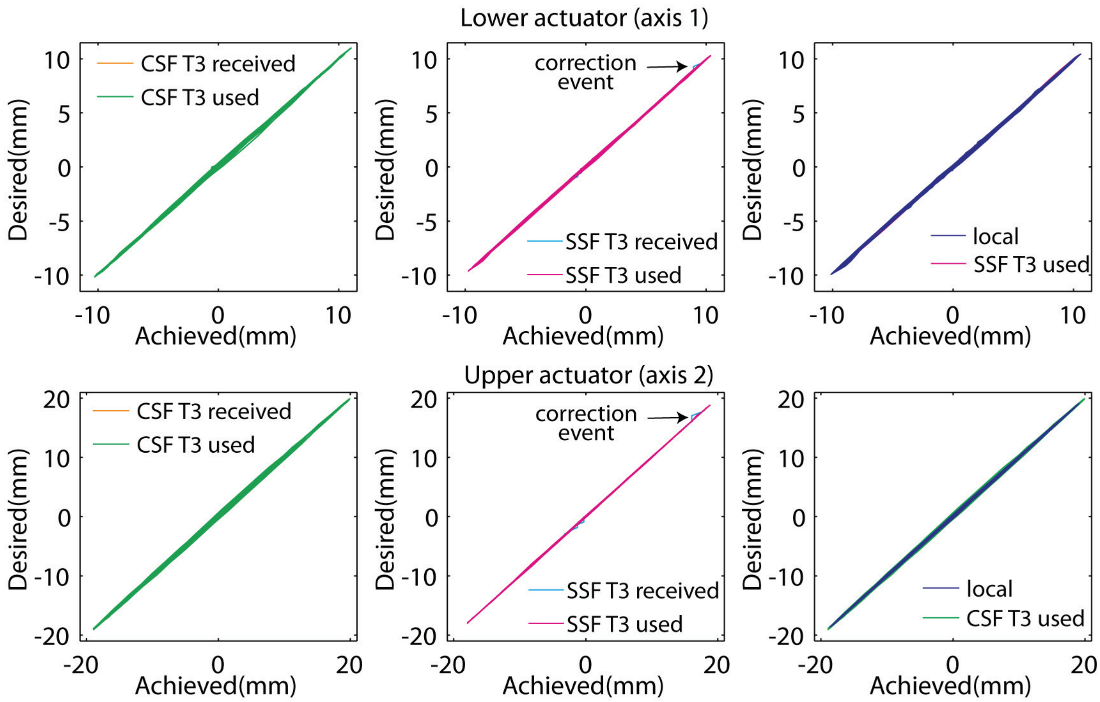

Figure 8. Error hysteresis between desired and achieved comparing local and real-time DHT

5. SUMMARY AND FUTURE DIRECTIONS

In this paper an example of a two-site multi-axis real-time distributed hybrid test has been presented conducted between Oxford and Bristol. Two real-time DHT tests have been compared to a local equivalent where both numerical and physical substructures are hosted in Oxford. In this case the test running at 20ms time-steps allowed a system with a maximum frequency of 5Hz and 5% first mode stiffness proportional damping to be tested. The physical part of the test consisted of 50% of the columns on the lower two floors of a three storey structure. Although the physical test setup displays significant coupling and slight nonlinearity due to specimen degradation the tests compare well with each other. The results demonstrate that the developed control systems are robust; stable, accurate and continuous real-time DHT can be achieved using realistic large scale physical substructures. The test is part of a larger series of several hundred real-time DHT conducted within UK-NEES.

While testing has been successful, allowing a wide range of challenging tests there are considerations for development. The first addresses the current dSpace controllers used. Testing revealed that when properly set up using the higher level data handler to act over existing network protocols, real-time communications between Internet distributed controller boards can be reliably achieved. However, test performance is reduced as computational load is increased, either due to additional substructures being tested, more data being saved or larger numerical models being used. The dSpace controller boards used are of much lower performance than current models. However, rather than directly replacing them it is worth investigating the use of alternative systems to change the fundamental hardware architecture of the system. It can be advantageous to run the test middleware and local controller board programs on a single real-time controller. Moving controller board functions on to a PC running in a dedicated real-time environment eliminates current board to PC saturation errors that can occur as the test increases in complexity or runs at much lower time-steps. This might be possible using real-time Linux and similar systems. The second consideration concerns network routing and test site topology. The current UK-NEES network runs via London increasing the network distance by several hundred kilometres. Although alternative more direct routes exist accessing them is a political issue and is controlled by regional operators. The best performance gains can be achieved by introducing a test node at London to host a client connecting to physical substructures at Bristol and Oxford.

The work conducted within UK-NEES proves that real-time DHT can be used as a tool to promote distributed collaboration. The developed testing system has the capacity to use the complimentary facilities of geographically distributed testing sites to conduct complex joint experiments in real-time. As a research tool the system can be used to plan experiments combining the expertise and the resources of UK-NEES nodes to explore seismic phenomena previously beyond the capability of either laboratory. In the case of UK-NEES the system may also be extended to for example, connect the Bristol shaking table to actuator arrays in Oxford. As a tool for supporting industry the system developed can improve the quality of tests used for qualification of new technologies. It can also act as a tool to quickly test the performance of new concepts and design implementations. In this way it may encourage the increased uptake of new technologies and allow verification in cases where numerical modelling alone is inadequate. The technologies developed are highly portable to other platforms. It is envisaged that other research centres can setup their own regional real-time DHT networks to extend national testing capabilities. Should network routing be more direct and improvements in delay compensation allow, real-time DHT between closely spaced national networks could become a reality.

REFERENCES

Blakeborough A., Williams M. S., Darby A. P. & Williams D. M. (2001) The development of real-time substructure testing. *Phil. Trans. Royal Soc. London, Series A*, **359**, 1869-1891.

Bonnet P.A. (2006) The development of multi-axis real-time substructure testing. DPhil Thesis, University of Oxford.

De la Flor G., Ojaghi M., Lamata Martínez I., Jirotka M., Williams M.S. & Blakeborough A. (*2010)* Reconfiguring practice: the interdependence of experimental procedure and computing infrastructure in distributed earthquake engineering. *Phil. Trans. Royal Soc.* London, Series A. 368, 4073-4088.

Foster I., Kesselman C. (2009) The Grid 2 Blueprint for a New Computing Infrastructure, Elsevier.

Mahin S. A., Shing P.B., Thewalt C. R. & Hanson R.D. (1989) Pseudodynamic test method current status and future directions. *Journal of Structural Engineering,* **115: 8.** Paper No. 23816.

Mosqueda G., Stojadinovic B., Hanley J., Sivaselvan M. & Reinhorn A. (2006). Fast hybrid simulation with geographically distributed substructures. *17th Analysis and Computation Specialty Conference*.

Mosqueda G., Stojadinovic B. & Mahin S.A. (2004) Geographically distributed continuous hybrid simulation. *Proc. The 13th World Conf. on Earthquake Eng*. Paper 959.

Nakashima M. (2008) Roles of large structural testing for the ad*a*vancement of earthquake engineering. *Proc. The 14th World Conf. on Earthquake Eng.* Keynote, K009.

Ojaghi M. (2011) Development of the UK-NEES test middleware, early tests, web services approach, network usage, security, and usability issues in distributed hybrid testing. OUEL Departmental Report No. 2322/11. University of Oxford.

Ojaghi, M. (2010) The Development of Real-Time Distributed Hybrid Testing for Earthquake Engineering. DPhil thesis, University of Oxford.

Ojaghi M., Lamata Martínez I., Williams M.S., Blakeborough A., Dietz M. & Crewe A. (2010a) Real-time hybrid testing in geographically distributed laboratories. *14th European Conference on Earthquake Engineering*. Paper 872.

Ojaghi M., Lamata Martínez I., Dietz M., Williams M.S., Blakeborough A., Crewe A., Taylor C., Madabhushi G., Haigh S. & Ali A. (2010b). UK-NEES - Distributed Hybrid Testing Between Bristol, Cambridge and Oxford Universities: Connecting Structural Dynamics Labs to a Geotechnical Centrifuge. *9th U.S. National and 10th Canadian Conference on Earthquake Engineering.* Paper 1024.

Schellenberg A., Kim H., Mahin S.A. & Fenves G.L. (2008) Environment independent implementation of a software framework for fast local and geographically distributed hybrid simulations. *Proc. The 14th World Conference on Earthquake Engineering*. Paper S16-02-005.

Spencer B. F., Finholt T., Foster I. Kesselman C., Beldica C., Futrelle J., Gullapalli S., Hubbard P., Liming L., Marcusiu D., Pearlman L., Severance C. & Yang G. (2004) Neesgrid: a distributed collaboratory for advanced earthquake engineering experiment and simulation. *Proc. The 13th World Conf. on Earthquake Eng.* Paper 1674.

Tsai K.C., Hsieh S.H., Yang Y.S. & Wang K.J. (2008) Networked hybrid simulation of large scale structures in Taiwan. Chp. 13 of Hybrid Simulation. Theory, Implementation and Applications. Saouma V. and Sivaselvan M. *Taylor & Francis*.