Ojaghi, M., Lamata Martínez, I., Dietz, M., Williams, M., Blakeborough, A., Crewe, A., Taylor, C., Madabhushi, G. S. P., & Haigh, S. K. (2018). Geographically distributed hybrid testing & collaboration between geotechnical centrifuge and structures laboratories. *Earthquake Engineering and Engineering Vibration*, 17(1), 53-71. <https://doi.org/10.1007/s11803-018-0425-1>

Peer reviewed version

Link to published version (if available):
[10.1007/s11803-018-0425-1](https://doi.org/10.1007/s11803-018-0425-1)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Springer at <https://link.springer.com/article/10.1007/s11803-018-0425-1>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Geographically Distributed Hybrid Testing & Collaboration between Geotechnical Centrifuge and Structures Laboratories

By Mobin Ojaghi^{1,*}, Ignacio Lamata Martínez^{1†}, Matt S. Dietz², Martin S. Williams¹, Anthony Blakeborough¹, Adam J. Crewe², Colin A. Taylor², S.P. Gopal Madabhushi³ and, Stuart K. Haigh³

¹ Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ, UK

² Department of Civil Engineering, University of Bristol, Bristol, BS8 1TR, UK

³ Department of Engineering, University of Cambridge, Cambridge, CB2 1PZ, UK

Distributed Hybrid Testing (DHT) is an experimental technique designed to capitalise on advances in modern networking infrastructure to overcome traditional laboratory capacity limitations. By coupling the heterogeneous test apparatus and computational resources of geographically distributed laboratories, DHT provides the means to take on complex, multi-disciplinary challenges with new forms of communication and collaboration.

To introduce the opportunity and practicability afforded by DHT, here an exemplar multi-site test is addressed in which a dedicated fibre network and suite of custom software is used to connect the geotechnical centrifuge at the University of Cambridge with a variety of structural dynamics loading apparatus at the University of Oxford and the University of Bristol. While centrifuge time-scaling prevents real-time rates of loading in this test, such experiments may be used to gain valuable insights into physical phenomena, test procedure and accuracy. These and other related experiments have led to the development of the real-time DHT technique and the creation of a flexible framework that aims to facilitate future distributed tests within the UK and beyond. As a further example, a real-time DHT experiment between structural labs using this framework for testing across the Internet is also presented.

Keywords: Centrifuge; Distributed; Dynamics; Earthquake; Geographically; Geotechnical; Hardware-in-the-Loop; Hybrid; Real-time; Testing

1. Introduction

To meet society's goals for ever improved safety and reliability of systems and infrastructure, there is a need for tools to optimise the use of existing resources and to enable enhanced capability for research into big scientific problems [1]. With widespread distribution of reliable high bandwidth computer networks, the way research is conducted is changing, creating new avenues for collaboration [2]. In tune with these ideas and developments, Distributed Hybrid Testing is one such tool, used for dynamic simulation. In DHT, technology is applied to allow traditional laboratory capacity limitations to be overcome so as to enable more complex experiments, beyond the capacity of any single laboratory, to be pursued. DHT is developed with earthquake engineering research in mind. In DHT, dedicated computer networks or the Internet are used to couple at rates up to and including real-time [3], experimental and computational resources across multiple geographically distributed laboratories to take part in a single dynamic experiment.

DHT promotes a collaborative culture, encouraging participants from diverse disciplines to share specialist technical knowledge and resources to tackle complex scientific problems. Within earthquake engineering and other disciplines DHT is extremely significant. With correct implementation there is real potential to conduct experiments that would otherwise not be feasible

* Correspondence to: Mobin Ojaghi, KBR Ltd., London Operating Centre, KT22 7NL, UK. E-mail: mobin@ojaghi.com

‡ Author acting in personal capacity. The authors remain solely responsible for this paper, which was produced independently, with no corporate involvement.

† Joint Research Centre, ELSA. Ispra, Italy.

due to sheer size, expense or lack of expertise at a single site. While benefits are clear, DHT presents significant challenges to be overcome for any successful implementation. These are not just technical (DHT is significantly more complex than single-site testing), but involve implementation of systems to support distributed collaboration.

This article represents the capstone synthesis of an innovative programme of research that has further developed and evaluated DHT in the experimental earthquake engineering field. In doing so, it has supported the innovative integration of structural and geotechnical engineering experimentation. The research has demonstrated the viability of DHT for particular classes of experiment and has identified and developed the technical network and control performance capabilities that must be in place to achieve this. As such, it is an important starting point for researchers considering possible implementation of DHT techniques in their own research. Through examples, including multi-site testing between geotechnical centrifuge and structures labs and two-site real-time DHT, we present experiences and achievements which have led to the development of real-time DHT and a flexible distributed hybrid testing framework. This is in the context of work from the UK and efforts made to make this complex technique more widely accessible.

2. DHT - a United Kingdom perspective

In 2006, development work began on UK-NEES (UK Network for Earthquake Engineering Simulation). A collaborative grid network, UK-NEES was setup with the aim of extending seismic and similar testing capabilities within the UK [4-6]. The UK-NEES network connected three of the main earthquake engineering research centres in the UK at Bristol, Cambridge and Oxford Universities. Each provides complimentary facilities for the network and together form a framework for distributed testing (Fig. 1).

Figure 1. UK-NEES distributed testing framework.

Bristol, the largest earthquake engineering research centre in the UK, hosts a six degree of freedom shaking table, an arrangement of approximately 200m² of reaction walls and a large array

of hydraulic actuators [7]. Oxford specialises in real-time hybrid testing and hosts a set of 6 high performance hydraulic actuators, 4 linear electrical actuators and a strong floor. Cambridge, a pioneering centrifuge testing laboratory, focuses on research into tunnels and foundations, notably hosting a 10m diameter beam centrifuge capable of achieving accelerations of approximately 125g at 4.125m with a 1.2 ton payload. Although, the network has access to a wide range of testing facilities, a wider aim is to expand the network to include other laboratories in the UK. As collaboration increases an eventual 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.

As with NEES [8] one of the main research themes within UK-NEES was the development of the DHT technique. Prior to and during the early stages of the UK-NEES project, several valuable demonstration tests had taken place internationally. These were to prove connections and also highlighted some of the inherent difficulties of the technique [9-14]. These tests were generally conducted at greatly extended times-scales and communication issues were reported, which tended to limit accuracy. A more detailed review of DHT tests is provided in [3,15].

An early aim of UK-NEES was therefore to enhance the reliability of the technique and provide useful insights to the testing community. In addition there was special focus on adapting the technique to tackle novel soil structure interaction problems - demonstrating distributed centrifuge testing and, to seek to implement DHT at faster rates up to and including real-time in challenging structural simulations. Extensive testing was pursued early on to explore the new distributed testing environment. Multiple communication approaches were investigated and a thorough understanding of the testing environment developed. Communications software development involved making the connections, using different network protocols and testing setups. Several alternative distributed control strategies were developed, using different data transfer approaches and delay compensation techniques [15]. One key series of tests included multi-site testing, connecting a foundation system in the Cambridge geotechnical centrifuge to structural systems at Bristol and Oxford [6]. These tests provided valuable insights. The lessons learnt were captured and used to develop the UK-NEES DHT testing architecture which was applied to enable a large series of tests demonstrating for the first time robust and realistic real-time DHT [15,16].

It became clear that distributed working is complex and test operators are highly occupied. When transitioning to distributed testing, careful consideration must be given to test usability, to increase operator awareness and reduce burden [2]. This was reflected in the software design and procedures developed for starting, running and stopping tests. To support testing, a variety of tools were investigated [17], tele-presence facilities installed [18,19] and work progressed on defining requirements for a shared data repository.

From 2010 work transitioned to the EU funded project SERIES (Seismic Engineering Research Infrastructures for European Synergies, [20]), where a decentralised repository was designed and implemented with the objective of sharing experimental data publicly from 22 leading earthquake engineering centres in Europe [21]. Work has since progressed to Celestina, a platform for the integration of hazard mitigation resources, which presents a means for sharing data between the EU (SERIES) and US (NEES) [22]. From a DHT perspective, some of the experiences of UK-NEES were made available to assist development of DHT within SERIES [23]. In recognition that several software platforms exist for DHT, the focus of the latest work has been to address the issue of standardisation and interoperability between different testing systems, defining specifications and a high level framework for conducting distributed experiments [24].

3. Roles and challenges of DHT in earthquake engineering

The main impetus behind single-site hybrid testing is to allow realistic simulation of seismic response. The hybrid nature means the experiment is split between coupled numerical and physical

parts. This is advantageous. Numerical modelling of earthquake engineering structures and components alone, though useful in design, has many limitations. Physical testing is essential where response phenomena are poorly understood or too difficult to model practically. While testing at small scale requires great care to ensure results are not erroneous [25], large or full scale testing is a must to ensure safe qualification of new technologies or design techniques [26-27].

However, the sheer size of structural systems and earthquake engineering components can mean large scale testing is prohibitively costly, time consuming to set up and experiment size is severely limited by actuation capacity and space.

Hybrid testing serves as an alternative approach for seismic simulation and advanced qualification testing when fully numerical simulation is inadequate and fully physical testing prohibitive. As only parts of the system are tested physically, larger structural systems and components may be tested without meeting single-site capacity limitations. The technique thus allows the realistic global response of the structural system to be captured with numerical modelling for the '*known*' (i.e. well-understood) parts and physical testing for the parts whose response characteristics are '*unknown*' or are the main focus of the experiment. Hybrid tests may be quasi-static or, as with shaking table testing, loading may be applied in real-time using dynamic actuators. In a real-time test, test duration equals real (earthquake) event duration, in order to capture rate dependent behaviour. Since larger scale specimens may be tested than with typical shaking table tests, time-scaling issues of smaller shaking tables may be mitigated. Hybrid testing therefore develops a synergy between fully physical and fully numerical simulation.

In the past two decades or so and in response to costly and devastating earthquakes, there has been significant investment worldwide (e.g. E-Defence [28], NEES [29], IEM [30], IIEES [31]) to upgrade or develop new large scale testing facilities. Interest in hybrid testing has increased with the realisation that these facilities are not big enough [32, 33]. Even large shaking table facilities, the largest of which is E-Defence, though invaluable, may only test medium size structures at full scale - most other facilities have significantly lower capacity. Real-time testing has become essential, with increasing adoption of rate dependent technologies such as isolation systems and energy dissipation devices to improve seismic performance [34, 35]. However, new earthquakes continue to expose limitations in current knowledge and practice. With the need for advances in earthquake engineering leading to the growing requirement for testing ever larger and more complex structural systems under realistic loading [32], resources at any one site may quickly become saturated.

Distributed hybrid testing serves to provide a solution to this capacity problem by extending the application of single-site hybrid testing, not only to allow bigger and larger scale experiments, but to promote sharing of specialist equipment and expertise. By applying actuators arrays, shaking tables, high performance computing or specialist equipment such as geotechnical centrifuges in a single experiment across multiple sites, a flexible framework for testing may be provided. This maximises the potential of the testing network to facilitate experiments that would have otherwise not been possible. Testing facilities may be brought online as required to conduct complex multi-disciplinary experiments at the frontiers of earthquake engineering for example, large scale soil structure interaction, attempting to address problems that may be overlooked by geotechnical or structural labs alone.

The technique is also significant for medium size testing facilities. These may form close relationships to better match some of the capabilities of larger labs. Labs in close enough proximity may conduct real-time DHT for example, to allow testing of large arrays of energy dissipaters, and their placement, whereas previous single-site testing would limit tests to just a few devices at a time [36]. The establishment of networks for dynamic simulation can allow an accessible route for industry to forge strong links with research labs. Hybrid simulation can be applied to assist design of large real-world structures in cases where seismic demand is higher than may be economically designed for with conventional techniques, or where levels of nonlinear response make conventional analysis methods uncertain.

3. a) Single-site vs. distributed hybrid testing

Fig. 2 illustrates a generalized hybrid test controller layout to introduce key components. In a single-site test typically inner and outer control loops are used to conduct testing. In a distributed test setting controllers can be heterogeneous, so each site requiring bespoke solutions to enable distributed communication. For example, a more typical setup is site 1 where distributed control data passes through a host PC to access the network, as enabled by a host PC communications program. An alternative less typical case is site 2 where the network interface is on board the outer loop controller. The advantage of site 2 is that network actions locally (i.e. send and receive) are guaranteed to take place automatically every time-step. In site 1, they can also occur every time-step but that depends on how host PC and controller board computation is balanced.

Figure 2. Generalized hybrid testing controller layout (site 1 top, site 2 bottom).

A typical single-site hybrid test applies substructuring [37] to split the system under test into numerical and physical parts, which are dynamically coupled via a transfer system of actuators and sensors. The test progresses in a time-stepping fashion, normally using a multi time-step strategy [38]. A numerical model which may be modal, implicit (using a predictor/compensator and correction algorithm) or explicit, is used to output the desired end of time-step position of the physical parts, at the start of the current time-step. These are smoothly and continuously actuated to (in displacement control) through the numerical model time-step at a control time-step. The achieved force at the end (usually) of the numerical model time-step is fed back to the numerical model to be used to calculate the next desired end of time-step displacement. Thus, the response of one part depends on the response of all the other parts. Smooth continuous loading ensures actuation is representative of the real loading event and artificial stress relaxation does not occur due to actuators holding until the next command [39].

A complication (more significant for real-time testing) arises due to the response characteristics and 'inner loop' control of hydraulic and electrical actuators. Often regarded as a phase delay between commanding and achieving a value, it depends on the performance of the actuator and the properties of the physical specimen and test rig. The delay will vary and an amplitude error may also exist, with compensation techniques used to overcome these [40, 15].

In a distributed environment, test complexity increases. A DHT testing system needs to facilitate reliable distributed communication, overcome data transfer delays and implement support systems for work in a distributed environment [3]. To enable testing, firstly, reliable communication

between distributed controllers must be established. To meet the strict timing requirements of testing, hybrid test control hardware operates in a ‘*hard*’ real-time environment. Here computational load is verified by the controller software prior to testing to ensure that all critical fixed time-stepping operations can complete on time. This is essential to ensure consistent actuation, and numerical model accuracy. This is as opposed to a ‘*soft*’ real-time environment where fixed time-step duration may in fact exhibit very slight variation in length, not measurable by the control software itself (e.g. [41, 42]). Currently, actuation control hardware is not designed with DHT in mind. Typically, test sites use legacy ‘*hard*’ real-time controller cards which require custom software to enable DHT networking operations or must be adapted to connect to networking hardware.

Once data leaves the ‘*hard*’ real-time environment at a single site, it is no longer afforded the same real-time priority as control data is accustomed to locally. It is essential therefore, that communication operations are balanced by the test set-up such that control data is available on the real-time controllers at each site on-time for accurate solution of numerical model equations and for actuation at fixed time-steps. Without this, data may be unduly delayed or lost, leading to test failure. For example, at a single site data transmitted between the controller board and a Network Interface Card (NIC) may do so in a non real-time environment, where it competes with other operations on the host computer. Significant delay may occur (leading to data effectively being lost). On the network, control data will usually share priority with other network users’ data during routing. Depending on the protocols used and transfer rate, data from each time-step may be queued for bulk transfer, data may be lost, corrupted, duplicated or possibly arrive out of order. To overcome these issues, and especially in the case where a host computer is used for network operations, host computer operations and data transfer must be optimised to ensure reliable communication at the rates required for the test. This in effect balances computational loads at each site and is tuned for the network. Furthermore, careful consideration must be made of DHT/network protocols and communication software operation. This also includes planning for safe handling of communication interrupts and approaches to manage, minimise or eliminate data arrival variations and possibilities of limited but not catastrophic data loss (including data arrival after it may be optimally used).

To maintain stable and accurate coupling, data needs to be made available at each site at regular intervals and in a controlled manner. However, while in single-site testing, data transfer between substructures is in effect instantaneous; in a distributed environment data transfer must contend with an additional variable delay on sending and receiving control data. This is mainly due to network latency and its variation (jitter). It is caused by the propagation time of information, data processing delays across routers and switches and may include packet queuing. Communication delay may also be dependent on the performance of local site computing hardware, communications program function and network protocols used. An important factor in communications delay is the test site topology and distributed control strategy chosen. In a multi-site test, the site hosting the main numerical model (*‘the client’*) is ideally placed centrally with the shortest network path to all other sites. The distributed test control strategy determines when data is transmitted between test sites and controls the synchronisation of distributed controller time-steps. While distributed controller time-stepping cannot be made to exactly coincide, the relative position of distributed time-steps may be controlled and may also be used to minimise communication delays [15].

The final essential part of a DHT testing system is the development of support systems to establish and manage interactions between distributed test operators and testing systems. Labs are often noisy and can be stressful working environments. Operators have to contend with control of multiple user interfaces and systems while ensuring that physical specimens are tested safely and correctly - especially in regards to samples that break. In moving to a distributed test setting, additional systems are overlaid and merged with existing ones, greatly increasing the chances of operator error. This is especially true in a multi-site test where the main test operator at the central

(‘client’) site needs to maintain awareness of both local and multiple distributed sub-structures. Therefore, a measure of a successful DHT testing system is how its design considers test usability. That is, how easy or practical it is to run a test and how effectively the test procedures and support systems meet the operators’ needs. Related to this is the operators’ experience of using the testing system and especially how they perceive the system e.g. as a tool to aid them or as a complex system that must be carefully managed. A good testing system has clear test procedures, enables safe operation and gives confidence to the operators by improving awareness without excess complexity. This allows operators to interact more easily and focus on test accuracy and the specimens being tested.

3. b) Distributed geotechnical 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. For testing of small scale structural models on shaking tables, scaling laws can mean that for some cases to achieve similitude between real systems and scale models, tests must be conducted faster than real-time [33]. For example, keeping acceleration un-scaled, mass is scaled and scale model earthquake duration may be considered as $L_R^{1/2}$ times real earthquake duration. L_R is the length ratio used and for a quarter scale model a 50s earthquake will last 25s. However, in general, for large or full scale testing of structural systems, experiments are either 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 because pore pressure generation is driven by inertial effects, and dissipation is therefore slowed to allow parity between the dynamic time scale and the consolidation time scale ($1/N^2$), [43]. Therefore, in experiments where physical parts are distributed to structural labs and are designed to study seismic soil structure interaction, these have to take place in real-time, at 1g. In this case, full scale tests are called for, but challenging, not just because of capacity issues but also due to difficulties in ensuring soil tested is adequately representative. Scaled model tests at 1g could be used, however, these will not represent true (full scale) soil stress states. Results from these tests would be evaluated in this context and depend on the scale factor used and how length scaling parameters are applied.

Distributed seismic centrifuge experiments are not practical but, if pore pressure generation is correctly applied, the resulting soil stress state is suitable for testing in situations where soil inertial effects are not significant. For example, harmonic loading due to waves or floating ice may be applied in an experiment involving a physical substructure of an offshore platform or wind turbine foundation. The experiment may be distributed with a physical structural unit e.g. damper and coupled together with a numerically modelled superstructure.

With this in mind, a centrifuge foundation system test coupled to a distributed structural system under seismic loading may still be simulated. The soil is not loaded with a short duration ground motion from a shaker inside the centrifuge. Instead, actuators located on the foundation system inside the centrifuge and on the distributed physical substructures apply seismic loads directly to the

physical parts and response is fed back. The soil may be prepared beforehand to simulate, for example, an upper bound strength. However, this neglects important features of seismic soil structure interaction, including inertial and rate effects on the soil, and the seismic wave propagation in the soil is not represented. Here the foundation system and the soil might represent a design resistance. Such tests may be an attempt to improve simulation rather than neglecting the foundation entirely, though other simpler methods may be appropriate.

4. Multi-site DHT with geotechnical centrifuge and structural labs

In this section we present experiments exploring the potential for multi-site DHT across both structural and geotechnical laboratories, at extended timescales.

A simplified experiment was devised as a development example (Fig. 3). The experiment was designed to test a multi-site distributed testing architecture connecting actuator arrays and physical substructures at Oxford and Bristol to actuators in the Cambridge centrifuge. It did 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. It was also used to explore the potential of distributed centrifuge soil structure interaction. It is part of a larger series of tests that were conducted to develop the real-time DHT technique.

Figure 3. Multi-site distributed hybrid test setup.

The test simulates the response of a linear five degree of freedom shear building model experiencing ground motion. It is represented by one main numerical model subject to lateral ground acceleration coupled to three physical substructures. Locally at degrees of freedom 2 and 3, the Oxford two storey column rig [38] using coupled 10kN actuators is placed. Distributed on the 1st degree of freedom, to allow direct comparison between the distributed sites, are both the Bristol column rig (with 50kN actuator) and the Y direction of the Cambridge XY shallow foundation pad rig resting atop a bed of loose dry sand. The foundation pad is controlled using two electrical motors (rated to 10kN) 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. A variety of simulations have been tested.

4. a) 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 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 internal control signals. The distributed architecture used for the three site test is shown in Fig. 4.

Figure 4. High level representation of the three site distributed architecture used.

The client server architecture shown was found to be the most effective [5,15]. This is because it inherently allows more control to the central test site operator as a client in initiating the test and this reduces the operational burden at the central test site. In this set-up the client can connect to the waiting servers when it is ready. This is in preference to the scenario where the central site acts as the single server waiting for connections from distributed clients and must be ready first.

In this view the Oxford node, which is most centrally located on the network and geographically, 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 and providing these as services for the Oxford node.

The architecture developed uses existing hardware and software and adapts these legacy systems for distributed use by incorporating changes in the control systems. The architecture uses a test version of the UK-NEES Distributed Hybrid Testing communications program, DHT, 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 institutional Internet connection (known as the JANET production network) may be used. Distributed testing poses unusual network requirements: network latency is of greater importance than bandwidth. The advantage of the Lightpath network is that packets have a fixed path to each site, they are not routed and do not compete with other data as on the production network, where routing occurs and the geographical path may change. Lightpath was installed to ensure that during a distributed test network usage fluctuations would not interfere with the test.

4. b) Local testing environments

Bristol and Oxford share similar testing hardware and software. They both use single-tasking processor boards allowing 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 Application Programming Interfaces. In testing, both Oxford and Bristol use dynamic actuators and have the capability to run real-time hybrid experiments. As real-time testing is not a priority, Cambridge uses high load capacity electric 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 strong-box. The Cambridge systems run LabView on a Windows (multi-tasking) environment to allow communication with a Computer Boards Analogue/Digital (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.

4. c) UK-NEES distributed hybrid testing program, DHT

These multi-site tests represented a culmination of preliminary development work. Extensive development had ensured that the distributed environment between Bristol and Oxford was well understood and stable communication at 20ms time-steps achieved; a marked improvement over initial numerical only testing where large delay fluctuations and delay up to 500ms would occur [17]. Part of the aim of this test was to extend distributed testing capabilities to Cambridge and to explore the issues experienced in the new multi-site test environment. Communication was achieved using the program DHT. The program aims to enable heterogeneous controller connections with little end user customisation. It is written to maximise efficiency and makes use of standard software libraries to transmit the required control signals between sites. The program and ancillary control software provide features both for background tasks related to control of distributed machines and foreground tasks related to human to human co-ordination. The version used for this test was adapted for multi-site testing with Cambridge, and a high level representation of the workflow is shown in Fig. 5.

Figure 5. The workflow used by the multi-site DHT program.

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 communication rate is governed by the slowest site - in this case Cambridge. Here TCP/IP (Transmission Control Protocol / Internet Protocol) was used for network communication. TCP/IP is the most commonly used method for Internet communication. Use of TCP all but guarantees that data is delivered, without duplication and in order. The program completes its operations loop as fast as possible sending or receiving the minimal information required. A synchronous mode of operation with TCP is used (though testing with mixed TCP/UDP (User Datagram Protocol) communication was also conducted). While this ensures robust communication, it can have severe consequences for test latency. This approach means that the client (communication program) on sending server commands can only continue its operation when the client receives server response data back. By waiting, the client may miss or can send the next time-step command much later than the processor board generates it. At the server, response data will not be sent as it becomes available but only after the receive operation completes and data arrives from the client. Therefore, tests proceed by selecting a suitable numerical model solution time-step to allow sufficient time for round-trip data transfer. In this case 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 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.

4. d) Multi-site test results and discussions

A series of hybrid tests with the centrifuge at 1g and 50g were conducted. Fig. 6 shows results from one seismic test at 50g. Here a linear modal model (fixed step, multi-rate with no predictor [15]) is used as the numerical substructure to provide command signals to the respective actuators as in Fig 3. In the example shown, the numerical model mass per storey is 89.1 tonnes and storey lateral stiffness is 42.1kN/mm. This results in numerical model natural frequencies of 0.98Hz, 2.87Hz, 4.53Hz, 5.82Hz and 6.64Hz and 5% first mode stiffness proportional damping is applied. The properties are chosen to provide representative dynamic motion but the focus is on test connections.

Figure 6. Sample results from time-scaled multi-site DHT (note Cambridge displacement and force scaled to client site scale. Achieved - Achv., Command - Com., Received - Recv.).

A third order time-step independent polynomial extrapolator/interpolator is used to smoothly actuate and compensate for local actuator delays at the client [38]. Delay compensation for the distributed actuators 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 second order polynomial interpolator – which earlier testing had found performed better in the distributed environment [15]. Second and third order polynomial interpolators perform similarly when sub-stepping actuation commands, if data is available in equally spaced time-steps at main calculation steps. However, in a distributed test if a main-step data point is received too early or late, due to communication fluctuations which are not corrected for, the main step data points are not equally spaced in time. In this case a higher order interpolation algorithm will lead to larger oscillations in the sub-stepped command. These oscillations introduce test errors which result in incorrect loading of specimens and can lead to instability. Second order interpolators are preferred, especially if no corrections are made to account for sub-step communication fluctuations.

The signal sent to Cambridge is absolute and offset so that the foundation pad may be actuated from the correct position. On arrival in Cambridge it is scaled (by the length scale 1/50 at 50g) to

convert to centrifuge scale. It is noted that applying $1/N$ scaling to the command is challenging at high N values and this resolution is possible through the stepper motor used. At Cambridge, on receiving commands, the local actuator controller smoothly actuates towards them using a feedback control loop. The measured forces are scaled up by 50^2 to return the force to full or client site scale, the achieved displacements are scaled up by 50. Since the test did not aim to represent a particular engineering problem but rather was aiming to evaluate DHT technology, on arrival at the client, the scaled up centrifuge force is reduced by a factor of ten. This reduction at the client is for convenience and does not adhere to centrifuge scaling rules. However, in a test designed to capture dynamic response with more accuracy, it is important to consider the relative scale of specimens inside and outside the centrifuge. In this case the forces from the foundation pad were reduced both for stability and to allow them to be on the same order as the Bristol actuator. Since delay compensation was not applied, test stability was achieved by ensuring that the physical part of the test did not greatly alter the overall dynamic response. The physical column at Bristol and the Cambridge foundation pad represent a small proportion of the overall storey load.

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 distributed sites. Therefore, test accuracy is limited. Minor data loss events would occur (leading to unwanted oscillations in sub-stepped command at Bristol), but due to the time-step chosen, not enough to significantly affect the test. The tests did however, indicate, that under the right test conditions a satisfactory foundation response may be represented.

5. Real-time DHT

The multi-site test above and other two-site experiments conducted led to the following general findings. Firstly, synchronous TCP communication leads to large communication delays. Asynchronous operation is superior as there is no waiting for data receipt before checking for availability of new data to send from the controller board. New data can be sent almost as soon as it becomes available. However, implementing asynchronous control with TCP is problematic if sending data at a high rate. For example, for efficiency, small data packets from multiple time-steps may be held to be sent together rather than individually [44]. TCP is popular for network communication due to its reliability with features including a capability to automatically resend lost packets. However, its data flow control features are not favourable for real-time control and may not be readily controlled by the DHT testing system. A simpler communications protocol is preferable and flow control may be customised if using a dedicated higher level DHT protocol.

It was further found that the dedicated Lightpath network did not perform better than the shared production network. This could be partly because the connection to Cambridge from Oxford was likely following a path via Warrington, a 1160 km round trip [15]. At the time a more direct route was not available. In comparison the production network has a roundtrip path of around 480 km via London. The roundtrip path to Bristol on the production network is around 700 km via London, on Lightpath it is more direct at around 360 km. Test performance was generally marginally better on the production network in two-site tests between Oxford and Bristol and some quantitative comparisons may be found in [15]. This is in spite of the fact that ping times were the same. This may indicate better hardware performance on the production network - it caters for more users.

While this result may be surprising especially since the dedicated network follows a shorter path in these tests, it is important not to draw broad conclusions. Network performance depends very much on how the network is set up. The Lightpath network was a software limited 100Mbps dedicated network, while the shared institutional network had a capacity of up to 10Gbps on the backbone at the time. However, this refers to bandwidth. Low latency (and jitter) is more critical for

real-time DHT. Conventionally most networks are set up for high bandwidth data transfer, low latency is not inherent but could be found if for example, higher performance network hardware is used. One additional benefit of a dedicated network is that it may be more readily optimised for low latency (though it wasn't necessarily in our case) by careful selection of hardware (including locally) and optimising how data operations are carried out, provided the resources are available.

Furthermore, once communications are appropriately set-up, most of the communications issues encountered were due to data transfer within the local computing environment. This is due to competing Host PC processes accessing the local controller board bus (one for communication via the DHT program and the other, data capture via the dSpace program Control desk) and due to the multi-tasking Host PC environment. Also, regardless of how well the local environment is set up, small fluctuations in data arrival time occur when data is received on to the fixed time-step hardware controllers. These fluctuations and the much less likely possibility of communication interrupts or data non-arrival need to be overcome or mitigated. Lastly, for fast tests up to real-time, the test must be set up to minimise delay. Single-site (local) delay compensation algorithms do not perform as required and improved delay compensation stability and accuracy is essential.

With these learning experiences, work continued between Bristol and Oxford to develop the real-time DHT architecture and in order to demonstrate that robust DHT at faster rates up to and including real-time is possible.

5. a) Components of a real-time DHT simulation system

In order to overcome the problems associated with earlier tests and to enable real-time DHT, the following test control layout was proposed [3,15], (Fig. 7). It is shown for two-site testing with a single distributed actuator. However, it is readily scalable to multiple sites with multiple actuators.

Note: Server operates at actuation time-step rate e.g. 1ms. Client operates at main numerical model time-step rate e.g. 20ms. Client data handler operates at actuation time-step rate e.g. 1ms. Rate data sent between sites depends on control strategy.

Figure 7. Scalable real-time DHT test layout.

Here, a generalized control layout is presented. However, the final test layout depends on the distributed control strategy applied. This determines how time-stepping is synchronised between local and distributed controller boards and when data is transmitted. Fixed-step time-stepping may begin on the client site or server site controller boards on receipt of the first command signal from the client communications program (IC-DHT [3,15]) on the host PC. All sites use the same main time-step size to solve numerical models and for actuation main step targets. The client site hosts

the main numerical model/substructure. This requires feedback control data from local and distributed substructures prior to the end of its current main step. It provides desired end of time-step actuation positions for local and distributed actuators. These (displacement control) commands may pass onto either local or distributed delay compensators. However, tests revealed that conducting distributed delay compensation at the client is simpler and less error prone. Distributed delay compensation relies on feedback of actuator achieved displacement corresponding to achieved force and a new type of delay compensator has been developed to enable accurate delay compensation of large variable delays at large time-steps [3,15].

Command data in the form of time-stamp and delay compensated target are then passed on to the network. Communication and optional control data saving is via the IC-DHT program. This uses the User Datagram Protocol (UDP) which is favoured when timely delivery of data is important [45]. It is applied in an asynchronous manner and new data transferred to and from the board (in a receive - write / read - send loop) practically as soon as it is available. Data is sent according to the control strategy and sending rate controlled by client and server time-stamps. Although with UDP data delivery is not guaranteed, testing on this network revealed no data delivery issues.

Locally data passes to an actuation sub-stepper for smooth continuous actuation. At distributed sites all incoming data is passed through a data handling protocol. This is to ensure that appropriate data is continuously available when required, to safely and accurately actuate and to provide data for numerical substructures. At a server, data handling occurs at actuation sub-steps. To deal with fluctuations in incoming data arrival times, new data may be held if arriving earlier than the end of the current actuation main step. If data arrives late, an accurate prediction of the expected command (using compensation algorithms) is made. This command prediction may be smoothly replaced with the received target if data arrives before the start of the next actuation main step. These additional predictions are possible due to the inherent accuracy of the received delay compensated command, and so do not introduce significant errors. Data handling prediction past one and half main time-steps will start to introduce more significant test errors. The data handler will provide sub-step actuation commands (at 1ms intervals) to the received or predicted target.

Concurrently the servers send back achieved actuation and force data to the client together with a time-stamp. This data arriving at the client passes through the client data handler. This data handler operates at actuation sub-steps but, provides data at the numerical model time-step. It uses the received time-stamp to determine control actions. If the expected time-stamp is received, the latest available data has arrived and is directly passed to the numerical substructure and delay compensation algorithms. If data does not arrive on-time, it will be considered late by one time-step. An accurate prediction of the expected achieved force and displacement is made based on previous incoming data, for use instead. In practice, since testing needs to select a time-step for which stable communication can be achieved, data is very rarely (this is not a usual test feature) not received by the end of the client time-step. At the server, data handling prediction events are more common but rarely applied past 5-6 actuation sub-steps after they are expected. Once stable communication is achieved, larger data arrival delays and loss at the server can be mitigated.

The control layout has been chosen as it is suitable for commonly used legacy controller boards that are installed on a host PC and do not have on-board NIC's (Fig. 2 - site 1). This is also the type available for testing between Bristol and Oxford. Network communication is therefore achieved through the host PC NIC's. However, the systems and control architecture presented are equally valid for testing with controller boards with on-board NIC's.

As it was known that the multi-tasking environment on the host PC caused much of the communication issues related to data loss and late arrival, changes were made. Firstly, data capture off the controller bus was minimised; in some cases data capture only occurred using the IC-DHT program. Secondly the host PC environment was optimised, removing unnecessary processes and enabling a soft-real time priority for the communication program on a dedicated processor. These

measures revealed that robust data transfer could be achieved within the operating range of the data handling protocol, and that acceptable test accuracy could be achieved, ensuring real-time actuation and solution of numerical model equations.

In the system implemented we rely on robust communication achieved between distributed hard real-time controllers running numerical and physical sub-structures connected together across a non real-time environment using a data handling protocol. The maximum numerical model solution frequency required for stable and accurate solution depends on the minimum time-step in which robust communication can be achieved. This can be determined by numerical only distributed tests. Extensive repeat testing revealed that once a time-step for stable communication is achieved, hard real-time testing could be conducted when required. Attempting smaller time-steps for control would also consistently fail. Therefore a successful test between distributed hard real-time controllers is achievable by balancing computational load on and off the controller by limiting or eliminating competing processes and minimising the rate and quantity of data sent to each site. The control signals sent between distributed substructures relies on the performance and reliability of the network. Testing should ideally be conducted with some oversight by network operators to avoid planned maintenance. While testing has been successful on the Internet, a dedicated network such as Lightpath provides additional reliability.

5. b) Real-time DHT demonstration test

As an example of the performance of the developed system and its potential in a relatively complex experiment, a multi-axis two-site real-time DHT experiment is presented in Fig. 8.

Figure 8. Two-site multi-axis real-time DHT demonstration test.

In this experiment the seismic response of a three storey shear building with rigid base and one lateral degree of freedom per floor is emulated. Two tests are planned, and the experiment is designed to capture the structural response with and without the inclusion of a shear type metallic energy dissipator on the first floor.

In this scenario, locally, three actuators are used to connect two separate physical substructures to the main numerical substructure in Oxford. However, since the other Oxford actuators could not be brought online at the time due to lab restrictions, one actuator and physical substructure is requested from Bristol to fulfil the experiment.

Details of the individual test rigs may be found in [3], and this demonstration test is designed to represent the general dynamic characteristics of earthquake engineering structural systems. The full model natural frequencies are estimated to be 1.16Hz, 3.26Hz and 4.71Hz and 5% first mode stiffness proportional damping is included. The main numerical model is linear and solved using 20ms time-steps, allowing a response of up to 5Hz to be accurately captured. All actuation main steps are 20ms and smooth 1ms sub-stepped commands are provided. As the column rigs are separate, to feedback the correct response force, the commands to the Oxford column rig actuators depend on the first floor response. An additional linear numerical substructure is used to represent the stiff dissipative device brace and the device desired deflection depends on the relative brace stiffness. Approximately half the stiffness of the lateral columns in the emulation are represented by physical columns and for compatibility with the model, the physical column feedback forces are multiplied by 48.

Results are presented in Fig. 9. In the first test the column rigs are tested together without the dissipator, and in the second test the dissipative device test rig is also brought online.

Figure 9. Two-site real-time DHT - 1 distributed and 3 local actuators - test with & without dissipator.

As shown, the inclusion of the device on the ground floor significantly reduces the storey drifts. Device hysteresis is significant. The response of the Bristol column is primarily linear. However, the Oxford column rig presents an unusual nonlinear response even under low displacements, this is

due to column degradation caused by prior tests close to yield. The hysteresis loops shown are all for the test with dissipator.

The Newmark explicit algorithm is used for testing here, providing the desired end of time-step response one time-step ahead. However, to interpolate smooth sub-step commands, a one time-step delay is introduced. Locally at Oxford, test control is achieved using a variable delay [46] second order time-step independent polynomial delay compensation algorithm [15, 38]. The 10kN column rig actuators exhibit a slightly varying actuation delay of around 12-14ms. The lower column rig actuator experiences coupling forces from the upper, and although minor additional outer loop feed-forward gain is added to the command, this does not significantly improve performance. Amplitude errors are not significant for the column rigs and the controller used adequately corrects for the dynamics of the transfer system. For the dissipative device test rig the stiff nonlinear device response means that both delays and amplitude errors vary considerably throughout testing due to hardening and degradation. A variable amplitude, variable delay second order polynomial delay compensation algorithm is used [15] to overcome this and control is achieved by feeding back the displacement across the device as measured by an encoder to the outer loop. With the device used in this test actuation delays were around 20-25ms.

To achieve distributed control a SSF (Server Start First) control strategy [15] is applied. The SSF approach synchronises local and distributed controller board time-stepping such that the server main-time steps begin and end before the client. The advantage of this approach is that overall test delay is minimised. Since the server actuation is ahead of the client, there is additional time to transmit the results back. The SSF control strategy chosen for this test applies fixed delay compensation. This was chosen for its computational efficiency and is suitable here since the distributed test rig at Bristol exhibits a relatively fixed actuation delay. Its operation is shown in Figure 10 at the client and server. A zoomed view is shown at the server in Fig. 11.

Figure 10. Control signals: SSF Client (Oxford) and Server (Bristol).

The test starts with the IC-DHT communications program at the client initiating the test following a command by the main test operator. IC-DHT sends the start test command to the server controller board, and after an initialise time counted by the program (12.5 ms in this test), it starts the client controller board. This time corresponds to the estimated fixed actuation delay. At the server on receiving the test start command, the data handling algorithm is used to predict end of time-step actuation commands (signal 1) and to provide sub-step actuation commands (signal 3). For this test, this is achieved by a second order fixed step polynomial algorithm. On test start at the

client a very accurate one time-step extrapolation is made using a large time-step prediction algorithm which uses input from a compensator predictor running in parallel to the test (signal 2 Oxford) [3, 15]. This compensator predictor uses knowledge of the numerical part of the test and prior force feedback. The client command is therefore sent two time-steps (40ms) ahead of when the server response is required back. On the server, the client command (signal 2) is used to update the original server end of time-step prediction. This is switched to immediately by the sub-stepper. Due to the relative accuracy of the server prediction, and as the more accurate target data arrives midway in the time-step, this does not have a significant impact on actuation quality. Small differences that may be observed in the command on switching target are smoothed by the high actuation rate of the actuator. The server sends back force displacement response data as soon as it receives command data. As the command data arrives at a time corresponding to the fixed actuation delay and the server is already running ahead of the client, the force fed back corresponds to the desired current end of time-step position at the client. As there is enough time to send this data to the client before the end of its current main time-step, the client will be able to use this data to accurately continue operation.

At the client and server the data handler on the board will manage incoming control data to assess its suitability for safe actuation. Within the tolerance of its own prediction algorithms, it will ensure that accurate end of time-step data is available for use on time by the controller board at each site, should data not be available on time. In these tests no client correction events (signal 5) are required for incoming data (signal 4).

Figure 11. Control signals: zoomed view at Server (test with dissipator).

Distributed control is accurate and repeatable, and physical substructures are accurately actuated at fixed time-steps. Fig. 12 presents various plots demonstrating test quality. These focus on the distributed part of the test. The top left plot shows that the desired response closely matches the achieved actuator response (measured at the client). This is confirmed by the top right synchronisation plot indicating little to no hysteresis between the desired and achieved response. Control errors may also be scrutinised by looking at the real or residual error defined here as the difference between desired and achieved response. The middle plots present the residual error for the distributed actuator in the two tests (with and without dissipative device). This is presented both relative to the desired response (left) and the time it occurs (right). Both tests indicate test errors are low and quite acceptable. The norm of the residual (the square root of the sum of the residual errors squared) is very low (zero is a perfect test). As a comparison, the residual error results for the comparable upper two storey column rig actuator using local actuation controllers is shown

(bottom). Distributed and local test errors are comparable, in-fact in this case distributed residual norms are lower than local ones. Differences in real errors for both local and distributed actuators in multiple tests are to be expected as a characteristic of testing. Actuation delay at Bristol is around 12-14ms and distributed test quality may be improved by slightly adjusting the initialisation time. Alternatively, variable delay SSF or CSF (Client Start First) strategies may be used [15]. The algorithms developed are equally valid for distributed non-linear physical substructures [3].

Figure 12. Comparison of distributed and local actuator actuation quality.

5. c) Setting up a DHT test

While DHT has significant benefits, it is significantly more complex than the local technique. In order to set up a distributed testing system, the following steps should be undertaken [15]. Establish network communication and determine network delays between sites. Establish communication between real-time controller boards. Determine the minimum time-step for reliable communication, selecting the optimum test site topology and control strategies to meet communication delay needs. If time-step is low enough consider real-time testing, if not and it is appropriate, use scaled-time testing. Implement a distributed test control strategy including distributed time-step synchronisation

and delay compensation. Implement adequate incoming data handling. Run software only and actuator only tests to better indicate extent of delays and effectiveness of controller implementation prior to main test. Consider applying additional tools to aid work in a distributed environment.

Once the testing system has been set up and communication and control issues resolved, where network delay allows, real-time DHT tests can be achieved that are on a par with local tests. Improvements in test quality can be achieved by improvements in delay compensation algorithms, inner loop control and actuation hardware, and this applies also to single-site testing. However, the current limitation in real-time DHT as compared to single-site real-time testing is the smallest numerical model time-step than can be achieved. Smaller time-steps allow higher frequency response to be more accurately captured but also require feedback at that rate. In real-time DHT the time-step is limited by packet queuing and serialisation delays and ultimately, the propagation time of information through a physical medium (around $2/3^{\text{rds}}$ of the speed of light). Hence the network distance is a key factor. A system which is capable of solving numerical models at 20ms time-steps is practical across large distances as we have demonstrated and can represent the dynamics of a system up to 5Hz with good accuracy. This is suitable for a broad range of earthquake engineering simulations. This limitation is not relevant for scaled-time DHT.

6. Conclusions

Distributed testing provides new and exciting possibilities to develop our scientific knowledge by making practical, experiments that have previously not been possible due to capacity limitations and a lack of expertise at a single site. This article presented groundbreaking developments in DHT representing a synthesis of experience from a comprehensive development programme. It presents the roles, challenges and opportunities of DHT and shares experiences through key test examples. It therefore acts as an important starting point for those wishing to apply DHT as a means to enable research and development.

Examples included early experimental work on coupling a foundation system inside a geotechnical centrifuge to structural systems at two geographically distributed structural dynamics labs. Centrifuge time-scaling incompatibilities can mean hybrid testing to capture real soil structure interaction effects of centrifuge foundation systems coupled to distributed structural elements under seismic loading has significant limitations. However, if the initial soil conditions are set up correctly, DHT under other dynamic load conditions such as waves or floating ice can be practical as these are not driven by soil inertial effects. The test, conducted at an extended time-scale, was also used to understand testing issues in the new distributed environment.

Presenting a real-time DHT system, we developed a relatively complex two-site real-time DHT experiment using three medium-scale physical substructures, demonstrating the influence on structural response with and without a physical energy dissipator. The testing system enabled high quality distributed testing on a par with existing local (single-site) test techniques. This was achieved through development of effective control strategies and advanced algorithms to overcome distributed communication delays and data arrival fluctuations. The technique developed is readily extendable to multiple sites and multiple distributed substructures.

The real-time DHT system has been implemented using commonly available ‘hard’ real-time controller boards, typical of the legacy controllers available at testing laboratories. This ensures the fixed time-stepping essential for accurate actuation. To allow robust communication between controller boards and local PC network interface, the testing environment has been optimised for communication. While this may be further improved by using proprietary software for real-time operating systems, extensive repeat testing revealed that real-time DHT could be achieved as required. In future the developed systems may be implemented on controller boards with on-board network interfaces, overcoming data transfer delays within the local test site and extending computational capacity.

The tests presented used both a dedicated private network which followed a fixed path and the shared institutional Internet connection. The testing system performed well on both. The advantage of a dedicated network is increased reliability. However, DHT should ideally have some oversight by network operators to ensure the network remains operational during testing. Improved reliability may also be achieved on the Internet by special priority routing - if granted by network operators.

A framework for hybrid testing has been presented in this article connecting the main earthquake engineering laboratories in the UK to conduct DHT. As testing develops this and similar networks will enable closer international collaboration and provide an accessible route for research institutes and industry to develop technology which can provide solutions to meet our infrastructure needs.

Acknowledgements

The authors thank the EPSRC for partially funding this work (particularly grant No's. EP/D079101/1 and EP/D080088/1). We also thank other members of the UK-NEES team, especially Arshad Ali, Javier Parra Fuente & Kashif Saleem.

References

1. Foster I. & Kesselman, C. (eds.) (2004) The grid 2 (second edition): blueprint for a new computing infrastructure. San Francisco, CA. Morgan Kaufmann.
2. De la Flor G., Ojaghi M., Lamata Martínez I., Jirotko 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.
3. Ojaghi M., Williams M.S., Dietz M.S., Blakeborough A. & Lamata Martínez I. (2013) Real-time distributed hybrid testing: coupling geographically distributed scientific equipment across the Internet to extend seismic testing capabilities. *Earthquake Eng. Struct. Dyn.* DOI: 10.1002/eqe.2385.
4. Ojaghi M., Parra-Fuente J., Saleem K., Blakeborough A. & Williams M.S. (2007) Grid based distributed hybrid testing. *Proc. of the UK e-Science All Hands Meeting*, Nottingham, UK, 190-196.
5. 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.
6. Ojaghi M., Lamata Martínez I., Dietz M., Williams M.S., Blakeborough A., Crewe A., Taylor C., Madabhushi G., Haigh S. & Ali A. (2010) UK-NEES - Distributed Hybrid Testing Centring Bristol, Cambridge and Oxford Universities: Connecting Structural Dynamics Labs to a Geotechnical Centrifuge. *9th U.S. National and 10th Canadian Conference on Earthquake Engineering, Toronto*. Paper 1024.
7. Dietz M.S., Dihoru L., Oddbjornsson O., Bocian M., Kashani M.M., Norman A.P., Crewe A.J., Macdonald J.H.G. & Taylor C.A. (2012) Earthquake and Large Structures Testing at the Bristol Laboratory for Advanced Dynamics Engineering - In Role of Seismic Testing Facilities in Performance-Based Earthquake Engineering. SERIES workshop. Fardis M. & Rakicevic Z. (eds.) Springer, DOI 10.1007/978-94-007-1977-4.
8. 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. 13th WCEE*. Vancouver. Paper 1674.
9. 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.
10. Kwon O. S., Nakata N., Elnashai A. & Spencer B. (2005) A framework for multi-site distributed simulation and application to complex structural systems. *Journal of Earthquake Engineering*, 9(5), 741-753.
11. Pan. P, Tomofuji H., Wang T., Nakashima M., Ohsaki M. & Mosalam K. (2006) Development of peer-to-peer (P2P) internet online hybrid test system. *Earthquake Eng. & Struct. Dyn.* 2006; 35:867-890.
12. Takahashi Y. & Fenves G.L. (2006) Software framework for distributed experimental-computational simulation of structural systems. *Earthquake Eng. & Struct. Dyn.* 2006;35:267-291.
13. Mosqueda G., Stojadinovic B., Hanley J., Sivaselvan M. & Reinhorn A. (2006) Fast hybrid simulation with geographically distributed substructures. *17th Analysis and Computation Specialty Conference*.
14. Wang K., Tsai K., Wang S., Cheng W. & Yang Y. (2007) ISEE: Internet-based simulation for earthquake engineering - part II: The application protocol approach. *Earthquake Eng. & Struct. Dyn.* 2007; 36:2307-2323.
15. Ojaghi M. (2010) The development of real-time distributed hybrid testing for earthquake engineering. DPhil. Thesis, Univ. of Oxford.
16. Ojaghi M., Lamata Martínez I., Williams M.S., Blakeborough A., Dietz M. & Crewe A. (2010) Real-time hybrid testing in geographically distributed laboratories. *14th European Conf. on Earthquake Eng.* Ohrid. Paper 872.

17. 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*. Oxford Uni.
18. Madabhushi S.P.G., Haigh S.K., Ali A., Williams M., Ojaghi M., Lamata I., Blakeborough T., Taylor C. & Dietz M. (2010) Distributed Testing of Soil-Structure Systems using Web-Based Applications. *7th International Conference on Physical Modelling in Geotechnics (ICPMG2010)*, Zurich.
19. Crewe A.J., Dietz M., Evans J. & Taylor C.A. (2008) UK-NEES: Summary of Activities at Bristol University. *Proc. The 14th World Conf. on Earthquake Eng. Beijing*. Paper S17-03-014.
20. Taucer F. & Apostolska R. (eds.) (2015) Experimental Research in Earthquake Engineering EU-SERIES Concluding Workshop. Springer DOI 10.1007/978-3-319-10136-1
21. Lamata Martínez I., Ioannidis I., Pegon P., Williams M. & Blakeborough A. (2014) Process and Future of Data Integration within the European Earthquake Engineering Laboratories. *J. Comput. Civ. Eng. ASCE*. Vol. 208 Issue 3.
22. Lamata Martínez I., Williams M. S., Dyke S., Krötzsch M. & Pegon P. (2017) Next directions in experimental data for seismic hazard mitigation, *Engineering Structures*. Vol. 136, Pages 535-546.
23. Dorka U.E., Santacana F.O., Williams M., Blakeborough A., Lamata I. Ojaghi M., Bursi O.S., Abbiati G., Mei Z., Bousias S., Kurc O. & Bahcecioglu T. (2013) Deliverable [D2.6] – Report on distributed testing and perspectives. *SERIES. Capacities Specific Programme Research Infrastructures Project No.: 227887*
24. Lamata Martínez I., Obón Santacana F., Williams M. S., Blakeborough A. & Dorka U. E. (2016) Celestina-Sim: A Framework to Support Distributed Testing and Service Integration in Earthquake Engineering. *J. Comput. Civ. Eng., ASCE* Vol. 30 Issue 1.
25. Williams M.S. & Blakeborough A. (2001) Laboratory testing of structures under dynamic loads: an introductory review. *Phil. Trans. Royal Soc. London, Series A*, 359, 1651-1669.
26. The Failure of an Offshore Platform, by R. G. Selby, F. J. Vecchio, and M. P. Collins (1997) *Concrete International* 19(8), pp. 28-35.
27. Taucer P. (2004) Cooperative advancements in seismic and dynamic experiments. Recent advances and future needs in experimental earthquake engineering. *CASCADE* report No. 7.
28. <http://www.bosai.go.jp/hyogo/ehyogo/index.html> (last accessed Mar. 2017)
29. <https://nees.org/> (last accessed Mar. 2017)
30. <http://english.iem.ac.cn/Index/index.html> (last accessed Mar. 2017)
31. <http://www.iiees.ac.ir/en/> (last accessed Mar. 2017)
32. Nakashima M. (2008) Roles of large structural testing for the advancement of earthquake engineering. *Proc. The 14th World Conf. on Earthquake Eng. Beijing*. Keynote, K009.
33. Rakicevic Z.T., Bogdanovic A. & Jurukovski D. (2012) Structural and Behaviour Constraints of Shaking Table Experiments- In Role of Seismic Testing Facilities in Performance-Based Earthquake Engineering. *SERIES* workshop. Fardis M. & Rakicevic Z. (eds.) Springer, DOI 10.1007/978-94-007-1977-4.
34. Highway Innovative Technology Evaluation Centre (HITEC) (1999) Technical Evaluation Report, Summary for the Evaluation Findings for the Testing of Seismic Isolation and Energy Dissipating Devices.
35. Clarke, C.S.J., Buchanan, R., Efthimiou, M., & Shaw, C., (2005) Structural platform solution for seismic arctic environments, Sakhalin II Offshore Facilities. 2005 OTC. 17378-MS.
36. Ojaghi M., Whittle J., Williams M.S. & Blakeborough (2010) Comparing shear type metallic and fluid viscous energy dissipation devices using real-time hybrid testing. *9th U.S. National and 10th Canadian Conference on Earthquake Engineering*, Toronto. Paper 898.
37. Blakeborough A., Williams M.S., Darby A.P., Williams D.M. (2001) The development of real-time substructure testing. *Phil. Trans. Roy. Soc., A359*, 1869-1891.
38. Bonnet P.A. (2006) The development of multi-axis real-time substructure testing. DPhil. Thesis. Univ. of Oxford.
39. Mahin S.A. & Shing P.B. (1985) Pseudodynamic method for seismic testing. *Journal of Structural Engineering*, Vol. 111, No. 7. Paper No. 19867.
40. Bonnet P.A., Williams M.S. & Blakeborough A. (2007) Compensation of actuator dynamics in real-time hybrid tests. *Proc. IMechE Part I: J. Systems and Control Engineering*. Vol. 221. Special issue paper. pg. 251–264.
41. McCrum, D.P. & Broderick, B.M. (2013) Evaluation of a substructured soft real-time hybrid test for performing seismic analysis of complex structural systems. *Computers and Structures*, Vol. 129, pp. 111-119.
42. Popovici K. & Mosterman P.J. (2012) Real-Time Simulation Technologies: Principles, Methodologies, and Applications. CRC Press
43. Madabhushi, S.P.G. (2015) Centrifuge modelling for Civil Engineers, CRC Press.
44. Braden R. (1989) RFC 1122. Requirements for Internet hosts -- communication layers. *Internet Engineering Task Force*.
45. Van Valzah R.A., Montgomery T.L. & Bowden E. (2009) Topics in high-performance messaging. *29West, Inc.*
46. Darby A. P., Williams M. S. & Blakeborough A. (2002) Stability and delay compensation for real-time substructure testing. *Journal Eng. Mech. ASCE*. Vol. 128, 2002, 12, 1276-1284.