Rittenbruch, Markus, Sorensen, Andrew, Donovan, Jared, Polson, Debra, Docherty, Michael, & Jones, Jeff
(2013)
The Cube: A very large-scale interactive engagement space.
In Horn, M S & Nacenta, M A (Eds.) *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces.*
Association for Computing Machinery, United States of America, pp. 1-10.

https://doi.org/10.1145/2512349.2512814

The Cube: A Very Large-scale Interactive Engagement Space

Markus Rittenbruch, Andrew Sorensen, Jared Donovan,
Debra Polson, Michael Docherty, Jeff Jones
Queensland University of Technology
2 George Street, Brisbane, 4001, QLD, Australia
{m.rittenbruch, a.sorensen, j.donovan, d.polson, m.docherty, ji.jones}@qut.edu.au

ABSTRACT
"The Cube" is a unique facility that combines 48 large multi-touch screens and very large-scale projection surfaces to form one of the world's largest interactive learning and engagement spaces. The Cube facility is part of the Queensland University of Technology's (QUT) newly established Science and Engineering Centre, designed to showcase QUT's teaching and research capabilities in the STEM (Science, Technology, Engineering, and Mathematics) disciplines. In this application paper we describe, the Cube, its technical capabilities, design rationale and practical day-to-day operations, supporting up to 70,000 visitors per week. Essential to the Cube's operation are five interactive applications designed and developed in tandem with the Cube's technical infrastructure. Each of the Cube's launch applications was designed and delivered by an independent team, while the overall vision of the Cube was shepherded by a small executive team. The diversity of design, implementation and integration approaches pursued by these five teams provides some insight into the challenges, and opportunities, presented when working with large distributed interaction technologies. We describe each of these applications in order to discuss the different challenges and user needs they address, which types of interactions they support and how they utilise the capabilities of the Cube facility.

Author Keywords
Multi-touch; Very large displays; Interactive wall displays

ACM Classification Keywords
H.5.2 [Information interfaces and presentation]: Graphical user interfaces (GUI)

INTRODUCTION
Interactive facilities that utilise large-scale multi-touch display surfaces are becoming increasingly common. However, the definition of what constitutes a large display surface varies widely. In the literature the notion of "large displays" ranges from single screens, such as large TVs [e.g. 4], to large rear-projected multi-touch wall displays [e.g. 5]. Due to advances in display technology there has been a noticeable shift from large, but relatively low-resolution applications, to setups that combine several interactive displays to create very high-resolution interactive surfaces, using either rear-projected [e.g. 1, 6] or LCD displays [3].

In this application paper we describe the design, implementation and end-user experience, of a unique, interactive exhibition and learning space deployed on a massive scale, named "The Cube". The Cube was commissioned at the Queensland University of Technology's (QUT) newly established Science and Engineering Centre in 2012/2013.

The QUT Science and Engineering Centre (SEC) has been established to foster research and education across the STEM disciplines. In order to support this activity the SEC identified visualisation, simulation and modelling across all areas of endeavour (science, engineering, technology, maths, society and culture), as an important research area. The new SEC precinct is perfectly positioned next to the Brisbane CBD, Botanical Gardens, and the popular South Bank precinct. The centre attracts a walking, transient, population of up to 70,000 people per week. This capacity to engage directly with a transient general population was seen as a powerful outreach opportunity for the University and has significantly impacted on the Cube's overall design and implementation.

The Cube was designed to be one of the major features of the SEC with a mandate to showcase QUT's teaching and research capabilities in the STEM (Science, Technology, Engineering, and Mathematics) disciplines. The Cube's giant interactive 'walls' are located on two levels of the SEC's open foyer area. The facility is open to the general public and the SEC receives a significant amount of thoroughfare from students, staff and visitors. The Cube is a demanding real world environment, with an operating period of at least 12 hours per day. A constant stream of tour groups, particularly from primary and high school outreach programs, also contributes significantly to the

Cube's utilisation. One of the challenges of the Cubes design was to implement a robust and scalable technical infrastructure supporting a wide range of applications suitable for the general public, whilst continuing to support significant ongoing research and education programs.

The Cube's technical infrastructure was developed in parallel with a set of five interactive "launch" applications that showcase the Cube's capabilities while providing a unique end-user experience.

In this application paper we will introduce and discuss the physical and technical platform, as well as describing each of the five applications developed for the Cube's launch. Our aim is to highlight the design rationale at different stages of the development process as well as to showcase how a range of different technologies and software frameworks were utilised to create engaging end-user experiences.

THE CUBE FACILITY

The physical and technical setup of the Cube is composed of 6 discrete zones, two of which can be combined into a mega-zone (1&2). The Cube's zones are designed to present a single continuous visual canvas, consisting of both multi-touch LCD displays as well as edge blended projection surfaces. In addition to the display surfaces the wall is constructed from a vast assemblage of network infrastructure, compute nodes, audio facilities and server infrastructure. We will present a brief overview of these infrastructure elements.

Physical and screen setup

The interior of the structure measures 14 meters long by 9 meters high, and each exterior wall measures 8.5 meters long by 4.5 meters high. The Cube structure is made up of 6 zones, as follows.

Zone 1 & 2

These zones each include 10 Multitouch panels in portrait orientation, 2 WUXGA projectors, 10 wide range loudspeakers (under each panel), dipole speakers for ambient/room audio and Gallagher T10 RFID readers.

The panels are arranged in a single row of 20 panels, bisected at the centre by an angle of 110 degrees.

The projectors are arranged 2 x 2, with a pixel dimension of 1920 x 1170px per wall. One computer provides signal to two projectors, however content can be scaled and projected to the entire projection if required. Each of the two zones contains a 10800 x 1920px multi-touch surface. Each pair of multi-touch panels is driven by a single SGI compute node. The total surface dimension of each zone is 6858 x 7788mm.

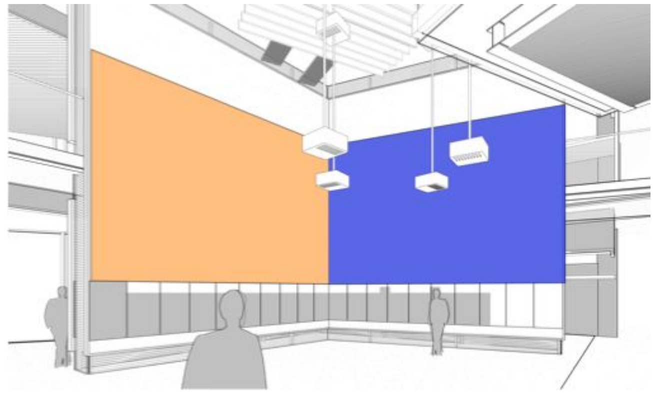

Figure 1. Zone 1 & 2 diagram

Zone 3 & 4

These zones each include 12 Multitouch panels in portrait orientation, 3 WUXGA projectors, 12 wide range loudspeakers (under each panel), 2 dipole speakers for ambient/room audio and 6 Gallagher T10 RFID readers.

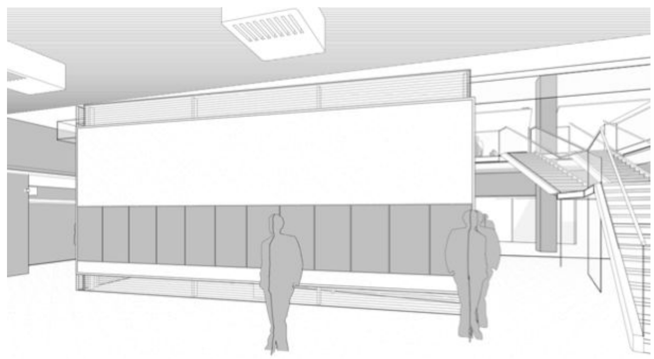

Figure 2. Zone 3 & 4 diagram

The projectors are arranged in a 3 x 1 format, with a pixel dimension of 5360 x 1114px per wall. A single computer provides signal to three projectors and alternate sources may be overlaid via the CorioMaster processor. The projectors are edge-blended to provide a single unified surface. Each of the two zones provides a 12960 x 1920px multi-touch surface. The total combined physical surface dimension of each zone is 8253 x 3374mm.

Zone 5 & 6

Zone 5 includes four Multitouch panels in landscape orientation, two WUXGA projectors, as well as four wide range loudspeakers (under each panel), two speakers for ambient/room audio, and two Gallagher T10 RFID readers.

The projectors are arranged in a 2 x 1 format, with a pixel dimension of 3030 x 1200px per wall. A single computer provides signal to two projectors and alternate sources may be overlaid via the CorioMaster processor. The zone contains a 3030 x 1200px multi-touch surface. The total physical surface dimension of the zone is 4866 x 1914mm.

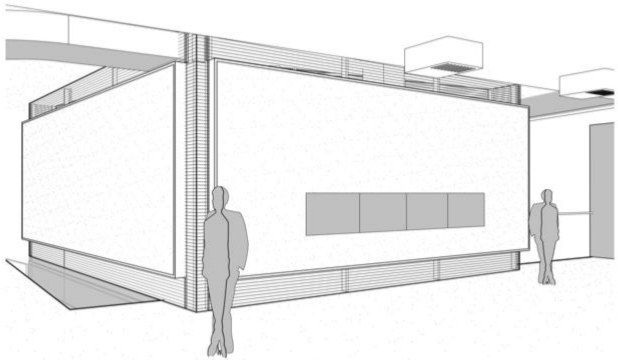

Figure 3. Zone 5 & 6 diagram

Zone 6 does not contain any multi-touch surfaces and is a pure projection surface consisting of two WUXGA projectors as well as four speakers for ambient audio.

Hardware setup

Components

The Cube is powered by a large number of computers, AV switching, audio, and image processing systems, all located in a dedicated server room on level 6 of the site.

The main technology features include:

- Graphics computing based on dual-socket, Windows 7 server-grade workstations with NVIDIA GTX690 display adapters (see spec. below),
- Windows Server 2008 or RedHat Linux Enterprise servers in physical or VM environments,
- Multiple, dedicated gigabit Ethernet networks,
- HDBaseT audio and video transport backbone,
- Broadcast class projection image processing,
- A MiFare 14443 Type A RFID sensor system,
- Custom-built dipole and line-array loudspeakers, provided per panel and throughout projection zones,
- Multi-channel audio processing (DSP).

Touch Panels

The 48 touch-panels installed throughout the Cube are MultiTaction MT553UTBs from MultiTouch Limited[1], Finland. The 55-inch panels are an optical (infra-red) based unit, with 32 cameras each, able to accept a virtually unlimited number of touch events. Each panel contains a dedicated Linux machine that processes and outputs touch events in the form of a TUIO stream at a rate of 100hz.

Server setup

Web, application, multi-user, and database servers complement the graphics nodes, with direct attached or SAN based storage available. Further, there is a dedicated (dark) fibre-optic tie to the central University network and border routers. Servers are either physical computers co-located with graphics nodes at the Cube, or are virtual machines hosted in the main University data centre.

Network

The data exchanged between nodes, servers and the Internet occurs over 3 networks: QUT LAN, Touch LAN, and the Interconnect LAN. Separate networks provide high performance and low latency, with little contention between data streams. The dedicated tie to the University data centres provides the Cube with almost direct attachment to the border routers (i.e. traffic does not enter the staff or student subnets). QUT has up to 10Gbits of bandwidth to the Internet.

DESIGN OF THE CUBE

The decision to incorporate a strong interaction and visualisation component into the new SEC building commenced in 2009. This early vision for the Cube included a public education program, visualisation research capacity, modelling and simulation, as well as the physical construction of the 'walls' and a collection of 'launch' applications. Extensive user group workshops and meetings were held to ensure concept design and design development incorporated fundamental requirements for visualisation technologies.

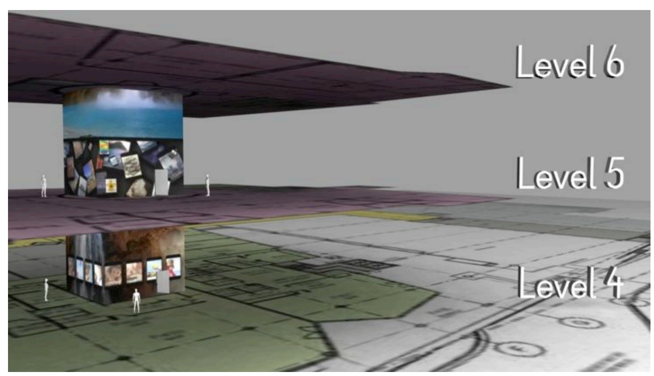

Figure 4. Early Cube design

In the early stages of planning the design context was defined by the anticipated use of the Cube, the physical layout of the building, the availability of technologies and the projected cost. One of the first design iterations consisted of a Cube[2] structure with continuous interactive display surfaces stretching over 2 levels (see Figure 4). This design was based on a rear-projected multi-touch setup. The original design had a range of issues. It posed structural problems (floating ceiling on Level 5), limited the resolution that could be achieved by exclusively using projection technology and, most importantly, did not constitute a good use of the available space.

[1] http://www.multitaction.com/products/ultra-thin-bezel/

[2] The name for this initial design was later used for the whole facility, despite the fact that the design evolved and the structure no longer resembled a Cube.

Subsequent designs were driven by a set of refined criteria:

- The **maximisation of surface area,** both in terms of physical space, resolution and interactive surface area.
- The anticipated **flow of crowds** through the facility, catering for large number of visitors.
- The suitability of the technical infrastructure to deal with **demanding high-frequency** use. This relates to the *graphics capabilities* to render complex content, the *multi-touch infrastructure* to deal with large number of touches and the *computing and network infrastructures* to allow for distributed execution of applications across a cluster of computing nodes.

These overarching criteria, and the requirements derived from the initial set of "launch" applications (see below) informed the eventual design of the facility. The final physical layout was significantly influenced by two factors. First, the availability of thin-bezel, LCD multi-touch technology and second the commissioning of a dedicated server room on Level 6. These factors supported relatively narrow walls, maximising both floor space and the availability of high-resolution interactive surfaces.

Decisions regarding computing infrastructure, in particular the specification of the machines driving the displays, were based on requirements derived from the most graphically demanding applications (The Reef and the Physics Playroom, see section "Cube Software"). Decisions centred on whether to support "gaming" graphics cards vs. "workstation" cards, the number of available cores, clock speeds and memory architecture. Performance requirements were balanced with maintainability, durability and supply factors. A further consideration for machine specification was an attempt to future proof, to the degree possible, the Cube's computing infrastructure.

The network infrastructure was designed to cope with the most demanding use scenarios, not only in terms of planned applications, but also in terms of future use. An important decision was the separation of touch and application data into separate physical networks. The rationale behind this decision was to allow for low latency syncing of applications across computing nodes as well dealing with an anticipated large number of touch events without mutual interference.

CUBE SOFTWARE

The design process for the Cube infrastructure was accompanied by the development of twelve separate interactive applications. Five of these initial 'launch' applications made it through to the Cube's final 'go-live' date in January 2012. The five launch applications are:

The Physics Playroom: a large virtual 'physics' simulation where people can explore a range of basic physical principles including, colour, light, sound, energy and motion.

The Virtual Reef: an education-oriented simulation of the Great Barrier reef, that displays marine life at scale spanning across the two largest Cube zones.

CubIT: a large-scale multi-user presentation and collaboration platform, which allows QUT staff and students to upload their own media content, allowing for shared interaction on any of the Cube's multi-touch zones.

ECOS: an interactive game-like application which uses interactive data and illustrations to help people understand how everyday energy use impacts people in a "green" building.

Community Science Wall (Flood wall): an interactive application that allows the public to connect with Queensland stories and experiences as told through a large-scale interactive map and timeline. For the Cube's launch, Community Science showcased Flood Wall, a virtual and interactive storytelling of the 2011 Brisbane floods.

These five 'launch' applications have been specifically designed to create engaging interactive experiences that let users and casual visitors interact with the Cube without requiring instruction or domain knowledge. Applications share the available wall surfaces and are deployed based on demand and scheduling. We will describe each application in turn.

The Physics Playroom

Overview

Tapping into our innate desire to construct and to deconstruct, the Physics Playroom is a space where people can come together to build (and destroy) each other's creations in a fun and engaging physics simulation. With a tie into QUT's first-year physics and game programming courses, the physics playroom is a collaborative, interactive and exploratory space, designed to support thousands of simultaneous user interactions.

Objects in the Physics Playroom are characterized by real-world properties, such as mass and friction. Additionally, physical laws such as gravity define the overall state of the room. Users interact with the playroom's artefacts in ways that allow experimentation with various physical principles including classical mechanics, fluid dynamics, sound and light. The playroom is fundamentally interconnected. A user may throw a block from one end of the wall to the other, knocking over another person's 'building' in the process. This interconnectedness pervades the Playroom and mandates that the Physics Playroom must be coordinated, as well as distributed.

Design

Developed for zones 3 and 4, the playroom is designed to run as a single unified interactive 3D virtual world. End-user interaction is largely centred on the manipulation of three 'physical systems'. Newtonian mechanics allow users to stack and throw blocks around the room; fluid dynamics allows users to draw sand paintings; and audio signal processing explores the sonic world of amplitude and frequency modulation.

Figure 5. The Physics Playroom

In order to support a diverse community of users, the playroom focuses on direct end-user interaction while providing a visually appealing 'canvas'. The playroom was designed to provide a visceral end-user experience by encouraging direct, fluid interactions with objects in the scene, as well as with other users in the ‘real world’.

Beyond its initial ‘visceral’ experience, the playroom encourages deeper reflection and understanding by providing a set of didactic resources designed to encourage a greater understanding of the mathematics that is central to the playrooms modelling and simulations. From simple linear motion, through to more complex dynamic systems, content is specifically designed to target a range of ages and skill levels.

The playroom is designed to handle up to fifty simultaneous users without suffering any noticeable performance degradation. At just over four users per 55inch panel this represents about the maximum possible physical utilisation of the space. Maintaining a 'silky smooth', visceral end-user experience, with fifty high school students intent on 'breaking' the system, was a primary design goal for the project.

Interactions & Engagement

The playrooms computational simulations are designed to respect a range of simple end-user touch interactions in-line with the playrooms goal of educating users about the underlying mathematical principles – throwing blocks, spinning cogs, flicking light switches, siphoning sand, turning dials, etc..

Motion: The playroom is full of wooden blocks of a variety of sizes and shapes. All of these wooden blocks react to human touch (dragging, flicking, swiping etc.) and to collisions between themselves and the environment. Blocks move in three dimensions and can be stacked vertically. Blocks behave according to standard classical mechanics.

Gravity: All of the blocks in the Playroom are affected by Gravity. Gravity can be changed to reflect the particular characteristics of each Planet in the Solar system. Changing gravity 'transports' the playroom to the surface of the chosen Planet. As well as planetary Gravity there is also a Zero Gravity option.

Fluid Dynamics: The Playroom incorporates two large 'paintings' that rotate through a selection of famous portraits. These portraits are drawn using a million particles, where each particle is coloured to match a particular image pixel (or group of pixels) location. The movement of these particles is then driven by a fluid dynamics simulation which is in turn driven by human touch. The effect is to produce a type of sand painting, allowing users to perturb the 'state' of the portrait.

Sound: A series of sinusoidal oscillators are used to introduce users to the physics of Sound. Wave concepts including wavelength, amplitude and frequency are introduced in an intuitive and direct manner. These oscillators are chained to introduce the principles of amplitude and frequency modulation. The general public can manipulate the amplitude and frequency of each oscillator using a selection of dials. Sound transformations are heard in real-time as well as being represented as a graphical audio signal.

Technology & Software

The physics playroom is a bespoke system developed using, Extempore, a new research programming language and environment being developed by Andrew Sorensen at QUT. Featuring a highly interactive development environment, and high performance programming language, Extempore provided a powerful environment for developing the Physics Playroom 'on-the-fly' [7].

The physics playroom was designed, in part, to showcase the Cube's technology, and as such was designed to utilise the available hardware to help demonstrate the Cube's potential to staff, students and visitors to the University. To meet this requirement the playroom was designed to run an interactive virtual world, at 60 frames per second, utilising the full 12960x3000 pixel resolution of the multiple-displays, while supporting 50,000 user events per second, across a distributed render farm of seven nodes. Most importantly the playroom is expected to run for a minimum of twelve hours per day, 365 days per year, and is a central feature in high profile University visits.

By leveraging Extempore's concurrency, distribution and coordination features, the playroom project was able to deliver a significant bespoke real-time distributed computational simulation platform.

CubIT

Overview & Purpose

The Cube flexible interaction and collaboration framework (CubIT) aims to make the Cube infrastructure accessible to

QUT’s staff and students by providing them with an easy and intuitive way to upload and showcase their own content. CubIT allows users to flexibly and spontaneously present and share their work, and collaborate with others on the Cube’s large multi-touch displays. While the design of CubIT shares similarities with some existing approaches of large-scale display interaction it differs from previous systems in a number of ways. Specifically, it is the combination of the following features that make the system unique:

- CubIT was specifically designed as a true multi-user system that allows multiple users to upload, and simultaneously interact with and share their own content on a large multi-touch enabled canvas.
- Users can log in and authenticate themselves at the display walls using RFID authentication. Each user has a workspace containing their content, which can be freely moved around the canvas. Each workspace contains interface mechanisms allowing users to manage their media content on the shared canvas.
- Interface mechanisms for uploading and sharing content have been designed to be as simple as possible, allowing for simple drag and drop interaction.
- CubIT has multiple user interfaces that serve different purposes. A *web-based interface* allows users to upload and edit content and change their user details. A *mobile interface* (phone or tablet) allows users to dynamically push content to the wall displays allowing them to interact remotely with the screen. And lastly, the *multi-touch interface* at the Cube wall(s) allows users to display, interact with and share user-generated content.
- CubIT uses an external object store (Redis) to dynamically maintain the state of the application across different interfaces. Changes to content (e.g. deletion of images) on each of the interfaces are dynamically represented on all of the other interfaces.

Interactions & Engagement

The CubIT multi-touch interface consists of a selection of bespoke widgets. The *user workspace handle* represents the user on the shared wall surface and gives them access to their own content, represented by thumbnails in a scrollable window. Content is displayed via *image, video and text widgets* which can be freely rotated, translated and scaled and provide controls appropriate to the content type. A *presentation widget* allows users to display stacks of images, videos and notes in a more convenient manner.

The multi-touch interface supports collaboration and sharing. Users can use their workspaces to freely share content. In order to copy content items between accounts users drag thumbnail representations of images, videos, notes or presentations into another user’s workspace.

Figure 6. CubIT in use

Technology & Software

The CubIT system consists of two main components. The *CubIT backend* (Ruby on Rails) manages all aspect related to content and user management, including content upload (images, videos, notes), the creation of custom presentations, content delivery and maintaining workspace, session and authentication state. The backend further provides two web interfaces, one aimed at users, to log into their system and maintain their content, and the other aimed at moderators, allowing them to assess content and suspend or delete user accounts if necessary. The *CubIT frontend* (Python/Kivy) manages touch interactions, widget display and syncing the canvas across a set of multi-touch screens and a cluster of display nodes using a peer-to-peer approach. Lastly, the *iPhone/iPad interface* communicates with both frontend and backend to allow for easy content access and upload.

CubIT requires a range of further services for its operation, including a TUIO multiplexer to merge separate input streams, a Redis server used for content and status synchronisation, and an RFID service that relays information from the Cube RFID readers and translates them into login information.

System use

CubIT was deployed in January 2013 and has currently over 450 registered users. The system is being used in a range of different ways including teaching (lectures, student presentations & exhibitions), event & conferences, visitors & demos and the engagement with the general public.

Community science wall

Overview & Purpose

The Community science wall is a large scale, multi-user map-based application that can be used to display user-generated as well as curated content in geographical context. It was designed for zones 3 & 4 of the Cube. The system was initially developed in response to the 2011 Brisbane and Queensland floods as a site for community engagement where people could learn more about the events of the floods as well as share their stories, pictures and videos of what happened.

As Figure 7 shows, the physical set-up of the wall consists of 12 multi-touch screens in portrait orientation on the lower portion of the wall and a projection display on the upper portion. The multi-touch displays allow users to interact with digital maps of the Brisbane area and informational layers of user-contributed, and curated content. The upper portion displays a looping multimedia presentation about the floods.

On the left edge of each multi-touch panel are a series of buttons users can press to change the rendering style of the maps and show additional layers of information about the floods. The maps can be displayed in 'satellite', 'terrain' and 'map' modes, allowing the user to navigate with the style of map that best suits their needs (e.g. whether finding a particular street, or looking for an area of parkland). The additional layers of information that can be shown include: an overlay of the 2011 peak flood level, shown as a blue outline on the map; user contributed 'flood stories' consisting of geo-tagged photos and videos; and a collection of planning concepts for the city from QUT Urban Design masters students.

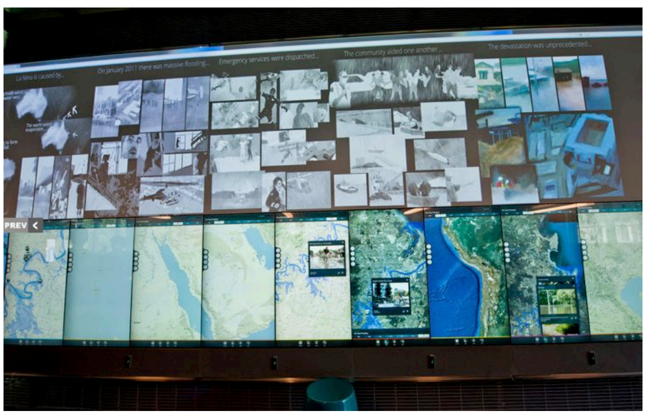

Figure 7. The Community Science / Flood Wall

Interactions / Engagement

The application is intended to appeal to a range of different visitor groups, from casual visitors, such as tourists and families to school groups and students engaged in directed learning activities. Because of this diversity in user groups, the design of the system needed to account for a range of levels of motivation and familiarity with the technology of multi-touch displays and interactive maps. This manifested in the design in several ways. To help users understand that the system can detect and respond to their touches, on-screen feedback is given by drawing small circle around each registered finger. This also gives more experienced users confidence that the system is responding correctly.

Another design strategy for supporting multiple levels of familiarity in the user population was to offer multiple ways of accessing the functionality. For example, the interactive maps can be zoomed in and out both through pinch gestures and with soft-buttons displayed on the screen. Similarly, locations can be navigated to both by dragging the map left and right and by accessing an alphabetically organized list of suburb names.

Providing quick access to specific suburbs was also an important design decision for supporting one of the common behaviours observed in users of the system – namely the desire of users to first navigate to the location of their home and zoom in on that. Independent of any desire to learn about floods, this seems to reflect a simple delight in seeing your own house presented in the interface, however, seeing the information from the flood wall in relation your own locality also provides a useful point of reference for people in understanding other information that is presented in the interface (such as flood levels) as well as allowing them to relate their own experiences of how the flood affected them.

The system is also designed to cater to a range of numbers of users from single people to groups of people interacting at different points along the wall. In the case of single users, the maps are synced across multiple screens and computers so that the system gives the appearance of a single large map. However, when additional users also touch the screen, the map 'breaks' so that each person has control of a separate set of screens. As more people interact with the wall, this breaking behaviour continues until each screen is showing its own independent map. After a period of inactivity, these 'breaks' heal again and the maps join back up into one.

Technology / Software

The system is technically notable because it is implemented as a web-application, the client-component of which runs in a standard web-browser (Google Chrome). Each of the 12 multi-touch monitors actually displays a single browser window running in full-screen mode. These are shared across six computers, each running two monitors. The user-interface of the system is implemented using the Google maps API, JavaScript and standard HTML/CSS.

TUIO touch events from the multi-touch displays are sent to a TUIO client running on the client machines, to drive the interactions on individual clients. Shared state is maintained across clients (for the syncing of screens, map mode, layer visibility) via a node.js server.

The system also consists of a web-application written in ASP.Net, which serves the client pages, layer contents and provides the suburb finding service. KML data can also be served to show other information such as the extent of flooding. The KML system is extensible for additional layers of information, which can be easily be added to the interface through the addition of a single line of code. For submission of user-generated content, there is also an inter-net facing web-form where users can upload content, which is then moderated by an administrator before being displayed on the wall.

Use & Lessons learnt
While the ‘breaking’ behaviour of the maps does to some extend solve the problem of fluidly supporting multiple numbers of users, it remains to be seen whether this is an optimal solution. One problem is that it is difficult to communicate the behaviour to users in an intuitive way. It is also difficult to encode a set of rules for the breaking behaviour that take account of edge-cases such as when a user is attempting to pinch zoom with two hands on adjacent screens (which the system would interpret as separate users).

The Virtual Reef

Interaction Design
The Virtual Reef is a simulated life-sized marine ecosystem expanding across two levels of the Cube (Zone 1 & 2). Multi-touch technology enables the user to manipulate, explore and interact with the reef world. The Virtual Reef was designed as an immersive virtual experience that engages and educates through sheer scale and novel interactions with a virtualised underwater environment.

The system’s interface was designed in an iterative manner involving various groups of High School students. Direct observation, video logs, speak aloud protocols and follow up questionnaires were used. An early prototype allowed users to move the camera into the 3D space. However, this was judged to be vertigo inducing for most participants. Thus, the display screens offer users a two dimensional wall as a barrier into an engaging world. There is point and hold, there is swipe to attract the fish; there are web base panels that appear and give you information on each of the species and more. The “glass wall of the aquarium” is the physical reference, but it is also a digital display surface.

Figure 8. The Virtual Reef

Engagement
The design of the Virtual Reef aims to achieve a number of Educational and Aesthetic objectives. Foremost, was the desire to engage with schools from grades 5 to 12 through interactive educational material and entertaining experiences. The animated behaviours developed for each species have been the result of thorough research, focusing on the animal’s role within the ecosystem. The user has the opportunity to not only observe and absorb the basic functionality of an ecosystem, but they are also able to witness behaviours observed by marine scientists themselves. The direct stimuli targeting curriculum-based learning objectives are available in the interface information pop-ups. These pop-ups are activated when any species is touched and also via a section at the bottom of each touch screen. Activity sheets are also provided to guide students and teachers through the relevant subjects, encouraging an exploration of the scene with an emphasis on necessary and relevant features.

A QR code containing links to more in-depth information is available to the users in the UI, mobilising educational material for further research.

Technology / Software
The Virtual Reef was developed using the Torque3D game engine, which supports 3D models in DTS, DIF and Collada formats. The 'reef' consists of modified Torque C++ engine code, torque script, various batch files for syncing the reef across the clients and server, and various art assets. The 'art' consists of AI fish meshes, and their behaviours, pre-animated sequences like whales, sharks, etc. It further contains the mesh for the reef landscape, dendra, sea grass and other animated and fixed flora. There are a range of sound files for the GUI, some animated sequences for the UI and general background sounds.

The fish AI is controlled in both torque script and cpp files. These files separately cover basic fish movement, the effect and range of influence that each fish has for avoiding the reef and other fish, and pre-set trigger data that ensures fish do not collide. There are also 'schooling' influence files relating to specific fish, and all fish have influence parameters for the screen touch events.

All the web pages are controlled through the Awesomium plugin. There is C++ code to control the view data, textures and fish models. Touch events are initialised for each client, then sent to the server. The individual clients handle their view of the reef, with the server managing the overall list of events and responses.

ECOS
The ECOS project is a playful interface that uses real-time weather data to simulate how a five-star energy building, in this case the new Science and Engineering Centre (SEC), operates in six different climate zones.

Design
The ECOS Project presents a simplified interface to the very complex domain of thermodynamic and climate modelling. From a mathematical perspective, the simulation can be divided into two models, which interact and compete for balance – the comfort of ECOS’ virtual denizens and the ecological and environmental health of the virtual world.

The final version of the ECOS Project (Figure 9) resolved to communicate one basic concern, that building energy consumption is determined by local climate conditions and as such, there are few climates on the globe that make it possible for buildings to be sustainable. The ideal user interaction order would be:

- Once the user activates a building, it detaches from the ECO-Source World as an ECOSphere displaying live weather conditions based on its original position on the ECO-Source World.
- The users must successfully negotiate the use of green technologies (solar, wind and gas) of each building to reduce the dependency on main (coal based) energy sources.
- The user must keep the internal temperature at a comfortable setting to maintain the productivity of the workers.
- If they are successful at balancing the use of green technologies with comfortable internal conditions of a single building, the ECOSphere will continue to appear clear and may even produce excess green energy.

Figure 9. Students interacting with the ECOS project on the Cube installation

Use & Lessons learnt

The user's choices can impact the amount of energy consumed by their building, the amount of energy generated by the building, and the comfort of the people inside the building. These choices are the internal temperature of the building, the internal humidity of the building, and the allocation of alternate energy sources. Users can allocate the distribution alternative energy sources between wind, solar, and on-site gas generation. The external weather conditions of each location affect the way each choice effects the energy usage and generation of the building.

Users receive feedback on two levels; individual feedback of a single ECOSphere, and the comparative feedback of considering multiple spheres together. We have observed that the feedback to the user is most useful in a comparative sense. For example, a user might modify and focus on the ECOSphere representing Beijing. Individually, the feedback they see represents that sphere's ability to generate energy from each allocated source, amount of energy used, and comfort of the people inside, based on temperature and humidity. But when they add another sphere to their session, for example Rio de Janeiro, the comparisons between the disparate climates do more to clarify the relationship between climates and energy efficiency than observing and modifying a single sphere.

LESSONS LEARNT

The Cube has been in operation since the start of 2013. Both the design process and the operations phase have provided us with a number of insights outlined below.

Interaction paradigms

The initial set of five applications support a range of interaction mechanisms which be categorised across a number of dimensions. First, *single-user interaction, shared experiences* (The Reef, ECOS, Community Science Wall) allow multiple users to simultaneously use an application and interact with its content through single-user interface elements. However, while the interaction is largely single-user focussed the experience can be shared by several users. *Second, multi-user interaction shared experiences* (The Physics playroom) allow multiple users to simultaneously engage with the same interface components (e.g. by throwing blocks to each other, collaborative building etc..). The first two categories both refer to users interacting with existing content directly on a multi-touch wall. The third category differs from this approach in that it focuses on user-generated content rather than a pre-designed content. *Multi-user interaction distributed shared experiences* (CubIT) allow multiple simultaneous users to interact with and contribute content through different user interfaces (web, mobile & multi-touch). The multi-touch interaction consists of interactions with both single and multi-user interface elements (e.g. share scaling, rotation, sharing across workspaces).

Each of the initial set of applications was built for a particular purpose, which is reflected in their unique interaction paradigms. The Reef, ECOS, and Community Science Wall all support shared interaction with a specific topic matter. The Physics playroom also does this, but adds playful multi-user interaction elements that allow users to jointly manipulate the digital environment. Lastly, CubIT differs in focus and predominantly supports the presentation and sharing of user-generated content, thus implementing a different user interface paradigm.

Multi-touch event handling

Due to the different requirements and interface paradigms outlined above, applications handle touch events in different ways. Most applications handle TUIO events locally (either per screen or per pair of screens attached to a computing node). By handling TUIO events 'locally' these applications are able to provide more responsive user feedback at the cost of increased application complexity.

By contrast, CubIT cannot solely rely on local touch information since it has interface elements (e.g. images, videos) can that can stretch across computing node boundaries requiring the multiplexing of multiple TUIO streams into a unified stream. The original CubIT design used a simple client-server approach to create a single large canvas. This lead to a unified, centralised, approach to event management. TUIO streams from each display were merged into a single multiplexed TUIO stream mapping the whole canvas. Due to latency and performance considerations this design was later changed to a peer-to-peer architecture where each display node would synchronise content with its neighbouring nodes. With regards to TUIO this meant that each node had a TUIO viewport that covered the two multi-touch displays attached to it as well as the two displays to the left and right, presented as a multiplexed stream. This effectively allowed users to scale and rotate objects across node and display boundaries.

Synchronisation

Due to the large number of displays and computing nodes in the Cube, the execution and synchronisation of content across multiple computing nodes was of critical importance. Existing solutions that support distributed rendering such as Chromium [2], VRJuggler and SAGE[3] fell short of our requirements. In large part this was due to the fact that these tools are designed to adapt existing, single user, environments. Our focus on bespoke, multi-user environments gave us the opportunity to build customized solutions.

This proved to be significant as many of the performance gains to be made in the Cube's clustered environment relate to end-user interaction and modelling rather than rendering per se. Indeed, distributed rendering was generally less problematic for Cube projects than managing distributed multi-user interaction.

Support structure and future projects

When the project transitioned into production stage at the start of 2013 the original development facility was turned into a production studio dedicated to supporting future projects. However, due to the demanding nature of the Cube environment the entry requirements for such projects were quite high (multi-node support, high reliability, etc.). While the production studio focused on preparing applications for this environment it was not designed to cater for a broader engagement with the QUT research community and a more exploratory use of the available technology. As a result QUT's Institute for Future Environments commissioned a "skunk works" facility in May 2013 that caters for a broader range of users and potential developers. Its activities range from teaching basic multi-touch development skills to supporting a wide range of interactive technologies needed by research projects.

[3] http://vrjuggler.org/ & http://sagecommons.org/

CONCLUSIONS

In this paper we discussed the Cube, a large-scale interactive learning and engagement spaces situated at QUT. The Cube is a demanding real-world environment that pushes the boundaries with regards to screen dimensions (both interactive and projected), resolution, number of distributed computing nodes and support for very large number of users. We have outlined the design and implementation of both the hardware and accompanying software projects in the hope that the lessons learnt can provide valuable for institutions planning to build similar facilities. The Cube is an active platform, that supports a wide range of further development and research projects. A host of new projects are underway to make use of this exciting large-scale infrastructure.

ACKNOWLEDGMENTS

We would like to thank the Cube project manager Gavin Winter for the design and documentation of the technical specification for the Cube and Bryce Christensen and Sherwin Huang for their contributions in the design and development of the Community Science wall. We further like to thank all the development teams of the applications presented here.

REFERENCES

1. Anshus, O., Stødle, D., Hagen, T., Fjukstad, B., Bjørndalen, J., Bongo, L., Liu, Y., and Tiede, L. NineYears of the Tromsø Display Wall. Available from: http://www.powerwall.mdx.ac.uk/papers/POWERWALL-Anshus.pdf. (2013)
2. Humphreys, G., Houston, M., Ng, R., Frank, R., Ahern, S., Kirchner, P.D., and Klosowski, J. Chromium: a stream-processing framework for interactive rendering on clusters. *ACM Trans. Graph. 21*, 3 (2002.), 693-702.
3. Kim, H. and Snow, S. Collaboration on a large-scale, multi-touch display: asynchronous interaction and multiple-input use. In *Proc. CSCW 2013 companion*, ACM Press (2013), 165-168.
4. Morris, M. Web on the wall: insights from a multimodal interaction elicitation study. In *Proc. ITS 2012*, ACM Press (2012), 95-104.
5. Peltonen, P., Kurvinen, E., Salovaara, A., Jacucci, G., Ilmonen, T., Evans, J., Oulasvirta, A., and Saarikko, P. It's Mine, Don't Touch!: interactions at a large multi-touch display in a city centre. In *Proc. CHI 2008*. ACM Press (2008), 1285-1294.
6. Qin, Y., Yu, C., Liu, J., Wang, Y., Shi, Y., Su, Z., and Shi, Y., uTable: a seamlessly tiled, very large interactive tabletop system. In *Proc. ITS 2011*, ACM Press (2011), 244-245.
7. Sorensen, A. and Gardner, H., Programming with time: cyber-physical programming with impromptu. In *Proc. OOPSLA 2010*, ACM Press (2010), 822-834