https://doi.org/10.2352/ISSN.2470-1173.2017.5.SD&A-105

Designing a Cloud-based 3D Visualization Engine for Smart Cities

Nicolas Holliman, Mark Turner, Stephen Dowsland, Richard Cloete, Thomas Picton
Digital Institute, School of Computing Science, Newcastle University, UK

Abstract

Urban data is being collected in increasing quantities as smart cities around the globe seek to understand and improve their operations and plan their growth. This trend is set to continue as UNICEF predicts that 75% of people will live in cities by the end of the 21st century. One of the aims of Smart City initiatives is to improve inclusivity by communicating more about the city to citizens and organizational stakeholders. The large amounts of data and computational cost of the calculations required to do this mean that cloud-based analytics and visualization is an attractive option as it can deliver results to virtually any client device.

In this article, we describe the design and implementation of the Urban Insight Cloud Engine (UICE) a pilot cloud-based 3D visualization system for Smart Cities that has been created using open software and data sets. This delivers a continuous live view of data collected by Urban Observatory sensors in the city of Newcastle-upon-Tyne, UK. We precede this with a discussion of our experience designing the facilities that exist, are being built, or are being planned to be built, to support our visualization research and production at Newcastle University.

Context

Like many cities, Newcastle-upon-Tyne is collecting measurements from new sources of data about the urban environment, this includes Internet of Things (IoT) and other sensing devices sited both internally and externally around the city. This data is being made publicly available via the Urban Observatory web portal [1]. However, accessing raw data on its own is not sufficient to meaningfully show the data being collected and to allow a range of users to explore what this data may mean for them.

Here we summarize our approach to visualizing this data. We describe the hardware facilities we are, and are planning to install, for visualization. In addition, we describe the software tools we are building to explore novel approaches to urban data visualization and give our first results from these tools.

Defining Visualization

The goal of data visualization is to support humans in understanding and acting upon data values, or analytics derived from data values. Fundamentally the process creates ideas in the mind using images, we believe there are three main roles for visualization:

Communication: presenting, describing and explaining ideas, examples of which might be education and training.
Confirmation: testing hypothesis, monitoring information streams, examples of which might be scientific results and internet of things streaming data.
Exploration: exploring and analyzing information seeking insight and answers to questions, for example policy related social data or the results of an engineering simulation.

All three roles can be supported interactively using computational methods, but they all originated before modern computing tools were widely available. The difference modern computing and display hardware makes is the speed at which large volumes of data can be acquired, processed and presented. There remains substantial research challenge in developing visualization methods that utilize computational tools to fully support these roles.

Experiences Designing Visualization Facilities

We have experience over twenty years of designing and running 3D visualization laboratories for a variety of purposes. Including, dedicated research facilities for novel hardware development and facilities to support visualization production and presentation by end-users in industrial (Sharp Laboratories of Europe) and academic (Durham, York, Newcastle) environments. Two of the key features in these facilities have emerged as being hardware flexibility and ease of use by end-users. Hardware flexibility often comes down to very specific requirements for interfacing to a range of devices. In this the WinTel PC platform has become dominant and is typically the first platform to be supported by new display and interaction devices. Ease of use for the end user we return to below.

The Core Decision Theatre

This facility we designed for Newcastle University's Digital Institute to support research and teaching in visualization. It also has a key role as a presentation system for external visitors, including businesses and the public as shown in Figure 1.
The main display hardware is a rear-projected display providing wide-screen presentation capability. In addition, there are moveable flat panel screens and the ability to setup the room with individual high performance teaching workstations, as shown in Figure 2.

Figure 1. *The Core Decision Theatre in use for a public presentation in S3D. Photograph courtesy of @Rich_K, http://richardkenworthy.com*

Figure 2. *The Core Decision Theatre setup for small group visualization teaching. Each workstation has a 4K monitor.*

The large screen presents excellent opportunities for displaying comparative versions of the same information and for presenting wide screen images at large scale. Additionally, the rear projection provides presenters with the ability to stand directly in front of the screen surface. One disadvantage of the screen is that it uses a pair of blended projectors, in many situations the blend area overlap is not visible, however, it does not support high contrast content well with visible blending artefacts when black regions of an image cross the image overlap. This can be masked in less demanding images by using a raised black level, or white/gray image background.

Distilled Experience

In designing several large-scale visualization facilities, we have concluded that there are some basic needs users of these facilities have. We summarize these requirements in five categories:

Ease of use: this equates to how easily and how quickly users can access the facilities and how easy it is to install and run the software tools they need. A simple test of a facility is can it run ubiquitous software, such as Microsoft PowerPoint, alongside any specialist tools needed.

Easy to maintain: hardware multi-projector systems, whether for stereo or for blended large screens, can be difficult to interface with and to keep aligned. If feasible we would always recommend using single projector display systems. End users will also have fewer difficulties connecting external equipment to these systems.

Standard aspect ratio: content production should drive major equipment decisions. Having non-standard aspect ratio screens makes content production and display a compromise. In all our current and planned systems we aim to standardize on 16:9 screen aspect ratios. Content authored on any screen can then be played back without borders or incorrect scaling on any other screen.

Bright, high resolution, high frame rate, 3D: Dim and low contrast displays with low refresh rate are painful to use for long periods of time. When making design tradeoffs a smaller image size with higher image quality should support users better than a larger image size with compromises. Brightness calculations can however be difficult as they must take into account ambient light as well as screen size.

Simple interaction devices: like all aspects of these facilities the temptation is to go for first impressions rather than long term usability. In decision theatres with a need to support extended interaction times the mouse is ubiquitous for a reason, it matches human fine motor control well, it is comfortable to use, and in versions like the Gyration™ mouse can be used in free space as easily as on a desk. For gesture tracking and wand-based interaction in free space users can tire quickly in longer visualization meetings, even though the initial impression of immersion is strong.

The above don't represent a recipe that will work for every visualization theatre, but to support a wide range of users our experience is they are all important points to consider at the design stage.

The Immersive Decision Theatre

Our latest facility implementation is the Immersive Decision Theatre at Newcastle University. The aim here is to create a room that provides an immersive display experience for groups of up to ten in big data applications. It also has the requirement that it is open to the building atrium and needs to be bright enough to be viewable in indirect daylight. The room is illustrated in Figure 3.

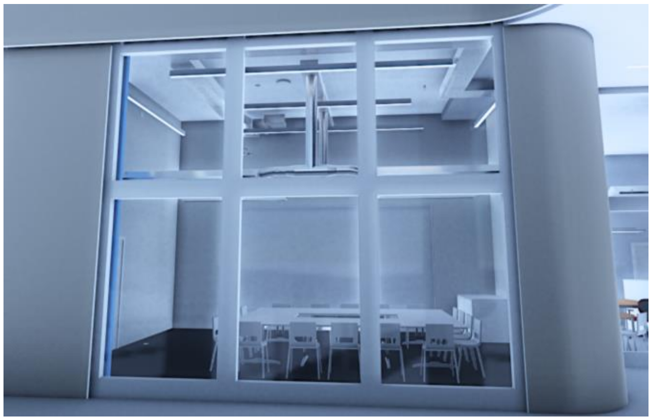

Figure 3. *Artists impression of the Immersive Decision Theatre in the new Urban Sciences Building at Newcastle University.*

The main screen display technology is a rear projected image onto a 16:9 aspect ratio screen with a physical width of just of 4m. Most users of the IDT will have at least a 60-degree horizontal field of view, which is the minimum specification for immersion in IMAX ™ theatres, as illustrated in Figure 4. In addition, we plan to mount a large touchscreen to the side wall of the theatre providing interaction and a second large display surface.

We will support some local graphics and compute resource and standard software applications for visualization and data presentation. But, as described below, we will use cloud computing for our major visualization research projects. A further key requirement is to be able to wirelessly connect portable devices to the main screen, this will be done directly via wireless display links in the room or indirectly via cloud-based applications.

Figure 4. *Artists impression of the Immersive Decision Theatre showing the user's view of the room.*

The Immersive Decision Theatre will be adjacent to a larger event space which is planned to support large group teaching and presentations with the same level of visual performance for audiences of up to one hundred.

The Team Decision Theatre

In the first stages of planning is a new type of space for visualization at Newcastle. The National Innovation Centre is a new building being designed for the Science Central site, its aim is to be a beacon for innovation work with businesses in the UK in the areas of Aging and Data Science. One way we plan to make the work in the building visible is to have the building represent externally what is happening inside it. A sketch of one concept that includes a large external screen is shown in Figure 5.

Figure 5*. Artists impression of one concept for the new NIC building showing an external screen for presenting research externally. Image courtesy of GSS Architecture Ltd.*

Within this building, the aim is to enable teams of innovators and technical specialists to come together and work on problems, while also including support for profile raising activities such as competitive data hackathons. To support these activities, we have designed a space to allow multiple teams to work in the same room on individual displays and then to bring their results back to the whole group via a central display system. A sketch of how this room might be laid out is shown in Figure 6.

With the use of some of the rear projection space for furniture storage we aim to be able to easily switch the use of the room from team based visualization for three teams of four to a group presentation format hosting up to about twenty viewers.

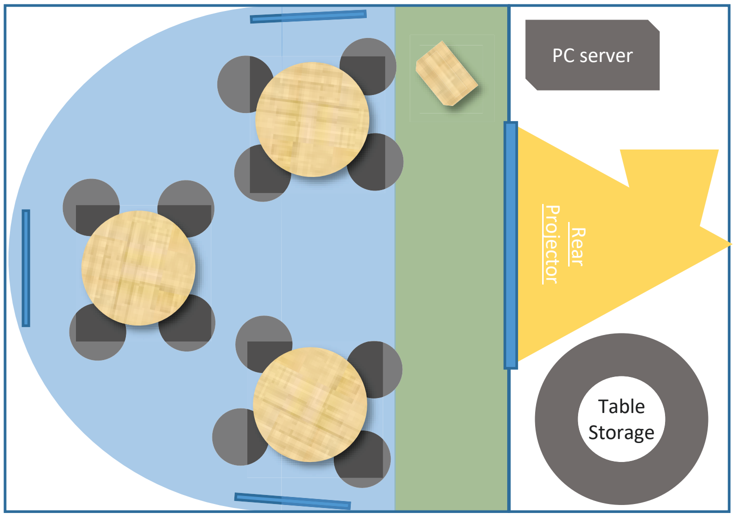

Figure 6. *Concept sketch for the Team Decision Theatre. Team screens are situated around the walls of the room and the main screen for group presentation is provided by a rear projection wall.*

In all the visualization theatres, we have designed a number of factors have raised questions from stakeholders. One of the most often repeated questions is why do we need space for rear projection, can't we instead use flat panel or blended projectors to reduce the projection room depth. Clearly this is a design decision based on space cost and display system cost as well as functional requirements. However, in our experience the walk-up nature and seamless image provided by rear projection works best in giving a high-quality feel to presentations.

Finally, in designing a visualization facility there are issues of user comfort that are not visualization technology related but are well worth some consideration. The air conditioning of the room is important for user comfort, particularly when larger groups need to stay awake in a darkened room during long presentations after lunch. In long development and testing sessions easy access to refreshments in the room is a continuing balance between the potential for equipment damage and the needs of the users to be comfortable. The flexibility, and design, of the furniture in the room is important. An ability to move and reconfigure tables and the ability for the chairs to work in standalone mode or with tables all help in creating a range of productive of working environments.

Cloud-based Visualization

The cloud presents new opportunities for visualization tools, it allows the use of scalable remote resources to support advanced techniques for computing with large data sets. As is demonstrated by commercial car sales software (Zerolight Ltd.) and GaaS gaming systems (NVIDIA), network latency is not a barrier to demanding real time computer graphics experiences in the cloud. The cloud also brings the benefits of any-device compatibility, click-to-run with no install time, and no need to continuously update your hardware or your own software patches.

For Newcastle's distributed visualization facilities, in the UK and overseas, the cloud has the potential to support high-quality visualization without replicating expensive local computing resource.

The Newcastle Urban Observatory

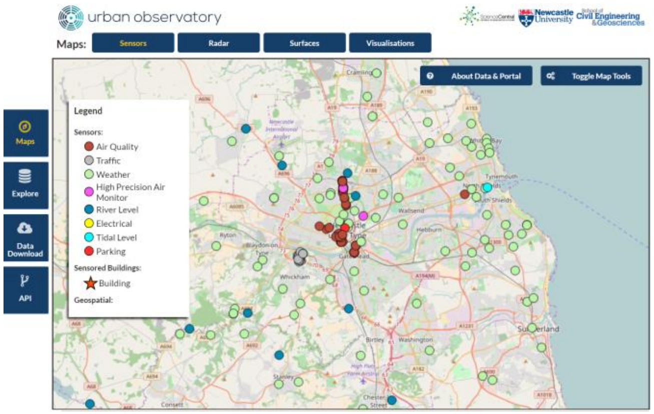

Figure 7. *The web page portal to the Newcastle Urban Observatory, providing interactive access to individual data sensors and their values as well as a scripting API interface.*

The Newcastle Urban Observatory [1] provides the largest set of publicly available real time urban data in the UK. It is updated with over one hundred thousand new readings a day.

It contains information from over fifty data types including environment, traffic, weather and power data, from sensors within the city and the local region. The data portal shown in Figure 7 provides interactive access to individual sensor readings and there is an API for direct script access to the underlying database.

Our interest here is in visualizing this data in 3D, using advanced rendering tools and novel visualization techniques that present the data and its analytics in new ways. We aim to use the power of cloud computing to help improve the inclusivity, the clarity and the scale of our visualizations. We detail these aims more below:

Inclusivity: by using cloud-computing to host our visualization systems we will allow the end-user to access them on any platform that can display images via a web interface. This removes the need for end-users to have expensive graphics hardware and should improve battery life for portable devices.

Clarity*:* by using cloud-computing we can chose to use more advanced graphics techniques and combine these with new psychophysically informed visualization designs to produce better quality visualizations of data and analytics results.

Scale*:* the ability to scale cloud-computing resource on-demand allows us to use much larger data sets than we might on personal devices. This enables both the 3D context model, e.g. a 3D map, and the data glyphs to be detailed and of large extent.

Image-based Visualization Architectures

This approach to visualization using cloud IaaS hardware splits the visualization pipeline so that the majority of the computation is undertaken cloud-side. Our objective in choosing this is to create a scalable system that can grow with user demands and problem size, and is not limited by client-side performance [2].

As shown in Figure 8 we place most of the computation associated with the visualization in the cloud, including the rendering. The cloud-side functions include reading raw data, filtering the raw data and mapping it to geometry before being passed to the renderer. The result is that there are no compute intensive modules running on the client.

Because interaction events need to travel to the cloud and the results back to the client there is the potential for noticeable latency to be added to every event. This may be an issue in hard real-time applications, although even there it is clear that cloud-based gaming is achievable. However in our current visualization applications, we don't have the same hard real-time requirements as in gaming and our experience to date is that this network latency is manageable.

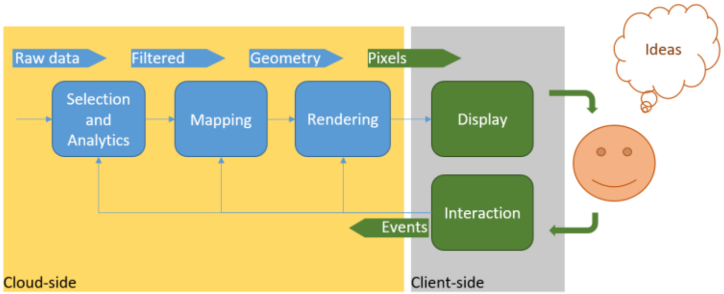

Figure 8. *An image-based architecture for visualization in the cloud. The division between client-side and cloud-side modules is chosen so that the client gathers interaction events, sends them to the cloud and receives back the pixels in the visualization image.*

The Urban Insight Cloud Engine

We designed the Urban Insight Cloud Engine (UICE) [3] to evaluate the system architecture ideas described above on cloud-based visualization. UICE takes a live feed from the Urban Observatory containing among other data both temperature and air quality (NO_2). It then renders a live view of these measures which is accessible via a standard web page on any system supporting a web browser.

Data Filtering

The live data feed from the Urban Observatory has a number of practical issues. Real-time sensing does not always return results, delays can occur at any point in the collection process at the sensor, in the wireless networks it connects to, in the intermediate processing of the sensor data and in the database it is stored in. As a result we can't rely on being able to request a single current reading and to have that reading instantaneously returned to us.

Instead we request data for the most recent one hour time interval and present an average value in the visualization. In the filtering stage we also restrict the range of the data requested to the geographic area we are interested in and, because at the moment Urban Observatory sensor locations don't include a height value, we add height above sea-level so that we can integrate them easily in our 3D visualization.

Mapping Data to Geometry

Our aim is to produce visualizations which show data within its 3D or 4D context, in this case located in a 3D representation of the city of Newcastle-upon-Tyne. To do this we need both a representation of the city and a compatible representation of the data values.

In our first version of UICE we use a topographic representation of the city, which excludes surface details such as buildings and trees, and consists of a 3D polygon mesh representing the underlying terrain. This was taken from the Environment Agency's 3D LIDAR scan of the whole of the United Kingdom. To provide further detail on this we added a texture map of the road network which was rendered from the Ordnance Survey's 2D vector map of the region. Additionally, we annotated this with street names using 3D text and model key buildings as height extruded geometry, see Figure 9.

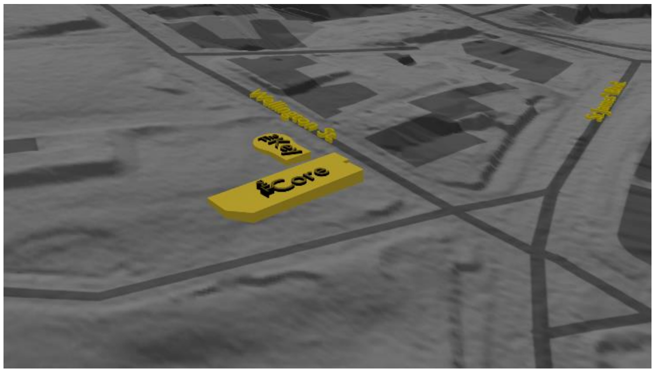

Figure 9. *The underlying 3D city model used in UICE rendering, this includes the 3D terrain mesh, road map texture, 3D road names and 3D representations of selected buildings.*

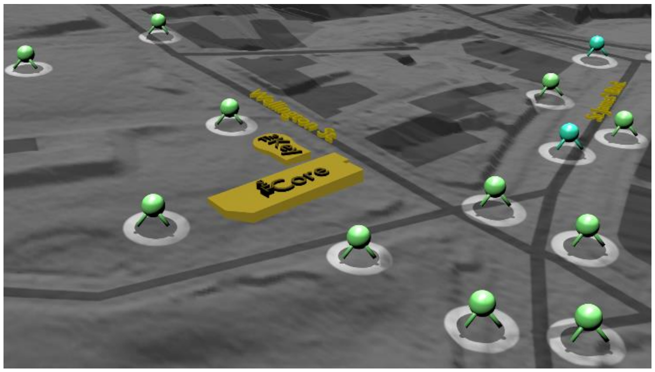

Figure 10. *An output from UICE showing the glyph design at each sensor location. Geometry is used to show the orientation of the sensors, colour to represent the values and a light/shadow combination local to the sensor to show the geographic location projected onto the 3D terrain, even from oblique and side-on camera angles.*

To create 3D representations of the sensors we created a 3D glyph that showed sensor value, geographic location and orientation, see Figure 10. This was designed to be placed over the topographic landscape described above, one glyph for each sensor, producing a simple point cloud representation of the data. Linking glyph colour to temperature was achieved using the UK Met Office temperature scale. In addition to the glyphs, we created a temperature scale widget in 3D, this currently shows the average temperature of all sensors in view over the last hour. Additional information is rendered as 2D floating text over the landscape.

Rendering

The modelling and rendering tool we chose to use was Blender, this provides a high-quality, physically-accurate rendering engine with advanced geometry modelling tools. It also allows us to replicate it on servers in the cloud without license restrictions.

We used Blender in offline UI mode to create models of the landscape and the fixed annotations. The process of converting the LIDAR scans and 2D maps from their original source data to a renderable 3D model was non-trivial, requiring careful coordinate system management and model scaling. LAStools was used in processing the Environment Agency LIDAR data into a polygon mesh, while QGIS allowed us to process the Ordnance Survey map data into an image we could use to texture map roads onto the mesh.

For the rendering stage of UICE we use Blender in command line mode and invoke it for each frame with the fixed 3D map model and a Python script that reads the live data from the Urban Observatory and creates the transient geometry for the sensor glyphs and the temperature scale. Our experience is that Blender running continuously in command line mode is unstable, so for each new frame we restart Blender and load in the base map model again. This adds a number of seconds overhead to each frame rendered.

Client-side functions

In the pilot version of the UICE system the client side is provided by a web page that supports a small range of interactions. The user can choose what type of data to display by selecting either temperature or air quality visualizations. Then for the chosen data type they can select a display format, from one of HD, 4K or S3D. They then have returned to them the latest rendering of the data values selected in the display format chosen.

A particular challenge is the difficulty of querying a web page to find its native display resolution. We therefore allow only the fixed size and aspect ratios above, supporting standard HD, UHD (4K) and HD above-below format 3D displays. On other size displays the image can be scaled with sometimes unpredictable results. However, the advantage of the chosen formats is that once displayed in full screen UICE can directly support consumer stereoscopic 3D TV and DLP projectors systems.

The web interface operates in a push mode, so that anyone viewing an image will automatically have it updated as the rendering system generates the next visualization from the latest data. Because we share a single image across all users we do not create a rendering bottleneck from multiple simultaneous rendering requests.

Performance Results

Our first implementation of the UICE system was implemented using Microsoft Azure cloud hardware. It used a single server for the cloud-side visualization computations and a second system to host the web pages that formed the user front-end. This generated outputs like those shown in Figure 11, updating roughly once a minute. Because a wider range of hardware choices were available in the USA both servers were sourced from the US East Coast pool.

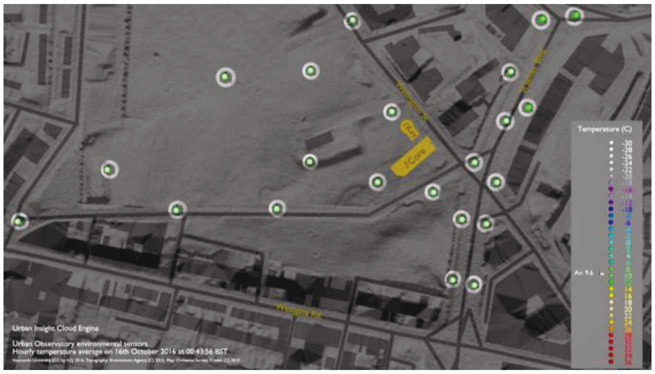

***Figure 11**. An output from UICE showing the temperature values from sensors around the Science Central site in Newcastle.*

The performance of Blender on a single computer is defined both by the hardware performance and the scene complexity. On the original server it took just under a minute to render the single view of the data shown in Figure 11.

In the second implementation of the system we parallelized the Blender implementation using a map-reduce framework. We scaled the system to use up to twenty compute nodes and this significantly improved the rendering speed we achieved, reducing it to about ten seconds per frame. However, as can be seen from Figure 12 the overheads of running Blender on a single image limit the number of nodes we can efficiently use to about eight.

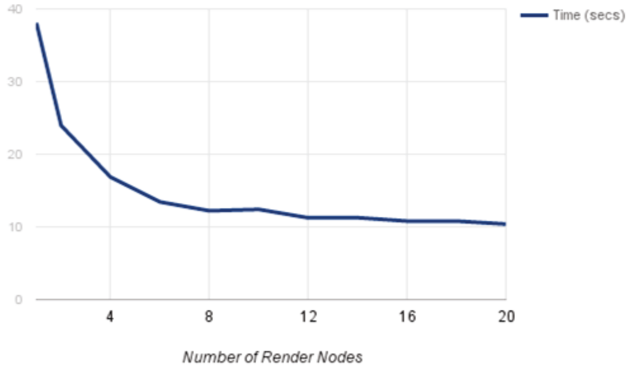

***Figure 12**. Results from the parallel version of UICE using twenty compute nodes from Newcastle's private cloud to render a single HD image.*

Future directions

We are working to provide a live view of sensor data from the city of Newcastle-upon-Tyne. This will also include a facility to view historic data from the Urban Observatory database. Our aim is to build a platform that allows us to experiment with novel high quality visualization techniques and with visualization system architectures.

The next stage in our experiments is to start to use a more realistic 3D model of the city as a backdrop for our data visualizations, see Figure 13. We believe this will prove more engaging to non-expert audiences but, as shown in Figure 14, it brings with it new challenges if we are to present clear informative visualizations. In addition, we plan to include support for new display devices, VR and AR headsets and 360 video presentations, while retaining our goals of inclusivity and cloud-side computation.

Acknowledgements

We thank the Science Central initiative and the SAGE Faculty at Newcastle University for supporting this work, also Microsoft for supporting the work through the Cloud Innovation Centre in collaboration with Newcastle City Council.

***Figure 13**. Aerial view of Virtual Newcastle Gateshead model rendered in Blender. (3D city mesh model courtesy of Northumbria University VRV group.)*

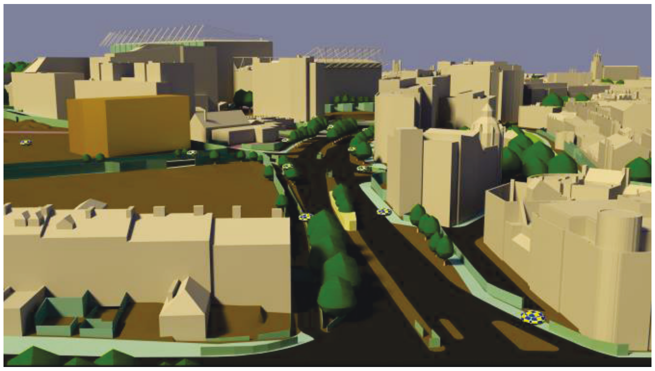

***Figure 14.** Sensor locations shown along St James's Boulevard in Newcastle-upon-Tyne. When displayed at natural heights the sensor glyphs are obscured in part by the trees and buildings in the area. How should we best represent data in crowded 3D environments? (3D city mesh model courtesy of Northumbria University VRV group.)*

References

[1] "The Urban Observatory", http://www.urbanobservatory.ac.uk, the School of Civil Engineering and Geosciences, Newcastle University, last accessed January 2017.

[2] N.S. Holliman, P. Watson, "Scalable Real-Time Visualization Using the Cloud", IEEE Cloud Computing, vol.2, no. 6, pp. 90-96, Nov.-Dec. 2015, doi:10.1109/MCC.2015.131

[3] "The Urban Insight Cloud Engine", http://portfolio.di-projects.net/, the Digital Institute, Newcastle University, UK, last accessed January 2017.