Voxelnet - An Agent Based System for Spatial Data Analytics

Charlotte Sennersten, Andrew Davie and Craig Lindley
Connection to the Physical World, Computational Intelligence and Decision Sciences
CSIRO Data 61
Sandy Bay, Australia
email: charlotte.sennersten@csiro.au, boo.davie@csiro.au, craig.lindley@csiro.au

***Abstract*—Voxelnet is a proposed voxel-based spatialised framework built upon internet communication and associated geospatial data standards. Each Voxelnet block has an individual IP address and functions as a computational agent. The internet of things (IoT) is also based upon objects having individual IP addresses on the internet so they can communicate using an individual unique identifier. IP addresses can roughly reveal where things are via geo locations. The internet, though, is a document based repository with 3D plug-ins that is not effective for sharing a unified world of interconnections from a volumetric point of view. The Voxelnet provides volumetric semantics on top of the current internet, within which measurements can be located, shared, interrelated and spatially analysed, supporting smart informed decisions in real and virtual worlds that account for spatial structure. It can also be used to interact with IoT systems using an integrated spatial representation.**

Keywords; Geospatial Data Systems; Block Models; Internet of Things; Data Sharing; Multi-Agent Systems.

I. INTRODUCTION

Increasing amounts of data are becoming available in all sectors. One area that is leading the growth and availability of large and heterogeneous datasets is in geosciences, and more broadly in the availability of different modes of remote sensing data in various spectral regions and resolutions. Point sensor networks are also becoming increasingly common, measuring factors such as climate (temperature, pressure, rainfall, wind direction and speed), river and water body temperatures, levels and flow rates, and traffic location and movement information. Concepts such as the IoT [1]-[3] propose the networked interconnection of an increasing variety of physical objects, ranging from sensors to consumer and industrial devices. The integration of these and other forms of sensors collecting data, using contemporary internet protocols and technologies, results in vast amounts of heterogeneous data, much of which is spatially specific or spatially located. However, current generic web tools are based upon a 2-dimensional (2D) text and image based presentation model, with few broadly available tools for analysing or even traversing or visualising large and complex data sets. To address this, this paper proposes the development of a *Voxelnet*, as an inter-networked system of volume elements that can provide an intrinsically 3-dimensional (3D) user interaction paradigm structured to readily provide visualisation of and access to spatial data sets. An active voxel system is proposed, where volume elements are also active computational agents that can process the data that they represent.

The Voxelnet is based upon points in space-time identified by [x, y, z, time], where time can be a specific time or a composite specification of times and/or time ranges. The Voxelnet space can be traversed (the spatial analog of browsing) using 3D interactive interfaces and systemised via a default index of 1 m³ cubes penetrating the whole digital world. The Voxelnet concept includes the composition or decomposition of blocks into smaller units down to an arbitrary resolution or up to a universal level providing a macro perspective. Every 1 m³ block also represents an agent with its own state that is able to react to and process data that enters and leaves the block, deciding what data is managed by the cube and what to do with it. The cubes will have their own IP addresses and can communicate with each other when changes occur to their current state. Hence, the Voxelnet readily lends itself to parallel computing and implement multi-block computations such as 3D cellular automata and finite element models.

This paper elaborates on various features of the Voxelnet concept. Section II proposes what a 3D indexation and annotation editor supporting the Voxelnet can look like, Section III reviews the IoT, Section IV discusses what is achievable and identifies some bottlenecks and issues, Section V discusses how this system can be implemented in a practical sense, Section VI summarises progress to date, and Section VII describes future work.

II. A 3D INDEXATION AND ANNOTATION EDITOR

The Voxelnet editor needs an annotation and indexation function supporting: 1) a functional communication architecture processing an individual spatialised IP address for each voxel, 2) agents accessed via their spatialised IP addresses, performing functions such as communicating what data they store (including 3D content such as 3D models with 3D sub-parts, 3D scanned objects, 3D environmental scans, data generated in a volume/area [x, y, z], etc.), 3) annotation and indexation functions so that 3D objects and related data can be created, changed, located, processed or displayed, 4) a 4D distributed data base management system so all data can be accessed and projected, and 5) scripting tools for active transformations and processing of block data.

In a dynamic world with huge amounts of data, in addition to capturing data we need to be able to display it as text or visually as 3D architectures, masses, areas and objects, and to represent where the location that it represents is in the real world. The vast amount of data needs to be categorised and indexed for use by researchers, industry and the public, allowing reuse of measurements, analyses and informed decisions that have been made. This includes data generated by agents in space, air, land, water and underground that perform distributed sensing and analytical functions while being contextualised via their location in the world. The Voxelnet with its indexation function will support researchers and engineers so they can focus on scientific challenges of analysis and discover new knowledge in cross disciplinarily correlations, rather than needing to handle details of data management.

III. THE IoT

The term "Internet of Things" was coined around 1999 by Kevin Ashton at Auto-ID Center mainly in reference to radio frequency identification (RFID) tags [3]. RFID tags are electronic transmitters that can provide identification and other data to a networked system. For the cargo industry this means that goods can be tracked anywhere at any time. In 2008 a group of companies launched the Protocol for Smart Objects (IPSO) alliance to promote the use of the internet protocol (IP) in networks of so-called smart or intelligent objects. Internet Protocol version 6 (IPv6) is the latest version of communication protocols providing an identification and location system for computers on networks [4]. As Leibson put it, "we could assign an IPv6 address to every atom on the surface of the earth, and still have enough addresses left to do another 100+ earths … if you take the surface of the earth as a perfect sphere and covered it with 1-layer-thick of atoms packed maximally close together." [5]. Even here the spatiality of the IoT starts to manifest.

The concept of the IoT, also referred to as the Industrial Internet, has expanded to encompass machine-to-machine internet connections in general. From a business perspective it needs to support machine and process-based analytics that are physics based, process deep domain expertise, are automated and are predictive. The analysis of physical machines and systems requires access to data from remote and centralised sources and visualisation in 3D and 2D graphical systems [1].

A. Unmanned Aerial Vehicles

As a caste study, Unmanned Aerial Vehicles (UAVs) flying and capturing data underground [6] will be considered as an example of active IoT nodes. A UAV is 'a thing' in IoT terms. The UAV has several sensors mounted on it and these are the things on the thing (the 'thing' concept is compositional: a bunch of things can constitute another thing). The sensors include video for navigation, inertial navigation sensors, and sonar distance sensors mounted on the UAV pointing towards surfaces in the environment that it flies by. Records of the physical location of the generation points of data includes complex relations like the thing (the UAV) located within 3D space with many other things (sensors) mounted to it directed towards many other things (walls, vehicles, people, etc.). The IoT universe is not merely concerned with individual things, but the compositional and spatial relationships among and between those things; locality is of fundamental interest. This can be achieved by associating every sensor value with a point in space, which can be specified in terms of a number of potential reference frames. For example, a sensor value might be represented in a location specified in a global reference frame (latitude, longitude and elevation), or by an orientation and displacement from a local coordinate frame centred on the UAV, which is in turn represented as a location and orientation in relation to the origin of a local mine (in this case) coordinate system, which has its own geometrical relationship to a wider area reference system such as latitude, longitude and elevation. In order to process all of this information, optimally in real time, data needs to be accessed in comprehensive and standardised formats, ideally contextualised based upon its meaning. Note however that this is not an argument for the semantic web in the common, ontology-based understanding of the term with all of its associated techniques. Rather, it is an argument for location and content-based indexing mechanisms. This is more of an argument for a denotational semantic web [7][8], where denotational meaning drives access. Remote control and navigation of UAVs flying in unknown spaces would benefit from a Voxelnet framework as a reference source of existing spatialised data, such as mine models and drill core data. 3D models of old mine voids provide an initial estimate of the likely structures that will be encountered. Drill core data can provide contextual information for the probabilistic data fusion of sensor information to make new measurements and maps of lithology and mineralogy [9].

The Voxelnet system provides a framework for the storage, use and reuse of data generated by the UAVs and other sensor systems. Drill core data is used routinely together with geological expertise, lithology, and contextual knowledge to create *block models*, which are voxel models of underground ore bodies, where features are represented for each voxel (which has dimensions on the order of 5 m on a side) such as specific gravity, hardness, grades of metals of interest, etc. [10]. In addition to this conventional mine analysis modelling, the Voxelnet also stores all data from the UAV. This includes data that can be processed to produce 3D models of mine void spaces, such as visual (video) data, LIDAR (laser scan data), sonar, optical or sonar flow, and inertial navigation data. All of these data can be retained in the spatial index system of the Voxelnet, together with all other forms of spatialised data gathered by other platforms, such as other UAVs or surface robots, manual survey instruments, and instrumentation built into mining vehicles such as load-haul-dump (LHD) machines, drill jumbos and drill platforms. All survey data can potentially add to mineralogical evaluations and estimates (e.g., of grades, percentages of target metal concentrations, with associated probabilities). But the totality of data can be used for potentially diverse purposes, such as relating vehicle performance and maintenance data back to spatial attributes of the working space (e.g., rock hardness, shape, gradients, sharpness of turns, stability), actual work records (distances and heights traversed, motor accelerations and decelerations, total vehicle accelerations and decelerations), and mean time between failure for vehicle parts, subsystems and systems. These data can feed into site and enterprise level operations analysis and optimisation. Not all of these data are or need to be spatially specific, but the availability of spatially indexed data supports forms of cross model analytics, such as the discovery of terrain regularities associated with temporal data features that in turn associate with maintenance trends.

The forms of spatialised data noted above could be (and currently are) recorded using methods that are not primarily indexed spatially, or if they are spatially indexed, the indices exist in highly localised systems. The Voxelnet generalises access to spatially indexed data, supporting a much greater variety of scales in analytic functions, and analytics across more diverse perspectives and interests than the design targets of existing tools.

IV. HOW CAN THE VOXELNET AND THE IoT CO-EXIST?

The Voxelnet uses IoT concepts in two ways: i) as a repository, spatial browsing and analytics system for IoT connected objects, and ii) each block behaves as an interconnected "thing" within the IoT. Where the IoT deals with actual physical objects and sensors, the Voxelnet can support virtual objects and sensors that may have been converted or derived from physical world data by agents within the Voxelnet infrastructure.

The IoT relies on IPv6 due to the huge number of individually accessible things (objects) that are expected to be part of it. Likewise, the Voxelnet has the same issue with the default 1 m^3 block index representing the world. The volume of the Earth is approximately 10^{21} m^3, so that is how many blocks that would be needed for a complete Voxelnet of the Earth. IPv6 supports about 3.4×10^{38} individual addresses and so could represent each individual block as an IP-addressable entity for the entire Earth and still have most of its addresses left over. Just considering the surface area of the Earth of 5.1×10^{14} m² and with 2000 m of height interest, then there are around 10^{18} blocks of 1 m^3 size.

However, the represented volume needs to accommodate the volumetric data available and context of how it is used. For example, geological block modelling used in geosciences seeks to achieve a block scale derived from the scale of available data samples together with considerations of the evidence that the data provides and a suitable degree of averaging for resource estimation purposes [11]. Raw data should be represented in the highest level of available detail. The proposition of a default 1 m^3 block is a human scale convenience for traversing the complete volumetric system, where specific functions may require the aggregation or decomposition of this scale to the degree needed. For example, underground mine planning typically uses block around a size of 5 m^3, while an autonomous vehicle that needs to analyse points in a point cloud to derive the location and classification of features and objects in its immediate operational environment might need a spatial index accuracy on the order of millimetres. The requirements of scale and the functional ability to transition to different scales impacts the design of the underlying computational and communications infrastructure, including storage requirements, network bandwidth, and processing time requirements (not considered in detail here).

This leads to questions for representing the world as a block model including how to meet requirements for: i) memory, ii) connectivity, iii) speed, iv) network bandwidth and v) processing (speed) related to a large number of addressable objects.

V. TECHNICAL APPROACH

Internet IP addresses currently using IPv6 have 128-bit addresses facilitating routing via octets representing hosts and subnet structures. A domain name or URL is a representation that is more understandable and memorable to a human, which maps onto a detailed numeric IP address. The semantics of an Internet URL or IP address is essentially a routing instruction identifying a specific machine within a network. This depends only upon the network topology and domain structure and has nothing to do with any other form of locational information, e.g., about location specific data.

The Voxelnet proposed here aims to create an alternative addressing system having semantics of spatial locations instead of routing pathways. This could sit on top of a transport layer protocol like IPv6, but it needs to provide users with spatial location, content, data modalities and feature information, rather than routing information. Hence, the semantics of a Voxelnet address is a 3D spatial location (within a designated geographic coordinate system), time or time range (default to latest), a data type selection specification, process instructions, etc. The list of data types and processing required includes specific spatial data sets (which could include temperature, surface type, pressure, gravity, magnetism, biomass, mineralogy, infrastructure, etc. The possible format for such a spatialised URL can be based upon existing standards for geospatial data systems, such as those specified by the Open Geospatial Consortium (OGC) [12], including netCDF, GeoSPARQL, Geography Markup Language Encoding Standard (GML), KML (formerly the Keyhole Markup Language), etc. However, currently there is no universal access tool that provides a higher level interface for access to the data mediated by these standards and hides source details from users.

Such a tool should have analogous seamless integration across multiple heterogeneous databases as provided by web browsers for data that is unified into a 2D text/image presentation paradigm. The Voxelnet is intended to provide this kind of access. This is fully in line with current broad standardisation efforts (e.g., [13]), but with an emphasis upon the creation of a unified user experience paradigm for 3D spatialised data and analytics. It can also be approached

 ISBN: 978-1-61208-462-6

as a user-oriented layer on top of ongoing initiatives to create a Spatial Identifier Reference Framework (SIRF) [14].

The Voxelnet is conceived as a system of agents, each of which is responsible for the following kinds of functions, including requested processing within the scope of its designated volume:

- Responsibility for management of a specified volume of information at a specific scale.
- Activated by and responds to Voxelnet requests falling within its volume of responsibility.
- Interpretation of the request.
- Assembly of a data package or service that satisfies the request.
- Conversion from stored coordinate format to coordinate format requested.
- Filtering by attribute value as requested.
- Aggregation and disaggregation functions associated with nested spatial structures. For example, the volume represented by an agent may exist within larger scale volumes represented by other agents, or may contain smaller scale volumes represented by other agents. A query might be satisfied by assembling a data package from several different scales representing varied resolution of sensor instruments and technologies.
- Streaming to participate in traversal interactions through virtual spaces that are constituted by many agents, e.g., to create a coherent virtual world experience assembled from many smaller scale spatial data elements.
- Implement computationally derived data modalities, either pre-computed or computed on demand. Examples of these might include finite element or cellular automata computations, e.g., to find a path through a landscape that minimises gradients or height changes, to predict flooding locations, for bushfire behaviour prediction, etc.).

On top of this network client functions must be provided to provide a coherent experience for users that hides the underlying standards and distributed repositories unless requested, and focuses upon the creation of a user task-oriented interaction, visualisation and comprehension paradigm.

VI. SUMMARY AND FUTURE WORK

The Voxelnet concept aims to provide seamless traversal through a 3D virtual space with content generated from diverse and heterogeneous data sources. Content may be multimodal, and can be computed from primary or other computational sources. Development of a demonstrator for this is work in progress that will support underground void mapping by UAVs as a first example use case. The Voxelnet concept is general and can be used for many applications.

REFERENCES

[1] P.C. Evans, and M. Annunziata, "Industrial Internet –Pushing the Boundaries of Minds and Machines", General Electric Co., November 2012.

[2] A. Noronha, R. Moriarty, K. O'Connell, and N. Villa, "Attaining IoT Value: How to move from Connecting Things to Capturing Insights – gain an edge by taking analytics to the edge", CISCO, 2014.

[3] K. Ashton, "That 'Internet of Things' Thing", RFID Journal, 2009. Available from: http://www.rfidjournal.com/articles/view?4986 [retrieved: January, 2016]

[4] A. Keranen, J. Arkko, S. Krishnan, and F. Garneij, "Toward Ipv6 world in mobile networks", The Advance to IPv6, Ericsson Review, 2/2011. Available from: http://www.ericsson.com/res/thecompany/docs/publications/ericsson_review/2011/IPv6-June-08-2011.pdf [retrieved: January, 2016]

[5] Quora, available at: https://www.quora.com/Is-Steve-Leibsons-statement-about-IPv6s-vast-address-range-true [retrieved: January, 2016]

[6] C. Sennersten et al., "Unmanned Aerial Robots for Remotely Operated and Autonomous Surveying in Inaccessible Underground Mine Voids", The Third International Future Mining Conference, November 2015, 101-08, Sydney, Australia.

[7] C. Sennersten, A. Morshed, M. Lochner, and C. Lindley, 2014, "Towards a Cloud-Based Architecture for 3D Object Comprehension in Cognitive Robotics", COGNITIVE 2014, The Sixth International Conference on Advanced Cognitive Technologies and Application, 25-29, May, pp. 220-225, 2014, Venice, Italy.

[8] M. Lochner, C. Sennersten, A. Morshed, and C. Lindley, 2014, "Modelling Spatial Understanding: Using Knowledge Representation to Enable Spatial Awareness in a Robotics Platform", COGNITIVE 2014, The Sixth International Conference on Advanced Cognitive Technologies and Application, 25-29 May, 2014, pp. 26-31, Venice, Italy.

[9] A. Rahman et al., 2015, "A Machine Learning Approach to Find Associations Between Imaging Features and XRF Signatures of Rocks in Underground Mines", IEEE Sensors 2015, Busan, South Korea, November 1-4, 2015.

[10] C. Lindley et al., 2015, "A multilayer 3D index tool for recursive block models supporting terrestrial and extraterrestrial mine planning" accepted for the Third International Future Mining Conference, 4-6 November 2015, pp. 289-296, Sydney, Australia.

[11] J.-P. Chiles, and P. Delfiner (2012). "Geostatistics: Modeling spatial uncertainty. 2nd Edn". New York: Wiley.

[12] http://www.opengeospatial.org/ [retrieved: January, 2016]

[13] http://www.anzlic.gov.au/foundation_spatial_data_framework [retrieved: January, 2016]

[14] http://portal.sirf.net/about-sirf [retrieved: January, 2016])