Visualisation

iVEC Interns 2013

Paul Bourke

iVEC @UWA

Introduction to visualisation

- The use of (advanced) computer graphics and algorithms to provide researchers with insights into their data.
- Finds application across a diverse range of disciplines.
- The common thread are the graphical techniques and algorithms.
- Very interesting field: requires a wide range of skills
  - computer programming
  - algorithms in computer graphics
  - mathematics
  - realtime / interactive APIs and technologies
  - human / computer interfaces
  - knowledge of human perception theory
  - creativity and design
- Will present a number of projects in 2011-2013 as a way of introducing various aspects of visualisation and the techniques involved.

Outcomes

- Informing researchers.
- Conveying research outcomes to peers.
- Educational resources for University courses, the general public, and school children.
- Exhibitions in museums and art galleries.
- Research outcomes for data visualisation include
  - Uncovering something new.
  - Understanding some aspect of the data faster.
  - Finding errors.

Visualisation @ iVEC

- Four 1/2 FTE funded positions: UWA - Curtin - CSIRO
- Budget to support visualisation activities of researchers at any of the iVEC partners.
- Compute infrastructure dedicated or optimised for challenging visualisation projects.
- Displays to support visualisation.

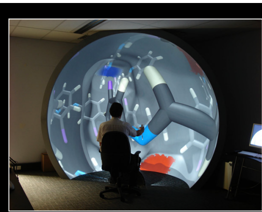
Curtin & ECU

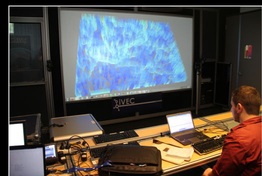
CSIRO


UWA

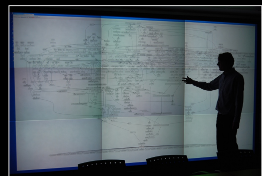
Murdoch


ECU

Visualisation @ iVEC

- Capture infrastructure:
  - Stereo3D video camera
  - High resolution video cameras
  - Specialist camera rigs
  - Structured light camera
  - 360 video camera
- Software tools and expertise.

Stereoscopic 3D cameras

High resolution video cameras

360 degree video capture

3D scanners

Hardware

- Three aspects of the human visual system not normally engaged
  - stereopsis
  - fidelity
  - field of view
- Displays that match the data topology
  - Magicplanet
- Other
  - Cylinder or partial cylinder
  - Head mounted displays
  - Tactile visualisation
  - Glasses free displays
  - Holography
  - Rapid prototyping

Stereopsis

- Exploiting our two eyes, the two views from which are responsible for the depth perception we experience.
- Easy to imagine that for geometrically complicated relationships that depth perception could be valuable.
- Lots of different methods for inducing the effect
  - Anaglyph
  - Polaroid, also referred to as passive stereo
  - LCD shutter glasses
  - Infitec glasses
  - Autostereoscopic (Glasses free), eg: barrier strip or lenticular
- In ALL cases there is only one thing going on ... a separate image needs to be presented to each eye. All the technology options are just a means to achieve that.
- Realtime or movie content needs to create the two views correctly.
- The iVEC facilities at Curtin and UWA have stereoscopic displays, large scale and television size.

Example: Sports Science

- Application of stereoscopic filming and capture in sports science.
- Providing realistic visual stimuli while tracking player behaviours/performance.

Courtesy Sports Science, UWA

Fidelity

- The human eye is capable of higher visual fidelity than the pixel size of a standard display at normal viewing distances.
- The real world we experience is continuous (except for the digital devices).
- Two solutions:
  - make the pixels smaller, eg: so called "retina" displays (too small)
  - stand further back
- Tiled displays allows one to present high volume of visual information at high fidelity.
- UWA, Curtin, and Murdoch have tiled displays.

Example: Tiled display at UWA

- Example below is 8x1920x1080 pixels = 32MPixels
- Upgraded to a version with very narrow bezels, still hard to get the same pixel count since the panels are only HD (6x1920x1080 = 12MPixels)
- Has space advantages over rear projection displays.

Example: Tiled display at Murdoch

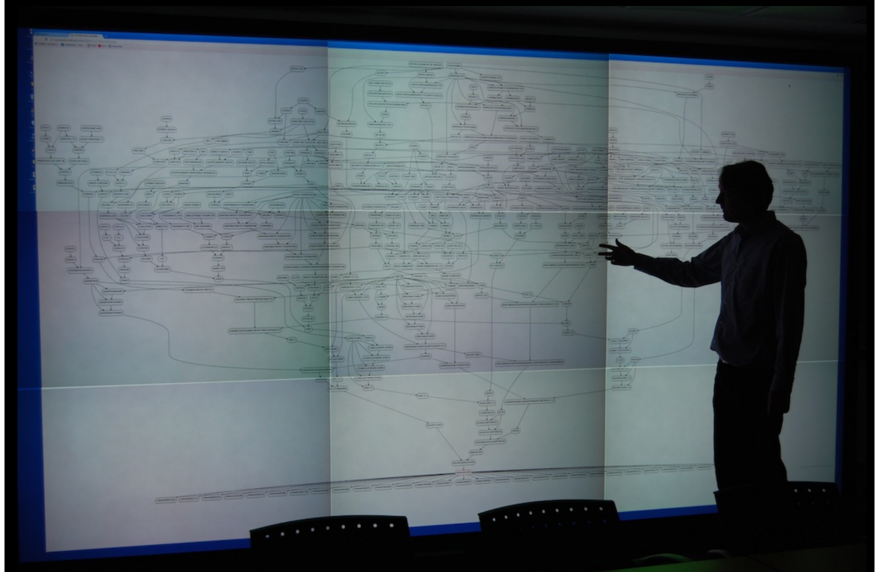

Visual field of view

- Usual display occupies perhaps 25 degrees of our horizontal field of view.
- But we are capable to almost 180 degrees horizontally and 120 degrees vertically.
- Historically planetariums have provided an immersive experience to the public. But orientated obviously in favour of astronomy.
- The iDome is a dome orientated to be more suited to visualisation and "front facing" experiences.

Example: AVIE

Magic planet

- Plane and sphere are topologically different, this means one cannot be mapped to another without distortion.
- Lots of ways of mapping one to the other to minimise some metric, may preserve area, angle, lengths, or combinations.
- Magic planet (others: Science on a Sphere, PufferSphere ...) can represent spherical data naturally.

Other senses ...

- Sonification: visualising using our sense of hearing.
- Classic examples:
  - Hospital pulse measurement: "The machine that goes PING".
  - Geiger counter.
- Two most common approaches are to map some variable to the waveforms (eg: amplitude or frequency modulation) or to map to instruments (eg: midi).
- Pure sonification can be difficult, often just sounds like noise or really bad music.
- More commonly used to accompany and reinforce the visuals.
- Good example is sonification of nuclear tests from 1945-1996 by Japanese artist Isao Hashimoto.
- Sonification of pulsars

Modulated white noise

Midi instrument, equal tempered scale

Sense of touch

- Force feedback has been used for some time to allow data to be “felt”, referred to as haptic technology.
- Commodity example is joysticks vibration in car driving games.
- Used extensively in remote surgery - force feedback scalpel.
- More recently it has been possible to make models that can then be explored physically.
- Exploring data in the same way as we explore objects in real life.

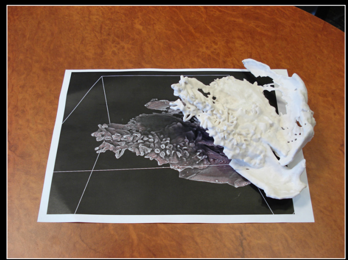

Sense of touch

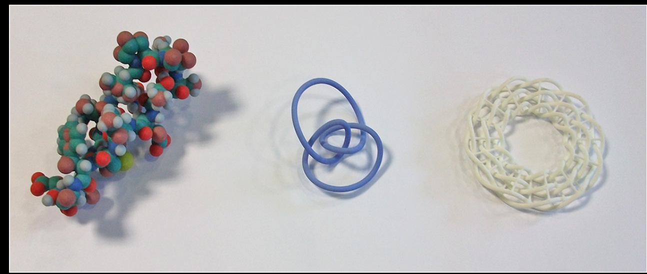

Peptide Knot theory Neuroscience

Illustrative vs data visualisation

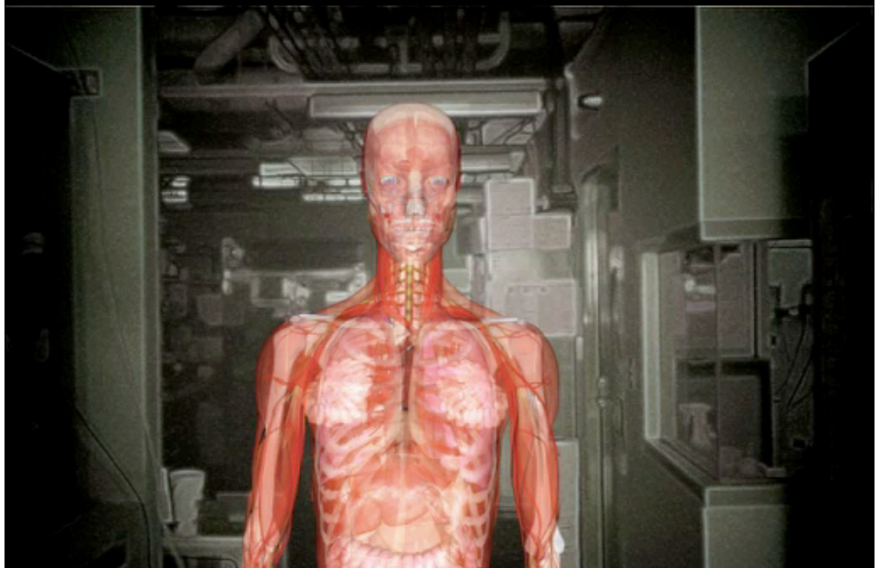

Courtesy Drew Berry

Data mapping

- Much of visualisation is about mapping data to graphical elements.
- Sometimes this mapping is obvious, other times there is some creative choices.

Courtesy Drew Whitehouse

Conventional mappings

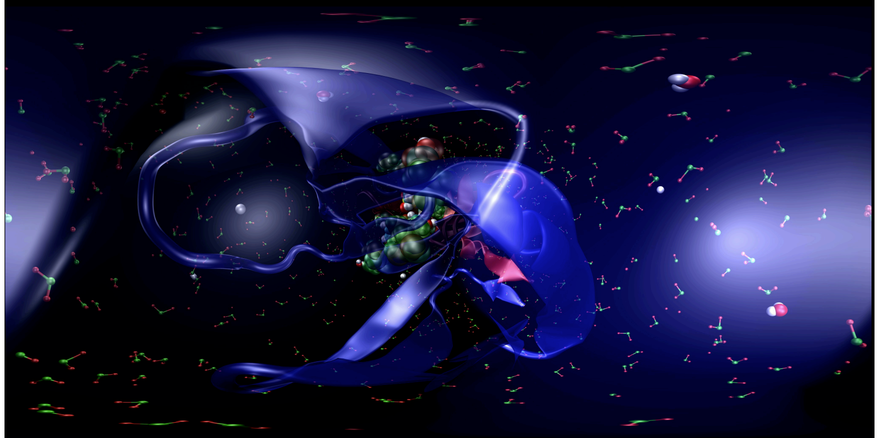

Abstract mappings

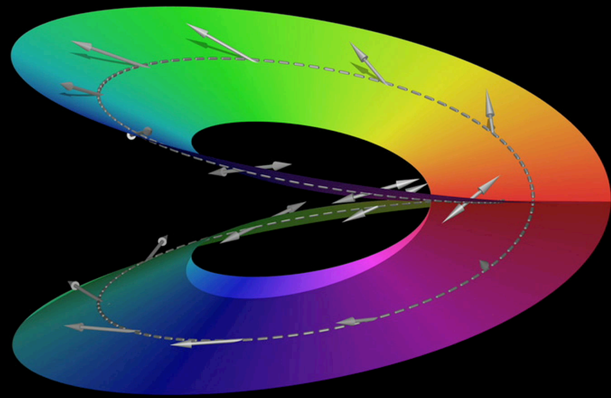

Pausiris mummy

Volumetric data

- A digital image contains some quantity sampled on a regular grid on a 2D plane.

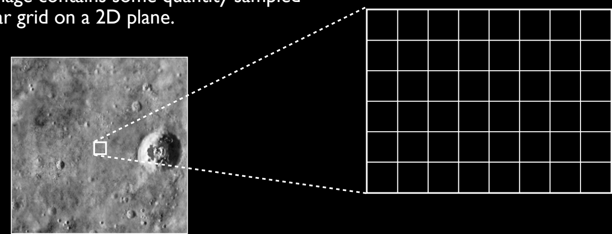

- In a volumetric dataset there is some quantity sampled on a regular 3D grid.

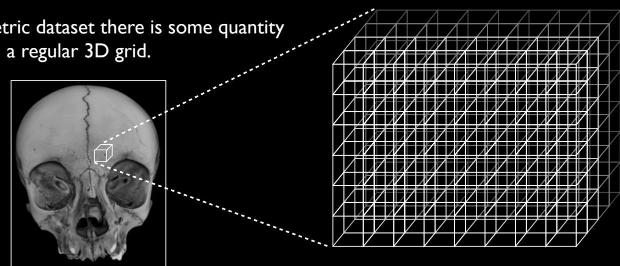

Volumetric data: terminology

- In a 2D image the fundamental unit of measure is a "pixel".
  The quantity represented by the image is sampled at each pixel.

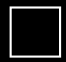

- In a volumetric dataset the fundamental unit of measure is a "voxel" (VOlume piXEL).
  The quantity represented by the volume is sampled at each voxel.

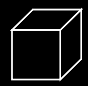

- The resolution of a 2D image is defined as the number of pixels horizontally and vertically. The resolution of a volumetric dataset is defined as the number of voxels in width, height, and depth.
- Image pixels are usually but not always square. Voxels are sometimes cubic (simulation)but generally not (experimental data), for slice based data the resolution within the slices is often very much greater than that between the slices.
  Note that some volumetric data (eg: finite element simulation) can have variable voxel sizes.
- Depends on who you talk to and their area of research but generally
  - A "small" volumetric dataset may be < 200 voxels on each side.
  - A volumetric dataset is considered "large" if it is > 1000 pixels on each side.
- Another important characteristic is the dynamic range of the data at each voxel. Most commonly a single byte, integer (2 or 4 bytes), or floating point. May even be vectors, multivariate, and so on.

Volumetric data

- Volumetric datasets have been a common data type in many areas of science for some time.
- Traditionally one thinks about medical data, for example MRI.
- Other scanning and 3D imaging technologies include CT (MicroCT) and CAT scans. There are many others.
- Volumetric data also arises from numerical simulations.
  Quite common in astronomy and engineering (finite element calculations).
- In scanned volumetric datasets the quantity per voxel depends on the scanning technology. For example: MRI essentially gives water content, CT gives density.
- For volumetric datasets derived from simulation there can be multiple variables per voxel.

Medical research (MRI)

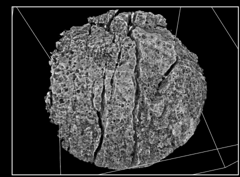

Geology (CT)

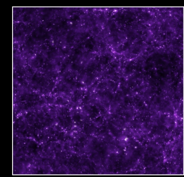

Physics (Simulation)

Volume visualisation

- The process of exploring and revealing the structure/interior of a volumetric dataset.
- The general approach involves a mapping between voxel values and colour/opacity.

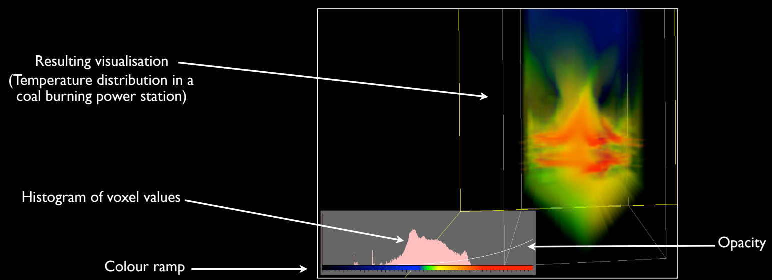

Example: MicroFossil

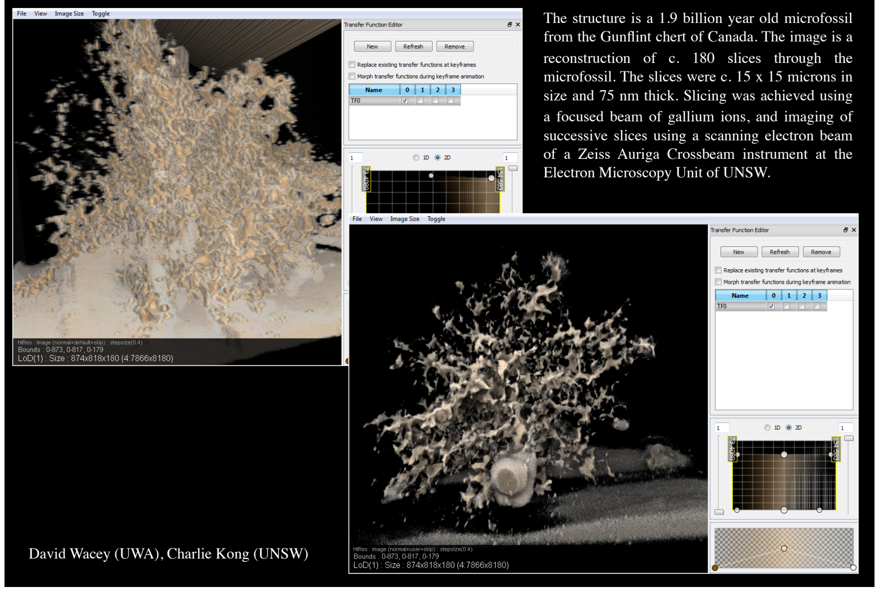

The structure is a 1.9 billion year old microfossil from the Gunflint chert of Canada. The image is a reconstruction of c. 180 slices through the microfossil. The slices were c. 15 x 15 microns in size and 75 nm thick. Slicing was achieved using a focused beam of gallium ions, and imaging of successive slices using a scanning electron beam of a Zeiss Auriga Crossbeam instrument at the Electron Microscopy Unit of UNSW.

David Wacey (UWA), Charlie Kong (UNSW)

Example: Western Gorilla (Male)

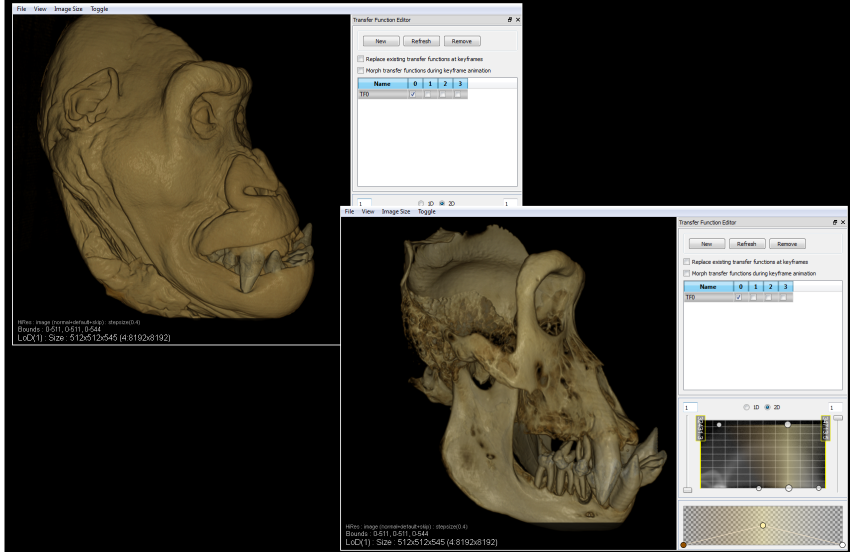

Graphics

- Visualisation has always been a demanding user of hardware accelerators, in particular graphics cards.
- While the gaming/entertainment industry might be the main driving force for improvements in the commodity market, there is still a high end graphics card product range dedicated to visualisation applications.
- Since visualisation often deals with large volumes of geometry (points, lines, polygons) there is a need for high performance geometry processing.
- Volume rendering is now almost always performed on the graphics card and the rendering is achieved using shaders that ray cast into the volume.
- Size of volume that can be handled interactively is usually limited by the amount of texture memory.

Example: Rabbits liver

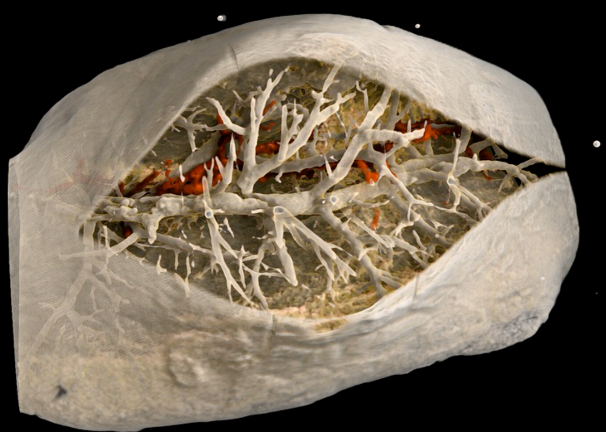

Courtesy Ajay Limaye, ANU

DARK

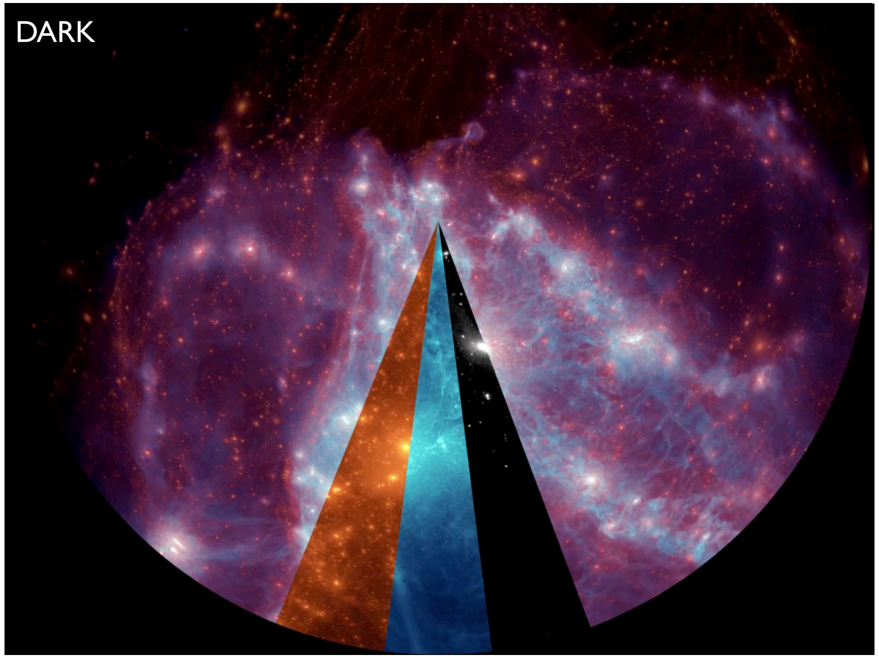

Visualisating astrophysical simulations

- DARK is a digital fulldome (planetarium) production explaining the nature of dark matter, the missing 80% of the mass of the Universe.
- http://darkthemovie.info
- Example where the visualisation is performed on the same supercomputer (epic) as where the simulations were performed.
- Three simulations were performed: cosmological large scale structure, small and large galaxy formation. In each case there are very large data files at each time step, for example the cosmology data is almost 2TB.
- Visualisation is not "arbitrary", characteristics of how the simulation is performed need to be used in the visualisation.
- For example: The simulation could be considered as a particle system but visualising as such would not be a satisfactory representation. Each point is actually a cloud, has a sphere of influence.

Smoothing kernel

- If it were just "points" it would be much simpler.
- 3D functions of radius, similar to a "point spread function" in optics. Note this is used within the simulation software so not an arbitrary choice for the visualisations.
- For particles without a smoothing kernel (eg: stars) a Gaussian is used which allows the same pipeline to be employed. Use a single standard deviation, star mass determines the amplitude.

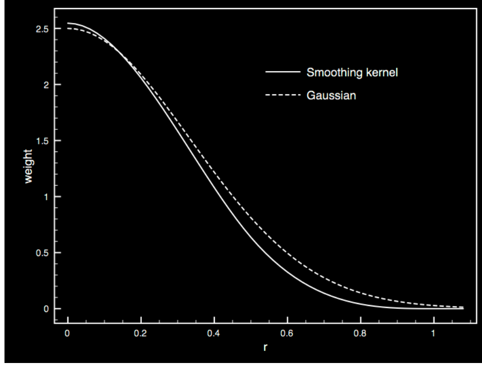

$$W(r) = \begin{cases} \dfrac{8\,(1 - 6\,(1 - \frac{r}{s})^2 + 6\,(1 - \frac{r}{s})^3)}{\pi s^3} & r < \frac{s}{2} \\ \dfrac{16\,(1 - \frac{r}{s})^3}{\pi s^3} & \frac{s}{2} \le r < s \\ 0 & r \ge s \end{cases}$$

Control over the time/quality trade-off

- Common requirements of very large datasets is to be able to render simply and fast, better quality as time becomes available.
- Be able to manage interactive performance as well as high quality rendered imagery.
- Implemented smoothing kernel by sampling (regular or stochastic) in 3D. Points are then projected onto plane, cylinder, or spherical surface. The image is then a histogram the projected points contribute their kernel weighted mass to.
- Advantage of being able to form image frames with speed/quality trade-off.

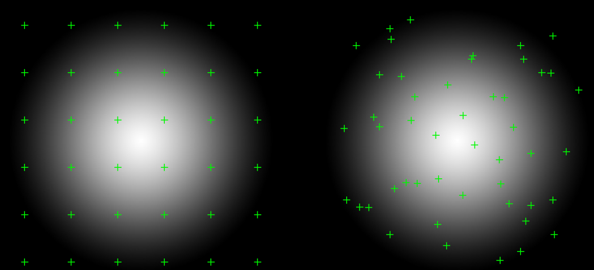

Uniform sampling

Random sampling

Data types and conversion

- The visualisation process often involves data conversion, either between file/data formats or to derive the quantities to be visualised.
- Often data storage format is ad hoc, that is, it is made up to suit the requirements at the time.
- However many disciplines have their own standard/favourite file format
  - DICOM in the previous case of medical 3D scanners.
  - FITS (Flexible Image Transport System), common in astronomy imaging.
- In this case the data is stored in HDF (Hierarchical Data Format).
- Many of these data formats are extremely complicated, for example, they can be self describing.
- Fortunately many come with an API to assist programmers read the data.
- Often the format used by the researcher or by existing code base is not efficient for visualisation. Or it may contain more information (variables) than required.

Parallel rendering

- Generally desirable for visualisation to be realtime and interactive.
- Many times due to either the volume of data or the visual fidelity required this is not possible.
- Often prohibitive to move the data to another system for visualisation.
- In these cases visualisation can occur "offline", that is, pre-rendered.
- The time required to do this can be reduced by employing parallel processing techniques.
- There are two general categories: data parallel or image space parallel.
- Data parallel renders N subsets of the data onto an image plane and then combines the image planes.
- Image space parallel subdivides the final image into sections and the data is rendered in parallel into each section.
- Combinations of the these are also possible.

Parallel rendering

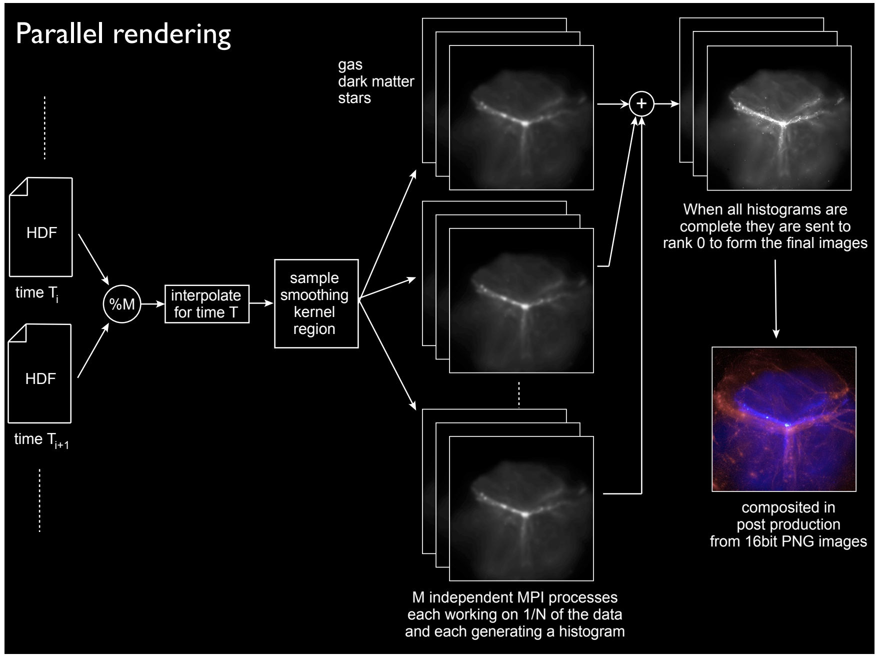

GIMIC

- Simulates the formation of a Dwarf Galaxy, similar to the Large Magellenic Cloud.
- The formation of these galaxies is a violent dynamic process.
- Dark Matters forms in filaments along which gas flows into the central disk where star formation occurs.
- Computed on cosma (Durham University).
  Used 32 CPUs, 92 hours (~3,000 CPU hours).
  Rendering performed on epic (iVEC).

COSMOS

- Simulation within a cubic region (periodic bounds) of the Universe just after the Big Bang.
- 600 million light years on each side of the cube.
- Shows dark matter collapsing over 14 billion years of cosmic time, forming filaments and collapsing haloes of the Cosmic Web.
- Note there is no smoothing kernel here, the images look smooth and continuous due to the 1 billion+ particles per time step.
- Even at 3Kx3K, if the whole dataset is in shot then on average there are over 100 points per pixel (if they were distributed uniformly).
- The final image is essentially a histogram formed on the projection plane.
- Original simulation computed on vayu (NCI).
  Used 1024 cores, 2.8TB RAM, took 19 hours (~20,000 CPU hours)
  Rendering performed on epic (iVEC).

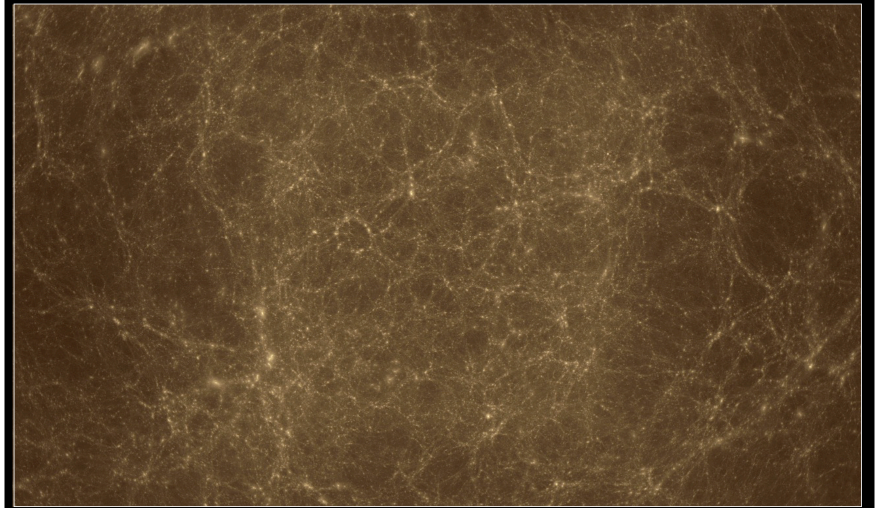

Engaging displays and other outcomes from visualisation

- DARK is an example of how visualisation outcomes can find application outside the original research aims.
- The engaging and informative visuals used in papers, presentation to peers, public outreach, and general education in schools.
- In the case of DARK the visualisation is presented on a hemispherical dome, a planetarium.
- An example of a display technology that leverages more of our visual system than a traditional flat display. It fills our entire peripheral vision.
- Requires rendering to other than standard perspective projections.

Projections

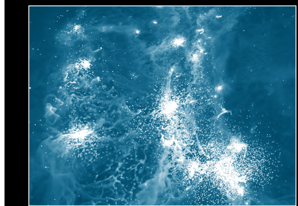

Orthographic

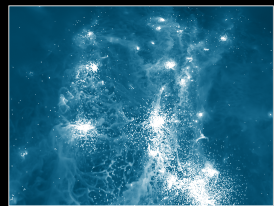

Perspective

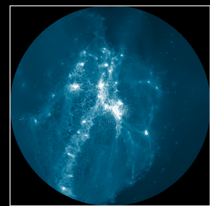

Fisheye

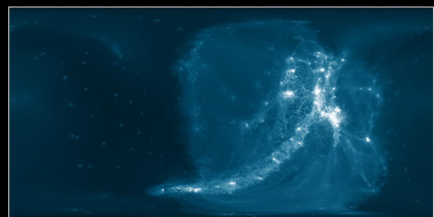

Spherical

Cylindrical

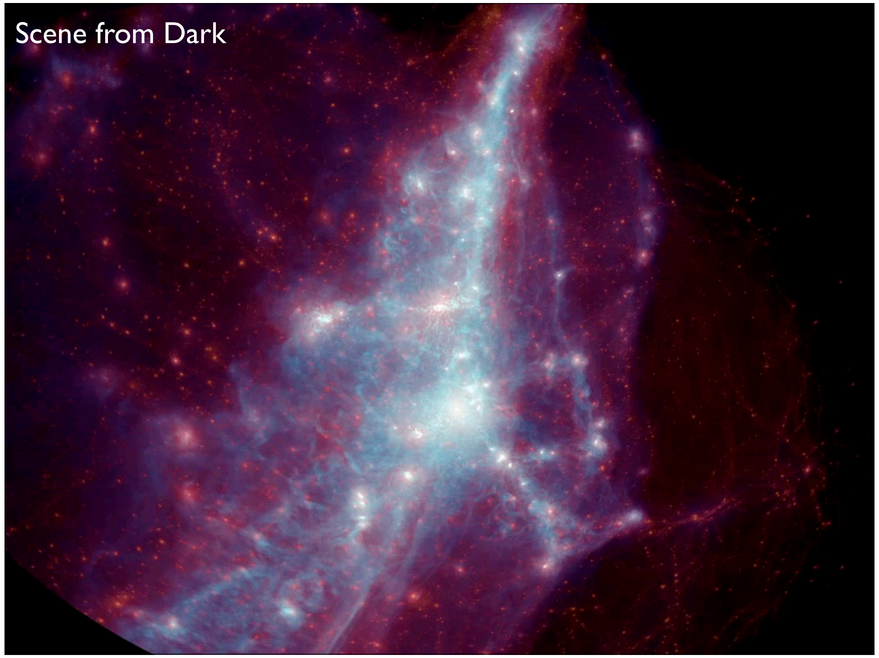

Novel Imaging

- Will consider the following
  - Photogrammetry, the automatic construction of 3D surface data from photographs.
  - Gigapixel photography
  - Photographic mosaics
  - High definition 360 panoramas
- Projects discussed here are largely based upon recordings in archaeology and heritage.
- Goal is to create
  - research objects that are richer than "normal" photographs
  - automatic creation of assets for virtual worlds and serious gaming

3D reconstruction: Rock art capture

- Photogrammetry, the automatic construction of 3D surface data from photographs.
- Historically the main application in mapping land elevation.
- Goal is to create
  - research objects that are richer than just photographs
  - assets for virtual worlds and serious gaming
- Wish to avoid any in-scene markers required by some solutions.
  Often impractical (access) or not allowed (heritage).
- Want to target automated approaches as much as possible.
  [Current site surveys recorded 100's of objects].


Photographs

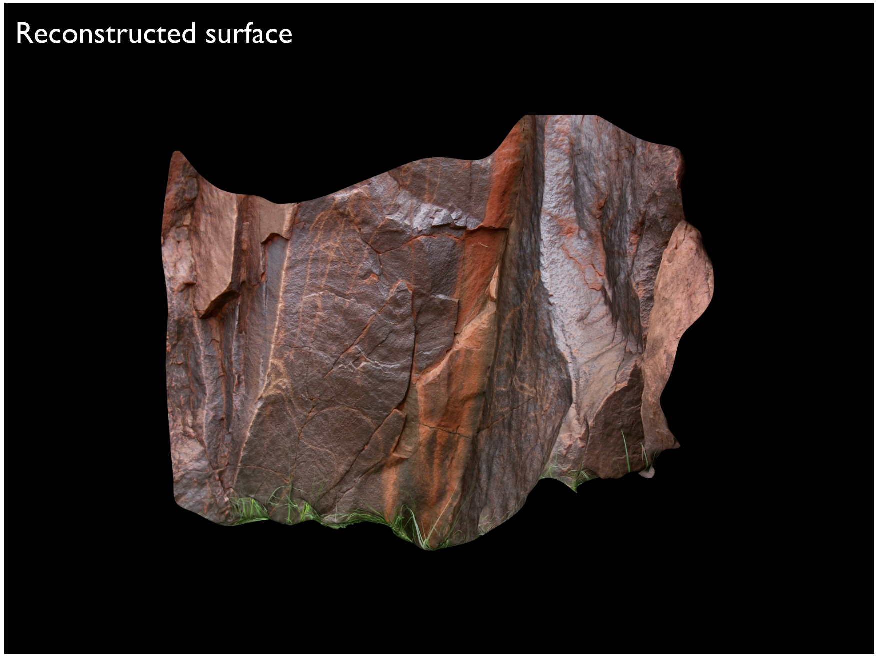
Reconstructed surface

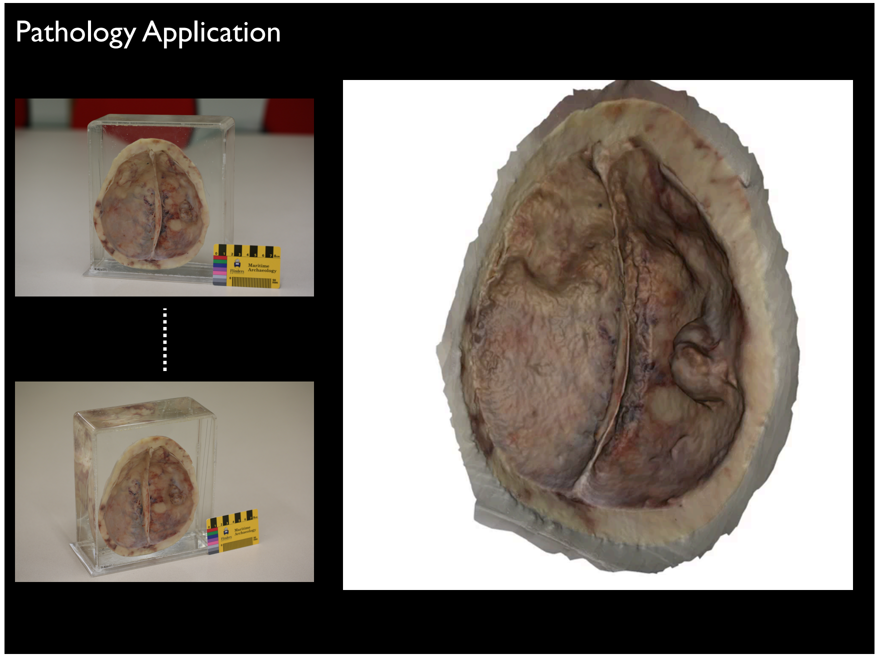
Pathology Application


Mining Application


Motivation: Artefacts in cultural heritage
Indigineous headdress

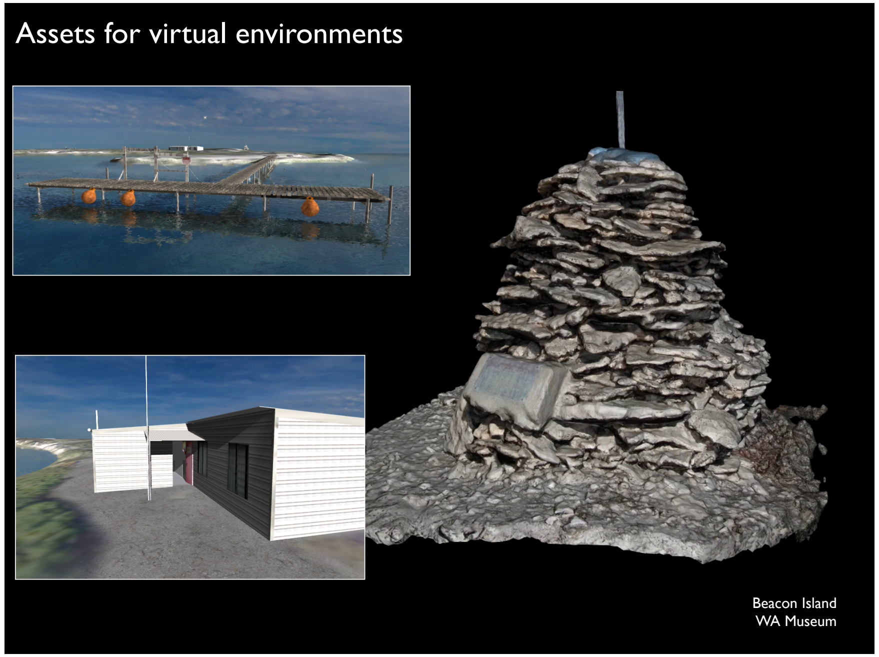
Assets for virtual environments
Beacon Island
WA Museum


Heritage preservation
Dragon Gardens, Hong Kong

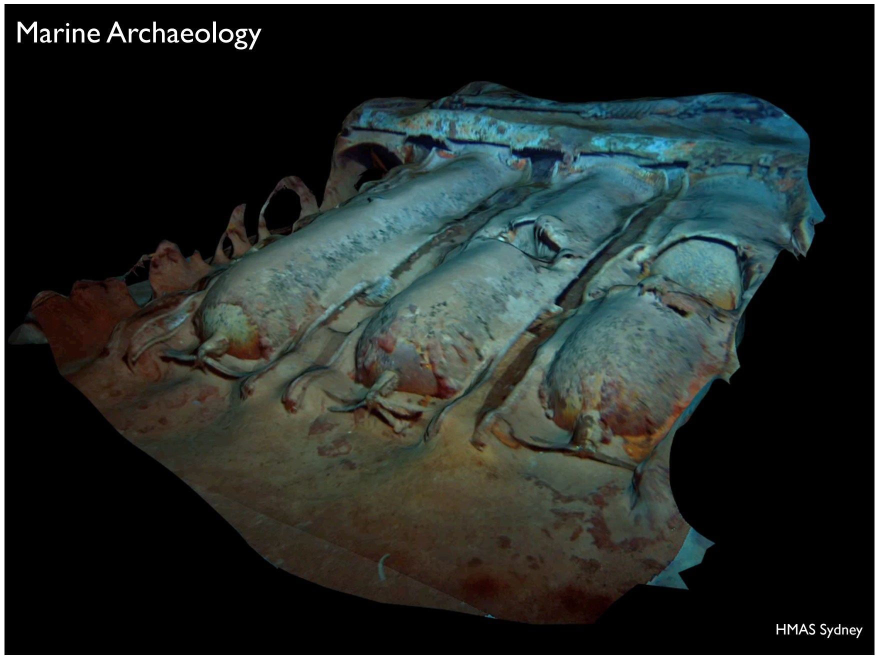
Marine Archaeology
HMAS Sydney

Gigapixel images

- Sensors in cameras are limited, what if one wanted a higher resolution image.
- Solution is to take many photographs in a grid and stitch them together.
- Solution is used in Astronomy (Hubble Space Telescope), Microscopy, Geology, etc.

40,000 by 20,000 pixels

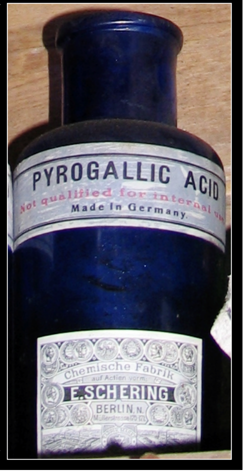

Hurleys darkroom, Mawsons hut (Antarctica), courtesy Peter Morse.

Example: ASKAP site

First ASKAP dish

Canon EOS 5D MkII camera and gigapan mount

ASKAP site, Boolardy

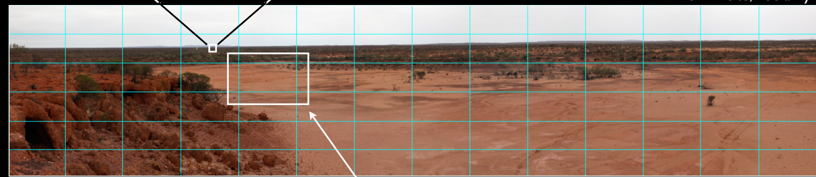

21 MPixels, Canon EOS 5D Mk11

Total: 1.5 GPixels

Example: Wanmanna (rock art site)

13 x 3 grid

40,000 x 10,000 pixels

15 x 4 grid

Single 10MPixel image

Example: microscopy

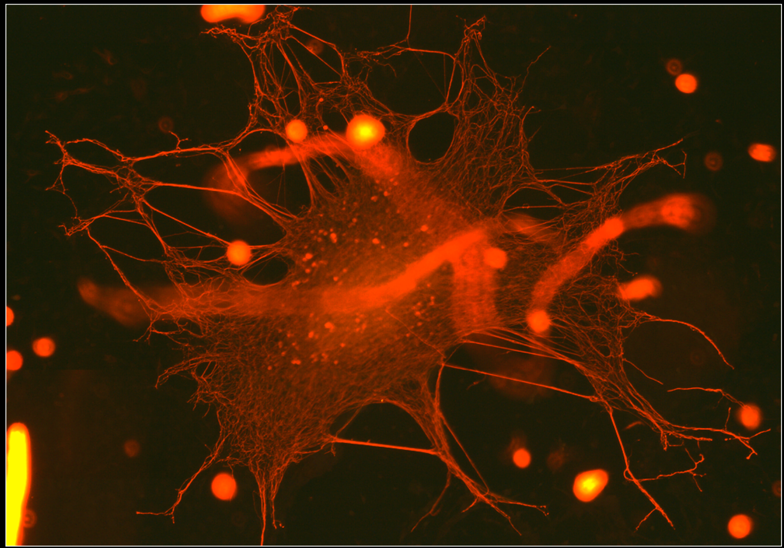

Image courtesy CMCA, UWA

11,000 x 8,000 pixels

Example: Hubble Deep Field

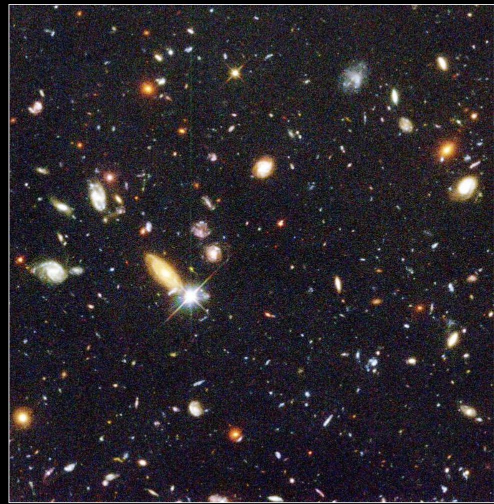

Hubble deep field, 340 images.

Data structures for visualisation

- An example of when one needs special data formats is for images so large they may not fit into memory.
- How does one interactively explore such high resolution images.
- Lots of approaches
  - spooling from disk as required
  - data subdivision: only see "close" data
  - automatic subsampling: easy for some datasets, harder for others. Data is resolved when time available.
  - precreated subdivision
- Often the data needs to be stored in such a way as to allow fast access for visualisation.

Example: volume data

- Volumetric data can be subsampled for fast interaction.
- Grids or octrees can be used to manage view pruning.
- Resolution determined as time is available, low resolution during interaction, higher resolution when interaction stops .... image gradually resolves.

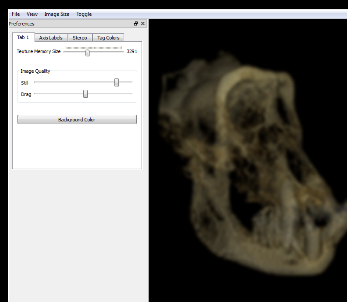

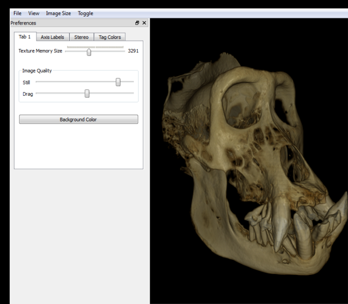

Example: Pyramidal tiff

- The tiles visible depends on where in the image one is exploring and the zoom level.
- In essence how Google Earth works.

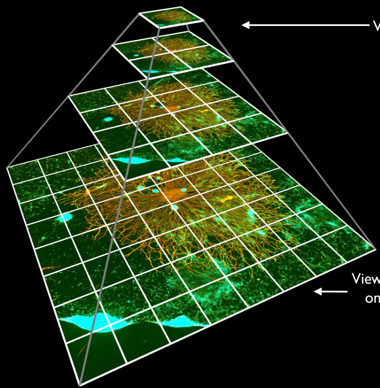

Photographic mosaics

- Very similar to gigapixel images except the photographs are generally more ad hoc.

Courtesy Ivan Zibra

Example

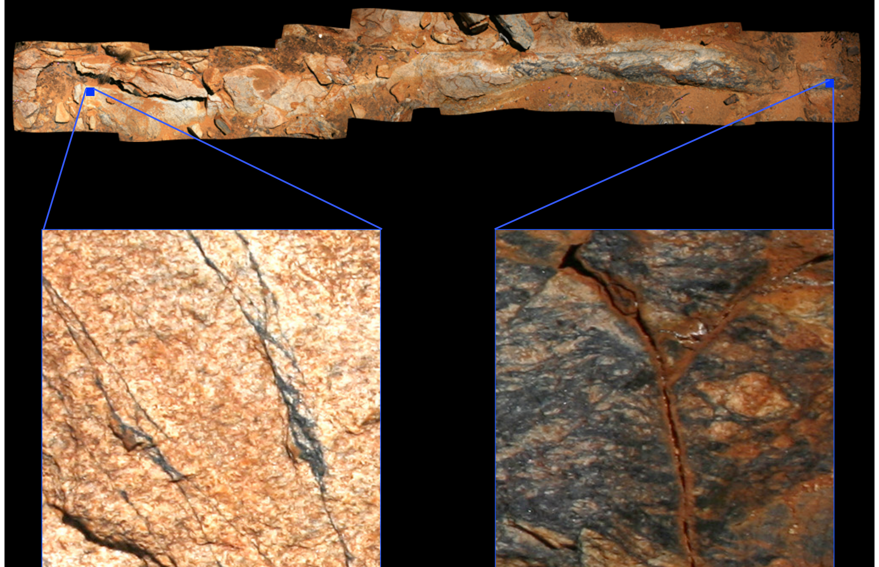

Courtesy Ivan Zibra

Virtual environments

- The types of imagery presented here is suited to creating virtual environments.
- These create a visualisation of the site and can be used as an interface to the research data.
- The 3D elements can be used to populate a 3D world.

Two examples: photographic tours and game engine example using Unity3D.

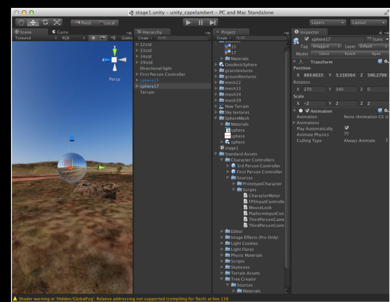

http://paulbourke.net/exhibition/Wanmanna/

Glyphs and information visualisation

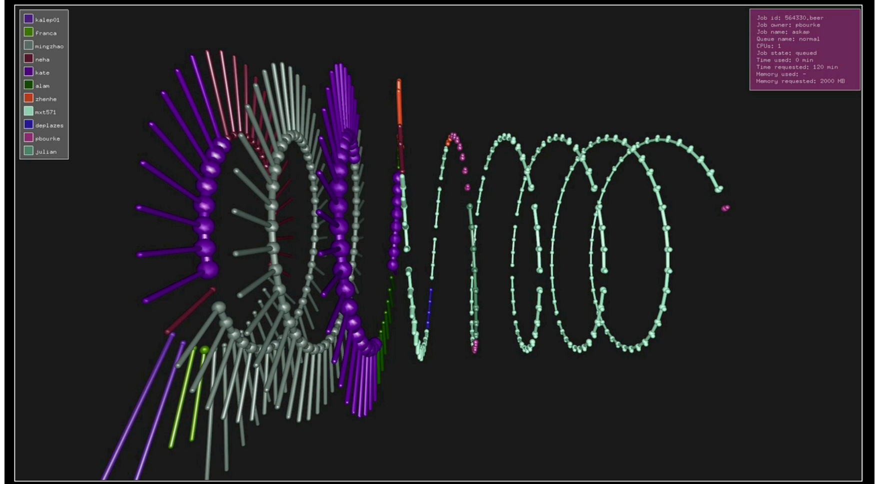

Mapping variables to geometry/graphics

- Visualisation is often about the mapping of variables to graphics or geometry.
- Scalars can be represented as colour.
- There is often a natural mapping between variables and geometry, for example: pretty obvious how one might visualise a surface.

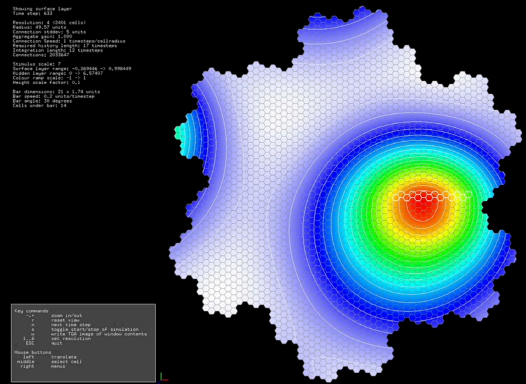

Glyphs

- Term given to the mapping of variables onto geometric objects.
- Ideally the structure/form of the glyph has some intuitive meaning.
- Common to map a quantity scalar onto the size of the glyph, obvious examples ...
  - Map direction onto arrows.
  - The strength of the direction (eg: velocity) onto length of arrow.
  - Scalars also mapped onto colour.
  ... and so on.
- Quite a common approach in flow visualisation.

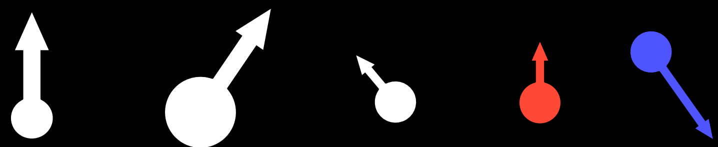

Glyphs

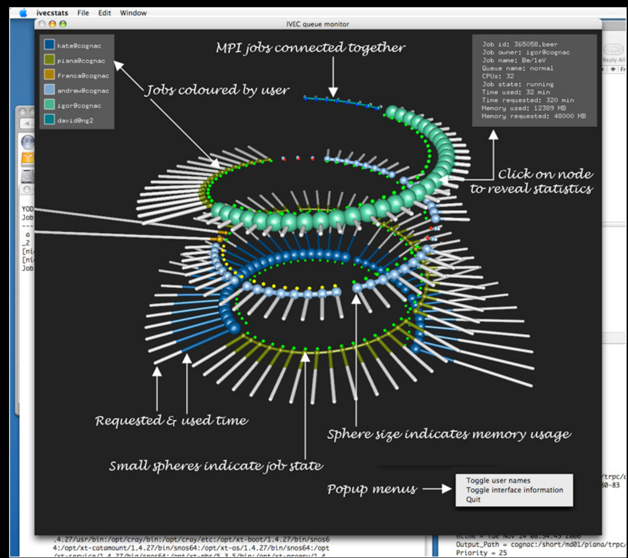

Example: Stock exchange

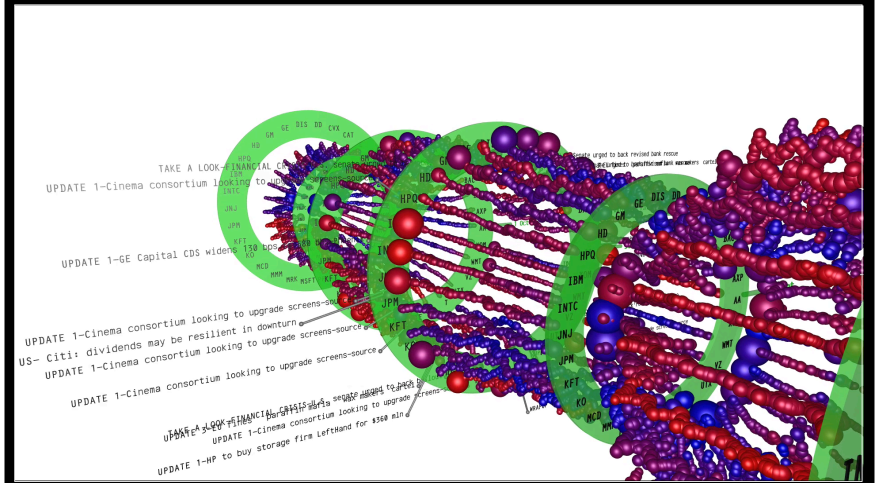

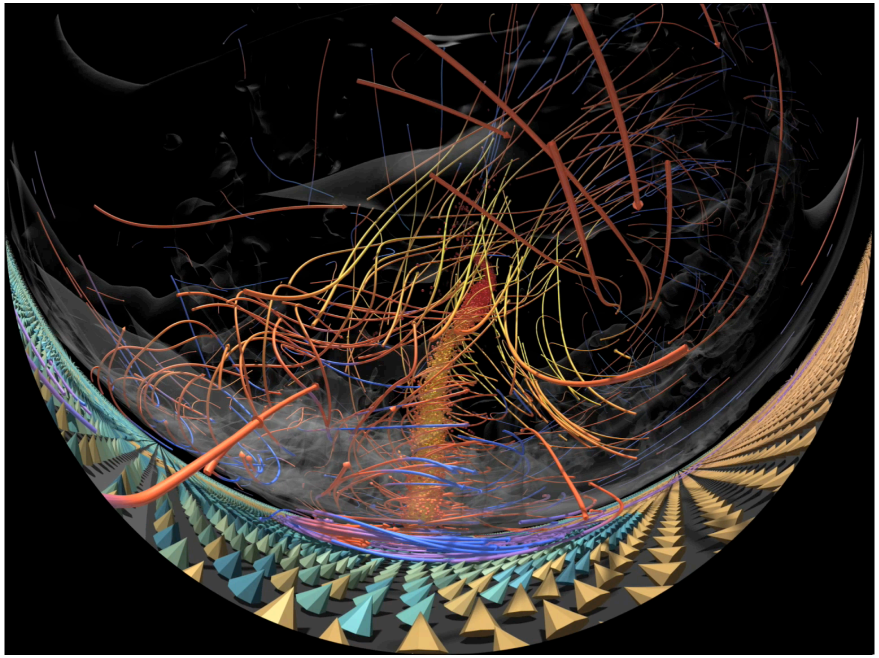

Mathematics

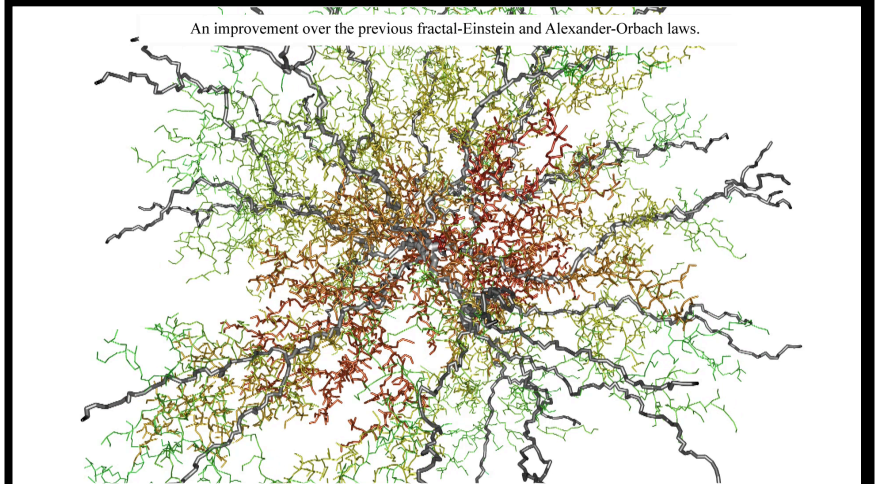

Courtesy Christophe Haynes and Anthony Roberts

Visualisation in Mathematics

- Widely used to visualise topology, higher dimensional space, knot theory, and so on.
- Not always so literal, for example to visualise abstract notions.

Minimal surfaces

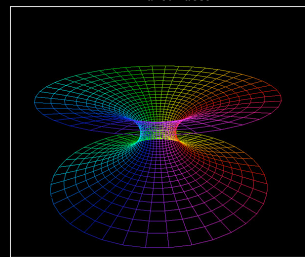

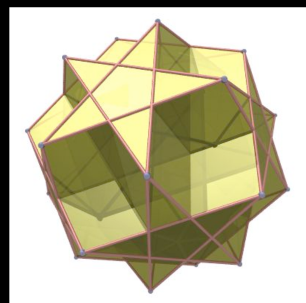

Polyhedra

Spherical harmonics

Examples

11 dimensional Calabi-Yau surface

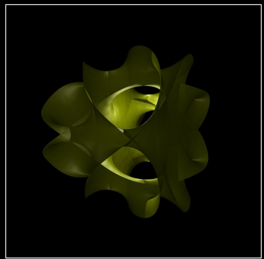

3D Truchet tiles

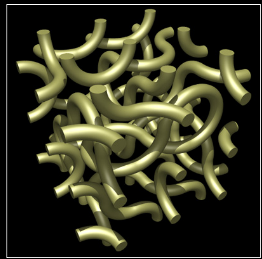

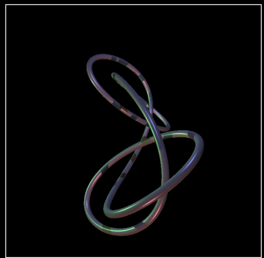

Knot theory

Borromean rings

Visualising connectivity in models of V1

Organisation of saturated coupling within and between local maps, and the approximate mirror symmetry of orientation preference in adjacent local maps.

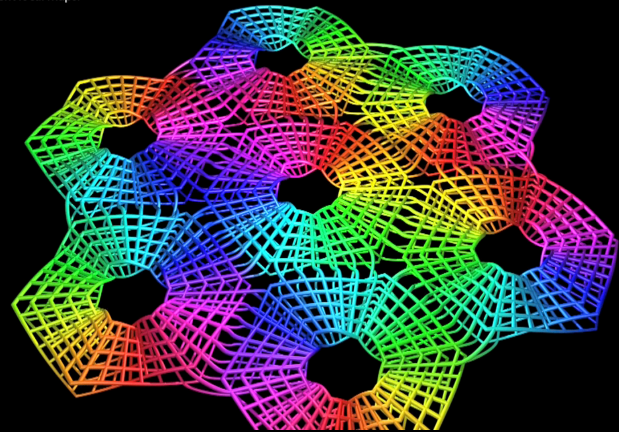

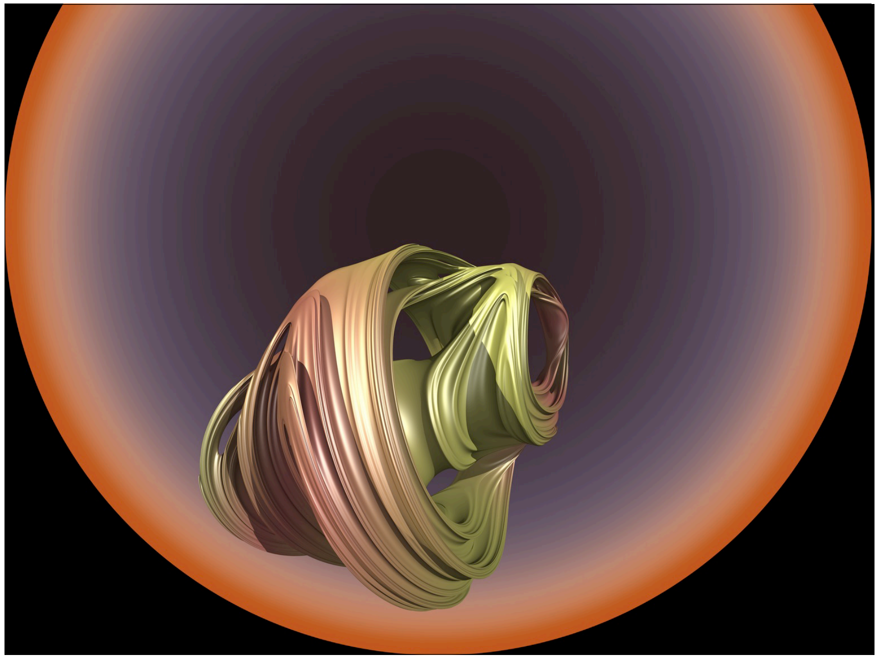

Visualisation in heritage

Courtesy Peter Morse

Virtual heritage

- Virtual heritage is conveying a sense of another place or time.
- Cultural heritage is conveying a sense of living in another culture, present or past.
- Often employs high resolution imaging, audio-visual, and immersive experiences.

360 x 150 degree video

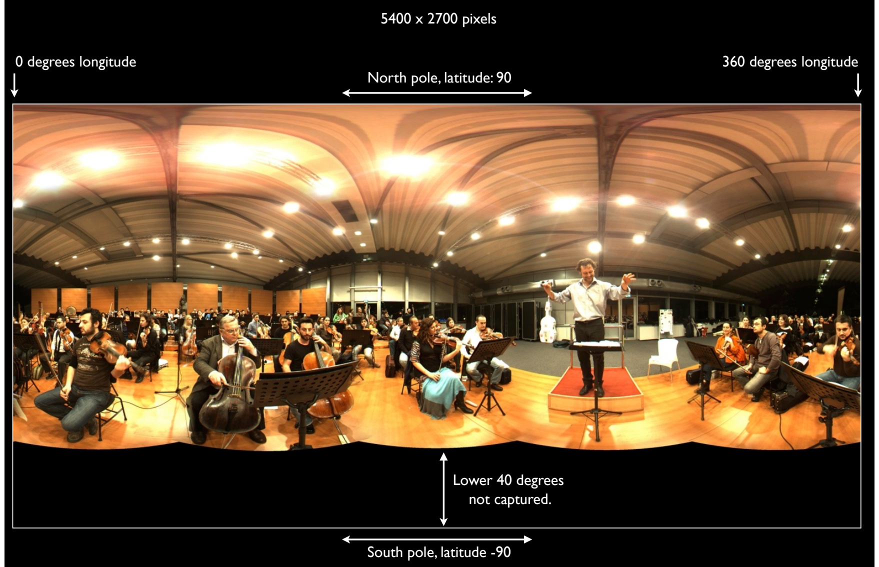

Cylindrical environments

Can derive cylindrical projection of any vertical field of view

Navigable movie in the iDome

Example: iJiao

- Showcase of cultural heritage of China.
- Venue: Main Gallery, Hong Kong Central Gallery.
- 360 degree video of various Taiping Qingjiao, also known as the Jiao festival.
- "The festivals, held throughout Hong Kong, appease the ghosts and give thanks to the deities for their protection. They take place every year or every five, eight, or ten years, depending on local customs. The religious rituals involved are meant to purge a community and prepare it for a new beginning."
  [Sarah Kenderdine]

Other applications

Other applications

Example: Ngintaka (Cultural heritage)

- One location for the Ngintaka story after the grind stone has been recovered is in a cave.
- The belly of Ngintaka.

360 x 150 degree video example

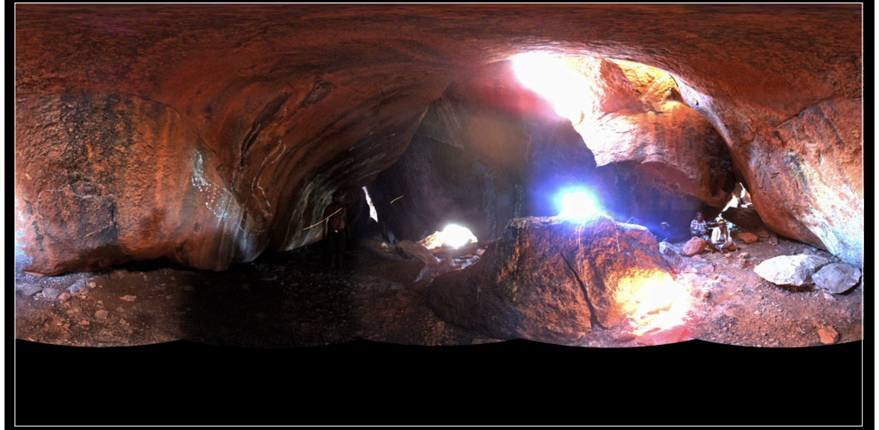

Navigable example in the iDome

Concluding comments

- Visualisation is used across a very wide range of disciplines.
- While visualisation is often associated with scientists and engineers who need to explore very large complicated data sets, increasingly many of the same techniques/algorithms are being used in the humanities.
- Visualisation has often employed custom hardware
  - historically always employed graphics cards in order to achieve real time performance
  - very often uses custom displays
- The outcomes range from
  - informing the researcher and conveying research to peers
  - to education in University courses, the general public, and school children
  - exhibitions in museums and art galleries