Glossary

Glossary of Graphics, Visualisation and High Performance Computing Terms

Mark W. Jones
Computer Science Department, University of Wales, Swansea
United Kingdom

adaptive - altering the focus of the computation in order to concentrate on parts that are more important to the solution. Computation takes place in areas of rapid change, leaving parts that contribute less significantly to be computed last or not at all. For example adaptive termination, adaptive rendering and adaptive re-meshing.

address space - memory within the computer that has contiguous addresses that are all visible. See also *virtual memory* and *shared memory*.

aliasing - a problem associated with the discrete nature of object edge representation. A binary decision occurs at each pixel – it is either inside the object or outside. This leads to a *staircase* effect or *jaggy* at the edge of objects. It is most noticeable during animation.

ambient light - the light that falls upon an object that is an amalgamation of light from all sources, for example, light sources and reflections.

anti-alias - the process of filtering an image so that sharp edges are softened by *averaging* neighbouring pixel intensities, in order to remove large intensity differences between neighbouring pixels.

area light source - a source of light that occupies some finite area in object space. See also *shadow rays*.

area sampling - values are evaluated at various points in space. In *unweighted* sampling, the sum of the samples is divided by the total number of samples. In *weighted* sampling values are multiplied by constants according to their position. Typically central values are given higher weightings in order to preserve their contributions.

backface culling - a method of reducing the number of polygons of a polygon mesh to be displayed, by removing those that face away from the viewer in a preprocessing step. Polygons facing away are selected by determining the dot product of the normal of each polygon with the view vector.

basic block - a self-contained part of a program that does not allow control to pass out of the block using jumps or loops.

bicubic surfaces - curved surfaces made up from parametric bivariate polynomial surface patches which are defined using three bivariate polynomials, one for each of x, y and z. Polynomials of various degrees can be used, but bicubic surfaces are made from cubic polynomials in both parameters.

bitmap - an image consisting of n pixels horizontally and m pixels vertically can be represented by $n \times m$ bits which determine whether a pixels is on (high intensity) or off (low intensity).

blobby models - surfaces created by isosurfaces of a certain threshold of field data, where each value has been calculated as the sum of the contribution of each primitive. The value contributed by each primitive is calculated according to its

strength and decay.

block - the act of suspending a process until some event occurs, for example a disk access.

boundary representation (b-rep) - the term given to the representation of a surface boundary using vertices, edges and faces to approximate the object to be represented. See also *polygonal meshes.*

bounding volume - a complex object is completely enclosed by a simpler object such as a sphere or box. Computation such as visibility takes place on the bounding volume and then only takes place on the more complex object if valid. See also *octree.*

broadcast - the method for communication in which a message is passed to all processes, each process responding as required.

bsp tree - subdivides space by successively dividing space into pairs of subspaces using a plane of arbitrary position and orientation. Each internal node has two pointers, one for each side of the plane. Leaf nodes contain those objects that occur in the subspace represented by that path through the tree.

bump mapping - the normals computed for some surfaces result in the surface looking very smooth. In order to introduce a roughness into the object, the computed normals are perturbed slightly in a random direction. The new normal is used to shade the object to give a stipple, or bumpy, effect.

classification (of elements) - Elements are classified according to position, value, and neighbouring values. In this way elements can be built into groups which represent objects upon which it is desirable for computation to take place collectively.

code partitioning - See *program partitioning.*

coherence - If there is a mapping between objects at a low level, and objects at a higher level, taking advantage of the coherence of the objects is achieved by using computation upon the groups to reduce the computation required for processing the objects in that group.

colour table - an array of (usually) RGBα tuples, where α is *opacity.* The array is indexed by the value of the objects to be displayed (for example pixels or voxels).

compositing - the process of merging together all the colours and opacities that contribute to one pixel of an image. The composition can take place in a back to front manner or a front to back manner.

computation-to-communication ratio - the ratio of time spent computing to time spent communicating. The ratio gives an indication of the effectiveness of the computational scheme being used.

constant shading - one normal is evaluated for a polygon to be shaded, and the pixels the polygon projects onto are all shaded the same intensity calculated by the shading model using some function of the normal.

convolution - a mathematical technique for combining two functions. In computer graphics its refers to the application of a function known as a filter function to the discrete function defining the image. Such filter functions can be defined to *anti-alias* an image, detect edges or enhance certain features.

CT scan - an image produced of a cross section of an object using a computed tomography (CT) scanner. The image describes the amount of X-ray absorbed by each of many discrete points on one slice plane taken through the object.

data distribution - the phase of computation during which the data is transferred to processors according to the *partitioning* strategy. The data should be distributed in such a way that the need for processes to request data is minimised, thus reducing inter-process communications.

data driven - a model of computation in which each process starts execution whenever it receives input data, or a change in input data. See also *demand driven.*

data flow - a model of computation in which the flow of data determines the order of computation. A process will execute only when it has the data available.

data parallelism - an operation can be applied to all members of the input data simultaneously.

deadlock - a problem that occurs when each process has been blocked, waiting for some other process to finish its calculation, which has also been blocked. See also *lock* and *mutual exclusion.*

decomposition - the act of dividing a problem into components that allow it to be distributed separately. Each component can then be calculated or acted upon separately. For **dynamic decomposition** the problem is adaptively divided during computation in response to changing conditions, usually to maintain a good *load balance.* In **static decomposition** the problem is divided before the start of computation, and the decomposition remains unchanged throughout. See also *data distribution,* and *partitioning.*

deformable models - object descriptions that not only represent the physical appearance of the object but also the physical characteristics of the object. Such objects are often described by particles connected by *springs* which allow the objects to deform under external forces, and revert to their original state using internal forces. Such deformable models are ideally suited to modelling elastic objects acted upon by physical forces such as gravity, friction, compressible force, extensible force and even heat.

demand driven - a model of computation in which each process computes only when its output is required. If it requires some input to perform its calculation, the process will request that input. See also *data driven.*

dependence - a process B is dependent upon a process A if the result of executing A and B changes according to the order in which they are executed. B is dependent upon A if it uses values calculated by A.

depth cueing - the attenuation of each pixel's colour in an image due to the distance of the object from the pixel using the view coordinates. The pixel's intensity, I is

$$I = s_0 I' + (1 - s_0) I_{dc}$$

$$s_0 = \frac{Max - z_0}{Max}$$

where Max is the maximum depth of any part of the object from the view plane, and z_0 is the depth of the part of the object from the current pixel. I_{dc} is the intensity of the depth cue colour and I' is the calculated intensity of the surface.

depth only shading - the distance z, is calculated from each pixel in the image to the closest object in the scene to be displayed. The intensity of the image pixel I is $I = 1 - \frac{z}{Z}$ where Z is the maximum distance any object can be from any pixel.

directional light source - see *parallel light source.*

distributed computer - a computer made up from many computers that may or may not be physically separated. One example would be that of using an network of computers to achieve some task.

distributed memory - memory that is physically distributed. This fact may or may not be hidden from the user, i.e. the user may only be aware of one contiguous *address space* or may be aware of many separate address spaces.

extended light source - see *area light source.*

farming - processors compute a number of jobs by obtaining a job from some scheduling process, computing the job, and then sending the results to some global receiving process. Upon completion of a job, a processor will obtain another job from the scheduling process.

flops - floating-point operations per second - a measure of processing power.

footprint - the area covered by the projection of an object when it is projected onto a plane. For example, a spherical voxel will have either an elliptical or circular footprint.

fork - the creation of a process that is a copy of its parent. See also *join* and *spawn.*

global illumination - the light falling upon any object in the scene is accurately calculated from the contribution of light from every other object in the scene.

Gouraud shading - a technique for determining the intensity of the shading for a polygon. The normals are calculated at the vertices of the polygon, and are determined along the edges of the polygon using linear interpolation. Intensities at the edges of the polygon are calculated by the shading model using the interpolated normal. Intensity within the polygon is calculated according to the linear interpolation of the intensities at the edges of the polygon.

gradient shading - a method of using a z-buffer to calculated surface normals which can then be used to compute pixel intensities using a shading model. The z-buffer is an array $z(x, y)$ of depths for each pixel (x, y). The gradient at (x, y) can be calculated using

$$\nabla z = \left(\frac{\partial z}{\partial x}, \frac{\partial z}{\partial y}, 1 \right)$$

At pixel (x, y) $\frac{\partial z}{\partial x}$ can be approximated by the backward difference

$$\delta_b = z(x, y) - z(x - 1, y)$$

by the forward difference

$$\delta_f = z(x + 1, y) - z(x, y)$$

by the central difference

$$\delta_c = \frac{1}{2}(z(x + 1, y) - z(x - 1, y))$$

or by a weighting of all three, and similarly for $\frac{\partial z}{\partial y}$.

granularity - the size of computation carried out between successive communication events. Course grained processes carry out a sizable computation between communications, whereas fine grained processes carry out very few operations. See also *computation-to-communication ratio.*

hidden surface - the parts of a surface that are not visible to the viewer because they are obscured by other parts. Hidden surface removal is the process of determining the hidden surface so that it may be removed from the viewing pipeline.

hypercube - a *topology* in which each *node* is at the vertex of an n-dimensional cube and is connected to its neighbouring vertex nodes. See also *mesh connected* and *processor array.*

image partitioning - the process of dividing an *image* in order for computation to be carried out separately upon each image partition. See also *partitioning.*

inbetweening - the process of interpolating objects between two key objects. If the two key objects (for example keyframes for images) are V_0 and V_k, the process of inbetweening produces $k-1$ interpolated objects – $V_1, \ldots, V_{k-1}$, which smoothly transform object V_0 into object V_k according to some criteria.

intensity - pixels, light sources and discrete points have intensity associated with them. It is usually either a grey-scale value representing brightness, or a triple (I_r, I_g, I_b) where I_λ represents the intensity at that wavelength (r = red, g = green, b =blue). Often an additional intensity (I_α) indicates the *opacity* of the point.

interpolation - the function of determining a new vector between two existing vectors using the existing vectors and the distance between them and the vector to be determined. For instance the vector could be as complex as a whole image, or as simple as a single real number. Interpolation can be linear, where the in-between vectors vary linearly according to distance, or some higher order function.

isosurface - a surface of equipotential. For a function $f : \Re^i \rightarrow \Re$, the τ-isosurface, S, is defined as

$$S = \{x^i : f(x^i) = \tau\}$$

Every point on the surface has a function value of τ, and no other points in the domain have a function value of τ. For $i = 2$, S would be referred to as isolines or contours.

join - the act of waiting for offspring that were *forked* to finish execution. See also *spawn.*

latency - the time to initiate an operation that is independent of the size of the operation. For example, the latency of a message-passing system is the time taken to deliver a message.

light source - the set of points in space that illuminate the surrounding space. A light source is usually associated with a triple (I_r, I_g, I_b) where I_λ ($\lambda = r, g, b$) is the intensity of light for that particular wavelength (r=red, g=green, b=blue).

linear speed-up - *speed-up* that is directly proportional to the number of processors used. See also *scalable.*

load balance - an indication of how well a job is distributed among available processors. Perfect load balance is achieved when all processors have exactly the same amount of work to do. It can be measured using the difference of the minimum amount of time a processor is active for and the maximum amount of time a processor is active for.

locality - an indication of the amount of computation that can be done by a processor using data held by, or close to, the processor. It can be measured by the number of non-local data requests.

lock - once achieved a process is given exclusive permission to use a resource. See

also *deadlock* and *mutual exclusion.*

mach bands - the term given to the discontinuous lighting effect observed on polygon boundaries in a polygon mesh description of a surface when a lighting model such as Gouraud shading is used. For the case of Gouraud shading, pixel intensities are C^0 continuous across polygon boundaries, but not C^1 continuous. The eye interprets this discontinuity in the form of *banding*. A higher order shading model such as Phong shading creates C^1 continuity which removes much of this effect.

marching cubes - a popular algorithm for determining isosurfaces in 3D. Function values are evaluated at regular discrete points to make a 3D grid of data values (voxels). Eight neighbouring voxels (four on one slice of data, and four aligned in an adjacent slice) make up a cube, with which the surface can be determined using look up tables for each of the possible cases. Since each of the eight voxels can be inside or outside of the surface, there exist 256 cube configurations. The look up table indicates the triangles to be added to a triangular mesh, with their vertices interpolated from known voxel positions and values.

mesh connected - a popular *topology* in which identical processors are laid out on a $m \times n$ grid, and each processor communicates with each of its four neighbours in the north, south, east and west directions. See also *processor array* and *hypercube.*

message passing - the method by which processors communicate with each other by sending discrete packets of data.

MIMD - Multiple Instruction, Multiple Data - many instructions are carried out simultaneously on many data sets. See also *SIMD.*

motion blur - the process of simulating the dynamics of moving objects by blurring them in the direction of motion. This occurs in the reality of photography or cinematography where the camera shutter is open for a finite length of time, and objects have a chance to move, and create multiple images (blur) on one frame. Simulating this from computer images, both stills and animations, adds to the realism of the image.

MRI scan - an image produced of a cross section of an object using a Magnetic Resonance Imaging (MRI) scanner. The image describes the disturbance of the magnetic field by each of many discrete points on one slice plane taken through the object.

multi-tasking - the method of executing many processes on a single processor. Each processor is given a fixed amount of time (time-slice) before it is swapped out, ready for the next process to be swapped in.

multi-processor - a machine which processors can execute simultaneously on different instruction streams and only have access to a single *address space.*

mutual exclusion - an event during which at most one process can execute. See also *lock* and *deadlock.*

node - a *processor* within a *topology.*

Nyquist rate - the lowest sampling rate that can allow a proper reconstruction of a signal. If the signal has a component with frequency f_λ which is the highest-frequency component, then the *Nyquist rate* is $2f_\lambda$. If less samples are used, the reconstructed signal may not be correct, and can represent high frequency components as a lower frequency signal. This misrepresentation is known as *aliasing.*

octree - a tree data structure where each node has up to eight children. It is most often used to partition 3D structures since the children of each node partition the space represented by the parent node. The process to be computed can be computed on a node. If successful the process is carried out on each of the eight children in turn. If unsuccessful, the part of the volume represented by that node can be ignored. For example – octrees are used as bounding volumes in ray tracing. If the ray being tested successfully intersects a node, it is tested against each of the children of that node, or if a leaf node, it tests for intersection with the objects within the volume bounded by the leaf. If the ray does not intersect the node, all the objects within that node are removed from the intersection computation.

opacity - usually a value between 0 and 1 which represents how dense an object is to light. If the opacity is 1 the object is opaque and obscures everything behind it, if the opacity is 0 it is completely transparent, and cannot be seen. Values between 0 and 1 represent, proportionally and linearly varying, the visibility of the object and objects behind it. See also *transparency*.

parallel balance point - the point at which increasing the number of processors available actually increases the time of the calculation. This could be as a result of increased communication cost or start up time.

parallel light source - a light source from which light falls on all objects as if from the same direction. It occurs when the light is infinitely far away from the scene.

parallel projection - the production of images wherein the projectors are all parallel to each other and they are defined by a single direction vector.

parallelisation - transforming a serial computation to a parallel computation.

particle systems - composed by a collection of particles that evolve over time. The particles have attributes associated with them such as mass, charge, or colour. The system evolves according to either random events, such as dying, introducing new particles and random motion, or by following equations for motion such as following trajectories according to initial velocity and gravity and following velocity vectors in numerical simulations.

partitioning - method for dividing objects amongst available processors. See also *decomposition*, *image partitioning*, *program partitioning* and *scene partitioning*.

perspective projection - the production of images using perspective. All the projectors pass through one point known as the centre of projection.

photorealistic - artificial images that closely resemble what you would see in reality, if positioned at the appropriate viewpoint in space and looking at the scene domain from the appropriate angle.

physically based modelling - creating an accurate behavioural model of an object based upon its physical properties and those of its surroundings. The motion and object shape are determined by the forces acting upon the object, and are calculated using interesting mathematics.

pixel - (picture element) is the name given to a value in a 2D image.

point light source - source of light in which all light rays radiate outwards from a single point. Light bulbs can be best approximated by point light sources.

polygonal meshes - a form of object representation. The object's surface is approximated by a set of planar polygons in which each edge is used by at most two polygons, and each polygon is a closed set of edges. A visualisation of the object

represented by the mesh is produced using any of a number of *rendering* techniques.

PRAM - parallel random access machine - a theoretical model of parallel computation in which finitely many processors operate on a single *address space synchronously.*

primary rays - originate from each pixel in the image in the direction of the view plane normal for **parallel projection**, or from the centre of projection in the direction of the pixel in **perspective projection**. The ray intersects surfaces that lie along the ray, and thus spawns *secondary rays* to compute the contributions to the pixel's intensity due to object *reflection*, *refraction* and *shadow*. These contributions are combined with the object's intensity to give the intensity of the pixel.

private memory - the memory available, and only available, to each processor.

process - a set of program instructions to be executed.

processor array - a computer composed of a regular grid of processors controlled by a single processor. Processor arrays are used to achieve data parallelism. See also *topology*, *hypercube* and *mesh connected.*

program partitioning - the method of distributing *basic blocks* amongst available processors. Once *dependency* has been calculated it is then possible to execute those blocks that can be executed in parallel. See also *partitioning.*

radiosity - models the lighting in a scene by using an analogy of thermal dynamics. The light in a closed scene follows the law of conservation of energy, and all light must be accounted for. The light at any surface is the sum contribution of all the light energy falling on that surface from all light sources and object reflections. Radiosity methods calculate the interaction of light in a view independent pre-processing step, and the resulting scene description can be rendered from this model using conventional techniques. Radiosity accurately models *global illumination* and removes the need for the *ambient lighting* term.

ray casting - the process of sending a primary ray from a point in space into a scene with direction. Upon encountering an object the ray terminates, and the intensity of the point is calculated according to the shading model.

ray tracing - is the process of sending a primary ray from a point in space into a scene with direction. Upon intersection with an object secondary rays are spawned in order to determine all contributions to the intensity of that point. Secondary rays determine contributions due to *reflection*, *refraction and transparency*, *shadows*, and *light sources.*

reconstruction - the process of determining an object from a partial description of the object. For examples, reconstruction of polygonal meshes from contours and reconstruction of objects from laser range data.

reflection rays - rays spawned from objects which are partial or total mirrors, in order to determine what is visible at the mirror. The rays are treated like primary rays, and may spawn other rays, including reflections rays, upon hitting an object. The origin of the ray is the intersection point with the object, and the angle of incidence is calculated using the direction of the incoming ray, and the surface normal, with the intersection point being the point of incidence.

rendering - the process of producing an accurate impression of a scene using any of the shading models, and any of the rendering techniques such as ray tracing or scan conversion.

resident set - that part of the process that is currently in main memory. See also

working set.

routing - the function of moving a message from its source to its destination.

scalable - to produce an increase in performance proportional to an increase in size. See also *linear speed-up* and *speed-up.*

scan conversion - the process of calculating pixel intensities from object descriptions in an ordered pixel by pixel, line by line way for each object.

scan line - one horizontal line in an image, derived from the line that a cathode ray tube sweeps (scans) across during anyone pass.

scene partitioning - the method of dividing objects in a scene and distributing them amongst processors. Each processor can operate upon the objects available within its private memory, but must communicate in order to fetch objects stored on other processors. See also *partitioning.*

scheduling - the process of deciding the order of process execution, and mapping those processes to available processors.

secondary rays - produced at the point of surface intersection during ray tracing. They are created in order to find the exact colour of the intersection point due to contributions from all objects that influence the light at that point, for example, light sources, shadows created by other objects and partial transparency of the object itself and objects that lie on the path between it and a light source. These rays are determined in the same way as primary rays except that the direction is some function of the view normal and surface normal, and the resulting colours contribute partially rather than totally to the final pixel colour.

segmentation - the process of determining similar objects with respect to certain defined parameters and grouping them together.

serialise - the process of placing processes in a sequence to be executed, even though they may be carried out concurrently.

shadow rays - rays spawned from objects in order to determine whether or not any objects occur on a path between that point and each light source which may place the point in shadow. The origin of each shadow ray is the intersection point with the object, and the directions are the directions to each light source. If the light source is an area light source, one method is to trace n rays to random points on the area, and if m rays get through, the fraction of light falling on the point from that light source is $\frac{m}{n}$.

shadows (umbra/penumbra) - created by objects obscuring light sources from an object. In the case of area light sources the shadow will be a transition from points that can see most of the light, to points that can see no light at all. Those that are in complete shadow are said to be covered by the umbra, and those that are in partial shadow are in the penumbra.

shared memory - memory that appears to the user to be a single *address space* that is available to a group of processes. See also *virtual memory.*

SIMD - Single Instruction Multiple Data - a single instruction stream is applied concurrently to multiple data sets. One instruction is processed by each processor on its data before passing onto the next instruction. See also *MIMD.*

spatial partitioning - the process of dividing space into smaller subspaces. The reasoning behind this is the fact that processes can be carried out on fewer objects in the particular subspace they relate to. See *octree* and *bsp tree* for examples of

spatial partitioning techniques.

spatial subdivision - see *spatial partitioning.*

spawn - the creation of a process by an existing process. See also *fork* and *join.*

specular reflection - the highlight that can be seen on a shiny surface from particular angles. Perfect mirrors reflect light in only the direction of the reflection according to the angle of incidence, whereas any other shiny object reflects light unequally in different directions, and so the highlight is seen from different angles.

speed-up - the ratio of two program execution times that indicates the increase in speed achieved. If a program takes time S_1 to run on one processor, and time S_2 on two processors, the speed up is $\frac{S_1}{S_2}$. See also *linear speed-up* and *scalable.*

splatting - the process of projecting voxels onto an image plane, throwing them against the image so that they *splat* their colour and opacities on the image. The distribution of a voxel's energy is calculated, and the energy is added to each pixel it projects onto by compositing the colour and opacity at that point in proportion to the energy.

spline - the name given to a line (in 2D) or a surface (in 3D) that has been produced using knots (control points). To describe an object, a small number of knots are stored which can then be used to generate a smooth surface or curve. Spline curves can either be *interpolated* from the knots, i.e. the surface passes through each knot, or they can be *approximated* to the knots, i.e. the curve passes close but not necessarily through the knots.

start up cost - the time to initiate any process. For example this could be the time to distribute the data to all processors.

stereo images - images that have been created from two images produced of the same scene at slightly differing angles and view points, in order to approximate the positions of the viewer's eyes. The images can be combined as one using one wavelength for one image, and a different wavelength for the second image (usually red and blue), and then viewing the composite through the appropriate glasses. Images can also be produced in the form of random dot stereograms, or by displaying the two images alternately, and relying upon additional hardware (glasses) to be able to synchronise the left eye with the left image, and the right eye with the right image.

subsampling - the process of obtaining values more coarsely than is usual, and interpolating the inbetween values. For example, instead of tracing a ray for each pixel in an image, rays are traced coarsely over the whole image with extra rays being traced in the regions of high change. Other pixels are interpolated from known pixels in the process known as *adaptive rendering.*

supersampling - the process of obtaining values more finely than is usual, and averaging neighbouring values to produce a more accurate sample of the point being determined. For example, a form of *anti-aliasing* - is to supersample the image, and then average neighbouring values to produce a smaller, but more accurate image.

surface normal - the vector which points from a particular point on the surface, outwards away from the object so that the angle between the normal and the tangent to the surface in any direction, is the same in every other direction.

synchronisation - causing some processes to suspend execution in order that all processes reach a given point in their execution at a given time.

threshold - a value which represents something interesting in a particular domain. For example, in 3D *isosurfacing* the threshold defines the surface in which we are interested in, and is calculated such that all points on the surface have a function value equal to the threshold.

topology - processors and their connections are laid out according to some topology. Topology more generally refers to a family of graphs and as such each processor is deemed to be a node, and each communication link is considered to be an edge. See also *mesh connected*, *hypercube* and *processor array*.

transparency - the inverse of *opacity*.

vector processor - a computer that performs operations on vectors (arrays) of data values.

vectorise - to convert a sequence of operations applied to data values to one instruction applied to a vector.

viewing pipeline - refers to the process of transforming an object description (usually a mesh of vertices) to screen coordinates. It typically involves the view transformation of the object coordinates, clipping the object to the visible volume, and projecting the object. Often additional steps for mapping these coordinates to the screen coordinates are included.

virtual memory - the user perceives virtual memory as addressable memory. Any data not available in main memory is retrieved from storage and placed in main memory. See also *address space*.

volume rendering - the process of obtaining images from three dimensional volume data by treating the data as cloudy material. Each value in the data has a colour and opacity assigned to it during a classification stage. Rays are cast for each pixel into the volume, and the colour and opacity are sampled at evenly distributed points along the ray. These samples are composited using standard techniques to produce the accumulative colour and opacity reaching the pixel. Images produced by this method usually have several differently colour semi-transparent surfaces overlaid on an opaque surface, for example, semi-transparent skin overlaid on opaque bone for an image of a CT scan.

volume visualisation - the general term given to the process of displaying three dimensional volume data, and encompasses the methods of isosurfacing, volume rendering and splatting.

voxel - (volume element) the name given to a value in 3D space.

voxelisation - the process of producing 3D data from an object description, such as polygonal meshes, so that each voxel has a value. The data can then be rendered using a volume visualisation technique to produce an image of the original object.

working set - the set of pages that have been referenced most recently for a process. See also *resident set*.

Author Index