VICTORIA UNIVERSITY OF WELLINGTON

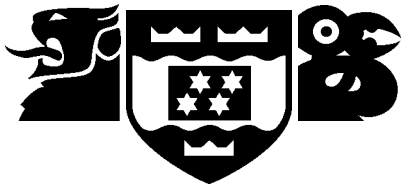

Department of Computer Science

PO Box 600
Wellington
New Zealand

Tel: +64 4 471 5328
Fax: +64 4 495 5232
Internet: Tech.Reports@comp.vuw.ac.nz

Simulated angiography

Peter Hall

Technical Report CS-TR-94/15
July 1994

Abstract

A method that simulates angiographic procedure is presented. The intended purpose is to provide simulated angiograms so that mechanisms for reconstruction the vasculature can be rigorously tested. Three elements are modelled; the vasculature and blood flow, the x-ray device, and the contrast agent and its injection. These modelling elements are specified by parameters supplied by the user, leading to a versatile simulation that captures many of the salient features of angiography. The simulation yields sets of time sequenced angiograms - an animation. Each set depicts flow of agent through the vasculature from one point of view. Such sets may be acquired either simultaneously, or through independent injection events. Results in the form of simulated angiograms are presented.

Author Information

Peter Hall is a lecturer at Victoria University. His principal research interests are scientific visualisation, and computer vision.

Simulated angiography

Peter Hall
Department of Computer Science,
P.O. Box 600,
Victoria University of Wellington,
Wellington,
New Zealand
Email: peter@comp.vuw.ac.nz

Abstract

A method that simulates angiographic procedure is presented. The intended purpose is to provide simulated angiograms so that mechanisms for reconstruction the vasculature can be rigorously tested. Three elements are modelled; the vasculature and blood flow, the x-ray device, and the contrast agent and its injection. These modelling elements are specified by parameters supplied by the user, leading to a versatile simulation that captures many of the salient features of angiography. The simulation yields sets of time sequenced angiograms - an animation. Each set depicts flow of agent through the vasculature from one point of view. Such sets may be acquired either simultaneously, or through independent injection events. Results in the form of simulated angiograms are presented.

1 Introduction

This paper shows how to simulate x-ray projections of blood vessel networks (*vasculature)*. X-ray projections of vasculature are called *angiograms*, and are routinely acquired for clinical use. In order to make the blood vessels visible a *contrast agent* - fluid that partially absorbs x-ray - is injected into the blood flow. Angiograms appear darker where the x-ray has travelled trough more of the contrast agent. The agent moves with the blood stream, driven by a pressure gradient composed of a cyclic part due to the heart beat, and a transient part due to the injection. A time-sequenced set of angiograms is acquired from a given point of view as the contrast agent courses through the vasculature. These *cineangiograms* demonstrate temporal information that is useful to clinicians. Often, sets of cineangiograms are acquired separated by an angle of about ninety degrees; these are *biplane cineangiograms*. At other times, the angiograms are separated by an angle of about four degrees, giving rise to *stereo cineangiograms*. Additional cineangiograms from different points of view may be acquired in either case. Distinct sets of cineangiograms may be captured simultaneously, or separately. An independent injection event is needed for each separate capture, and the x-ray device may be moved into a new position for each injection. The simulation to be presented can reproduce all of these features.

The motivation for the simulation arose in the context of reconstructing vasculature from biplane cineangiograms. We have developed a novel mathematical system that learns the information it needs for reconstruction, Hall and McGregor (1993). The reconstruction method needs rigorous practical testing; producing a simulation is one step in our test strategy. Other systems designed to reconstruct vasculature have also used computer simulation for test purposes. The principal

contributions made here are (1) the introduction of a modelling scheme capable of representing complete vasculature, (2) the animation of contrast agent as it flows through the vascular network, and (3) a rendering mechanism that is inexpensive but effective.

Background, in section 2, precedes the method, given in section 3. Animation implies the production of time sequenced images, and such images are shown as results, section 4. The paper concludes in section 5.

2 Background

The notion of producing x-ray simulations is not new. Accurate simulation of the x-ray process is possible using radiation transport theory. General methods for a solution are given by Lewis and Miller (1984). A method for x-ray simulation by ray-tracing voxel data sets is presented by Siddon (1985), and for ray-tracing constructive solid geometry models by Li and Williamson (1992).

The geometric modelling primitive used here is a generalised cylinder. In a computer graphics context Bronsvoort and Klok (1985) produced a mechanism for ray-tracing the surfaces of generalised cylinders. This algorithm may be adapted to ray-trace the interiors of the cylinders. However, ray tracing such objects is very inefficient. In response to this Bronsvoort (1992) published a mechanism for rendering generalised cylinders by point sampling its surface. Unfortunately this rendering method is not suitable here, for it models reflected rather than transmitted radiation. The rendering method used here point samples the interior volume of generalised cylinders.

Specific to angiograms, simulations vary in their aim as well as in degree of detail. Barba *et al.* (1987) seek to reconstruct vessel lumen. They model the interior boundary of vessels. These boundaries may be crescent-shaped so that internal blockages can be modelled. Their simulations seem to model these blockages by removing portions from the simulated projection rather than model them in three dimensions. Gaussian noise is used to degrade individual projections. Angiograms of single vessels are simulated by Kitamura *et al.* (1988) in order to test vessel tracking algorithms; tracking operates on a pair of angiograms to produce a trajectory in three dimensions. Rolland and Puff (1993) also model single vessels but depict more than one such vessel in any given angiogram. Tree structures of vessels are modelled by Fessler and Macovski (1991) for testing reconstruction from Magnetic Resonance Angiograms. These authors allow for effects such as point blurring, and noise. None of these simulations approach the issue of animating blood flow.

Modelling the flow of blood through a vessel network is a very complicated task, texts such as Fung (1984) make this clear. Difficulties include; the viscoelastic properties of blood, the elasticity of vessel walls, the time-varying pressure gradients, branching structures, and turbulent flow. Particular models seem to be designed to allow researchers to investigate specific aspects of flow. Models of blood flow along a single vessel with branches do exist, see Anliker *et al.* (1978), and models of branches exist, Resch and Küster (1994). Pulsatile flow in a single vessel has also been the subject of mathematical investigation, see He *et al.* 1993. Melchior *et al.* (1992) model the whole vascular system, but not in a form useful here. These authors also provide a useful overview of the large body of literature in the field. The task here is not only to simulate flow but to animate it, and no literature relating to that task has been found by the author.

3 Method

Despite the many simplifying assumptions made here, we have produced a model that simulates many features of angiography. A detailed description of the method by which flow is simulated and animated, including explicit reference to assumptions made, can be found in Hall (1994b). Rendering issues are discussed, in detail, in Hall (1994a). The techniques developed in those papers to produce are put to use here.

The simulation models: (1) the x-ray device, (2) vasculature and blood flow through it, and (3) the contrast agent. Each will be discussed in turn. Animation is a requirement of the simulation, and an outline for its control appears in section 3.4.

3.1 The x-ray device

The x-ray device consists of one or more gantries, as in figure 1. Each gantry consists of an x-ray source and an image plane. An x-ray source is modelled as a point emitter, radiating isotropically. The image plane has particular position, orientation, and size. The image plane is decomposed into pixels.

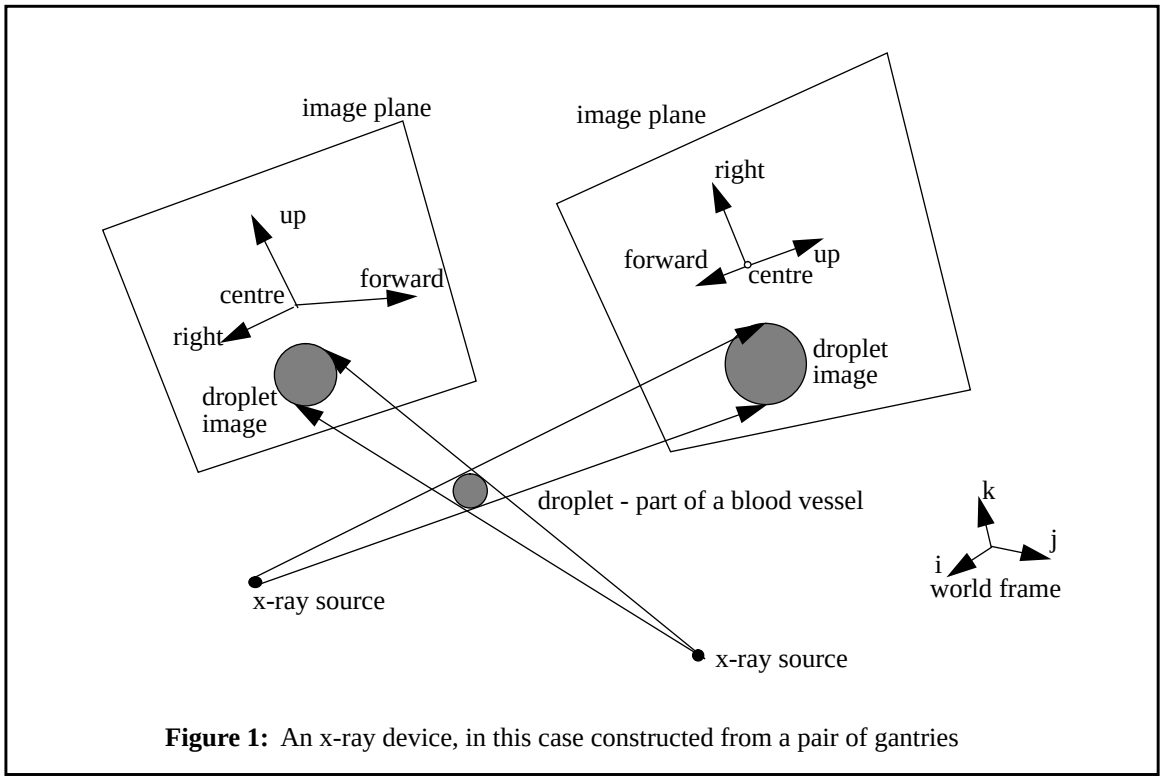

Figure 1: An x-ray device, in this case constructed from a pair of gantries

The geometry of each gantry is fixed, and the geometrical relation between gantries is also fixed. Should the x-ray device be moved then each of gantries moves also; the x-ray device is a rigid object. The x-ray device simultaneously acquires images from each gantry.

Image acquisition is based on sampling points in three dimensions, sample points appear on a regular, cubic lattice. By decomposing the vessels into a series of small cones and sampling only

in the bounding box of each cone, the number of points sampled can be greatly reduced. Each sample point is tested against the object being rendered. If the point is found to be interior to the object then it is considered to be the centre of a small but finite droplet of absorbing material, and the droplet is projected onto the image plane. Pixels in the image plane accumulate path-length between it and the x-ray source. An image of a droplet makes an additive contribution to the path-length value of every pixel is covers. Once all objects have been projected the accumulated path-lengths are used to compute the fraction of radiation that is transmitted.

Suppose a droplet of radius R projects onto an image plane, with centre (u,v) and radius r. A pixel at (i,j) is covered by the droplet if $|(u,v) - (i,j)| < r$. For a single droplet, the path length of a ray is defined to be the product of the length of the ray through the droplet, x, and the attenuation coefficient of the droplet, μ. The total path length, p(i,j), of the pixel is the sum of such products;

$$p(ij) = \sum_k x_k \mu_k \tag{1}$$

where k ranges over every droplet that covers the pixel. The distance a ray travels through a droplet can be found by elementary geometry; $x = 2(R^2 + b^2)^{1/2}$, where b is the closest distance that the ray passes to the centre of the droplet. The distance b is related to the distance $|(u,v) - (i,j)|$ by a scale factor, R/r, so a look up table may be used to compute a suitable value of s. Alternatively, a simpler function can be used as the convolution kernel of the sample point on the image plane. The final intensity is the pixel, I(i,j) is computed by

$$I(i,j) = I_0 e^{-p(i,j)} \tag{2}$$

where I_0 is the initial intensity of the radiation.

Notice that this method accounts for attenuation at a single wavelength only. Effects such as scattering and refraction are ignored. The principal advantage of this method is its time cost in producing multiple images; any point needs to be tested for membership to the object being rendered once only - as many projections as needed can the be produced at little extra cost. This contrasts with ray tracing where each projection must be independently rendered.

In order to acquire a time sequenced set of images the x-ray device must be operated at regular time intervals. The simulation is paused to permit capture of a frame, motion blur is not modelled. The mechanics of pausing the simulation is discussed in section 3.4, below.

3.2 The vasculature

Vasculature are modelled from a collection of vessel segments and furcations. Both vessel segments and furcations are constructed from tubes. A tube has geometrical properties, such as shape, and physical properties such as conductance to flow. A complete vasculature is depicted in figure 2. The model vasculature must contain open ends, this allows fluid to flow through it. There is cyclic flow because the heart is not modelled directly, though pressure fluctuations that arise from it are represented. The component parts of the model vasculature are explained next.

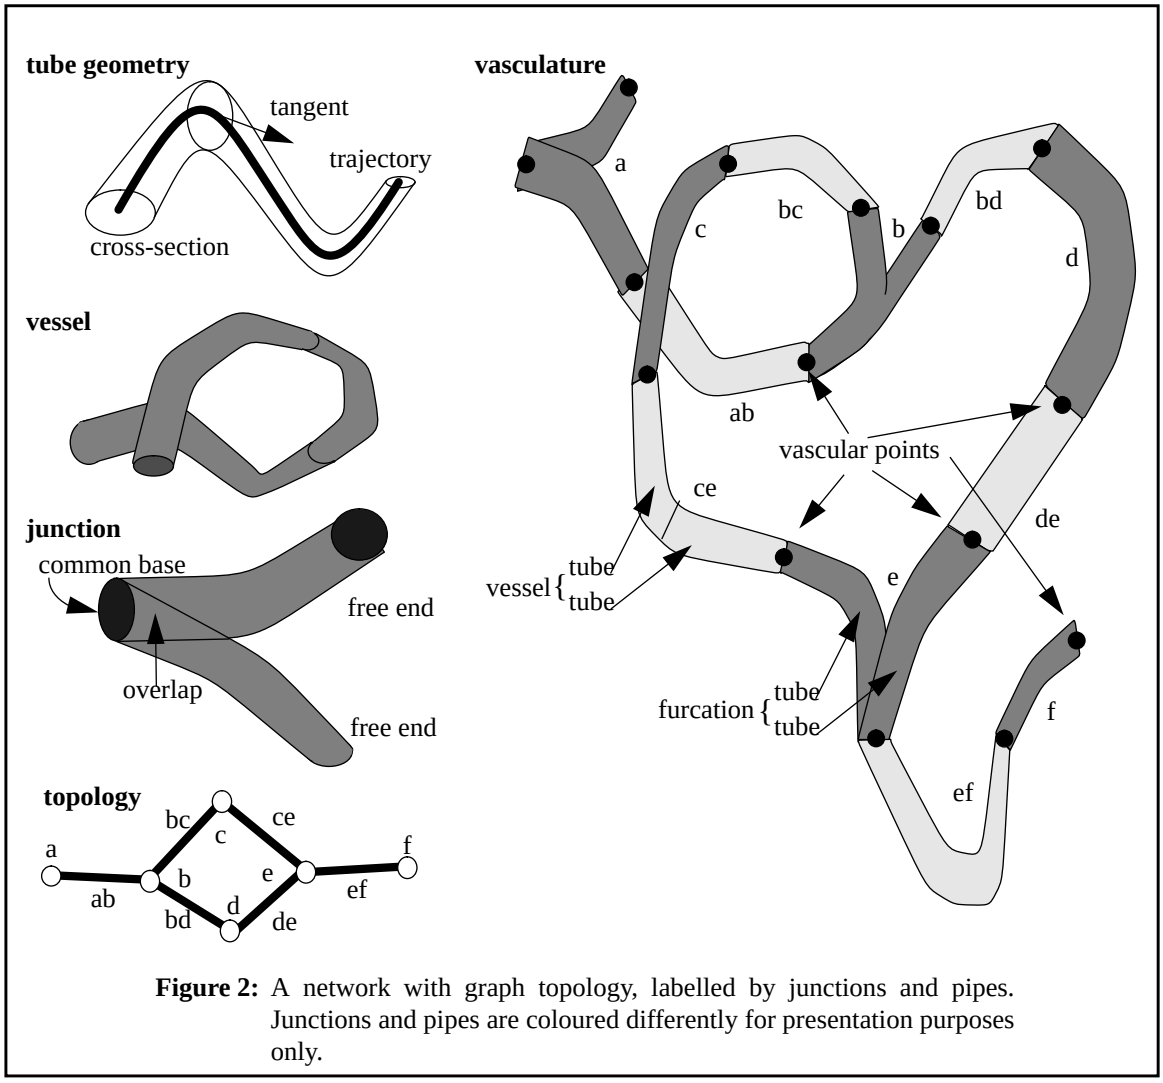

Figure 2: A network with graph topology, labelled by junctions and pipes. Junctions and pipes are coloured differently for presentation purposes only.

Tubes are the basic modelling primitive. The shape of a tube is described by a circular cross section swept over a parametrically defined cubic space curve; this is a restricted form of generalised cylinder. Radius of cross section is given at either end of the space curve and linearly interpolated in between. The cross section is always perpendicular to the tangent of the trajectory. To ensure a linear variation of radius with physical distance the space curve is reparameterised in terms of arc length, s.

Tubes are combined end-to end to create vessels. The position and slope of abutting tubes must match, and radii of cross-section must match also. Furcations are formed by arranging a set of tubes so that they all share a common base. Ends of tubes in junctions that are not part of the common base are free ends. At the common base the position and tangent of tubes must match, and the radii of cross-section must match also. In general tubes in a furcation will overlap to create a common volume. This common volume effects flow only to a small extent and so can be ignored. Consequently, vessels are regarded as tubes in series, and furcations are regarded as tubes in parallel.

The topology of a particular vasculature is represented by a graph of nodes and arcs; edges are labelled by vessels and nodes are labelled by furcations. Vessels and furcations are used to describe the geometry of a vasculature. Tubes in a vessel abut with tubes in a furcation in a manner analogous to the way tubes in a vessel abut - the same continuity conditions are enforced. The graph used as a topology base is undirected because the direction of flow along any vessel is decided by the instantaneous pressure distribution, and so is subject to change.

Conductance to flow is defined for a straight tube of constant cross-section as

$$C = \frac{r^4 \pi}{8L\eta} \tag{3}$$

where r is the radius of cross-section, L is the length of the tube, and η is blood viscosity. Conductance is estimated by decomposing the tube into a series of small tubes, each of constant cross-section. Iterative improvement upon the estimate can be provided by decomposing at successively finer scales until a pair of consecutive estimate match sufficiently well. The resulting estimate ignores any curves in the tube and, more importantly, assumes that blood viscosity is a constant. The rate of flow, Q, of blood along a uniform tube of circular cross section is

$$Q = C\Delta p \tag{4}$$

where Δp is the pressure differential over the tube, and C is its conductance as defined in equation 3. The conductance of a tube is a constant of the simulation.

In this simulation, pressure is defined on the medial axis of the tube, at either end. When the tube is part of a junction such a point will be called a vascular point. There is a single vascular point at the common base of every furcation, and a distinct vascular point as each of its free ends. Not every vascular point has an abutting vessel, blood flows into or out from the simulated system at such points.

The pressure at the set of vascular points in a vascular network defines a pressure distribution that is sufficient to compute the flow rate everywhere in the system. Subject to certain assumptions, the pressure distribution is readily computable given the pressure value at a few vascular points. There are two principal assumptions. One is that there is no energy lost into a redundant form as the blood flows through any vessel. This assumption permits a system of simultaneous linear equations to be set up and solved. Another assumption is that vascular points for which there are no abutting pipe are at zero pressure, unless specifically given otherwise. Also, pressures are permitted to vary in time, but this variation is sampled only when a frame is due to be captured. This can result in severe aliasing of the pressure curve, for which the simplest remedy is an increased frame rate.

Suppose we enumerate the vascular points, 1 to N, then we can write

$$p_i = h_i \tag{5}$$

where p_i is pressure at the i^{th} vascular point, and h_i is the value either given or assumed. Pressures are given as time varying, either cyclic due to heart beat, transitory due to injection of agent, or an additive mixture of these. At vascular points where pressure is neither given nor assumed we can appeal to conservation of volume and write

$$\Sigma_j (p_j - p_i)C_{ij} = 0 \quad (6)$$

in which C_{ij} is the conductance of the tube or vessel that connects point i to point j, and Σ_j accounts for every point distinct from the i^{th}. If there is no connection then $C_{ij} = 0$. These expressions (equation 5 and equation 6) define a set of N simultaneous linear equations in N unknowns, and so are soluble via standard methods (see Press *et al.* 1988). The solution yields a pressure distribution. Having computed a pressure distribution the computation of rate of flow uses equation 4 in each tube. The rate of flow in each of the component tubes of a given vessel must be equal.

3.3 The contrast agent

Contrast agent is used to make flow visible. Contrast agent is injected into the blood stream at a particular location. As the contrast agent courses through the vasculature it splits, and sometimes merges, at furcations. This splitting and merging mixes the contrast agent, changing its concentration.

In the simulation the contrast agent is represented by a bolus. A bolus is a section of blood that contains a specified concentration of contrast agent,; the agent is homogeneously distributed throughout the bolus. The concentration of contrast agent is assumed to change no property other than the attenuation coefficient of droplets making up the bolus. Thus, the μ_k of equation 1 depends upon concentration. Here the dependency is modelled as a linear relation

$$\mu_k = M\kappa \quad (7)$$

where M is the attenuation coefficient for a droplet at unit concentration, and κ is the concentration of the bolus that the droplet is in.

Bolii are made to appear as if they move by moving the boundaries between them. Such a boundaries will be called a face. Faces are advected by the flow. At furcations a single face may divide into two or more faces, or many faces may combine into a single face; concentration level may or may not change with such event. Assumptions relevant to faces are: there is discontinuity in concentration level at a face, faces are impervious, faces are planar, they are cross-sections. Each face has five variables associated with it. These are: the tube it is currently in, its position inside that tube, the concentration of agent behind it (with respect to the direction of flow), the time instant at which the face is at the position, and a time period. The latter two are used to control animation.

Advecting faces by forward integration is difficult because tubes are permitted to taper. Instead, preservation of volume is used directly, as follows. The volume of fluid, $V_{required}$, that a face sweeps in a tube in a time, dt, is just

$$V_{required} = Qdt \quad (8)$$

where Q is average rate of flow. For a tube that tapers linearly the volume captured between positions s_{old} and s_{new} is

$$V_{captured} = \pi\left(r^2 + \frac{(\Delta r)^2}{3}\right)|s_{new} - s_{old}| \quad (9)$$

where r is the minimum radius at the old at new positions, and Δr is the absolute change in radius. A solution to advection requires a value for s_{new} be computed. This is done by iterating positions of s_{new} until $V_{captured}$ is pleasingly close to $V_{required}$. Depending on the direction of flow, starting values for s_{new} will be either 0 or L, the length of the tube.

The computation can be used to determine whether a face has advected out of the current tube. If the initial estimate of $V_{captured}$ is less than $V_{required}$ then there is no solution; the face has spilled from the tube. The time taken, δt, to reach the spill is computable. This is just

$$\delta t = \frac{V_{captured}}{Q} \tag{10}$$

this time period is used to control the simulation.

Spills are most interesting at vascular points that correspond to the common base in some junction. It is at such points that concentration may change. Appeal to preservation of number of dye particles leads yields the governing equation for concentration change;

$$\kappa_{out} = \frac{\sum_{in} |\kappa_{in} Q_{in}|}{\sum_{out} |Q_{out}|} \tag{11}$$

κ_{out} is the concentration in all tubes whose flow is directed away from the spill point; $\Sigma_{out} |Q_{out}|$ is the sum of the outward flow rates. κ_{in} is the concentration in a tube whose flow, rate Q_{in}, is directed in to the spill point, Σ_{in} sums over all input tubes. All of the terms in this equation have a time dependency. Because pressure are sampled only when a frame is captured, flow rates are constant between frames and hence constant in equation 11 for a given frame. This simplification cannot be made for the concentration functions. The solution adopted is to accumulate the concentration function in each tube at every vascular point. Whenever a spill occurs from a tube, at a network point, the corresponding concentration function is updated by the addition of a semi-infinite step function. The edge of the this step function occurs at the time instant of the spill. Before then, the concentration function is unchanged, after then the concentration is taken to be the height of the step; which is the concentration value. This remains at a constant level, until the next change that arises from a spill.

3.4 Animation

Controlling the faces as they move through the vascular network is a complex task. Spill events must be accounted for (these change the number of heads), and animation frames must be captured (so the simulation must be paused at regular intervals).

The proposed solution used three lists of faces, and two cycles as part of its control strategy. The three lists are called the ‘active list’, the ‘done list’ and the ‘part done list’. The active list is a list of faces that require advection. Faces on the active list are ordered by time, earlier faces first. The done list is a list of faces whose advection is complete, for a given frame. The part-done list is a

list of faces whose advection is only partially complete, for a given frame. The advection of a face may be partially complete if the face spills from its current junction or pipe. Spills from one tube to another within the same pipe are not handled by the control strategy being outlined, rather, the simulation continues to advect the face in its new tube. The overall strategy is depicted as a flow chart in figure 3.

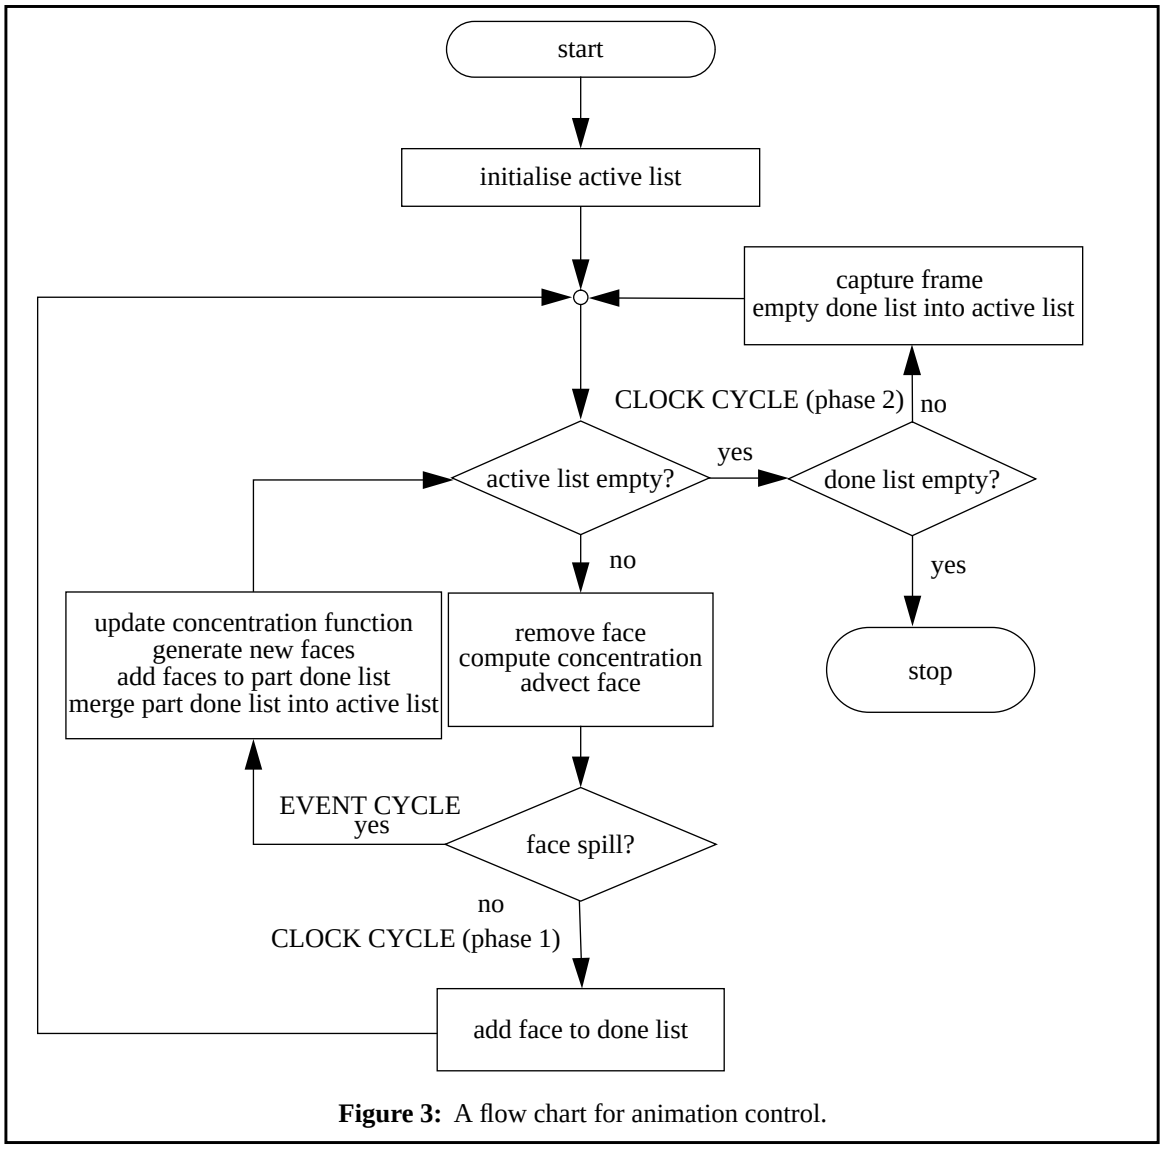

Figure 3: A flow chart for animation control.

To begin the animation, an active list is constructed. One way to do this is simply to specify the faces on the list; that is supply an initial state for the vasculature. In such a case, the time period of each face will be set to be the time interval between frames. Also, in such a case the time location, t, of each face is set to zero, say. Another way to generate the list is to simulate injection of fluid. To inject is rather more difficult, and is discussed below. For the immediate discussion, we assume that a set of faces exist within the vasculature, and that these faces are recorded on the active list.

The event cycle handles spill events. During this cycle faces are removed form the active list, and

faces appear on the part done list. However, the transfer is not one-to-one because a face that spills generates one or more new faces, and it is these that appear on the part done list. Also, faces that spill from tubes that have no abutting vessel are lost by the simulation; they have leaked from the vasculature. In the reverse phase of the cycle, faces are removed from the part done list and merged into the active list, the time ordering of the active list is maintained. Care is taken to ensure that faces that appear in the same place of the network, at the same time, are not duplicated; this process destroys faces.

The clock cycle accounts for the clock events needed to regulate animation. During this cycle faces are transferred from the active list to the done list. Transfer occurs only after they have been advected without spill. This phase of the cycle stops when there are no more faces on the active list, the simulation has then been moved forward from one frame to the next. If the done list is empty as well then all faces will have leaked out of the network, and the animation halts. Otherwise, all faces on the done list are in position for a frame of animation to be captured. Once a frame is captured, the done list is emptied into the active list, and the cycle begins again.

Both of these cycles are part of a single strategy. This removes faces from the front of the active list. At attempt is made to advect it for the time period, Δt, associated with it. The simulation then places faces on either the part done list or done list, as appropriate. The time consumed by that head during its advection is compute. If the face has spilled then the a time δt as in equation 10 has been consumed ($\delta t < \Delta t$). Otherwise the full time period associated with the face has been consumed. In either case its time instant is incremented by that the time consumed, and its time period decremented by the same amount. When the active list and the part done list is empty the done list is checked. If it not empty then an animation frame can be captured, otherwise the simulation should halt. Alternatively, the simulation may be halted after a given number of frames.

Injection, which is the simulation of a syringe injection, must be explained. Injection is defined to take place at a vascular point. The volume and concentration of dye injected must be given, this is the syringe. Any transitory pressure associated with the injection must be given also. Two sets of faces are released into the fluid, one to mark the start of the injection of dye, and one set to mark its completion.

Initially, the ‘starting’ faces are injected. The volume of dye that can be injected in the time leading to the next frame is computed. If this amount is less than the volume currently in the syringe then that amount is deducted from the syringe volume. This deduction of volume in the syringe continues at the start of each new frame until the amount of dye that needed to be injected is greater than the volume currently in the syringe. Then the time to the next frame must be divided into two parts, that for which dye is injected, and that for which it is not. The animation proceeds as normal over the first time period, the trailing faces are injected, and the animation proceeds once more. So, two complete iterations of the clock cycle are required before next frame is rendered.

4 Results

The model shown resembles a network of vessels that might appear in the brain of a human - the cerebral vasculature. The ‘circle of Willis’, and its associated arteries are depicted, see Newton and Potts (1974) for example. This structure is shown from the both front and left side, both static images. These were acquired simultaneously.

To obtain animated images two injection events were used - one to acquire each set of images. In both cases, the left carotid artery has been injected with an dye that partially absorbs x-rays. The carotid appears alone in frames number 1. The dye flows through the vasculature, and out via the left pericallosall and the left middle cerebral arteries. The first of these appears alone in frames numbers 3, but with the middle cerebral in frames 4. These frame numbers imply, at best, a loose correspondence between frames, this is a feature or real angiography. Also, the fact that x-rays are partially absorbed leads to a gray level that depends on the amount of material the x-ray has traversed, again a feature of real angiography.

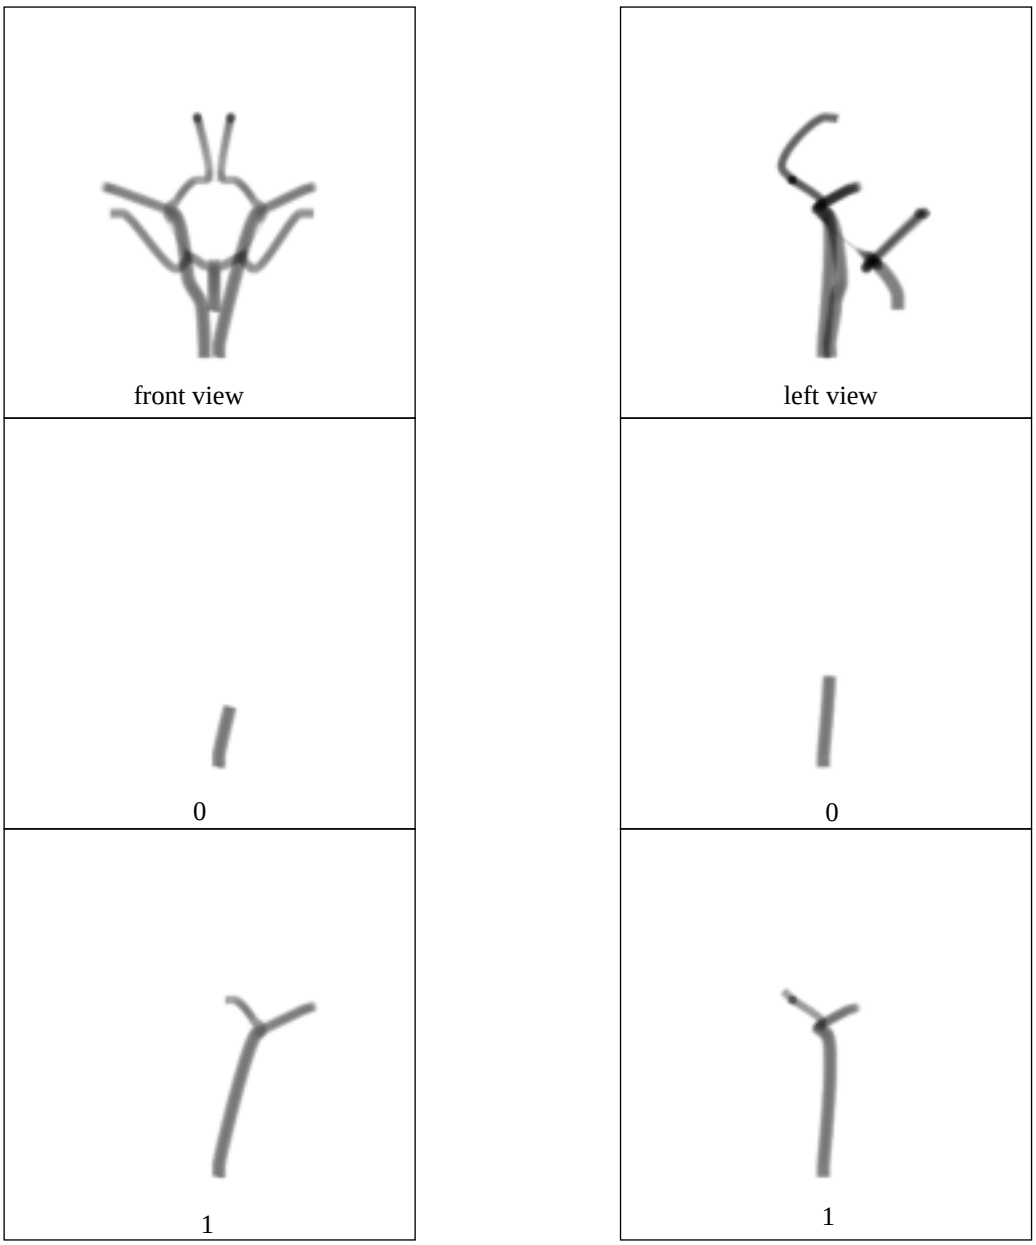

Figure 4: Flow through a subsystem in simulated network of blood vessels. The figure continues...

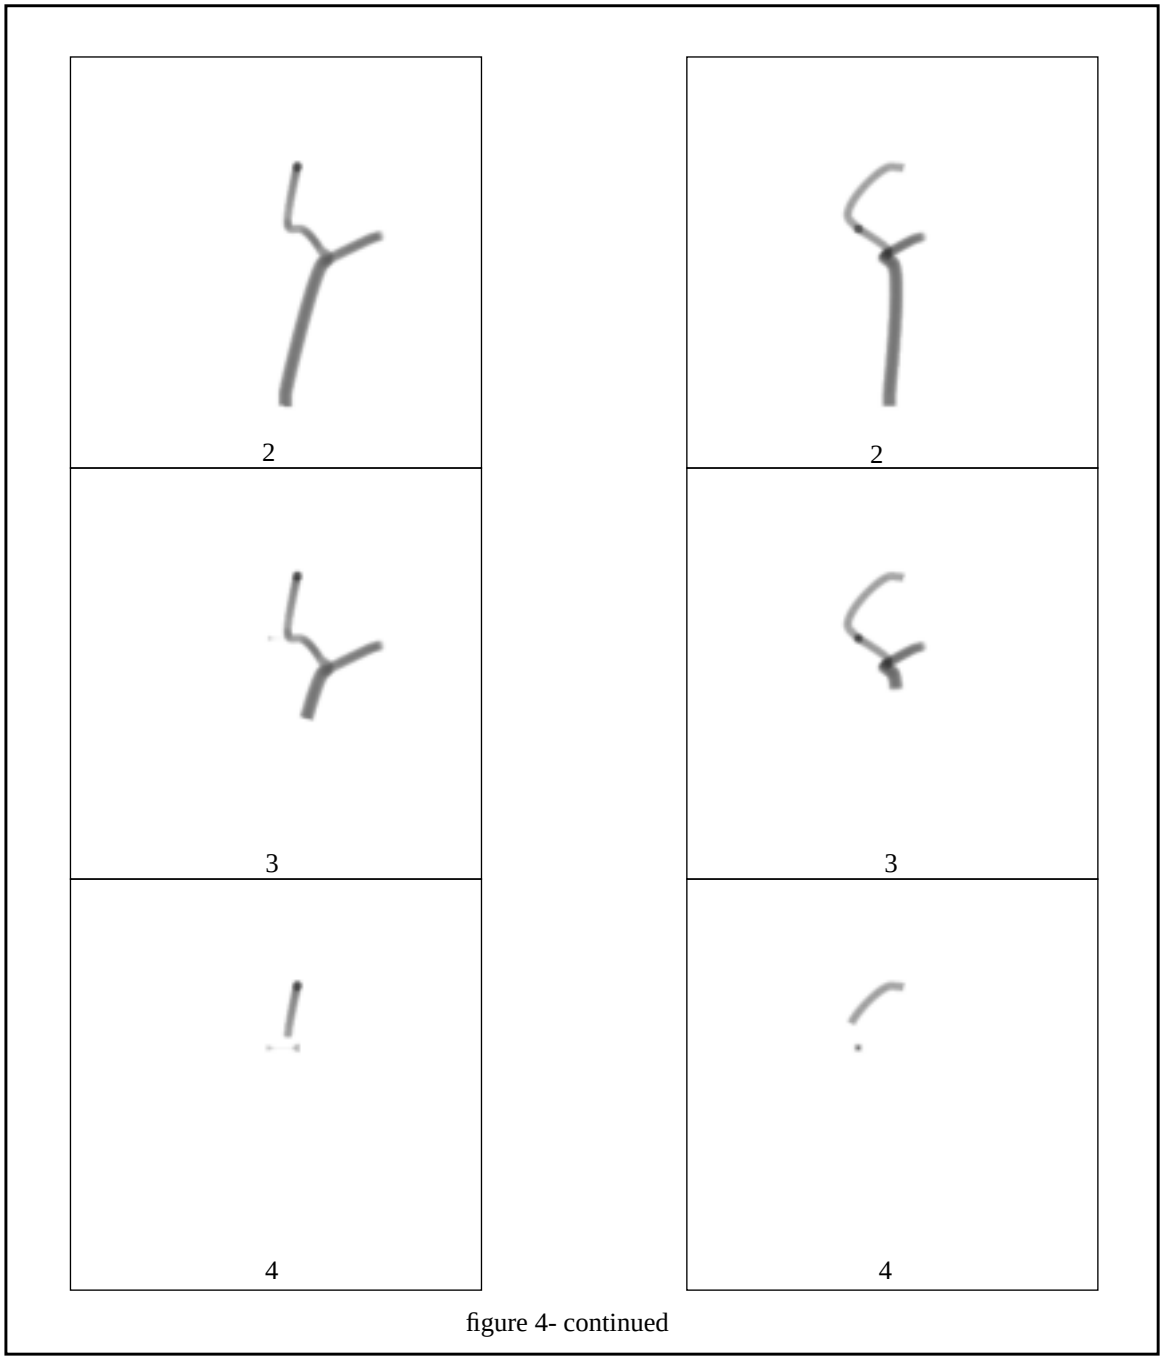

figure 4- continued

Additional results are found in Hall (1994b). These show that the changes in concentration at furcations follow observed trends. In particular, if contrast agent is injected with uniform concentration in time into a delay network, then its variance in time at the outlet quickly reaches a peak before slowly decaying. The same results can be used to show that the simulation conserves volume to an accuracy of a few percent. Also, pressure distribution is strongly effected by the distribution of cross-sections, as expected since conductance to flow depends on the fourth power of radius.

5 Conclusion

At the outset, the salient features of angiography that we wished our simulation to reproduce were as follows.

* The fact that the gray level value at a pixel in an angiogram depends upon the amount of contrast agent that between the pixel and x-ray source. This may be used to estimate the angle that a section of blood vessel presents to the image plane.

* The fact that contrast agent flows through the vasculature and that its trajectory is sampled via cineangiograms. This temporal information could prove very useful when reconstructing from only a few angiograms.

* We wanted to acquire simulated angiograms from many points of view, simultaneously. This is useful because some angiogram machines work that way. Additionally, we believe that it will provide a measure of the information that is lost when angiograms from separate points of view are acquired at separate times, as is most often the case.

* We wanted to be able to commence testing our reconstruction strategies at any desired stage. This allows us to develop our system in a systematic and rigorous manner.

* We wanted the simulation to use models of vasculature in the form produced by the reconstruction unit. This means that our rendering system will provide an immediate mechanism for visualising the results of any reconstruction.

Our simulation fulfils all of these wishes. We know of no other simulation capable of reproducing these features.

There are many unsatisfactory aspects to the simulation. In particular, x-ray simulation does not account for scattering, and blood flow simulation is very simplistic - it cannot be regarded as a physiological model. Improvement to either of these would be at the expense of computational efficiency. We believe that the model is adequate for our purposes, which is reconstruction of vasculature from angiograms. Consequently, we defer any improvement unless it proves necessary.

6 References

Anliker, M., Stettler, J.C., Niederer, P., and Holenstein, R. (1978) Prediction of shape changes of propagating flow and pressure pulses in human arteries. *In The arterial system. Dynamics, control theory and regulation. Bauer, R.D., and Busse R. (Eds.). Springer-Verlag, Berlin*

Barba, J., Fenster P., Suadiaz, M., Wong, K.K., Herrold, E.E., Kenet, R.O., and Borer J.S. (1987): 3D reconstruction of arterial lumen, *Proc. Comput. Assisted Radiology. Berlin, '87*, pp 769-778

Bronsvoort, W., and Klok, F. (1985) Ray tracing generalised cylinders. *ACM Trans. Graph 4(4):* pp 291-303

Bronsvoort, W. (1992) A surface-scanning algorithm for displaying generalised cylinders. *The Visual Computer 8(3)*: pp 162-170

Fessler, J.A., and Macovski, A. (1991) Object-based 3D reconstruction of arterial trees from magnetic resonance angiograms. *IEEE Trans. Med. Imag.* 10(1): pp 25-39

Fung Y.C. (1984) *Biodynamics: circulation*: Springer-Verlag, New York, USA

Hall, P.M., and McGregor, J.J. (1993): A graph based model of a collection of physical vasculature, *Proc. Digital Image Computing Techniques and Applications,* Sydney 93, pp 414-421

Hall, P.M. (1994a) Implicit volume rendering of generalised cylinders. *Tech. Report CS-TR-94-10*. Department of Computer Science, Victoria University of Wellington, Wellington, New Zealand, Internet Tech.Reports@comp.vuw.ac.nz

Hall P.M. (1994b) Simulating and animating fluid flow through a network of tubes," *Tech. Report CS-TR-94-13*. Department of Computer Science, Victoria University of Wellington, Wellington, New Zealand, Internet Tech.Reports@comp.vuw.ac.nz

He, X. Ku, D.N., and Moore, J. (1993) Simple calculation of the velocity profiles for pulsatile flow in a blood vessel using mathematica *Annals of Biomedical Engineering*, 21 (1993): 45-49

Kitamura, K., Tobis, J.M., and Sklansky, J. (1988): Estimating the 3D skeletons and transverse area of coronary arteries from biplane angiograms. *IEEE Tans. Med. Imag. 7(3),* pp 173-187

Lewis, E.E., and Miller, W.F. (1984) Computational methods of neutron transport. *Wiley, New York.*

Li, Z., and Williamson J.F. (1992) Volume-based geometric modeling for radiation transport calculations. *Med. Phys. 19(3):* pp 667-677

Melchior, F.M., Srinivasan, R.S., and Charles, J.B. (1992) Mathematical modeling of human cardiovacular system for simulation of orthostatic response. *Am. J. Physiol* 22(6), (Heart Circ. Physiol. 31): H1920-H1933

Newton, T.H., and Potts D.G. (1974) *Radiology of the skull and brain*: Volume 2 / Book 2. Mosby, Saint Louis

Press, W.H., Flannery, B.P., Teukolsky, S.A., and Vetterling, W.T. (1988) *Numerical recipes in C, The art of scientific computing*. Cambridge University Press, Cambridge, England.

Rech M., and Küster, U. (1994) On viscoelastic fluid flow simulation using the finite volume method. ECCOMAS, 1994.

Rolland, J.P., and Puff, D. (1993) Angiogram simulation software documentation. *Tech. Report. TR93-018* University of North Carolina at Chapel-Hill, Computer Science Department,

Siddon, R.L. (1985) Fast calculation of the exact radiological path for a three-dimensional array. *Med. Phys. 12:* pp 252-255