Year Two Progress Report Three-dimensional Imaging of Drill Core Samples using Synchrotron Computed Microtomography

Brent Lindquist - Principle Investigator
Sangmoon Lee - Graduate Student

State University of New York, Stony Brook, NY 11794-3600

1 Introduction

Fundamental theories of rock structure are limited by the absence of high resolution, pore level, three dimensional images which could be used for statistical analysis. The ability to produce such images in a non-destructive manner would also allow for repeated measurements of dynamic processes such as fluid motion which could be correlated to the medium properties.

One aspect of this funded work is the production of cross sectional images of rock drill core samples with one micron resolution. This work is being done by Keith Jones and Per Spanne of Brookhaven National Laboratory using beam lines X-17 and X-26 of the National Synchrotron Light Source. Jones's annual report will provide a description of the progress in this area.

The second aspect of this work is the geometrical and statistical analysis of the pore structure of the three dimensional images provided. This aspect is being carried out at SUNY-Stony Brook; we report here on the progress that has been accomplished to date. Year one progress was discussed fully in the previous progress report (DOE/ER/14261-1). Year one effort focused on the problem of voxel-by-voxel pore/grain identification of the tomographic images (image filtering), and on analysis of the filtered images via 2-point correlation structure.

2 Year Two Progress

Year two effort focussed on continued improvement of the image filtering algorithm, on completion of medial axis software to investigate the geometrical

Figure 1: Example attenuation histogram (unfiltered) for high porosity Berea sandstone.

structure of the “spine” of the void pathways, and on coupling of the results to three dimensional visualization software for routine viewing of output.

2.1 Data Filtering

A tomographic image slice consists of a rectangular array of x-ray attenuation values, each value corresponding to a finite volume cube (voxel) of the sample. Ideally the attenuation values are proportional to the average material density in any voxel; and a plot of the attenuation values will give a good image of the density variation in the sample. Plate 1 presents a typical x-ray attenuation image for a slice of high porosity Berea sandstone. Due to the nature of X-ray tomography, and finite volume effects, the distinction between void space and material is not sharp. The attenuation histogram for a typical binary material consists of an overlapping bivariate distribution as illustrated in Figure 1 (corresponding to the image of Plate 1). The peak at higher attenuation values is associated with material while the peak at lower attenuation values is associated with void space.

PLATE 1: Example x-ray attenuation image for a typical slice
of the high porosity berea sandstone.

CORR_DAT

For many quantitative analyses of tomographic images it is necessary to have a well defined, self consistent method of identifying each voxel as either pore or grain. Due to the finite size of each voxel, some voxels contain a pore/grain edge; such voxels are denoted as partial voxels. If edge positions were resolved on a sub voxel level, a reasonable algorithm would be to identify those voxels containing 50% or greater grain (void) volume to be grain (pore) voxels. Since tomography provides no sub-voxel resolution, other methods are required.

A traditional approach involves the choice of a cut-off attenuation value lying somewhere between the peaks in Figure 1; any voxel with attenuation value greater (less) than the cut-off is identified as grain (pore). In general, choice of this cut-off is arbitrary, and has little physical basis. We have developed an algorithm which incorporates additional knowledge of the curvature of the spatial distribution of attenuation values in order to guide the choice of cut-off. The procedure is outlined below.

We define the curvature κ of any scalar function $I(x, y)$ as

$$\kappa(x, y) \equiv -\nabla^2 I(x, y);$$

thus the curvature is negative (positive) for increasing (decreasing) slopes. The computation of the Laplacian is implemented by standard, second order finite differences. In order to avoid directional bias, the Laplacian for any voxel is computed twice, once using the 5-point finite difference stencil oriented along the natural $x - y$ axes of the tomographic image, and again using the 90-degree rotated stencil. The two values of Laplacian obtained for each voxel are averaged together to provide the final measure of curvature. The histogram of voxel curvatures obtained for the entire sample of Plate 1 is shown in Figure 2b; a similar histogram for only those voxels contained in a region of air exterior to the sample proper is displayed in Figure 2a. In both plots, the bin at $\kappa = 0$ has been removed, as it contains a large δ -function spike of data counts.

We proceed based upon the following assumptions:

- - A region of physically constant density produces a curvature histogram consisting of a spike of zero curvature, superimposed upon a background Gaussian distribution of noise (tomographic reconstruction noise, slight physical density variations)

Figure 2: Curvature histograms for (a) air region exterior to berea sample, (b) full berea sample.

- – The gaussian noise distribution in a constant density grain region has the same standard deviation as in the constant density air region. This is particularly true if the gaussian distribution is largely due to reconstruction noise, and not due to sizable density differences in grain material. We believe this to be the case in the samples studied here.
- – The full sample curvature histogram will therefore contain the union of a zero-curvature δ -function spike, a gaussian distribution of noise contributed from uniform tomographic reconstruction noise, and non-gaussian ‘tails’ of negative and positive curvature signals coming from partial voxels containing grain/pore edges.

We isolate this sample of edge voxels in the following manner. The δ -function spike is removed from both histograms, and a gaussian fit performed on the air distribution (Figure 2a). The amplitude of this fit is scaled to the full sample distribution (Figure 2b). Thus for each value κ_i of curvature in the full sample histogram, we obtain a fraction f_i of actual counts that are predicted as due to gaussian noise, with the remaining fraction $1 - f_i$

PLATE 2: Voxels identified as edge voxels based upon curvature in the spatial distribution of attenuation values.

Figure 3: Attenuation histograms for full data set, positive and negative curvature voxels.

attributed to edge voxels. For a voxel n in the full data sample having curvature κ_i , a random number ξ_n is then chosen (from a flat random number generator) in the range $(0,1)$. If the $\xi_n > f_i$, the voxel is identified as being non-gaussian, ie. a candidate for an edge voxel. This procedure will produce the statistically correct number of non-gaussian (edge) voxels, but does misidentify an unknown fraction. We reduce this misidentified fraction by requiring that single, spatially isolated (from each other) edge voxels not be produced in this procedure. Since the object of our procedure is to produce a subset of voxels from which to choose an cut-off attenuation value, some contamination of misidentified voxels is not necessarily harmful. The subset of voxels thus identified as being edge voxels are plotted in Plate 2. Voxels having negative (positive) curvature are colored red (blue). We shall term these negative (positive) curvature voxels rising (falling) edge voxels.

We now concentrate on the two subsets of identified rising and falling edge voxels. The histograms of attenuation values associated with a) the full data sample and the identified sets of b) falling edge voxels and c) rising edge voxels are displayed in Figure 3. Mirroring the full data set, the falling edge attenuation histogram contains two peaks, corresponding respectively

to falling edge voxels in the pore and grain sets of voxels. An analogous statement holds for the rising edge attenuation histogram. The falling edge histogram contains a well defined separation between the pore and grain peaks, the separation for the rising edge histogram is less well defined. Further the inter-peak minimum attenuation values for the three distributions do not coincide. In the three data sets examined, we have found the minimum for the falling (rising) edge histogram to lie below (above) the full data set minimum.

We choose the inter-peak minimum in the falling edge attenuation histogram as the best demarked cut-off between a subset of air and grain voxels and apply this cutoff to the *entire* data set; any voxel with attenuation less (greater) than this is identified as pore (grain). The subsequent attenuation image obtained is illustrated in Plate 3. Comparison with Plate 1 reveals some blurring of grain boundaries, where grains are close, within a voxel or two, together. Such voxels correspond to rising edge voxels with rather high attenuation values.

We attempt to identify this subset of rising edge voxels in the following manner. From the rising edge histogram, we determine an attenuation window. The lower attenuation value of the window is chosen as the inter-peak minimum of the rising edge histogram. The upper attenuation value of the window is chosen to lie on the left rising slope of the right hand peak of the rising edge histogram. (Rising edge voxels with attenuation values larger than the right hand peak value tend to correspond to density changes in the grains themselves, rather than voxels sandwiched between close grains.) Currently this upper attenuation value is chosen by trial and error. It would be preferable to develop a more physically based choice. Any grain voxel in Plate 3 that is identified as having i) negative curvature and ii) attenuation value within this window is reidentified as a pore voxel.

The resultant filtered image, with voxel-by-voxel pore/grain identification is shown in Plate 4a compared with an image (Plate 4b) prepared by the "standard" procedure of choosing a single cut-off attenuation from the full data set inter-peak minimum. It is apparent that, for grains that are close together, better grain-grain separation is achieved with our algorithm, than by the conventional one.

PLATE 3: Image obtained after cut-off determined from the attenuation histogram of falling edge voxels is applied to data.

Plate 4: Binary pore/grain image obtained from (a) our filtering procedure (b) a standard, single attenuation cut-off procedure.

(p)

(p)

2.2 Medial Axis Analysis

The medial axis of any volume filling object is intuitively the “spine” of the object running along its geometric middle. It defines the inherent shape of the object. The usefulness of medial axis analysis to the void distribution in a porous medium is immediately apparent. Knowledge of the spine of the void space gives basic information concerning the geometry of the void space such as the average pore length, the average pore connectivity, the percentage of void space cut-off from communication. The medial axis of a two dimensional object consists of a union of lines and points; of a three dimensional object: surfaces, lines and points. We shall show that in constructing the medial axis, some volume information can also be associated with each elementary segment of the medial axis. We discuss the medial axis analysis in the context of tomography; a continuum description is also possible and can be inferred from our discussion.

Consider a three dimensional tomographic image consisting of a cubic array of voxels, each of which has been identified as either pore or grain. We desire to construct the medial axis of the void space of the sample, defined as the union of the pore voxels lying within a polygonal surface that defines the extent of the sample. (Tomographic images typically consist of an external region of air often used for calibration purposes.) The medial axis is constructed according to the following iterative procedure:

- initial step: Label all voxels in the exterior region (outside the image-defining surface) with integer -1 . Label all grain voxels with integer 0 . All pore voxels are initially unlabeled.
- k 'th iteration step ($k = 1, 2, 3, \dots$): Identify all unlabeled pore voxels that lie next to a voxel having label $k - 1$. Label each such voxel with the integer k . Any such voxel bounded on two *opposite* sides by $k - 1$ labeled voxels is further identified as a medial axis voxel.

If we view this procedure as one of “burning” the pore voxels, with the fire starting simultaneously in the “layer” of voxels lying next to the grains, and progressing into unburnt voxels at a rate of one layer per “timestep”, then the medial axis voxels are those voxels in which any component of the fire’s direction is extinguished.

Labeling exterior region voxels by -1 , grains by 0 , and starting the algorithm at $k = 1$, ensures that exterior region voxels do not affect the burning

procedure in any fashion. The slightly ambiguous phrasing “that lie next to” appearing in the iterative step of the algorithm, implies two variations of the algorithm. The identification of “lying next to” as “sharing a common side with” means that only the six neighbors of a three dimensional voxel are considered in the algorithm. The identification of “lying next to” as “sharing any common boundary point” means that all twenty-six neighbors of a three dimensional voxel are considered in the algorithm. The quantitative differences in the resultant medial axis construction produced by these two algorithms are minor; for subjective reasons, we prefer the latter algorithm and have used it exclusively.

The result of the algorithm is a set of identified medial axis voxels, each having an integer label k (a “burn” number). A voxel’s burn number label has physical interpretation, it is the radius (in units of number of voxels) of the largest sphere centered at that voxel, which just fits inside the void space. Thus, in addition to obtaining the geometrical portrait of the spine of the void space, we have some measure of the smallest cross sectional area of the void space at each point along the spine.

We have developed the details of this medial axis algorithm during year two and implemented it in software. The results to date of this investigation are discussed in §2.4

2.3 Visualization Graphics

Part of this funded work involves utilization of graphics software for visualization of the three dimensional images. This has been accomplished at no cost to this grant through use of existing collaborations and public domain software. There are three fundamental graphics paradigms for building/displaying three dimensional objects: mesh rendering, surface rendering and volume rendering. Although mesh rendering is the fastest, the absence of “solid” objects makes this method unappealing and was not considered. A public domain surface rendering code, GEOMVIEW, and a voxel based volume rendering code, VolVis, were considered as state-of-the-art examples of each of these two paradigms. The fundamental building object in GEOMVIEW is a surface bordered by a polygon. This surface can be colored and solid objects built from unions of surfaces. The fundamental object in VolVis is a voxel, a three dimensional rectangular volume. Each voxel can be colored, and voxels stacked together to build larger three dimensional objects.

The voxel based VolVis code is very natural for visualizing tomographic images, as the data comes inherently in voxels. This enables relatively compact input data sets (for an image of $n \times n \times n$ voxels with intensity values stored as characters (256 colors), n^3 bytes of memory are needed). For surface based rendering, a naive rendering utilizing 6 rectangular sides for each voxel, requires listing 4 corner point coordinates and a color for each side, leading to repetitive listing of points ($13 \times n^3$ bytes of memory for displaying n^3 voxels). Although this can be reduced by judicious merging of those unions of sides that form larger rectangles, reductions only of the order of a factor of two are typical. While practical image sizes of $50 \times 50 \times 25$ voxel renderings were possible using VolVis (image refresh time approximately 30 seconds on SGI T400 Indigo workstations), rendering times for GEOMVIEW images of the same size required 3 to 4 minutes and substantially larger memory sizes.

2.4 Initial Results from Medial Axis Investigation

Three data sets of rock core images are being used for analysis. The data sets are summarized in Table 1. The Berea sandstone represents an industry standard rock; the chalk sample illustrates a sample that is difficult to work with using invasive laboratory techniques such as microsanding/microslicing; the glass bead sample provides a realization of a non-overlapping sphere model, a model which has been extensively studied.

rock type	voxels/slice	slices	voxel resolution	slice separation
Berea Sandstone	541×541	56	$5\mu m$	$5\mu m$
Danish Chalk	401×401	22	$5\mu m$	$5\mu m$
100 micron glass beads	601×601	33	$5\mu m$	$5\mu m$

Table 1: Summary of Data Sets

Plates 5, 6 and 7, illustrate, respectively, the medial axes computed for these samples. For simplicity of presentation, two dimensional slices defined by each of the principle planes (xy , xz , yz) in the sample are presented. Grain voxels are black, exterior and non-medial axis void voxels are white, medial axis voxels appear in a spectrum of colors, each color corresponding

PLATE 5: Illustration of medial axis computed for berea sample.
Shown are xy (upper left), xz (lower left), and yz
(right) median slices through the sample.

PLATE 6: Illustration of medial axis computed for chalk sample.
Shown are xy (upper left), xz (lower left) and yz
(right) median slices through the sample.

PLATE 7: Illustration of medial axis computed for glass bead sample.
Shown are xy (upper left), xz (lower left), and yz (right)
median slices through the sample.

to a different integer burn timestep number (red indicates burn number 1, green is burn number 2, dark blue is burn number 3, etc.)

The chalk sample consists of narrow pore spaces occupying a small volume fraction. Nowhere does the medial axis lie further than three voxels (measured center-to-center) from a grain voxel. The berea sample consists of relatively large void spaces separating relatively irregularly shaped grains. The medial axis is nowhere further than 11 voxels (center-to-center) distance from a grain voxel. Due to the uniform size of the glass beads, the medial axis is nowhere further than 7 voxels from any grain voxel.

Figure 4a plots, on a log-linear scale, a histogram (open squares) of the burn numbers obtained in all pore voxels during the construction of the pore medial axis. Superimposed on the histogram is a least-squares straight line fit. The agreement is excellent, indicating an exponential distribution

$$\rho(k) = e^{(-k/\lambda+b)}, \quad k > 0, \quad k \text{ integer},$$

of burn numbers in the Berea sample. The least squares fit values for λ and b and standard deviations thereof are given in Table 2. The same histogram

Rock Type	Distribution Type	λ	σ_λ	b	σ_b
Berea Sandstone (all pore voxels)	exponential	2.49	0.06	5.99	0.06
Berea Sandstone (pore medial axis)	exponential	2.42	0.06	5.62	0.06
Danish Chalk (all pore voxels)	exponential	0.71	0.03	5.57	0.14
Danish Chalk (pore medial axis)	exponential	0.58	0.03	5.83	0.19
100 micron glass beads (all pore voxels)	normal	2.42	0.02	5.92	0.04
100 micron glass beads (pore media axis)	normal	2.49	0.02	5.54	0.04

Table 2: Least squares fit predictions for sample burn number distributions

(filled squares) and fit is also shown for the burn numbers of only those pore voxels identified as medial axis. As evidenced in Table 2, the exponential distribution of this subset of burn numbers is essentially identical to the full sample.

Figure 4a also shows the histogram for the chalk sample. Although only three burn numbers are present in the chalk data, it appears that the burn number distribution is also exponential, decaying much more rapidly ($\lambda = 0.71$) than the berea sample ($\lambda = 2.49$). There also appears to be a

Figure 4: (a) Log-Linear scale histogram of burn numbers for berea and chalk samples. (b) Log-Square scale histogram of burn numbers for bead sample.

Figure 5: Histograms of number of medial axes occurring versus medial axis length for the three data samples considered.

slight difference between the burn number distribution for the subset medial axis voxels as compared with the full pore voxel sample.

In contrast, the burn number distribution for the glass bead sample appears to be normally distributed

$$\rho(k) = e^{(-0.5(k/\lambda)^2 + b)}, \quad k > 0, \quad k \text{ integer},$$

as shown in Figure 4b which plots the histogram on a log-square scale. The distribution of the subset of medial axis voxels appears the same as for the full pore voxel set.

Of interest is the connectivity of the void paths. One simple measure of this is the number of disconnected medial axes found in the sample. The distribution of number of medial axes found versus length of medial axis (as measured by number of voxels composing the medial axis) for the three data sets is shown in Figure 5. In all cases the data seem to suggest a composite distribution consisting of a power-law distribution of lower length medial axes

$$n(l) = b_1 l^{-a}, \quad 0 < l < l_0, \quad l \text{ integer},$$

followed by an essentially constant distribution of medial axes of intermediate lengths

$$n(l) = b_2, \quad l_0 < l < l_1, \quad l \text{ integer}.$$

In the berea and beads sample there is also a single very large interconnected void space, not present in the chalk sample.

3 Proposed Work - Year Three

We intend to explore the range of quantitative information available from the medial axis data, both in a statistical and geometrical manner. As illustrated in Figures 4 and 5, intriguing distribution data is already emerging from this work. Clearly the non-overlapping glass bead model predicts the incorrect burn number distribution (normal instead of exponential) when compared with the berea and chalk samples, though it is quite similar in medial axis size distribution to the rock samples. It would be interesting to test the medial axis features of an overlapping sphere model. Geometrically related questions related to the void pathways will be addressed; these include such questions as: what fraction of the void space is closed to flow (ie. inaccessible from

surface pores), measurements of the tortuosity of the pathways, what is the distribution of branching ratios for the pathways. One intriguing possibility is whether one could characterize the geometry of a "typical" segment of the medial axis, and thus construct "realistic" rock samples by hooking together such typical segments.

It would be very interesting to attempt to predict flow properties such as average permeability from information on about the medial axis geometry.

In addition to exploring these issues, we intend to provide the software created in this effort for use by others.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DATE

FILMED

5/25/94

END