



N° d'ordre : 3847

THÈSE
PRÉSENTÉE À
L'UNIVERSITÉ BORDEAUX 1

ÉCOLE DOCTORALE DES SCIENCES CHIMIQUES

Par Erwan PLOUGONVEN
POUR OBTENIR LE GRADE DE
DOCTEUR

SPÉCIALITÉ : Science des matériaux

Lien entre la microstructure des matériaux poreux et leur perméabilité


Mise en évidence des paramètres géométriques et topologiques influant sur les propriétés de transport par analyses d'images microtomographiques

Thèse dirigée par Dominique BERNARD

Soutenue le : 6 octobre 2009

Devant la commission d'examen formée de :

M. Dominique BERNARD	Directeur de recherche - ICMCB (CNRS)	<i>Directeur de thèse</i>
M. Michel COUPRIE	Professeur - ESIEE Paris	<i>Examinateur</i>
M. Christian GERMAIN	Professeur - ENITAB	<i>Examinateur</i>
Mme. Daniela LAVRIC	Ingénieur de recherche - Corning	<i>Examinateur</i>
M. Peter LEE	Professeur - Imperial College	<i>Rapporteur</i>
M. Mario MAGLIONE	Directeur de recherche - ICMCB (CNRS)	<i>Président du jury</i>
M. Jean-Michel MISSIAEN	Maître de conférence - SIMAP (INPG)	<i>Rapporteur</i>



Contents

- Introduction** **7**

- 1 Background** **9**
 - 1.1 Porous materials 11
 - 1.2 Acquisition of the microgeometry 11
 - 1.2.1 Tomography 12
 - 1.2.2 Physical principles of X-ray absorption contrast tomography 13
 - 1.2.3 Principles tomographic reconstruction 15
 - 1.2.4 Image artefacts in tomography 16
 - 1.2.5 Direct calculation of permeability 16
 - 1.2.6 Pore space characterisation 18
 - 1.3 Digital image processing 18
 - 1.3.1 Image binarisation 19
 - 1.3.2 Object segmentation 19
 - 1.3.3 Pore structure characterisation 21
 - 1.4 Pore space partitioning 21
 - 1.4.1 The distance transform 22

- 2 Pore positioning** **25**
 - 2.1 Direct use of the distance map 27
 - 2.1.1 Maxima dynamics 27
 - 2.1.2 The component tree 29
 - 2.1.3 Limitations 30
 - 2.2 Skeletonisation 30
 - 2.2.1 Definitions 31
 - 2.2.2 Homotopic thinning algorithm 32
 - 2.2.3 Boundary conditions 33

2.3	Digitisation artefact removal prior to skeletonisation	34
2.3.1	Initial methods	35
2.3.2	Digitisation artefact detection	37
2.3.2.1	0D and 2D artefacts	37
2.3.2.2	1D artefacts	38
2.3.2.3	Conclusion	40
2.4	Skeleton/graph conversion: topological classification	44
2.4.1	Surfaces in the skeleton	44
2.4.2	Topological classification of the pixels of the skeleton	45
2.4.3	Topology preservation during conversion	47
2.5	Graph post-processing	48
2.5.1	Types of node cluster merging	48
2.5.1.1	Classical methods	48
2.5.1.2	Determination of the common volume between intersecting balls	51
2.5.2	Node insertion	53
	Conclusion	54
3	Pore delimitation	57
3.1	Region-based methods	60
3.1.1	Region growing	60
3.1.2	The topological watershed on the distance map	63
3.1.2.1	Algorithm	63
3.1.2.2	Sequence of constructible pixels to raise	63
3.1.2.3	Eliminating oversegmentation	65
3.1.3	Modifications to the watershed approach	65
3.1.4	Generating thin separations	65
3.1.5	Seed maximisation	65
3.1.6	Conclusion	67
3.2	Separating planes and surfaces	67
3.2.1	Separating planes	67
3.2.2	Separating surfaces	69
3.2.3	Limitations	72
3.3	Validation	75
3.3.1	Image generation	75
3.3.2	Procedure	76
3.3.3	Relabelling for similar partitions	78
3.3.4	Results	78
3.3.4.1	Computation times	81

Conclusion	81
4 Applications of pore-space decomposition	83
4.1 Evolution of permeability with geometry	85
4.1.1 Presentation	85
4.1.1.1 Sintering experiment	85
4.1.1.2 Precipitation simulation	85
4.1.2 Material characterisation	86
4.1.2.1 Permeability	86
4.1.2.2 Geometric characteristics using the decomposition	89
4.1.2.3 Geometric characteristics without decomposition	92
4.1.2.4 Conclusion	92
4.2 Pore network models	93
4.2.1 Merging non-pairwise connected pores	94
4.2.2 Inserting interstitial pores	96
4.2.3 Inserting special branches in the network	97
Conclusion	101



Introduction

Porous materials are studied in a wide variety of fields. Soils are examined in earth science for water retention; reservoir rocks in petrophysics for hydrocarbon transport; pulp in the paper industry for its mechanical strength; alloys in metallurgy for solidification, and many more.

In all these applications, macroscopic transport properties, such as conductivity or permeability, are of prime interest. It is understood that these characteristics are strongly dependent on the structure of the porosity, not only its amount, but the way it is arranged. The fact that a link exists between the intrinsic structure of the porosity and the extrinsic properties of the porous material has motivated much work on accurate characterisation of the pore space.

A basic requirement for this characterisation is to be able to properly visualise the object of study. Although the acquisition of the structure is a wide area in itself, the geometric complexity of the 3D object that is the pore space is extremely difficult to grasp, and the logic one can have in characterising 2D structures is not applicable to 3D.

Another requirement is the ability to acquire a correct representation of the porosity. A 2D method or 3D reconstruction by serial sectioning is inadequate for such applications. An efficient and non-destructive way to acquire the porous structure is X-ray microtomography, a non-destructive imaging technique with micrometer precision, and thus capable of correctly transcribing the information at the pore scale of many porous materials.

Obtaining the microstructure is a necessary step, but methods need to be put in place to define descriptors that characterise and quantify information from the imaged materials. The work presented in this thesis begins at this stage.

Porous media characterisation often relate to features of individual pores. In practical cases the individual pores that constitute the porosity are not easily identifiable. In fact, they are not even easily definable. One of our objectives is to propose a definition and establish a reasonable decomposition method, taking into account existing works.

The process of decomposition allows a pore mapping of the porosity, along with geometric descriptors, such as volume and connexity, associated to each. This rich data set is somehow linked to macroscopic properties like permeability. A second part of this work will attempt to understand this link through the use

of statistical tools, such as variance analysis or regression. A strong correlation between certain geometric descriptors and a macroscopic property, permeability in our case, will validate our decomposition method.

Another objective is to use this decomposition to generate an accurate pore network model that will serve for flow simulations, in order to allow accurate and rapid prediction of macroscopic transport properties.



1 Background

Contents

1.1	Porous materials	11
1.2	Acquisition of the microgeometry	11
1.2.1	Tomography	12
1.2.2	Physical principles of X-ray absorption contrast tomography	13
1.2.3	Principles tomographic reconstruction	15
1.2.4	Image artefacts in tomography	16
1.2.5	Direct calculation of permeability	16
1.2.6	Pore space characterisation	18
1.3	Digital image processing	18
1.3.1	Image binarisation	19
1.3.2	Object segmentation	19
1.3.3	Pore structure characterisation	21
1.4	Pore space partitioning	21
1.4.1	The distance transform	22

1.1. Porous materials

A percolating porous material is a solid punctured by voids, through which, when connected to the exterior, a fluid or gas can flow. The set of voids is called the porosity. The understanding of fluid flow and transport phenomena is of great interest in many fields of science and affects numerous types of materials: sandstones and other rocks in petrology and geophysics for enhanced oil recovery and CO₂ sequestration, soils in agriculture for water retention and transport of nutrients or contaminants, pellets and thin layers in chemistry for better catalyst supports and chromatographic adsorbents, mortars and cements in civil engineering for durable materials, paper for ink imbibition, bone in medicine for more robust bone implants, ceramics in the automotive industry for efficient particle filters. Examples are even found in the food industry: porous apple tissue for fruit postharvest quality and shelflife [Mendoza et al., 2007], or the breaded coating to optimise frying times for tasty chicken nuggets [Adedeji and Ngadi, 2008].

The macroscopic transport properties of a porous material are naturally correlated to the structure of the porosity, but in a more complex way than a simple relation with its amount. For instance, Frost [1989] showed how two granular materials, prepared differently but with equal porosity, showed different mechanical responses. The geometric structure of the pore space play a crucial role in macroscopic properties, prompting the need for accurate characterisation of this structure, which in turn implies proper visualisation of the pore space.

1.2. Acquisition of the microgeometry

In regards to prediction of transport properties, simulations are traditionally performed on synthetic models adjusted by external measures such as mercury intrusion curves. These models tend to oversimplify the geometry of the porosity, and the preferred solution is to directly acquire the microstructure of the pore space.

Two-dimensional acquisition methods, such as micrographs or SEM (Scanning Electron Microscope) images have been largely studied before X-ray tomography could reach micron (or sub-micron) resolutions, but a 3D representation of the pore space must be obtained, because simpler 2D views of the porosity does not provide a realistic representation of the complexity of the porosity [Al-Raoush and Willson, 2005]. Konstandopoulos [2003] uses statistical tools presented by Kikkinides and Burganos [2000] to reconstruct a three-dimensional structure from a SEM scan that has the same porosity and pore size. Muntean et al. [2003] and Lymberopoulos and Payatakes [1992] use serial section data reconstruction: a series of micrographs are acquired, each after a polishing of the sample of a few microns, then a 3D image is reconstructed by realigning the sections and interpolating the space in between. This has the disadvantages of presenting inaccuracies due to the (tedious) polishing phase and the interpolation between micrographs, of having a significant difference in horizontal and vertical resolutions (preventing non-isotropic measure-

ments) and of being a destructive method [Lindquist et al., 2000]. X-ray microtomography does not have these shortcomings.

1.2.1 Tomography

Tomography is the non-invasive process of generating quantitative data of the cross-sections of a material, from a series of transmission or reflective data acquired from different directions. The type of tomography depends on the physical excitation measured. Here are a few examples:

Ultrasonics PhotoAcoustic Tomography (PAT), used in the medical field, detects the ultrasonic emissions of tissue when heated by short light pulses [Wang, 2008; Xu and Wang, 2006].

Magnetic fields In radiology, Magnetic Resonance Imaging (MRI) or Nuclear Magnetic Resonance Imaging (NMRI) uses the property of the hydrogen nuclei (protons) of water molecules to align with a strong magnetic field. A radio frequency electromagnetic field transfers its energy to the water which, when the field is turned off, emits a resonance radio frequency detected by the scanner [Ganssen et al., 1981].

Electric fields Several electrical measurements on the periphery of the object can be made when an alternating current is applied. Common examples include Electrical Resistance Tomography (ERT) [Bolton and Primrose, 2005], Electrical Impedance Tomography (EIT) [Holder, 2004] and Electrical Capacitance Tomography (ECT) [Williams and Beck, 1995].

Electrons Electron Tomography is an extension of transmission electron microscopy, where a beam of electrons is transmitted through the sample to form an image. The most common type of imaging used for electron tomography is bright field imaging, where the occlusion and absorption of electrons define the contrast in the image [Midgley and Dunin-Borkowski, 2009].

Gamma rays For Positron emission tomography (PET), a nuclear medicine imaging technique, a positron-emitting radionucleotide (tracer) is introduced in the body, and through short life and beta decay, emits a positron that annihilates with a nearby electron, thus generating a pair of gamma photons [Phelps et al., 1975]. For Single Photon Emission Computed Tomography (SPECT), the tracer directly emits gamma photons [Jaszczak, 2006]. For Gamma-Ray Tomography (GRT), in material science, an exterior (high-energy and near-monochromatic) source of gamma rays can be used (when ionising radiation does not degrade the sample), and attenuation in the sample of those rays are measured by the detectors [Cesareo et al., 1999; Patel and Thorat, 2008].

X-rays The process, abusively called Computed Tomography (CT) or Computed Axial Tomography (CAT), is similar to GRT except that the radiations transmitted through the sample are X-rays.

In material science, CT and GRT are most commonly used for non-destructive imaging (GRT is mostly used for high-speed phenomena or highly absorbent materials). With recent progress in this field, CT can attain very fine resolutions, down to a micron or less per pixel, which is why the term microtomography or microCT is often used. This thesis will focus on X-ray microtomography, and more precisely absorption

contrast X-ray microtomography[†].

1.2.2 Physical principles of X-ray absorption contrast tomography

The theory of X-ray tomographic imaging was first applied for clinical purposes with the invention of the X-ray computed tomographic scanner by Godfrey Hounsfield, for which he received a Nobel Prize in 1972, shared with Allan Cormack who found some of the reconstruction algorithms.

The acquisition process, based on X-ray radiography, is an application of the Beer-Lambert law, or attenuation law, which states there is a logarithmic dependence between the absorption coefficient of a material and the transmission of light:

$$I_1 = I_0 e^{-\mu l}$$

where I_0 and I_1 are respectively the incident and detected light intensity, μ is the attenuation coefficient of the material, and l the width of the traversed material. Obviously, the value of μ is not uniform inside the sample, therefore the sum of all attenuation coefficients along the light path must be considered:

$$I_1 = I_0 e^{-\int_{path} \mu(x) dx}$$

where $path$ is the path the light takes through the material, and x the coordinate along this path. This line integral transformation was introduced by Johann Radon who gave it his name, and demonstrated the reversibility of the process [Radon, 1917]. Figure 1.1 illustrates this for one path.




Figure 1.1: Illustration of the attenuation law, which defines a logarithmic dependence between transmission of light and attenuation of the material. In a material with a homogenous attenuation coefficient, transmitted intensity depends on the width l of the material. In a real material, however, as in figure (b), the sum of attenuation coefficients along the light path must be considered.

The coefficient of attenuation is dependent on the material density ρ and atomic number Z , and the photon energy E . In the photoelectric domain, i.e. for energies below 200 keV, the following equation is

[†]With a coherent X-ray beam, the phase shifting of the rays through the sample can be imaged, this is called phase contrast tomography, a process pioneered by Ando and Hosoya [1972].

used to describe this relation [Attix et al., 1968]:

$$\frac{\mu}{\rho} = K \frac{Z^4}{E^3}$$

where K is a constant.

A radiography for a given material, often referred to as a projection, is a 2D map of attenuations through distinct paths. Two types of paths are shown in figure 1.2, depending on the X-ray source. If a laboratory scanner is used[†], a conical X-ray beam is created by an electron beam hitting a micron-thick target metal layer with high Z and ρ , such as Tungsten or Molybdenum (the loss of kinetic energy of the electron translates into radiation energy, a small percentage in the form of X-rays). If synchrotron radiation is used[‡], then the light is well collimated, making the rays nearly parallel. The bright and coherent X-rays of synchrotron radiation is generated by the acceleration of relativistic electrons travelling inside the storage ring. This acceleration can be achieved by an insertion device (ID), an array of magnets which create a spatially periodic field through which the electron passes. Figure 1.3 shows examples of both systems (the circular storage ring is shown in 1.3(b)). For demanding applications, microtomography at a synchrotron facility allows finer resolutions (below $0.3\mu\text{m}$ per pixel at the ESRF), much brighter light for imaging materials of higher density or higher atomic number, and better quantification due to monochromatic light [Baruchel et al., 2000].




Figure 1.2: Two types of paths the lights can take: figure (a) shows a cone beam, where all paths originate from a focus point and thus describe a cone, whereas figure (b) presents a setup that is found at synchrotrons, where all paths are parallel.


Finally, detection of the attenuated rays depends on energy range, resolution, and frame rate. Earlier detectors use a gas ionisation chamber [Sauli, 1998], while more common and recent systems used in microtomography are scintillation detectors, where a fluorescent material (e.g. Caesium Iodide) converts x-ray photons into visible light, detected by a CCD camera [Cendre et al., 1999].

[†] few companies produce microCT scanners, among which are Skyscan, Phoenix|x-ray, Toshiba, and PreXion.

[‡] There exists about 50 synchrotron radiation facilities worldwide, a list can be found at lightsources.org.



(a) The phoenix x-ray Nanotom scanner



(b) The European Synchrotron Radiation Facility

Figure 1.3: Example of two different types of light sources for microtomography. Figure (a) shows a laboratory microCT scanner, a machine 2 meters long, and figure (b) shows an aerial view of the ESRF, where the storage ring of the synchrotron facility is about 270 meters wide. The choice of the light source depends on the application.

1.2.3 Principles tomographic reconstruction

A single projection provides information of the attenuation from one point of view, which is insufficient to deduce the complete attenuation map of the material. In order to reconstruct the original map, a series of projections is taken from different points of view. Typically, the object is set on a rotating platform and rotated by small increments on 360° (only 180° is sufficient for parallel beam, since the projection at an angle of $180 + \theta$ is a mirror of the projection at an angle of θ). In reality, many more projections are needed for an accurate reconstruction (it depends of course on the number of pixels in the reconstructed image [Deans, 1983]).

Several reconstruction methods exist. Fourier reconstruction methods make direct use of the projection-slice or Fourier slice theorem, which is a simply demonstrated equality between a projection and a cross-section in the Fourier domain. If an infinite number of projections are measured, then the entire Fourier transform of the object can be deduced, and therefore an inverse Fourier transform produces the attenuation map. Errors of this method derive from the finite amount of projections along with the necessary polar-to-cartesian resampling in the Fourier space. A more numerically accurate and faster method, still derived from the projection slice theorem, is called Filtered Back-Projection (FBP), developed by Ramachandran and Lakshminarayanan [1971] and generalised by Shepp and Logan Jr. [1974], and is the most common reconstruction method. It compensates for the finite amount of information of the slices in the Fourier space by filtering each projection (several filters exist, such as the Hanning filter, Butterworth filter, Weiner filter, etc.) before performing an inverse Fourier transform or backprojection (so called because resembling a smearing of the projection onto the final image) in the direction corresponding to the projection angle. A comprehensive review of this method can be found in Chapter 3 of [Kak and Slaney, 1988].

A completely different type of reconstruction is Algebraic Reconstruction Technique (ART) [Natterer, 2001] which uses an analytical definition of Radon inverse transform. It defines linear system $\vec{p} = M\vec{x}$, where \vec{p} is the projection, \vec{x} the image pixels, and M the contribution matrix of the projection points on the image pixels, and inverts it in a least square sense due to noise.

Note that these methods can be used for parallel beam as well as cone beam geometries, the equations being written accordingly[†].

1.2.4 Image artefacts in tomography

X-ray absorption tomography is affected by specific artefacts. A few examples are listed:

Beam hardening When a polychromatic x-ray beam traverses a material, low energy photons are absorbed more quickly, resulting in an attenuated flux with a higher mean energy (i.e. towards the hard x-rays, hence the name). This implies that the longer the beam traverses an absorbant material, the harder the x-rays become, therefore the less it interacts with the material, making the centre of a sample appear less absorbant in the reconstruction. This effect is called the *cupping effect*. For the same reason streaks may appear behind highly absorbant materials during the backprojection. These artefacts can be reduced by placing a filter (such as a 0.1mm thick copper foil) between the source and the sample, in order to pre-harden the beam.

Hot points There is a probability for an x-ray photon to traverse the scintillator and directly hit the CCD camera (if such is the acquisition system), resulting in a very high intensity photon on a particular radiogram. The backprojection of that radiogram will result in a pixel-wide streak across the reconstructed volume. A conditional median filtering can remove these hot points.

Ring artefacts When the detector response is not linearly related to the received light intensity (this can be caused, for instance, by dead, hot, or stuck pixels of the CCD sensor), ring artefacts appear in the reconstruction. This artefact is a circular arc centered on the rotation axis of the system. Such artefacts are illustrated in figure 1.4, along with its correction, made prior to the backprojection with the method from Bernard and Chirazi [2006].

X-ray microtomography is a powerful tool that allows to accurately visualise the structure of a porous material, but for a quantitative characterisation of its transport properties, additional tools, such as numerical modelling, are required.

1.2.5 Direct calculation of permeability

The flow through a porous material can be characterised by its permeability, a measure indicating the ease a fluid has of permeating through. Henry Darcy discovered in 1856 that the flow of a fluid through a porous medium (sand, in this case) was related to the pressure gradient by the following relation [Darcy,

[†]A widely-used reconstruction method for cone-beam light source is the Feldkamp algorithm, which is a weighted filtered backprojection [Feldkamp et al., 1984]




Figure 1.4: Example of a common type of artefact found in x-ray tomography, called ring artefacts. Figure (a) shows a cross-section, on the plane perpendicular to the rotation axis, of the reconstructed volume (a porous crinoidal limestone) without correction. Figure (b) shows the effect of the ring artefact filter on the reconstruction.

1856]:

$$q = \frac{-\kappa}{\mu} \nabla P$$

where q is the flow, μ is the fluid viscosity, ∇P is the pressure gradient, and κ the permeability, expressed later on as a symmetric second-order tensor for anisotropic materials.

Permeability can be measured experimentally (although only a one-dimensional version of the permeability can be obtained), or computed by using the pore-scale geometry of a material (over a representative volume), by determining the flow field of the fluid in the porosity.

The flow through a porous material can be described by the Navier-Stokes equations, or if only creeping flow is considered, the simplified Stokes equations (that assumes negligible inertial forces). These non-linear equations are an application of Newton's second law, conservation of momentum (along with conservation of mass in these equations), written for an arbitrary control volume:

$$\begin{cases} \mu \nabla^2 \vec{V} - \vec{\nabla} P = 0 \\ \vec{\nabla} \cdot \vec{V} = 0 \end{cases}$$

The complete permeability tensor can be obtained by solving the closure problem associated with the Stokes equations (chapter 4 of Whitaker [1999]). The result is a flow field (and a pressure field), which once averaged over a representative volume provides the components of the permeability tensor.

Micro-tomography gives access to the geometry at the pore-scale, on which the permeability tensor can be computed, for instance by a finite volume technique (an illustration is shown in figure 1.5) based directly on the pixels of the fluid phase from the tomographic acquisitions [Anguy et al., 1994; Bernard, 1995; Bernard and Vignoles, 2000; Bernard et al., 2005b; Mercet, 2000].

Numerical approximation of this non-linear system is computationally intensive, it demands a number of unknown proportional to the number of pixels of the volume under examination, and convergence of the solution can take tens of thousands of iterations. Resolutions in microtomography can reach 2000^3




Figure 1.5: Illustration of a simple geometry used to solve the Stokes equation. The geometry consists of a cylinder in which a ball was placed at the centre, partly obstructing the flow. The system consists of cubic finite volumes one fiftieth the diameter of the central portion of the cylinder. The calculated horizontal flow is represented as a set of vectors.

pixels, but it would be totally unreasonable to solve the Stokes equations on such a large system, in terms of necessary memory and computation. This is the principal reason for much ongoing work on pore space characterisation and generation of simplified but representative models.

1.2.6 Pore space characterisation

Mitchell and Soga [1993] used the term *fabric* for describing particle arrangements and pore spaces in the study of soil behaviour, and Oda [1972] showed the influence, for granular materials, of fabric on compression mechanisms, but for a better understanding of the microscale processes involved in macroscopic properties, local or pore-level characteristics should be examined. The term pore is introduced cautiously because its definition is the first difficulty in this thesis. The intuitive idea we have of the shape of a pore (a spherical opening potentially connected to other pores by cylindrical channels) can hardly be projected onto a real porous material. Yet qualifiers for what we call pores do exist.

Dullien [1991] defines a pore as "part of the pore space bounded by the solid and planes erected where the hydraulic radius is minimal". The idea is repeated in many other publications, sometimes claiming the definite and unambiguous quality of this definition [Baldwin et al., 1996; Sederman et al., 1997]. As we show in chapter 3, this definition is unfortunately not so unambiguous and is inapplicable to real porous materials.

For this reason, one objective of this thesis is to survey available methods for defining and identifying pores. All the developed procedures use concepts from the field of digital image processing.

1.3. Digital image processing

An image can be loosely defined as a representation, whether in a continuous or discrete space, on a square or hexagonal grid. The notion is not limited to two dimensions, as it merely depends on the

data to be represented. We limit our work to the 3D case and on the cubic grid because of the method of acquisition, X-ray tomography. We refer to the images obtained as applications from \mathbb{Z}^3 to \mathbb{R} , and elements as pixels (the portmanteau word *voxel* is often used for 3D images).

There are two classical representations of a pixel. They can be seen as 0-dimensional objects, or points, defined on the 3D cubic grid, or as unit cubes centred on those points. Although both representations have their qualities, we will prefer the latter, as tomography computes an attenuation coefficient on a certain volume. The attenuation coefficient is computed as a floating-point variable, therefore defining a greyscale attenuation map of the porous material. Because of noise and sub-pixel features, pixels are not directly defined as contained in the pore space or the solid phase. Subsequently, a binarisation process must first be performed.

1.3.1 Image binarisation

Image binarisation is a wide subject in itself, and can occupy an entire thesis for an exhaustive survey of existing methods. Recent reviews can be conferred [Pham et al., 2000; Wirjadi, 2007]. It consists in assigning every pixel of an image to either the object (or foreground) or the background. Among the most common segmentation algorithms, we can cite simple thresholding, edge detection, active contours, or indicator kriging. The work presented here begins once the images have been binarised into the object, which is the porosity, and the background, i.e. the solid. The method is based on anisotropic diffusion, in order to enhance edges, and followed by thresholding [Mrázek and Navara, 2001; Perona and Malik, 1990].

1.3.2 Object segmentation

Binarising the images, i.e. identifying which pixels belong to the void phase, is one step of pore space characterisation. Of course, global quantities such as the amount of porosity can be directly computed, but additional analysis is critical. As a reference in the study of porous media states [Dullien, 1991]: "It is now well understood that transport and capillary phenomena in chaotic pore spaces cannot be explained, let alone predicted, in terms of a single parameter, such as the volume fraction of the void phase. Instead, detailed descriptions of the pore space must be sought that reflect information on local variations in the geometry of pore channels as well as information on the degree of pore space connectivity."

The general problem of object segmentation in images is another wide subject, and we will focus only on partitioning of porous materials. The approaches belong to either of two general classes, depending on the object of interest: when the material allows it, the solid phase can be examined and separated into physically meaningful elements, e.g. for bead packings or sand beds [Garboczi, 2002; Thompson et al., 2006]. In these cases, where the object of study is composed of physically well-defined and distinct entities, efficient image processing methods have been established.

The second class concerns the separation of the porous phase into elements (pores) that do not stem from distinct and naturally separable entities (except in special cases, as in figure 1.6), which are generally

named pores, voids, chambers or nodes. Many authors use the terms *pore size* or *pore distribution*, implicitly implying that individual pores have been identified, which is not the case in most porous media, as in figure 1.7 for example.

Classical methods to introduce pore-level details into numerical models of mass transfer consist in generating artificial networks of pores with volumes and connections adjusted by measures from mercury porosimetry or nitrogen gas adsorption. The generated networks unfortunately provide a oversimplified view of the pore structure, that does not account for all the connectivity or spatial variations inside the material. This type of approach becomes irrelevant if it is possible to directly retrieve the pore structure, through non-destructive 3D imaging techniques. The problems is then shifted to finding a decomposition of the pore space into individual pores, separated by contact surfaces called necks or throats, and forming a pore network that accurately represents the porous phase.

Note that this model is one of two generally accepted view of the pore space, where the void phase is separated into pores having thin contacts between them. Another view is called the chamber-and-throat concept [Constantinides and Payatakes, 1991; Dullien and Dhawan, 1974; Li et al., 1986; Tsakiroglou and Payatakes, 1990, 1991; Wardlaw et al., 1987], where the pores are separated by throats having a certain volume and the property of a *relatively constant cross-sectional area*, assuming that the material allows such an idealisation [Lymberopoulos and Payatakes, 1992]. The two views are related, and we can suppose that a bijection exists between the two (thin throat can be dilated to generate a volumetric throat, and vice-versa). In any case, all methods to partition the void space starts with the use of morphological tools such as homotopic thinning, skeletonising, erosion or medial axis.






Figure 1.6: Illustration of an idealised porous material. The material was produced by inserting PMMA (Polymethyl methacrylate) beads, a porogenous organic material removed by heat treatment, inside a calcite matrix. In this model sample, pores are easily identifiable.



(a) Surface representation of the pore space of a Fontainebleau sandstone



(b) Fluid phase of a quenched Aluminium-Copper alloy [Bernard et al., 2005b]

Figure 1.7: Examples of the porosity in real porous media, representing a volume of about 1cm^3 .

1.3.3 Pore structure characterisation

Many statistical values can be extracted from the information given by the pore structure. Whether or not they are relevant or influent in transport properties such as permeability is what this thesis will assess.

Directly from the pore space, pixel count can provide approximations to porosity ratio, surface porosity, specific surface area or burn number distribution [Lindquist et al., 1996], i.e. the distribution of the distance of the pixels of the porosity to the solid phase. Topological information can include connectivity, genus (linked to the amount of different paths through the material), and can infer measures of tortuosity [Quiblier, 1984]. Finally, from the decomposed set of pores, a wealth of quantities can be determined: pore volume and throat surface area distributions, pore-throat correlations, pore and pore throat shape factors, etc.

Before we can assess the pertinence of any of these quantities, we must verify that the image processing scheme we use to obtain them yields as clear a description of the pore structure as possible.

1.4. Pore space partitioning

The next two chapters present a critical analysis of various existing decomposition methods, along with improvements we have brought. From this study we define the requirements and qualities of a pore space partitioning.

In terms of decomposition, not only for pores but interconnected elements in general (grains or particles), two fundamental steps can be distinguished: positioning and delimitation. The former, presented in chapter 2, handles the number and spatial location of the elements in the medium, while the latter, described in chapter 3, determines their volumes. These two steps are not systematically independent,

especially for pore space decomposition, and pore number and position cannot be known accurately before the delimitation is completed. Usually a set of markers serves as initialisation for the delimitation, but there is not a perfect correspondence between them and the pores, and they can be further modified in the delimitation process. Finally, we attempt to assess the robustness of our decomposition method.

As a foreword to these image processing tools, we present the concept of distance transform, fundamental in almost all existing decomposition methods and often necessary for both positioning and delimitation.

1.4.1 The distance transform

In a binary image, each pixel belongs either to the object, X , or the background \bar{X} . In pore space partitioning, the object is the porous phase, and the background, the solid. The purpose of the distance transform is to assign each pixel $p \in X$ its distance to \bar{X} , that is the distance $d(p, \bar{X}) = \min_{p' \in \bar{X}} \|\vec{pp}'\|$. The resulting image is called a distance map or distance field. Marcke [2008] refers to it as an aperture map, while Shin [2002] calls this distance the *burn number* and many others name the transformation a *burn* or *grassfire* algorithm [Atwood et al., 2004; Lindquist et al., 1996; Thompson et al., 2006], or erosional analysis [Bakke and Øren, 1997].

Nearly every existing decomposition method uses this distance transform step, even implicitly: Al-Kharusi and Blunt [2007] uses the notion of maximal ball, which is based upon a distance transform. Some exceptions include Delerue [2001] who does not compute directly the distance map, but a Voronoï diagram that can lead to the distance map with an additional scan of the image. Aste et al. [2005] does not even perform a segmentation, but uses a priori information on the shape of the elements to partition (this was possible only because the material was a ball packing and the partitioning was applied to the solid phase, i.e. the balls).

An important parameter in the distance transform is the norm used to compute the distance. Most of the previously quoted articles use a non-euclidean norm[†], either the Manhattan (L_1) or chessboard (L_∞) norm, for alleged reasons of computational cost [Al-Raoush and Willson, 2005]. This strongly accentuates grid orientations, as illustrated in figure 1.8. As a reminder, if $\vec{p} = (x, y, z) \in \mathbb{R}^3$, then the L_1 , L_2 (Euclidean) and L_∞ norms are respectively : $\|\vec{p}\|_1 = |x| + |y| + |z|$, $\|\vec{p}\|_2 = \|\vec{p}\| = \sqrt{x^2 + y^2 + z^2}$, and $\|\vec{p}\|_\infty = \max(x, y, z)$.

Much work has been done on Euclidean distance transforms and fast linear algorithms (algorithms that have a linear relation between the image size and computation time) for the Euclidean distance transformation are well established [Cuisenaire and Macq, 1999; Hirata, 1996; Meijster et al., 2000; Saito and Toriwaki, 1994]. For all subsequent methods presented, unless stated otherwise, we have implemented and used the exact Euclidean distance transform as described by Hirata [1996].

[†]Burn, grass-fire or erosion algorithms use non-Euclidean norms

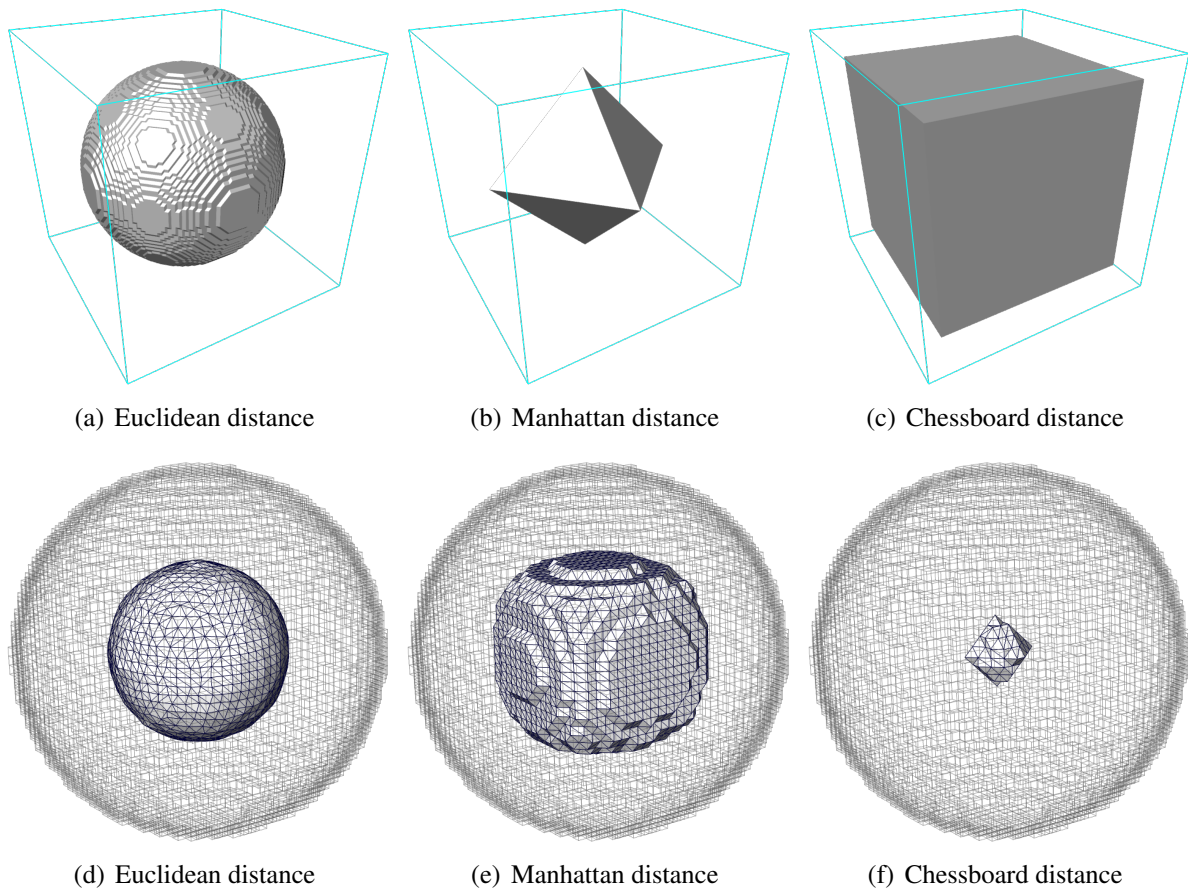


Figure 1.8: Illustration of three norms used in image analysis through isosurfaces : (a)-(c) show isosurfaces at an isodistance of 40 pixels from the 3D image centre using three different norms. (d)-(f) show the same process inside a digitised ball, with its outlying pixels outlined, at a distance of 10 pixels from the outside. (b) and (e) use the Manhattan distance, favouring the directions of the reference axes, and (c) and (f) use the chessboard distance, which favours the diagonals.



2 Pore positioning

Contents

2.1	Direct use of the distance map	27
2.1.1	Maxima dynamics	27
2.1.2	The component tree	29
2.1.3	Limitations	30
2.2	Skeletonisation	30
2.2.1	Definitions	31
2.2.2	Homotopic thinning algorithm	32
2.2.3	Boundary conditions	33
2.3	Digitisation artefact removal prior to skeletonisation	34
2.3.1	Initial methods	35
2.3.2	Digitisation artefact detection	37
2.4	Skeleton/graph conversion: topological classification	44
2.4.1	Surfaces in the skeleton	44
2.4.2	Topological classification of the pixels of the skeleton	45
2.4.3	Topology preservation during conversion	47
2.5	Graph post-processing	48
2.5.1	Types of node cluster merging	48
2.5.2	Node insertion	53
Conclusion		54

This chapter provides the means for positioning markers in the porosity, which will effectively define distinct elements in the pore space that can be called pores. Several methods are reviewed, beginning with the most direct ones in section 2.1. Section 2.2 describes the topological approach, leading to processing difficulties dealt with in section 2.3 by a novel and efficient method. Finally sections 2.4 and 2.5 present the processing for obtaining a representative graph of the pore structure.

2.1. Direct use of the distance map

The distance map can be used as such to compute a pore decomposition [Al-Raoush and Willson, 2005; Atwood et al., 2004], with a method called watershed and presented in section 3.1. This transformation, usually computed on the inverse of the distance map, considers the input image as a height map or relief (the value of each pixel of X is called its height), and separate the minima, or basins, of this topography. Figure 2.1 shows a simple 2D example of two overlapping circular pores. Of course, when considering the distance transform of an object as the relief for this type of algorithm, the distance map must first be inverted, thus transforming the maxima of the distance map into minima.

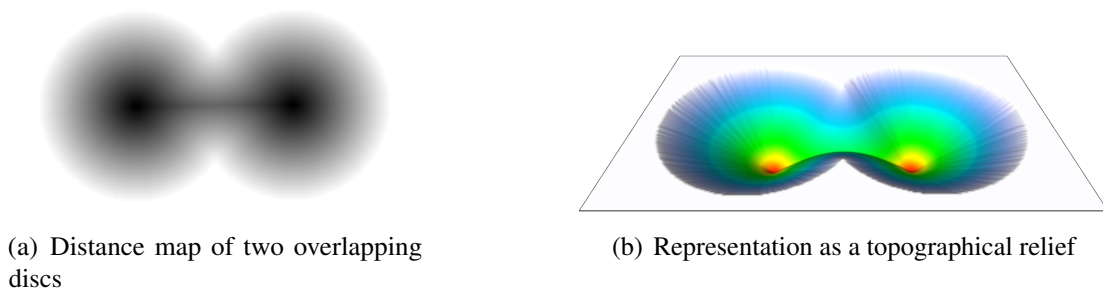


Figure 2.1: Representation of a distance map as a relief on a 2D example: (a) shows the Euclidean distance map of two overlapping discs, and (b) is the same data but inverted and represented as a height map (the lowest pixels are those at the greatest distance from the disc exterior). This representation is easily extendable to 3D, although its visualisation would be less trivial.

The first difficulty to consider is the overabundance of maxima in a distance map of a typical pore space image: this implies a merging process [Thompson et al., 2006]. Saadaftar et al. [2006], for example, defines a merging condition according to the distance between maxima and their heights, i.e. their distance from the solid. A similar type of merging, used by Marcke [2008] and Tariel [2008], consists in examining the height difference between the maxima and the contacts separating the regions. This condition, verified during the delimitation step, is related to the notion of maxima dynamics [Bertrand, 2005; Soille, 1999], described in the next section.

2.1.1 Maxima dynamics

The dynamic of a maximum is the smallest difference in height between it and the lowest pixel of a path joining it with another maximum. The implementation of this notion in 3D is not trivial. There are


generally many paths from one maximum to another, and the difficulty resides in finding one with the highest minimum height. This concept is illustrated in 2D in figure 2.2, where the greyscale image is represented as a relief, and paths between the maxima are drawn.

The objective of maxima dynamic calculation is to select those with a sufficiently high dynamic because the others are considered insignificant[†]. This notion is identical to the measure of region significance in Frucci and di Baja [2008].


An efficient way to determine the dynamic of a maximum is to use the component tree of the distance map.




(a) A 2D image with three regional maxima.




(b) The image as a Digital Elevation Model. From this point of view, the lowest regions are defined as maxima.



(c) An example of a path between the two top maxima. This is not the one with lowest pass value.



(d) A path with lowest pass elevation.



(e) The altitude difference between the pass and the maximum, defining the dynamic.


Figure 2.2: 2D Illustration of the notion of maxima dynamic. (a) shows the initial height map, and (b) a representation as a relief. (c) shows one among many paths between the two maxima of the image and (d) shows one whose pass is lowest. The difference in elevation between the maximum and its lowest pass, shown in (e) for the red region, defines its dynamic.

[†]since no delimitation has taken place, there is no merging performed, only a removal of some maxima. The end result of the decomposition, however, produces the same number of elements.


2.1.2 The component tree

Implementation of the tree construction algorithm is based on the works of Najman and Couprie [2004], but applied to images with non-integer values. Let $\mathcal{F} : \mathbb{N}^3 \rightarrow \mathbb{R}$ be an image, and $k \in \mathbb{R}$ a height. In our case, \mathcal{F} is the distance map. We define the level k component of \mathcal{F} as the set \mathcal{F}_k of pixels with a height greater than or equal to k . Furthermore, we define the level k^+ component as the set \mathcal{F}_k^+ of pixels with a grey level strictly greater than k .


Level k components can include either zero, one, or more than one level k^+ component, and are respectively called leaf, branch and node components. Since $k \in \mathbb{R}$, there can exist an infinite number of branch components in \mathcal{F} , although we consider only the branch components for a level $k \in \mathbb{R}$ such that $\mathcal{F}_k \neq \mathcal{F}_k^+$. Since the branch components give no additional information on the topology of \mathcal{F} , the vertices of the component tree are defined by the non-branch components, and edges are created between 2 components \mathcal{C}_1 and \mathcal{C}_2 if $\mathcal{C}_1 \subset \mathcal{C}_2$ and if $\exists \mathcal{C}' : \mathcal{C}_1 \subset \mathcal{C}' \subset \mathcal{C}_2 \Rightarrow \mathcal{C}'$ is a non-branch component, and we say $\mathcal{C}_1 = \text{Father}(\mathcal{C}_2)$. An example is presented in figure 2.3: figures (c) to (f) show the level k components of \mathcal{F} for the significant values of k .




(a) Initial image \mathcal{F} . The numbers are the grey levels.




(b) Component tree of \mathcal{F} . The top component, the root of the tree, is a branch component, preserved for illustration purposes. The leaves represent the maxima in the image.




(c) \mathcal{F}_1



(d) \mathcal{F}_2



(e) \mathcal{F}_3



(f) \mathcal{F}_4

Figure 2.3: Illustration of the component tree along with all the level k components.

Once the component tree is calculated, it becomes very easy to determine the dynamic of a minimum: it corresponds to the difference in altitude between the leaf in the tree that represents that maximum, and the height of its lowest ancestor in the tree having more than one son.

2.1.3 Limitations

The component tree is a powerful tool for image segmentation and various morphological operators: Vachier [1995] uses the component tree to define *extinction functions* to merge extrema before a watershed transform. Unfortunately, on the subject of pore space decomposition, this tool is insufficient to account for the complex topology of the porosity. The fact that a maximum (significant or not) marks the position of a pore is true, but the reverse, a pore is always marked by a maximum, is not. Figure 2.4 illustrates such a problem: using the minima alone will not detect pores at intersections where the void space is no more extended than in its neighbourhood, therefore too few markers will be generated and the pore space will not be properly represented.




Figure 2.4: A 2D example of a pore space that presents a problem when marking the pores only with the high-dynamic maxima of the distance map: in (a) circles completely contained in the porosity are drawn to show that no local maxima exist where the pore paths intersect, therefore no marker will be placed near the central triangular area. (b) shows a graph that would accurately represent this configuration, where the vertices would define the pores, and the branches the different connections between them. The dotted lines show a centred decomposition based on these vertices.

In fact, this problem is recurrent to all methods that do not include topological information in their selection of markers. The most efficient way of retrieving the topological information of an object is through skeletonisation.

2.2. Skeletonisation

Intuitively, the skeleton of an object is a thin simplified representation that preserves topology. It was first introduced by Blum [1961] and has been extensively studied. We provide a general definition for the skeleton of an n -dimensional object $X \subset \mathbb{R}^n$ ($n \in \mathbb{N}$):

Skeleton of X Subset \mathcal{K} of dimension at most $n - 1$ homotopic to X .

Homotopy of X and Y There exists a continuous deformation from one to the other.

Note that this loose definition does not imply uniqueness nor centredness of the skeleton. In a topological space, the skeleton is defined as equivalent to the medial axis, and very often medial axis transform and skeleton are used interchangeably [Al-Raoush and Willson, 2005; Al-Raoush et al., 2003;

Liang et al., 2000; Lindquist et al., 1996, 2000; Prodanović et al., 2006; Shin et al., 2005; Thovert et al., 1993]. Lieutier [2003] proves the homotopy of the medial axis with the object X , but only provided X is an open subset of \mathbb{R}^n .

The accepted definition for the medial axis is as follows:

Medial axis of X Centres of the maximal n -balls of X .

Maximal n -ball of X n -ball \mathcal{B}^n contained in X such that no other n -ball contained in X contains \mathcal{B}^n .

n -ball Interior of a sphere \mathcal{S}^{n-1} . [Conway and Sloane, 1992]

Sphere \mathcal{S}^{n-1} Set of all points in \mathbb{R}^n at a given distance from a point defined as its centre. [Hocking and Young, 1988]

In discrete spaces, however, the equivalence between medial axis and skeleton does not hold [Thompson et al., 2006], as the medial axis is not guaranteed to preserve topology (see figure 1 in [Couprie et al., 2007]).

A classical method for image skeletonisation is homotopic thinning, which consists in removing pixels of the object (provided the topology is preserved) until obtention of the skeleton (by removing, we mean that we assign a pixel $p \in X$ to the background \overline{X}). Whether the result is irreducible (no more homotopic thinning can be applied: the result is called an ultimate skeleton), or centred, or contains extremities etc. depends on the thinning scheme used, which is why our definition of a skeleton is so loose.

In the following section, we present the classical framework of digital topology for defining the notions used for homotopic thinning.

2.2.1 Definitions

Let I be a binary image, i.e. an application $\mathbb{Z}^k \rightarrow \{0, 1\}$. We recall the necessary definitions, taken from [Bertrand and Couprie, 2007], in the case $k = 3$. A pixel p is defined by its indices $(x, y, z) \in \mathbb{Z}^3$, and its intensity is noted $I(p)$. We present three typical neighbourhood relations Γ_6, Γ_{18} and Γ_{26} , illustrated in figure 2.5, and defined for a pixel p by $\Gamma_6(p) = \{p' \mid \|\vec{pp}'\|_1 \leq 1\}$, $\Gamma_{18}(p) = \{p' \mid \|\vec{pp}'\|_2 \leq \sqrt{2}\}$, $\Gamma_{26}(p) = \{p' \mid \|\vec{pp}'\|_\infty \leq 1\}$, and denote by $\Gamma_n^*(p)$ the set $\Gamma_n(p) \setminus \{p\}$ for $n = \{6, 18, 26\}$. A pixel p' is said to be n -adjacent to p if $p' \in \Gamma_n^*(p)$.

Let $X \subset \mathbb{Z}^3$, \overline{X} the complementary set of X , and $\{p, p'\} \subset X$. We say that there is a n -path between p and p' in X if there exists a sequence of pixels $\{p_1, p_2, \dots, p_m\} \subset X$ such that $p_1 = p, p_m = p'$, and $(\forall i \in [1; m - 1]), p_i$ is n -adjacent to p_{i+1} . We say that two pixels p and p' of X are n -connected in X if there exists an n -path between p and p' in X , and that a non-empty subset $Y \subset X$ is an n -connected component of X if $(\forall (p, p') \in Y \times Y), p$ and p' are n -connected and are not n -adjacent to any pixel of $X \setminus Y$.

The set of all n -connected components of X is written $C_n(X)$. A subset $Y \subset X$ is n -adjacent to a pixel $p \in X$ if there exists a pixel $p' \in Y$ n -adjacent to p . The set of all n -connected components of X n -adjacent to a pixel p is written $C_n^p(X)$.

We also introduce the notion of geodesic neighbourhoods. For a pixel $p \in X$, we call a *geodesic*




Figure 2.5: The three neighbourhood relations for 3D discrete images defined on a cubic grid. (a) shows the 6-neighbours (in transparent) of the central blue pixel. (b) shows the previous neighbourhood in dark grey, and the 12 additional pixels that together comprise the 18-neighbourhood. Finally (c) shows the 8 additional pixels that make up the 26-neighbourhood of the central pixel.

n -neighbourhood of p in X of order k the set $\Gamma_n^k(p, X)$ defined recursively by $\Gamma_n^1(p, X) = \Gamma_n^*(p) \cap X$, and $\Gamma_n^k(p, X) = \cup\{\Gamma_n^*(p') \cap \Gamma_{26}^*(p) \cap X, p' \in \Gamma_n^{k-1}(p, X)\}$.

In order to preserve the Jordan-Brouwer theorem [Spanier, 1966], different adjacency relations must be used when considering the object and the background [Kong and Rosenfeld, 1989]. The two are respectively noted n and \bar{n} , and the compatible relations are either 6 and 18, or 6 and 26. A distinction needs to be made when 6-connectivity is used with 18-connectivity, in such case we use the notation 6^+ . The *geodesic neighbourhoods* $G_n(p, X)$ are defined as $G_6(p, X) = \Gamma_6^2(p, X)$, $G_{6^+}(p, X) = \Gamma_6^3(p, X)$, $G_{18}(p, X) = \Gamma_{18}^2(p, X)$, and $G_{26}(p, X) = \Gamma_{26}^1(p, X)$.

Finally, we can define the *connexity numbers* $T_n(p, X)$, for a pixel p in a set $X \subset \mathbb{Z}^3$, as $T_n(p, X) = |C_n[G_n(p, X)]|$. These number allow to topologically characterise a pixel of an object [Bertrand, 1994; Bertrand and Malandain, 1994; Malandain et al., 1993].

The central idea behind many homotopic thinning algorithms is the notion of simplicity. A pixel p of an object X is said to be simple if it can be removed without changing the topology of the object. Using the above definitions, a pixel is n -simple if and only if $T_n(p, X) = 1$ and $T_{\bar{n}}(p, \bar{X}) = 1$.

Simplicity is a local characterisation, which is of prime importance when considering the computational cost of an algorithm, especially on large 3D microtomographic images. In the rest of this thesis we consider 26-adjacency for the object and 6-adjacency for the background, because 26-connected skeletons are thinnest and therefore contain a minimum number of pixels.

2.2.2 Homotopic thinning algorithm

There are two types of homotopic thinning methods : sequential and parallel thinning. Sequential thinning removes pixels one after the other, while the latter and more difficult type removes more than one simultaneously. Parallel algorithms have the advantage of not requiring a distance map to obtain a

properly centred result. Many have used a parallel algorithm of Lee et al. [1994][†], although it has been shown to be based on inexact assertions [Bertrand and Malandain, 1995].

There exists a vast and ongoing literature on skeletonisation algorithms, and their review is beyond the scope of this thesis. For the avid reader, a recent publication on the subject [Bertrand and Couprie, 2006] presents a general framework for describing classes of parallel skeletonising algorithms.

The work presented here will use a computationally efficient sequential algorithm, based on distance-ordered homotopic thinning [Pudney, 1998], that produces centred skeletons when used with the Euclidean distance map.

A sequential skeletonisation algorithm can be easily described. The process consists in iteratively removing simple pixels of the object. This means that at each iteration, a pixel is checked for simplicity and is removed if the test succeeds. The process is ended when no more pixels can be removed, and this depends on the properties the skeleton should have (conserve extremities, surfaces, contain the medial axis). In order to obtain a centred skeleton, the Euclidean distance map of the object is used: it serves to define processing priorities of pixels [Bertrand and Couprie, 2007]: the simple pixels with smallest distance are removed first. To handle the priorities, the algorithm uses a red-black tree [Bayer, 1972], with each element, sorted by distance in the tree, containing a queue of pixels to process. The pixels at the object border are processed first and their neighbours, which are not yet in the tree, are appended in the queue corresponding to its distance. At each iteration, the first pixel of the first queue of the sorted tree is processed and its neighbours added to the tree if necessary.

Figure 2.6 shows a few examples of skeletons. The two first ones are centred using the above algorithm. The last one has been obtained by processing each pixel in a random order, leading to a much more complex but thin structure (which is still homotopic to the original object)[‡].

2.2.3 Boundary conditions

We are interested in transport phenomena inside the material, therefore the skeleton must preserve any connection with the exterior of the image. If we suppose that the exterior is part of the solid phase, then the homotopic thinning will erode the border pixels, thus removing all connections to the border. To avoid this, we provide two different boundary conditions, depending on the application. The first one supposes that the exterior is part of the porosity, leaving branches in the skeleton having an extremity in contact with the image border. The second is a periodic condition: two opposite sides of the bounding box of the image are in contact. The geometric space is periodic in all three directions. The skeleton defined in the latter geometry can be used to build an equivalent network model for permeability calculations [Combaret,

[†][Al-Raoush and Willson, 2005; Al-Raoush et al., 2003; Arns et al., 2004; Lindquist et al., 1996, 2000; Mendoza et al., 2007; Neethirajan et al., 2008; Peth et al., 2008; Prodanović et al., 2006; Shin, 2002; Shin et al., 2005; Sok et al., 2002; Venkatarangan, 2000]

[‡]Note that this is an ultimate skeleton, even though there are two dangling ends connecting disordered groups of pixels. This is the classical problem of the difference between contractibility and collapsibility [Zeeman, 1963], with the house of Bing being the most famous example (although in this case the object can still be collapsible, but not with the proposed algorithm [Passat et al., 2008]).




Figure 2.6: Illustration of a few results from the described skeletonisation algorithm: (a) shows the object to skeletonise, (b)-(d) are examples of skeletons that can be obtained.

to appear].

2.3. Digitisation artefact removal prior to skeletonisation

A well-known problem with skeletonisation is its sensitivity to small features. Indeed, a one-pixel cavity in the object is just as sure to lead to a surface in the skeleton as a hundred-pixel cavity. The skeleton is only concerned with preserving the topology, no matter the thickness of the features it describes. These small features can be generated during the binarisation process because of remaining noise from the acquisition or from the tomographic reconstruction algorithm, and consist mainly in misclassified pixels (a pixel of the solid phase in the middle of the porosity, or vice-versa), referred to as surface noise in [Venkatarangan, 2000], for instance. Another significant possibility that is seldom described is the effect of digitisation of physical object, which either have features finer than the size of a pixel, or features that are not aligned with the pixel grid, which is nearly always the case. Figure 2.7 shows this effect on an image of a simple mathematical object made of two balls having a point contact. Shin [2002] refers to this as a *finite voxel size effect* among the digitisation errors, and refers to a particular type of digitisation artefact as a *punch-through artefact*.

An intuitive assumption is that low image resolution (i.e. the pixel size compared to the size of the represented objects) is to account for these types of artefacts, and the finer the resolution, the more these effects tend to disappear. What happens is actually the contrary, as is demonstrated in figure 2.8.

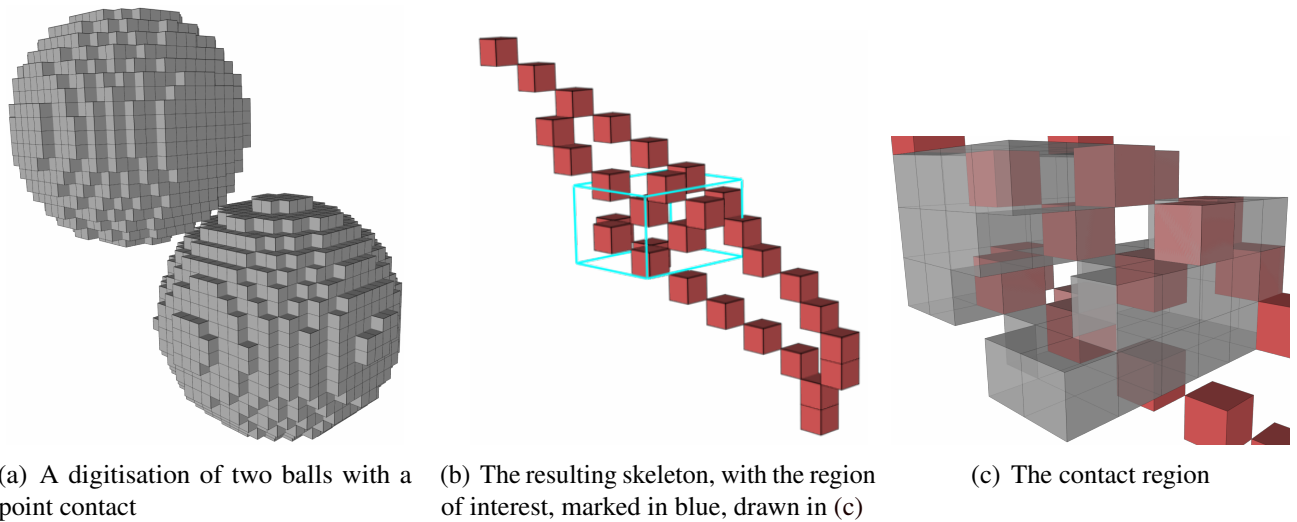


Figure 2.7: *Illustration of artefacts from the digitisation process. In this case, two digitised balls with a point contact (a) generate several branches in the skeleton (b). The region around the contact (c) shows numerous distinct contacts between the pixels of the two balls.*

It is therefore necessary to pre-process the segmented images before skeletonisation.

2.3.1 Initial methods

Shin [2002] uses a custom morphological operator called a majority filter. It is applied to the pixels of \bar{X} (i.e. the solid phase), and consists in reassigning a background pixel to X if a majority of its 26-neighbourhood belongs to X . This filter is equivalent to a median filter exclusively applied to pixels of \bar{X} . Although it significantly reduces the amount of artefacts, it has two drawbacks. Firstly, it modifies far more pixels than necessary: more than 95% of the modified pixels are nowhere near digitisation artefacts. Secondly, and more significantly, Shin shows a configuration where an artefact is not removed by his filter.

A more aggressive method could simply be the median filter, used by Venkatarangan [2000], which modifies around 50% more than the majority filter but is able to remove more problematic configurations. Yet again, we can find situations where this filter is insufficient, for instance in figure 2.9.

A first personal attempt was to use a morphological closing. This is a dilation followed by an erosion of the object[†]. In effect, it produces similar results to the majority filter, modifying around 30% more pixels, and still unable to handle all configurations (e.g. figure 2.11). A morphological closing followed by an opening can actually remove all artefacts, but the number of modifications to the object is completely unreasonable (considering the four images of figure 2.8, an average of 3.65 times more pixels are modified than with the majority filter).

These observations prompted the need for a proper and novel approach to handling these artefacts,

[†]A dilation assigns to X all pixels in its neighbourhood, and similarly an erosion assigns to the background \bar{X} all pixels neighbour to \bar{X} .




Figure 2.8: Illustration of the effect of resolution on digitisation artefacts on an image of a set of 14 balls extracted from a random close packing (with a 5% overlap). The skeletons, superimposed with the centres of the balls, show many erroneous edges and even surfaces. Contrary to popular belief, the higher resolution and thus the more pixels, the higher likelihood of digitisation artefacts to appear.




Figure 2.9: An illustration of how the median filter is unable to properly handle certain configurations containing digitisation artefacts. (a) presents two balls of a random packing in near contact. (b) shows the contact, in blue, between the balls (the darker left ball is shown in transparent), which is a set of 6 pixel corners. These 6 points should not represent a physical contact and be considered as artefacts. (c) shows the modifications, in red, made by the median filter, and the modified contacts between the balls.

and thus our solution consists in directly identifying the responsible pixels.

2.3.2 Digitisation artefact detection

The basic assumption guiding this method is that small features of the object below the scale of the pixel cannot be considered reliable, and therefore the parts of the skeleton induced by these pixels are considered as artefacts.

The described process is actually a subclass of well-composedness [Latecki et al., 1995; Latecki, 1997], where the purpose is to obtain an object having the same topology no matter the adjacency relation chosen for the object (6-, 18- or 26-adjacency). We say subclass because this method modifies only part of those modified by Latecki, those that generate artefacts in the skeleton.

A natural classification of the artefact types is by their dimensionality. Since the skeleton of a 3D object is at most 2-dimensional, there can therefore be 0D (points), 1D (branches), and 2D (surfaces) artefacts.

2.3.2.1 0D and 2D artefacts

Handling the 0D artefacts are trivial, as they results from small clusters of pore pixels in the solid phase. They are easily removed by selecting all connected components of the porous phase that contain only a few pixels. These artefacts, as much as any isolated void inside the solid phase, are irrelevant in the study of transport phenoma. Therefore all connected components of the porosity completely surrounded by the solid phase are removed. Note that this removes more than simply the 0D artefacts, as one can imagine for instance a closed loop-shaped cavity inside the solid phase, described by an isolated circle in

the resulting skeleton.

2D artefacts, or surface-type artefacts, can occur for two reasons: either because of noise from the acquisition, creating isolated pixels of the solid phase in the porosity (the dual of the 0D artefacts), or pixels on the solid interface that are disconnected from the rest of the solid when using 6-connectivity, as illustrated in figure 2.10. In the same manner, these artefacts are removed by assigning the small connected components of the solid phase to the porosity [Lindquist et al., 2000].




Figure 2.10: Example of the reason of the appearance of 2D artefacts in the skeleton: (a) shows a portion of the ball packing used previously, at the interface between two balls, (b) shows the pixels lying in between the balls, with the two highlighted in red that cause the surface in the skeleton, because they are not 6-connected to the others. (c) shows the resulting surface encompassing these two pixels, as shown in (d). The surface is much wider than those two pixels because the skeletonisation process produces a centred skeleton.

2.3.2.2 1D artefacts

The troublesome artefacts result in erroneous branches in the skeleton. We can distinguish two subclasses in the branch-type artefacts: we refer to them as corner-wise and edge-wise.

The easiest to remove are the former, an example of which is presented in figure 2.11. We call them corner-wise because they originate from a connection between two pixels of X that are corner-adjacent, i.e. 26-adjacent and not 18-adjacent, such that no other pixel of X is 26-adjacent to both. This configuration can be tested locally by finding the 2x2x2 pixel configuration shown in figure 2.12, and to avoid the artefact, either one or both (for purposes of symmetry) pixels of X can be set to \bar{X} .


The most difficult type of artefacts are edge-wise branch-type artefacts: they appear when portions of the solid phase are connected by an edge or a series of pixel edges, as seen in figure 2.13.

To define these artefacts, we assume that a pixel of \mathbb{Z}^3 is the union of the unit cube, 6 unit faces, 12 unit edges, and 8 vertices. We use the following definitions:

Edge adjacency Two pixels are edge-adjacent if they are 18-adjacent and not 6-adjacent, i.e. if they share exactly one common edge.

Weak edge An edge of a pixel $p \in X$ is *weak* if it is contained in exactly two edge-adjacent pixels of X

Strong edge An edge of a pixel $p \in X$ is *strong* if it is contained in at least three pixels of X



(a) A portion of the porous phase, with a salient shape (b) The resulting branch in the skeleton (c) The pixels of the porous phase that cause the artefact

Figure 2.11: Illustration of the corner-wise branch-type artefacts: (a) shows a porous phase with a sharp convex feature, (b) shows the edge created in the skeleton, passing through the highlighted pixels from (c).

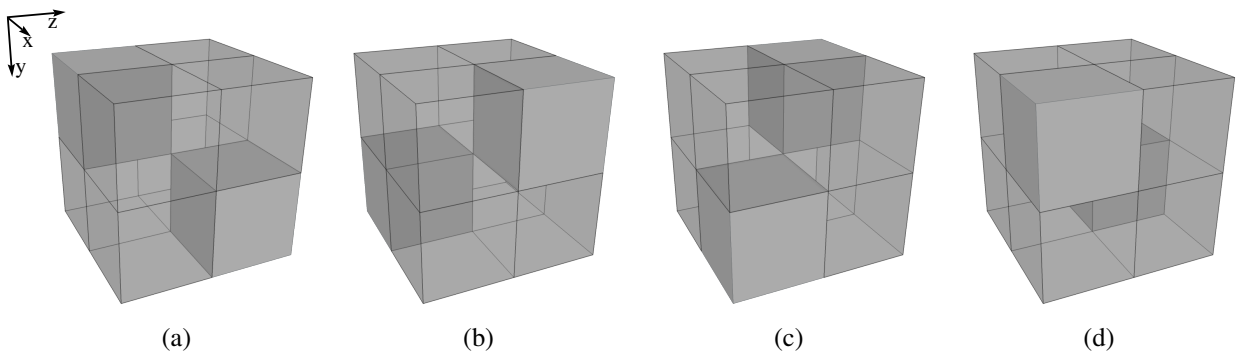
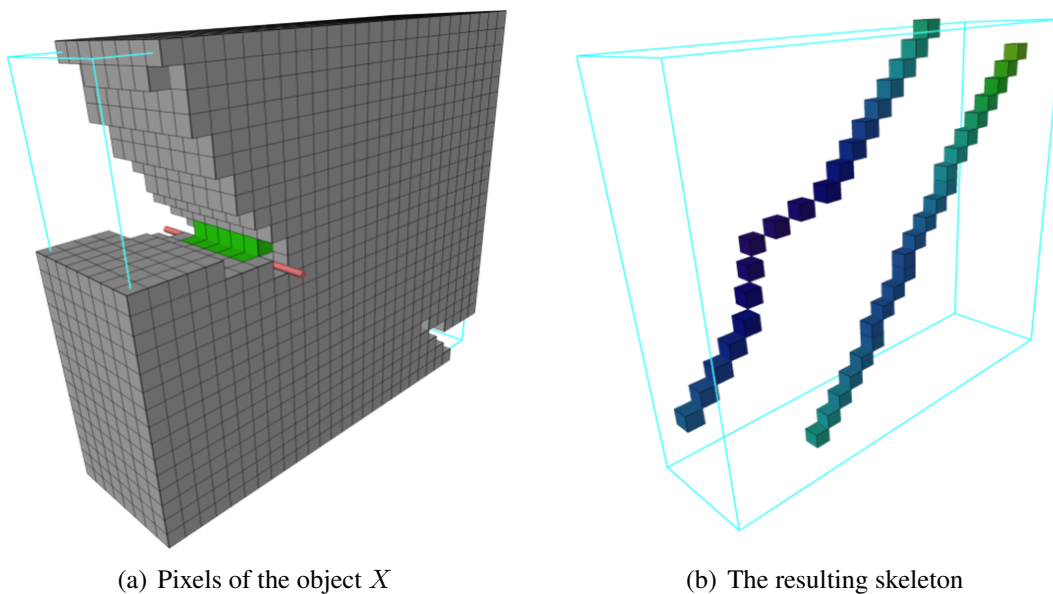


Figure 2.12: Local $2 \times 2 \times 2$ pixel configuration, under the four possible rotations, to test for the appearance of corner-wise branch-type artefacts. The filled-in cubes represent pixels of X , while the transparent ones belong to the background.



(a) Pixels of the object X


(b) The resulting skeleton

Figure 2.13: Example of an edge-wise branch-type artefact: the top skeleton branch in (b) results from the fact that the green pixels of the object in (a) are edge-connected, or that there is a tunnel in between the two balls, represented by the pink line.

Edge chain An edge chain is a union of edges $E = \{e_1, e_2, \dots, e_n\}$ such that $\forall i \in [1; n - 1], e_i$ and e_{i+1} share a common vertex.

Extremities of an edge chain The extremities of an edge chain is the set of vertices such that there exists only one edge containing each. In the case of a loop, this can be the empty set.

The property can be presented as follows: a chain of weak edges having no extremity contained in a strong edge is a nonlocal branch-type digitisation artefact. This property, illustrated in figure 2.14, needs to detect weak and strong edges, drawn in red and blue respectively in the figure. Figure 2.14(c) shows a chain of weak edges having one extremity connected to a strong edge, therefore not satisfying the condition.



(a) A configuration generating an artefact

(b) Pixels removed after processing

(c) A configuration that does not generate an artefact

Figure 2.14: Illustration of the removal of edge-wise branch-type artefacts: (a) shows the pixels of the object generating the artefact, with the detected strong and weak edges (highlighted in blue and red respectively). Since the chain of weak edges satisfy the artefact condition, our method removes the adjacent pixels, as shown in (b). Figure (c), on the other hand shows a case where a chain of weak edges does not satisfy the condition of having its extremities unconnected to a strong edge, therefore no pixels are removed in this case.

In order to present the removal algorithm, we need to define neighbourhood relations for pixel edges e . We note $\Gamma(e)$ the set of edges containing at least one vertex of e , and $\Gamma^*(e)$ the set $\Gamma(e) \setminus e$. These neighbourhoods can be divided into two subsets $\Gamma^1(e)$ and $\Gamma^2(e)$, depending on which vertex is considered. Figure 2.15 illustrates these neighbourhoods. The removal algorithm is written on page 41. It requires two additional functions: `edge(pi, pj)` to define an edge from two edge-adjacent pixels, and `pixelsContaining(e)` to find the four pixels that contain an edge.

2.3.2.3 Conclusion

This new scheme removes, by construction, the pixels responsible for artefacts in the skeleton, and only those pixels. It is designed to minimise the number of modifications to the image, and handles all possible artefacts by using an exhaustive classification. Table 2.1 shows the number of pixels modified with this method on the four images shown in figure 2.8. As a comparison, the majority filter and the morphological opening/closing is also present. In the best case scenario, the majority filter still modifies 28 times as many pixels.




Figure 2.15: Neighbourhood of an edge $e = \{v_1, v_2\}$, for the three possible orientations. Each neighbourhood can be divided in two half-neighbourhoods Γ^1 and Γ^2 depending on which vertex the edges share.

Input : Binary image $I : \mathbb{Z}^3 \rightarrow \{0, 1\}$, with $X = \{p | I(p) = 1\}$
Data : E_w, E_s set of weak and strong edges, $\{e, e', e''\}$ edges, i an integer, and $\{p, p', p_1, p_2\}$ pixels.
Output : Modified binary image I

begin

```

     $E_w \leftarrow \emptyset$ 
     $E_s \leftarrow \emptyset$ 
    foreach  $\{p, p'\} \subseteq X | (\Gamma_6^*(p) \cap \Gamma_6^*(p') \neq \emptyset)$  do
         $\{p_1, p_2\} \leftarrow \Gamma_6^*(p) \cap \Gamma_6^*(p')$ 
         $e \leftarrow \text{edge}(p, p')$ 
        if  $\{p_1, p_2\} \subseteq \bar{X}$  then  $E_w \leftarrow E_w \cup e$ 
        else  $E_s \leftarrow E_s \cup e$ 
    foreach  $e \in E_w$  do
         $i \leftarrow 0$ 
         $e' \leftarrow e$ 
        if  $(|\Gamma^1(e) \cap E_w| = 0) \wedge (|\Gamma^1(e) \cap E_s| \neq 0)$  then  $i \leftarrow 2$ 
        else if  $(|\Gamma^2(e) \cap E_w| = 0) \wedge (|\Gamma^2(e) \cap E_s| \neq 0)$  then  $i \leftarrow 1$ 
        while  $i \neq 0$  do
             $E_w \leftarrow e'$ 
            if  $\exists! e'' \in \Gamma^i(e') \cap E_w$  then
                if  $e' \in \Gamma^1(e'')$  then  $i \leftarrow 2$ 
                else  $i \leftarrow 1$ 
             $e' \leftarrow e''$ 
            else  $i \leftarrow 0$ 
        foreach  $e \in E_w$  do
             $Y \leftarrow \text{pixelsContaining}(e)$ 
             $X \leftarrow X \setminus Y$ 

```

end

Algorithm 1: Edge-wise branch-type artefact removal algorithm

Figure 2.16 plots the percentage of pixels modified versus the pore radius. Results for the median filtering is also plotted, but using the secondary axis. Finally, in order to show the influence of these modifications on a macroscopic property, we have computed the permeability of synthesised ball packings before and after our modifications, and plotted the error as a function of ball radius in figure 2.17. With sufficient resolution, we observe that the error quickly drops to around 1%, which is acceptable but not negligible nonetheless.

	2.8(a)	2.8(b)	2.8(c)	2.8(d)
Majority	668	995	1364	2537
O.C.	2416	3632	5103	9364
D.D.A.S.	24	24	34	64

Table 2.1: Number of modifications of the different digitisation artefact removal methods on the image shown in figure 2.8. The second line is the sequence of morphological opening and closing, and the third is our novel Direct Digitisation Artefact Suppression method.




Figure 2.16: Percentage of pixels of the image modified from the preprocessing as a function of the pore radius, measured on the same random ball packing digitised at varying resolutions. Our method modifies less than 1% of pixels of the image, even for small pores, whereas the median modifies an average of almost 14 times as many.




Figure 2.17: Error percentage of our artefact removal compared to the original image, on the computed intrinsic permeability as a function of the pore radius in pixels, using again the digitised ball packing, with a 10% overlap to avoid singularity effects.

Note that generation of certain artefacts, namely the 1D type, can be avoided by performing an initial skeletonisation with the 6-adjacency, followed by a second skeletonisation with 26-adjacency [Sheppard et al., 2005]. Figure 2.18 shows that similar results are obtained, but only if an ultimate skeleton is computed. If dangling ends should not be trimmed, then a 6-connected skeleton keeps too many pixels.




Figure 2.18: *Effect of using the 2 connectivities for avoiding digitisation artefacts, applied to the object of figure 2.8(d). (a) shows the skeleton after our digitisation artefact correction, while (b) shows the skeleton obtained from skeletonising the 6-connected skeleton. Note that these two results are ultimate skeletons, that is all dangling ends are trimmed. (c) shows the untrimmed skeleton after our correction, while (d) shows how skeletonising while keeping the dangling ends in 6-connectedness preserves too many pixels (the subsequent 26-connected skeletonisation will thin the object, but not trim it).*


2.4. Skeleton/graph conversion: topological classification

We have implemented a skeletonisation algorithm, and developed a pre-processing scheme for removing all unwanted information from the skeleton. We obtain a centred Euclidean skeleton, with the intent to classify pixels of the skeleton into branch intersection and branch types, effectively representing markers for the pores and marker contacts respectively. This supposes that the skeleton is curvilinear, or 1-dimensional. Our digital artefact suppression method removes artefacts that would have generated surfaces in the skeleton, but surfaces can appear for other reasons.

2.4.1 Surfaces in the skeleton

Normally if the skeleton of a pore space contains a surface this would mean that a disconnected solid particle is *floating* inside the porosity, which is physically unrealistic[†]. However, surfaces can appear at the image boundary due to cropping effects.

Microtomographic acquisitions are often cropped to focus on the interior of the material, and this cropping can cut particles located at the edges, leaving portions of the solid phase disconnected in the image because they are connected outside this region of interest. The pore space surrounds it, and because we keep the connection with the image boundary, by the method described previously, an open surface appears around this disconnected portion of a particle, the border of this open surface being on the border of the image. Figure 2.19 presents such a situation. Note that if we consider periodic boundary conditions, this cropping effect results in closed surfaces [Chaussard et al., 2008].



(a) A cropped image of a glass bead packing

(b) A surface view of the solid phase

(c) The skeleton of the porosity

Figure 2.19: Example of a non-curvilinear skeleton: (a) shows the cropped region of a glass bead packing and (b) shows a 3D view of the solid phase, which clearly shows a cut bead near the lower corner of the image. This result in an open surface in the skeleton, shown in (c)

This problem is most often avoided by preprocessing the image: Lindquist and Venkatarangan [1999] simply removes the pixels of the solid phase that created the surfaces, which in some cases can result in a

[†]It is however possible that a microtomogram depicts such a situation due to resolution or thresholding limitations, outside of situations considered as artefacts.

significant modification of the pore space.

A idea for solving this problem is to define pixels of this surface that can be removed (deletable) in a second skeletonisation process, and others that cannot (anchors). Figure 2.20 illustrates this. By construction, this process alters as little topological information as possible for obtaining a curvilinear skeleton.




Figure 2.20: Illustration of the idea of a second skeletonisation for removing the surfaces in a skeleton: (a) shows a portion of a skeleton that contains a surface in contact with the image boundary. On the intersection of the surface and the border, anchor points can be defined (circled in the figure), and a second skeletonisation, starting from the border, can iteratively remove points from the surface until only a curvilinear shape remains, as in (b).

This type of approach requires a topological classification of the pixels of the skeleton, which is not a trivial problem.

2.4.2 Topological classification of the pixels of the skeleton

Pothuaud et al. [2000] and Liang et al. [2000] use similar custom classification methods, based on an order relationship in 26-neighbourhoods (or the 3D extension of λ -adjacency in the latter). We prefer to use the classification method of Bertrand [1994]; Bertrand and Malandain [1994]; Malandain et al. [1993] because of its sound topological framework.

Three quantities are necessary to classify pixels of a skeleton: the number of neighbours in X for a pixel p , or $|\Gamma_{26}^*(p) \cap X|$, and two connexity numbers $T_{26}(p, X)$ and $T_6(p, \bar{X})$ defined in 2.2.1 page 31.

The initial classification procedure is shown in table 2.4.2, which introduces temporary labels for misclassified pixels: border, surface and junction. Two more passes on the pixels are necessary. Firstly, all pixels marked as surface are relabelled to nodes if their neighbourhoods do not contain any surface pixels. This situation can arise in complex junctions, such as the one illustrated in figure 2.21. The colour code represents the topological classification proposed in Malandain et al. [1993], where according to the definition, the turquoise pixel is a surface pixel. Locally, this pixel satisfies the condition for being a surface pixel, but the fact that none of the adjacent pixels are surface pixels means that this surface is entirely composed of one pixel, which is the result of a misclassification. Secondly, the border and junction pixels are relabelled to surface if their neighbourhood does not contain any branch pixels, otherwise they become nodes. This second scan removes misclassifications in tortuous surfaces and reclassifies surface junctions.

$T_6(p, \bar{X})$	$T_{26}(p, X)$	$ \Gamma_{26}^* \cap X $	Label
0	1	*	Interior pixel (no occurrence in the skeleton)
1	0	*	Node (Isolated pixel)
1	1	1	Node (End of a Branch)
1	1	>1	<i>Border</i>
1	2	2	Branch
1	2	>2	Node (Branch junction)
1	>2	*	Node (Branch junction)
2	1	*	<i>Surface</i>
2	>1	*	<i>Junction</i> (Branch-surface junction)
>2	*	*	<i>Junction</i> (Branch-surface or surface-surface junction)

Table 2.2: Table summarising how the pixels of the skeleton are classified for the graph conversion. Labels in italics are subject to post-processing to avoid misclassification.




Figure 2.21: Example of initial misclassification of pixels of a skeleton: yellow pixels are border pixels, blue are branch junctions, green are branch pixels, the turquoise a surface pixel, and the pink a branch-surface junction (beneath which there is no pixel).

This topological classification of the pixels of a skeleton allows us to identify surfaces, which are connected components of pixels labelled as surface pixels, and propose a new method for handling them. If the surface is closed (when periodic boundary conditions are used), the surface is *pierced* by removing the surface pixel closest to the solid phase. If the surface is open, we choose the pixel on the surface boundary farthest from the solid and anchor it, all others marked as deletable. A second skeletonisation then iteratively removes simple points, starting from the surface border. An example, taken from the skeleton of the pore space of a particle packing, is presented in figure 2.22.




Figure 2.22: Portion of a skeleton of the pore space in a particle packing. The open surface, shown in figure (a) is in contact with the image border along an edge of the bounding box. After processing the surface, the pixels in (b) remain.


Once the skeleton contains no more surfaces, this classification is identical to a simple neighbour counting [Shin, 2002]. When all pixels are classified either as node or branch, it is a simple matter to convert the information to a graph. Firstly, all nodes are defined as vertices, and secondly, branch pixels are *followed* from one node pixel to another in order to define the edges of the graph. This new graph structure in fact contains more information than just vertex and edge sets, that will be useful for post-processing and delimitation. For instance, with each vertex v is associated the value $d(v)$, which is the distance of v (or the corresponding node pixel) to the solid phase.

With this classification, as can be seen in figure 2.21, several node pixels can be adjacent to each other. These node pixels clearly identify only one branch intersection, but topologically speaking, this intersection cannot be associated to strictly one pixel. Therefore connected components of node pixel are merged into one marker[†]. To avoid any loss of information, the vertices of the graph are not internally merged into one vertex, but are assigned an identical label l that identifies this merging. In the rest of this thesis, we refer to a marker as a cluster of vertices (vertices with identical labels) in the graph.

2.4.3 Topology preservation during conversion

The assumption that regions defined by nodes clusters are all homotopic to a ball, i.e. contain no holes or cavities, is not necessarily true. If we consider the 2D example of figure 2.23, the branch consisting of one pixel has both endpoints belong to the same cluster, therefore the edge itself is considered as being part of this cluster, which contains a hole.

[†]Pothuaud et al. [2000] Calls this deletion of the nil length skeleton graph branches.



(a) Skeleton before conversion to a graph. The pixels in grey are classified as curve junctions. (b) Graph transformation, the 3 nodes belonging to the same cluster (c) Non-homotopic representation of the cluster

Figure 2.23: 2D illustration of how a node cluster can contain holes, and thus not necessarily be homotopic to a ball. The centre edge path contains only one intermediate pixel and joins to curve junction pixels to be merged. If the cluster was represented as a single vertex, the topology changes.

Fortunately, because of the digitisation artefacts removal applied to the image before skeletonisation, such situations never arise. This does not remove the possibility that cycles appear in clusters during the node merging post-process, although the probability of this phenomenon occurring is low.

2.5. Graph post-processing

The graph produces an initial approximation of the pore network that we wish to obtain. It provides the essential topological information of the pore space, but geometrical features also need to be taken into account. A pore network is a graph that identifies each vertex to a pore, an object described in geometric terms: a local void, surrounded by the solid phase, separated by constrictions etc. Figure 2.24 shows a portion of the skeleton of a ball packing, where three distinct branch intersections appear around the ball centre. From the previous conversion process, three clusters are defined inside the sphere, each defining a separate marker for the delimitation. In fact, it is unlikely that the branches of the skeleton all intersect at the pore centres, resulting in a greater number of markers than required, prompting the need for a merging method.


2.5.1 Types of node cluster merging

2.5.1.1 Classical methods


The process of merging markers (whether obtained from the maxima of the distance map or the branch intersections of the skeleton or medial axis) is a typical procedure in pore decomposition[†]. The merging condition must consider the extent to which they belong to the same *opening* in the porosity. This latter property can be characterised by the height of the markers[‡] (or if the medial axis was used to find these,

[†]We assume the merging of connected node pixels from the graph conversion has been previously applied.


[‡]For more readability, we refer to the distance to the solid phase as the height.



(a) Portion of a ball packing, with the centre ball highlighted



(b) Portion of the skeleton inside the centre ball



(c) Portion of the converted graph

Figure 2.24: *Illustration of the need for a merging process: (a) shows a portion of a random ball packing, the centre ball being highlighted. (b) shows the part of the skeleton inside the centre ball, and (c) shows the subsequent graph conversion. The three connected components of node pixels are merged during the conversion process into three distinct markers, but a second merging process is required for node clusters not directly adjacent, but sufficiently close in regards to the surrounding geometry.*

the radii of their corresponding maximal balls), and either the distance between them, or the size of the constriction between them (determined by the minimum height of the branch connecting them). A list of the merging methods found in the literature is presented in table 2.3.

Method	Paper
$d_1(C(m_1, m_2)) \leq \alpha \min(d_1(m_1), d_1(m_2))$	Marcke [2008]
$\ \vec{m_1 m_2}\ _\infty < d_\infty(m_1) + d_\infty(m_2)$	Al-Raoush and Alshibli [2006]; Al-Raoush and Willson [2005]
$\ \vec{m_1 m_2}\ _\infty \leq \min(d_\infty(m_1), d_\infty(m_2))$	Shin [2002]; Shin et al. [2005]
$\ \vec{m_1 m_2}\ _\infty < \max(d_\infty(m_1), d_\infty(m_2))$	Lindquist et al. [2000]; Thompson et al. [2006]
$\ \vec{m_1 m_2}\ _\infty < \max(d_\infty(m_1), d_\infty(m_2)) + 2$	Prodanović et al. [2006]
$\ \vec{m_1 m_2}\ _2 < \lfloor d_2(m_1) \rfloor + \lfloor d_2(m_2) \rfloor$	Al-Kharusi and Blunt [2007]; Silin et al. [2004]
$\ \vec{m_1 m_2}\ _2 < \max(d_2(m_1), d_2(m_2))$	Delerue [2001]
$\ \vec{m_1 m_2}\ _2 < \max(d_2(m_1), d_2(m_2)) + 4$	Saadaftar et al. [2006]
$f\left(\frac{\min(d_2(m_1), d_2(m_2))}{d_2(C(m_1, m_2))}, \frac{\ \vec{m_1 m_2}\ _2}{d_2(C(m_1, m_2))}\right)$	Sheppard et al. [2005]

Table 2.3: Merging methods, where m_i are the markers, and $d(x)$ the distance from the solid phase. We denote by $C(m_i, m_j)$ the pixel farthest from the solid and lying on a separation between the delimited regions marked by m_i and m_j .

Omitting Marcke’s method, in which a definite percentage α is not set, and Sheppard’s approach, which uses an unexplicited non-linear function for an iterative merging (where the merging of two markers is dependent on the entire porosity structure), there are two classical types of merging, summarised as marker inclusion and ball intersection. We denote by $\mathcal{B}(m, d(m))$ the ball centred on m with a radius of $d(m)$ the height of m , and consider two markers m_i and m_j . Marker inclusion checks whether m_i is contained in $\mathcal{B}(m_j, d(m_j))$ (or vice versa), and ball intersection checks if $\mathcal{B}(m_i, d(m_i)) \cap \mathcal{B}(m_j, d(m_j))$. Figure 2.25 illustrates these two merging methods in 2D. Note that for any merging method, if the markers m_i and m_j are defined by vertex clusters $\{v_i^1, v_i^2, \dots, v_i^n\}$ and $\{v_j^1, v_j^2, \dots, v_j^m\}$, then the condition must be verified for at least one pair of vertices v_i^x and v_j^y .




Figure 2.25: 2D illustration of two common merging methods (shown with the use of the Euclidean norm), in a figurative pore space. The solid phase is shown in grey, and the markers m_1 and m_2 are the branch intersections from the skeleton. (a) is an example of marker inclusion: m_2 is included in $\mathcal{B}(m_1, d(m_1))$. (b) shows ball intersection: the set $\mathcal{B}(m_1, d(m_1)) \cap \mathcal{B}(m_2, d(m_2))$ is non-empty. Note that the former condition is more strict than the latter, as marker inclusion implies ball intersection.

Some authors modify this merging process by adding a delta in the equation, to account for surface irregularity (the circularity residue R_c in [Saadafar et al., 2006] for instance). Because it is a constant value, it will favour merging of small regions. It can also be argued that the ball intersection condition is too loose and will merge too many markers together, and in relatively high-porosity unstructured media, it will result in one big pore comprising a vast majority of the pore space. In the end, there is no clear answer as to which method is optimal in real porous media, as the decomposition into pores is not unique [Prodanović et al., 2006], although the merging can have a significant (and non-deterministic) impact on the results [Marcke, 2008].

What is clear in these methods is that they try to estimate the amount of redundant *pore information* identified by each marker. If two markers share a given amount of information, then we can assume they identify the same pore. If this is indeed the underlying idea guiding a merging process, and if a strong hypothesis is that *pore information* can be characterised by the maximal balls $\mathcal{B}(m, d(m))$, then we propose that an effective way of determining this redundancy is to directly compute the ratio of overlapping volume between balls.

2.5.1.2 Determination of the common volume between intersecting balls

We consider the markers m_1 and m_2 , and we note $r_1 = d(m_1)$, $r_2 = d(m_2)$ and $l = \|\overrightarrow{m_1 m_2}\|$. We want to determine the volume $\mathcal{B}(m_1, r_1) \cap \mathcal{B}(m_2, r_2)$. In the following, we suppose that $|r_1 - r_2| < l < r_1 + r_2$. The first part of the inequality means that one maximal ball is not contained in the other. The second part simply means that there is an overlap. It is interesting to note that, by definition, if $l < r_1 + r_2$ then $|r_1 - r_2| \leq l$. If this were not the case, that would mean that either $r_1 \neq d(m_1)$ or $r_2 \neq d(m_2)$.


Figure 2.26 shows the values needed for calculating the intersection volume: (a) illustrates a cross-section along the line $(m_1 m_2)$, and (b) shows this intersection, which corresponds to a union of two spherical caps, \mathcal{V}_1 being that of $\mathcal{B}(m_1, r_1)$. The point p is the orthogonal projection on $(m_1 m_2)$ of a , the intersection between the circles $\mathcal{C}(m_1, r_1)$ and $\mathcal{C}(m_2, r_2)$ defined by the ball boundaries on the cross-section. Finally, 2.26(c) represents a ball $\mathcal{B}(m, r)$ and an associated spherical cap, with $h = r - \|\overrightarrow{m p}\|$.

To find \mathcal{V}_1 , we need the length h , which is equal to $r_1 - \|\overrightarrow{m_1 p}\|$, and $\|\overrightarrow{m_1 p}\| = r_1 \cos \theta_1$. $\cos \theta_1$ is found with Al-Kashi's theorem on the triangle $(m_1 a m_2)$: $r_2^2 = r_1^2 + l^2 - 2r_1 l \cos \theta_1 \Leftrightarrow \cos \theta_1 = \frac{r_1^2 + l^2 - r_2^2}{2r_1 l}$.

Finally, the volume of the spherical cap shown in figure 2.26(c) can be written $\frac{\pi h^2}{3}(h - 3r)$. Consequently, the overlapping volume $\mathcal{V}_1 + \mathcal{V}_2 = \mathcal{V}(\mathcal{B}(m_i, r_i) \cap \mathcal{B}(m_j, r_j))$ can be written $\frac{\pi h^2}{3}(h - 3r_1) + \frac{\pi h^2}{3}(h - 3r_2) = \frac{\pi h^2}{3}(2h - 3(r_1 + r_2))$.

The condition should depend on the ratio between the volumes of the intersection and the smallest of the two balls, because if we consider balls with a large variation in radius, the ratio with the largest one can become insignificant. The condition then depends on the percentage of overlapping volume, and is written


$$\frac{\mathcal{V}(\mathcal{B}(m_i, r_i) \cap \mathcal{B}(m_j, r_j))}{\min_{k \in \{i, j\}} \mathcal{V}(\mathcal{B}(m_k, r_k))} < \alpha \implies m_i \text{ and } m_j \text{ are merged}$$




(a) Cross-section of the maximal balls along the line joining their centres

(b) The common volume between the balls

(c) A spherical cap



(d) Illustration of 2 overlapping balls



(e) Their common volume

Figure 2.26: Diagrams for calculating the overlapping volume of intersecting balls.

where α is the threshold percentage, the guiding parameter of this merging process. The choice of this threshold percentage remains subjective [Prodanović et al., 2006] and shows the difficulty of identifying pores in a porous structure.

Tests have shown that a percentage too low (0-10 %) can lead in highly porous structures to a snowballing effect [Sheppard et al., 2005], or excessive merging, and too high a value leads to an unapplicable merging condition. A percentage around 30% produces acceptable decompositions. Figure 2.27 illustrates some overlapping ratios for equal-sized balls.




Figure 2.27: Illustration of overlapping volume ratios between two balls of equal radius. The threshold value for when the corresponding nodes should be merged remains an adhoc parameter.

2.5.2 Node insertion

It was said on page 30 that using the maxima of the distance map alone was insufficient to account for all the regions in the pore space, making it preferable to integrate topological information through the use of skeletonisation. Unfortunately, in some cases, the skeleton can also prove insufficient, where a significant maximum in the distance map is not marked by a branch intersection in the skeleton. Baldwin et al. [1996] illustrates this problem with a chain of 3 pores. Figure 2.28 shows this situation. The centre element is traversed by a branch of the skeleton, with no associated node pixel, even though the constriction between the centre and right discs is as significant as the others. Such type of situations have been shown to exist in real and simulated porous media [Ioannidis et al., 1997].

It is therefore appropriate to insert markers where such openings appear along branches of the skeleton. These openings are represented by local maxima along the branches, but to avoid the same problem of excessive marking as mentioned in section 2.1 page 27, we only insert the significant local maxima, measured by their dynamics. Maxima dynamics are defined in the same way on the skeleton, but are computed much more easily because on a 1-dimensional object that is the skeleton branch. A low value between 1 and 3 is sufficient for avoiding insignificant maxima.




Figure 2.28: 2D example of how the use of the vertices of the skeleton as markers for the watershed can create an insufficient number of pores. Figure (a) shows the object to partition, in black, (b) its skeleton, which has only two vertices on the centerline. Subsequently, only two pores are separated in the centre, when it might be more pertinent to have defined a third.

The node insertion and merging post processings are two non-commutative operations, as markers created along edges can satisfy the merging condition with neighbouring markers, as shown in figure 2.29. The depicted situation shows the need to apply node insertion prior to merging.

Conclusion

We have compared various methods for positioning markers to define pores in a pore space, and shown how both geometrical and topological information need to be examined for a reasonable marker positioning. Using the maxima of distance map considers only geometrical *openings* in the porosity, which proves insufficient in some cases, therefore a skeletonisation procedure was developed. The skeleton contains all the topological information of the pore space, even the one due to digitisation artefacts. We have presented a novel and efficient method to handle all types of artefacts, so as to only conserve the pertinent data. The skeleton, converted to a graph, provides marker locations at its vertices. It was also shown that this topological information alone was insufficient to properly represent the pore space, and thus post-processings were added to account for geometric features, among which a new node merging method that computes the amount of redundancy in the vertices. The method is almost entirely automatic, the only parameter that can be varied is the overlap volume percentage for the merging condition, as no methodology exists to assess the quality of such a choice.

Definition of these markers is a starting point for a decomposition method to effectively assign every pixel of an object to an element, or potential pore. The next step of delimitation is described in the next chapter.




Figure 2.29: Example in a random ball packing of the non-commutativity of the graph post-processing operations. (a) shows a portion of a digitised ball packing, highlighting a specific ball, and (b) showing the portion of the skeleton inside it. (c) shows the ball defined by the marker generated inside this ball after application of the merging method, and (d) shows the new marker created from the node insertion post-processing, which almost entirely covers the previous marker. This shows the need to apply the merging method last.



3 Pore delimitation

Contents

3.1	Region-based methods	60
3.1.1	Region growing	60
3.1.2	The topological watershed on the distance map	63
3.1.3	Modifications to the watershed approach	65
3.1.4	Generating thin separations	65
3.1.5	Seed maximisation	65
3.1.6	Conclusion	67
3.2	Separating planes and surfaces	67
3.2.1	Separating planes	67
3.2.2	Separating surfaces	69
3.2.3	Limitations	72
3.3	Validation	75
3.3.1	Image generation	75
3.3.2	Procedure	76
3.3.3	Relabelling for similar partitions	78
3.3.4	Results	78
	Conclusion	81

Pore volume delimitation is a particular example of the problem of object segmentation in image analysis. The general goal of segmentation is to assign a label to every pixel in the image. Pixels with the same label share a common visual characteristic (intensity, texture, color, etc.), and are separated from other labels by discontinuities in that characteristic [Stockman and Shapiro, 2001]. In the case of pore space decomposition, we consider the object X of the image that is the porosity (all other pixels are assigned a distinct label identifying \bar{X}), and its partitioning into n elements or pores P_i , with $i \in [1, n]$, the problem can be written as follows:

- $\bigcup_1^n P_i = X$
- $\forall i \in [1, n]$, P_i is a 6-connected region
- $\forall (i, j) \in [1, n]^2$, $i \neq j \Rightarrow P_i \cap P_j = \emptyset$

Pore space decomposition is different from typical image segmentation in several key aspects. First of all, the images are in three dimensions: this excludes certain segmentation methods which are specific to two-dimensional images, such as snakes, geodesic active contours, and any curve propagation technique. Secondly, the non-binarised image contains no visual cues (i.e. discontinuities of a given visual characteristic) as to where partitions are located nor which pixels belong in the same region, every pixel of the void phase having a constant intensity, making a majority of typical segmentation methods, such as clustering or histogram-based thresholding, irrelevant, and others need to be adapted.

Furthermore, the problem defined above requires more constraints for a representative pore decomposition, the first of which is the use of the markers defined in the previous chapter. This implies that n is the number of markers, and for $i \in [1, n]$, m_i must be contained in P_i .

Two main approaches to pore delimitation are found: region-based and boundary-based methods[†]. The first method iteratively assigns pixels to a given region, starting from the markers and spreading outwards until all pixels of the object are assigned. This naturally creates separations between regions. The second method directly constructs these separations, by building planes or surfaces where a hydraulic minimum is assumed to be, naturally delimiting pore volumes.

Section 3.1 and 3.2 presents these two types of methods, and section 3.3 attempts to determine their robustness through application on a series of synthesised images.

[†]Also called pore-body detection and throat construction, respectively [Prodanović et al., 2006].

3.1. Region-based methods

3.1.1 Region growing

Region growing relies on a simple concept to partition an image. Beginning at the markers, neighbouring pixels with similar properties are grouped together. The regions, initially defined by the markers, expand until the entire space is covered [Pratt, 2001]. Because in the case of pore decomposition there is no measure of similarity between pixels of X , the determination of the groups will depend solely on the propagation, and the logical method for controlling it is to define a propagation speed for each pixel.

With a constant propagation velocity field, the result is equivalent to the skeleton by influence zone [Delerue, 2001]. The logical choice for the propagation speed is the distance map. The farthest a pixel is from the solid phase, the higher priority it is given for being assigned to a neighbouring marker. Internally, this means that at each iteration, a pixel farthest from the solid and having unassigned neighbours has these neighbours assigned his label.

Figure 3.1 illustrates this method on a simple geometry, and shows several intermediate stages. With the processing priorities given by the distance map, growth of the upper region does not begin immediately, since the grey level of its pixels are inferior to those of the lower region. It allows to position the separation closer to constrictions in the image.

Because the distance map is used, the subsequent separations should not exhibit preferential orientations due to the grid. This is unfortunately not the case, as is shown in figure 3.2. Two pores have been extracted from a decomposition of the interstitial volume of a glass bead packing. The contact between the two forms a type of 3D wedge parallel to the grid, clearly irrelevant to the surrounding geometry.

The reason for this is illustrated in figure 3.3. When two regions come into contact in the diagonal directions of the cubic grid, they will extend side by side in parallel but opposite directions, thus creating this wedge-type artefact.

Another type of region-based segmentation method is the watershed. Originally proposed by Lantuéjoul [1978], the watershed transform is a common segmentation approach in greyscale mathematical morphology. Typically defined in flooding semantics [Vincent and Soille, 1991], the principle is to consider the image as a relief (c.f. figure 2.2 on page 28), and find the crest lines between basins.

The classical watershed algorithm is actually very similar in its implementation with region growing. During the *flooding*, the unlabelled pixels are examined: if all pixels of their neighbourhood contains either unlabelled pixels or pixels with the same label, then it is assigned it, and its unlabelled neighbours are added to the queue. Other methods of computing a watershed exist. A recent algorithm that has been proven to conserve the most significant contours of the original image is called the topological watershed [Bertrand, 2005]. This algorithm is based on a rigorous framework for defining and characterising the watershed [Coupré and Bertrand, 1997] and is described in the following section.




Figure 3.1: Illustration of the process of region growth. Figure (a) shows the object to decompose, made up of two balls, (b) shows a cross-section of the distance map used as propagation velocity field for the region growing. Figure (c) shows the initial state t_0 of the process, and (d) through (f) show different intermediate steps: in the beginning the white element has priority since the values of the neighbouring pixels of the region are greater. At some point, the dark element begins growing, and the two propagation fronts meet at the centre of the neck between the two balls. Figure (g) shows the final state at time t_∞ , i.e. when all pixels have been assigned a label.




Figure 3.2: Illustration of a grid artefact appearing from the use of region growing method. Figure 3.2(a) shows two adjacent pores obtained from a region growing delimitation (the pore are in the interstitial volume of a bead packing), and figure 3.2(b) shows the white pore, with its faces in contact with the other pore coloured in grey. A wedge artefact is clearly apparent and adversely affects the representativeness of the delimitation.




Figure 3.3: *Illustration of the reason for the wedge effect in region growth. If two regions arrive in contact diagonally, each will grow parallel to each other in the 3 directions of the grid axes.*

3.1.2 The topological watershed on the distance map

This algorithm consists in a sequential raising of the pixel grey levels until convergence. A pixel of level k is raised (i.e. their grey level k is increased) on the condition that it does not alter the topology of the higher cross-section \mathcal{F}_k^+ (c.f. section 2.1.2 page 29). The only way it can alter the topology is either to create a new connected component in \mathcal{F}_k^+ , or merge components. Therefore it can be raised only if it is adjacent to exactly one component. Such pixels are called *constructible*. The process is reiterated until no more pixels are constructible. It can be noted, for instance, that in figure 2.3 page 29 all pixels having a grey level of 2 in \mathcal{F} are not constructible because they are neighbour to two connected components in \mathcal{F}_2^+ , which is equal to \mathcal{F}_3 . The result is a set of regional maxima separated by pixel-thin crest lines. These lines define the topological watershed.

3.1.2.1 Algorithm

Before any constructible pixels can be raised, the component tree is constructed. The implemented method uses the works of Najman and Couprie [2004] and Couprie et al. [2005]. Figure 3.4 illustrates this process. For purposes of relevance, the example of the figure has the structure of a distance map. The method initially constructs a tree structure containing the branch components, as in figure (b), and a post-process produces the component tree represented in figure (d) by removing those components. Of course, the component membership of the pixels must be updated: this is illustrated in figure (c), where the grey levels of the pixels have taken the value of the component to which it belongs.

If a pixel p , belonging to the component $\mathcal{C}(p)$ and with intensity k , is constructible, then all points in \mathcal{F}_k^+ in the neighbourhood of p belong to the same connected component in \mathcal{F}_k^+ . This does not always mean that those neighbours belong to the same component in \mathcal{T} , but rather that the components to which they belong have a common ancestor in \mathcal{T} that is a descendant of $\mathcal{C}(p)$. Therefore, if *one* common ancestor is a descendant of $\mathcal{C}(p)$, then the lowest common ancestor is a descendant of $\mathcal{C}(p)$. Checking the constructibility of a point p is therefore summarised by verifying that the lowest common ancestor of the components $\{\mathcal{C}(p'), p' \text{ is a neighbour of } p\}$ is a descendant of $\mathcal{C}(p)$. Raising the pixel p thus means changing its membership $\mathcal{C}(p)$ to the lowest common ancestor. Figure 3.5 illustrates this process with the central pixel, that has a grey level of 0, which is constructible.

The lowest common ancestor lookup is one of the fundamental problems in algorithmics for tree structure manipulation. Since this lookup must be made at least once each time a point is checked for constructibility, it should be done in minimal time. This problem has been extensively studied, and with a linear preprocessing of the tree, the lookup can be made in constant time [Bender and Farach-Colton, 2000].

3.1.2.2 Sequence of constructible pixels to raise

A watershed of an image is rarely unique (if there are regional maxima, then there is more than one path a crest line can take). A very important factor in all watershed algorithms, in the same way as for




Figure 3.4: Illustration of the component tree on a 2D distance map \mathcal{F} . The initial tree structure contains the branch components, which are not essential in characterising the topography of the image. The structure is then reduced to \mathcal{T} , and the image can be modified to \mathcal{F}' without any loss of information.




Figure 3.5: Relation between the lowest common ancestor lookup in the component tree, and the number of connected components in \mathcal{F}_0^+ : (a) is the original image, and (b) its component tree. When considering the central pixel p with a grey level of 0, its neighbours with a strictly higher grey level (its upper and lower neighbours p' and p'' , with intensities of 4 and 3 respectively) do not belong to the same component in the tree, nevertheless they belong to the same connected component in \mathcal{F}_0^+ (shown in figure (c)): this is translated in the tree by having the lowest common ancestor of $\mathcal{C}(p')$ and $\mathcal{C}(p'')$ be a descendant of $\mathcal{C}(p)$, i.e. the component with a value of 1.

skeletonisation algorithms, is the order in which the pixels are processed. To ensure that all pixels will be raised at most once, those having the highest grey level are processed first (c.f. [Couprie et al., 2001] for more details).

To ensure this order, the pixels to process are first added in a hierarchical queue [Meyer, 1990], i.e. a queue that sorts the pixels according to their grey level. It is called a hierarchical queue because if more than one pixel has the same value, then the first one added will be processed first (First In First Out principle). When a pixel is to be processed, the one with the maximum value that was first added to the queue is selected.

The queue is first filled with all points neighbouring the leaf components of \mathcal{T} , or the regional maxima. The pixels are removed from the queue and processed one by one, and whenever a pixel is found to be constructible and thus raised, its neighbours are added to the queue. The system converges when the queue is empty, and the result is the upper homotopic kernel. The pixels p of the crest line of the watershed are those such that $\mathcal{C}(p)$ is not a leaf component.

3.1.2.3 Eliminating oversegmentation

The question of oversegmentation has already been discussed in section 2.1. If the watershed was applied on the distance map, every local maximum would define a separate region. The task of selecting a proper set of markers that was done in the previous chapter can nonetheless be injected into the calculation of the watershed. All markers will define new maxima in the distance map, while any unmarked maxima are lowered to the level of their component mappings.

3.1.3 Modifications to the watershed approach

3.1.4 Generating thin separations

Watershed-type methods commonly define crest lines (or surfaces, in 3D) that are one-pixel thick, meaning that not all pixels of X are assigned a pore P_i . This problem can be easily circumvented by arbitrarily assigning to a neighbouring basin the pixel of a watershed. The implemented solution actually used the region growing approach on the watershed pixels. This hybrid method, in essence, assigns watershed pixels to the neighbouring pore P_i with highest value in the distance map.

3.1.5 Seed maximisation

An important difficulty in the combination of the watershed of the distance map and markers defined by independent information (the skeleton) is occasional incompatibilities. The markers do not necessarily fall within distinct maxima in the distance map. If several were to share the same maximum (or at least belong to the influence zone of that maximum), then the marker with initial highest value will *encroach* on the adjacent pores.

A solution for reducing this type of complication is to increase the size of each region prior to the watershed, so as to begin placing the separation appropriately. We have stated in section 2.5.1.2 page 51 that the balls $\mathcal{B}(m_i, d(m_i))$ contained information regarding their corresponding pores P_i , therefore it is natural to assume that all pixels belonging distinctly in these balls also belong to P_i . Figure 3.6 illustrates the process on two pores extracted from the interstice between glass beads.




Figure 3.6: *Illustration of the effect of seed maximisation when two markers share the same basin in the distance map. The pore space is the interstice in a bead packing, a subvolume and its skeleton was extracted, shown in (a) and two subsequent pores considered, shown in (b). When the node pixels are used as markers for the watershed, this results in one pore invading most of the surrounding volume, leaving the adjacent marker in near contact with their separation, as shown in (c). The purpose of seed maximisation is to prelabel all pixels of the balls $\mathcal{B}(m, d(m))$ defined by the markers: (d) shows these pixels for the green pore, along with the preformed beginning of the separation between the two. (e) shows the complete separation after the watershed: the blue pore preserves more of the local geometry it was initially representing.*

It was also proposed to use, for the relief, a combination of the distance map and the geodesic distance from the markers, so as to privilege the geodesic distance when close to the markers, and fade to the distance map when moving away. In effect, this alternative would present similarities with the previous solution, in the sense that it would effectively assign by default pixels within a certain range of the markers, but that range would be constant, because of the geodesic distance. In the selected method, that range

depends on the local geometry for each markers since the balls $\mathcal{B}(m, d(m))$ are used.

3.1.6 Conclusion

Region-based segmentation methods are a natural choice for the problem of pore separation. Unfortunately, region propagations cannot be stopped by discontinuities in pixel color or texture because none exist. Therefore simple region growing, with propagation speed dependent on the distance from the solid, creates significant angular wedges between the partition, despite the use of the Euclidean distance map, rendering the method unsuitable. A better delimitation is provided by watershed-type algorithms. The topological watershed was implemented because it was proven to obtain the most significant crest lines in a relief. Note that watersheds are defined for the entire image, but our implementations only applies it to the pore space.

Contrary to simple region growth, separations from this watershed algorithm are much more representative of the local geometry. One little inconvenience is that the separation is one-pixel thick, while our problem is to assign *every* pixel of the pore space to a pore. Therefore a post-process is defined to assign these watershed pixels to an adjacent pore.

Finally, the problem of occasional incompatibilities between the markers and the distance map is presented. We reduce this effect by pre-labelling all the pixels in the balls $\mathcal{B}(m, d(m))$.

From this decomposition method, it is a simple matter of examining each pixel neighbourhood to identify the actual separations between elements, but the inverse approach, i.e. constructing the separations without identifying the volumes, is in theory just as reasonable.

3.2. Separating planes and surfaces

The second approach directly constructs the separations between the pores, naturally decomposing the porosity into distinct regions. It assumes that the number and approximate positions of every separation is known. Just as the region growth method needs a set of markers to identify the regions to separate, throat construction methods need throat markers, which are given by the skeleton of the pore space. Every throat construction algorithm developed assumes that a throat has a corresponding branch in the skeleton[†].

Two methods were tested. The first and earliest method is based on a strict definition of a pore, while the second is more flexible and better accounts for the difficulties of complex porous media.

3.2.1 Separating planes

This first method is based on a rigorous application of Dullien's definition for pores: "portions of the void space confined by solid surfaces and planes erected where the hydraulic radius of the pore space exhibits local minima" [Dullien, 1991]. The skeleton of the pore space, used in the previous chapter to

[†]The reciprocal, however, is not assumed.

identify the elements to decompose, can also provide the locations of the separations. If the markers are placed at the openings inside the pore space, then there can only be constrictions connecting them, and such connections are represented by the branches of the skeleton. The problem is then to best position a separating plane along each branch[†].

A plane can be defined by two parameters: an origin o and a normal \vec{N} . The normal is accepted to correspond to the local tangent along the skeleton branch. The determination of the origin can differ. For instance, Bakke and Øren [1997] build an approximate planar separation centred on each pixel of the branch, and select the plane origin according to the separation with minimal average radius. Another method is simply to use the branch pixel with lowest value $d(p)$ [Liang et al., 2000], which is equivalent for simple geometries. For purposes of speed, we have implemented the latter method. The plane normal \vec{N} is then defined by the two neighbouring branch pixels p_i and p_j , i.e. $\vec{N} = \overrightarrow{p_i p_j}$.

The plane is defined as a connected set of pixels in the object. To find these pixels, the algorithm initiates at the plane origin, examines the origin's neighbours, selects those that are *close* to the plane (if a pixel is considered as a unit size cube, then a pixel is selected as belonging to the separating plane if the cube that represents it intersects the plane), then examines their neighbours iteratively, in the same way as in a region growth, only constrained to the plane. It converges when the solid completely surrounds the planar surface.

This approach functions only on very simple cases. Firstly, digitisation can cause errors in the defining parameters of the plane. It often occurs that at least two pixels have the minimum value along a branch. A solution to this is to use the barycentre of the branch pixel with minimum value as plane origin. The plane normal can also be affected by digitisation, since it is defined by the local tangent along the branch, i.e. the two neighbouring pixels, leaving a limited number of possible orientations, dependent on the cubic grid. A way to reduce this dependence is to use more than the direct neighbours, but several pixels on either side (with the previous modification, i.e. the use of the minima barycentre, these neighbouring pixels are chosen as closest to the plane origin). The linear regression of these pixels define the plane normal. The number of points on which to perform the linear regression should be big enough to smooth out staircase effects, but small enough to keep the branch curvature. Figure 3.7 shows this process on a branch of the skeleton of the ball packing shown in figure 2.8(c) page 36.

We observe that these few improvements are insufficient in certain cases, where planes spread through the entire image volume, intersecting other planes and *cutting* neighbouring pores in two. To circumvent this, Liang et al. [2000] proposes to remove the largest overlapping planes. A very simple way to implement this is to perform a first calculation of the planes only to determine their sizes, then order them, begin with the smallest one, and forbidding the traversal of one plane by another.

Figure 3.7 shows that the planar condition is too strict (another example is presented in figure 3.10(b)), and this method, even on relatively simple geometries, is not well adapted to detect constrictions. For this reason, more recent work have proposed methods that are not limited to planes.

[†]Of course, if a branch connects two vertices that have been merged, i.e. labelled as belonging to the same element, then no separation should be created.




Figure 3.7: Illustration of the separating planes method on a contact between two balls taken from the packing of figure 2.8(c). (a) shows two balls with portion of the skeleton. (b) shows the generated plane along the branch connecting these two balls, with the bottom one displayed. The local tangent along the branch defines a poorly oriented separating plane, an effect reduced by smoothing the tangent along the branch, resulting in (c).

3.2.2 Separating surfaces

A constriction between two pores can be seen as a minimal surface normal to the branch of the skeleton that represents their contact. An efficient way to determine this surface is to begin by locating its boundary, which is a one-dimensional loop on the solid-fluid interface that winds (once) around the corresponding skeleton branch. A strong supposition for justifying this method is that the set of pixels of the interface that are closest to the branch and that constitute a loop surrounding it is the boundary of the minimal surface. In this case, the problem consists in finding a loop of the closest solid pixels around the skeleton branch.

Lindquist and Venkatarangan [1999] dilates the skeleton branch, which becomes a curved cylinder, until it comes in contact with pixels of the solid phase. When a closed loop of contact points encircles the cylinder, the boundary is obtained [Lindquist et al., 2000; Neethirajan et al., 2008]. This algorithm implicitly uses the L_1 or L_∞ norm to find the closest pixels, depending on the connectivity used for the dilation.

The solution we have tested is based on a faster wedge-based method [Shin, 2002], that first specifically searches for the closest pixels in angular wedges around the branch. The wedges are defined by a cylindrical coordinate system oriented by the tangent \vec{N} along the branch. We select the origin of the system o in the same way as in 3.2.1, as well as smoothing the tangent \vec{N} . In each cylindrical wedge k ($k \in [1; N]$) of angle $\frac{2\pi}{N}$, illustrated in figure 3.8, the closest solid pixels p_k is found (with the Euclidean norm L_2).

The k pixels $P = \{p_1, p_2, \dots, p_k\}$ gives a rough lower bound approximation of the separating surface perimeter, where each pixel p_i is the closest solid pixel from o . The complete border is a 6-connected set of




Figure 3.8: Illustration of an angular wedge around a skeleton branch, in which to search for the closest solid pixel.

solid pixels containing P . Therefore the second step is to *fill the gaps* in between each pair $(p_i, (p_i + 1) \% k)$ for $i \in [0; k - 1]$. This is easily performed by a breadth-first search in the graph defined by the pixels of the solid at the interface with the pore space and the 6-adjacency relation[†]. To simplify this refinement step, Shin et al. [2005] first computes the *local grain boundary*, or the set of solid pixels on the interface that lie closest to this skeleton branch, and restricts his boundary search to this set. This condition is sometimes too restrictive [Prodanović et al., 2006], therefore our method does not precompute any local grain boundary beforehand. The result is an extended set L of k' pixels forming a 6-connected loop around the given throat of the pore space. The actual separating surface is the set of k' (non-coplanar) triangles formed by L and their centre g , i.e. the set $\{(g, p_1, p_2), (g, p_2, p_3), \dots, (g, p_{k'-1}, p_{k'}), (g, p_{k'}, p_1)\}$. Figure 3.9 and 3.10 present this method and compare it with the separating planes.





Figure 3.9: Illustration of the separating surface method for one branch through a circular constriction: (a) shows the edge connected to node 0, the initial plane used to guide the surface, and the arrow on that plane define the different sectors in which to look for pixels of the solid phase. (b) shows the resulting refined surface. (c) is the final separation, and as a comparison, (d) shows the separation with the planes method.


To avoid having thick separations, for every pixels of the void phase that intersect a triangle of the separation between elements i and j , it is directly assigned to one or the other according to the side on which lies its centre, allowing every pixel to be assigned a region.

This throat construction algorithm performs much better than the restrictive planar method, the result approaches those obtained with a watershed separation, but there are still fundamental limitations to this type of approach.


[†]Since each edge of such graph has a constant cost, there is no need for a more complex search algorithm such as Dijkstra's algorithm.



(a) The object to decompose



(b) The separation with the planes method



(c) The separation with the surfaces method

Figure 3.10: Comparison between the two methods that explicitly define the separations: (a) presents an object composed of 23 intersecting balls, from which the skeleton and a set of 23 markers were determined. (b) shows the delimitation result from plane construction. We can clearly see the misorientations of two planes, which are much larger than the others because they penetrate an adjacent ball. (c) shows the result with non-planar surface construction, which proves more accurate.


3.2.3 Limitations

Three main points should be addressed regarding throat construction. The first is its complexity: it has not been mentioned but each branch is considered independently and it can occur that surfaces cross or intersect, in different configurations described in [Prodanović et al., 2006]. It then requires a throat correction algorithm that iteratively deletes certain throats judged unable to properly separate two elements. Although irrelevant for simple geometries, throat construction methods are also applied to 50% porosity unstructured media, making efficient correction functions essential.


If we assume that separations are quasi-planar, then all throat construction methods should produce acceptable results (planar separations failed to convince because of our inability to determine the best parameters o and especially \vec{N}). All examples presented consider random ball packings, an object that contains only planar separations. Wedge-based throat construction can produce more practical results for non-planar separations, but is still unable to handle truly curved separations, as is illustrated in figure 3.11. It shows an object containing two elements, a bean-shaped elements that curves around a ball. The separation, properly identified with the watershed method, is not well approximated with neither planar nor wedge-based throat constructions.

The third point, and probably the most problematic, is the fact that throat construction is unable to handle all situations, even in what is assumed as simple geometries such as ball packings! An example is presented in figure 3.12. The constructed throats will surround only a fraction of the pixel faces shared by the two balls, thus leaving unseparated connections.


If we recall the method for handling the digitisation artefacts of Sheppard et al. [2005], which consisted in first using 6-connectivity to skeletonise the object, and subsequently using 26-connectivity, then more branches are preserved between balls. But, as figure 3.13 illustrates, this is still insufficient. On the other hand, this issue is resolved if the 6-connected skeleton is used to define the branches.




(a) A digitised object consisting of two elements



(b) Watershed separation



(c) Separating plane



(d) Separating surface

Figure 3.11: Illustration of one limitation of throat construction methods. (a) presents an simple set of two elements in a curved contact. (b) shows the result of the watershed, accurately following the crest line of the distance map. (c) shows the planar throat construction, creating a separation, although positioned as well as possible, that slices through the bean-shaped element. (d) shows the result with wedge-based throat construction, it accurately locates the border of the separation, but the triangulation is too coarse to produce a truly curved surface, resulting in a V-shaped deviation.

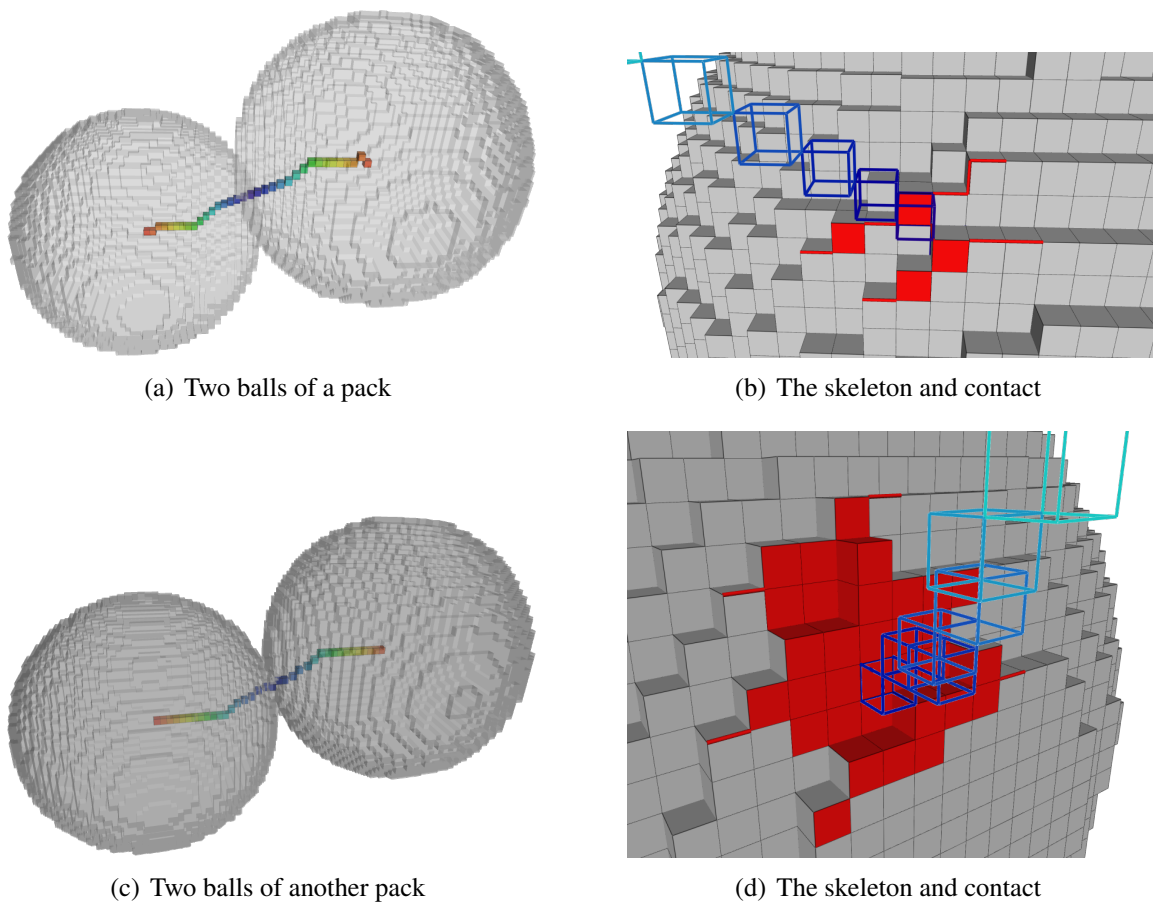


Figure 3.12: Two examples where the separations are more complex than simply connected pixel faces, and where the throat construction methods fail! (b) and (c) show the connected sets of pixel edges and faces, in red, shared by the two balls, and represented by a single branch in the skeleton.

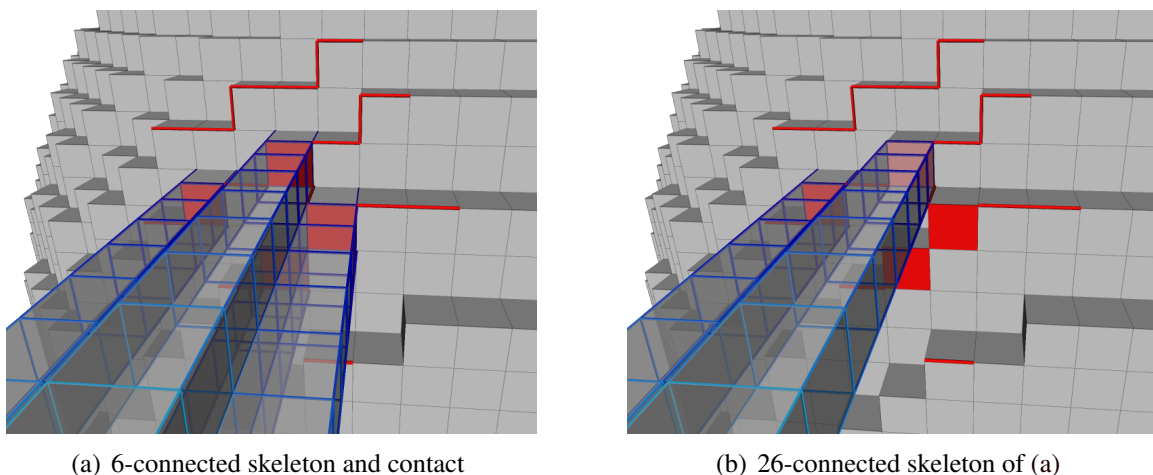


Figure 3.13: Effect of using 2 connectivities for skeletonising. (a) presents the 6-connected skeleton of the two balls from figure 3.12(a) and its contact, in red. The contact set contains additional edges compared to figure 3.12(b) because the digitisation artifact correction has not been applied. This skeleton contains four 6-connected branches representing the contacts between the two balls, three of which are 26-connected. Therefore, subsequent skeletonisation using 26-connectivity result in only two branches, shown in (b). Throat construction on this skeleton will also miss contacts.

3.3. Validation

In order to test the robustness of delimitation methods, we need an object for which a correct decomposition is already known. Several authors have used regular ball packings (cubic or hexagonal), but these regular structures are too idealised to represent real porous materials. Because ball packings are useful in the sense that the elements that comprise the set are simply defined, we decided to use random close packings, where the object to decompose is of course the balls (we assume the solid phase is the space between the balls). The objective is to verify the similarity between pixel assignment in the decomposition and in the image generation.

Although better adapted than regular packings, random ball packings are still simple geometric constructs, in which ball contacts are, mathematically, points. This translates in a digitised version as a significant constriction, becoming ever so simple to detect the higher the resolution. Since we want to determine the robustness of a decomposition method, the constriction should be made less obvious to detect. To this end, we progressively dilate the balls, by increasing their radii, creating a certain amount of overlap. At each stage, a decomposition is performed, and its accuracy is verified.

With the choice of geometry made, the first stage is to generate the digitised object on which to apply a decomposition.

3.3.1 Image generation

Figure 3.14 shows the random close packing used, and various stages of radius dilation. Since the packing is already almost 64% dense, increasing the radii by 20% already presents very severe topological changes.




Figure 3.14: *The random ball packing used to generate images for the validation. 1707 balls were taken from a random close packing, and the radii uniformly dilated to decrease the distinction between the individual elements.*

At each stage of dilation, the object B defined by a set of balls $B = \mathcal{B}_i = \mathcal{B}(c_i, r_i)$ is digitised into an image of a specified dimension, where a pixel p (defined as a unit cube) has a zero value if a majority of its volume is not contained in B , i.e. if $V(p \cap B) < \frac{V(p)}{2}$. Otherwise, it is assigned the value i of the

ball that contains the largest volume of p , or $\max_i V(\mathcal{B}_i \cap p)$. If this maximum is not unique (for instance if p is contained entirely in several balls), then we determine on which side the pixel centre lies in the ball intersection. If we take 2 balls, \mathcal{B}_1 and \mathcal{B}_2 , that contain an equal volume of p , it will be assigned the value 1 if it is contained in the spherical cap \mathcal{V}_2 (refer to figure 2.26(b) page 52), and vice versa. In the cases of more than two balls, they are taken two by two, one ball being discarded at each stage.

Note that this method is overtly complicated in the case of balls of equal radius, but it can be applied to any distribution of radii. Furthermore, if all radii are dilated enough for the set B to cover the entire digitised volume, then our method is equivalent to creating a digitised Laguerre tessellation (a weighted generalisation of the Voronoi tessellation).

Using this method, we directly assign each pixel of the object with a given value, thus naturally providing a reference decomposition, shown in figure 3.15.

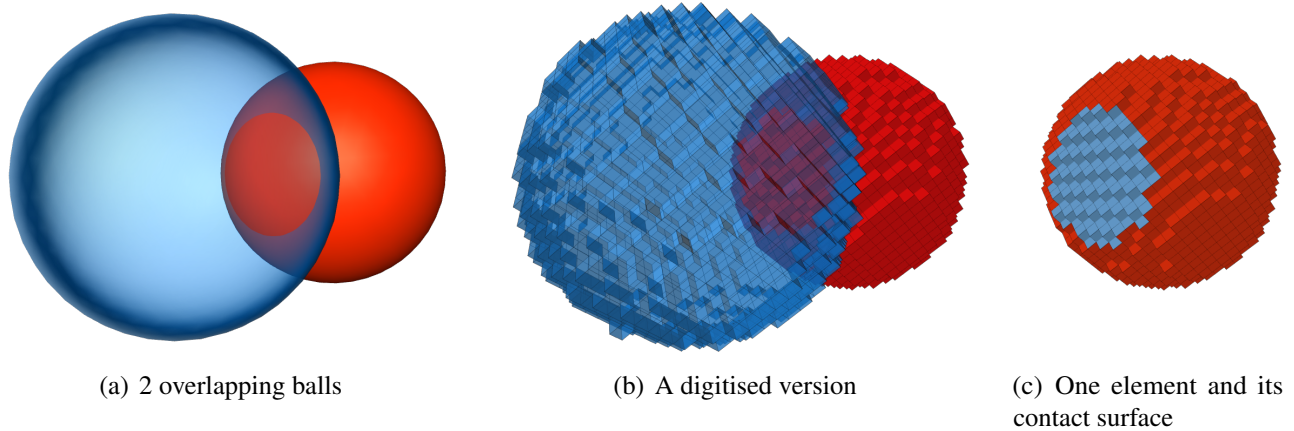


Figure 3.15: *Illustration of an optimal partitioning between two balls. (a) shows two overlapping balls (one is shown in transparent for display purposes). (b) shows a generated image, where each pixel is assigned either the background or one of the two balls during digitisation, thus defining by construction a partitioning. (c) shows the red element along with the throat surface. Results of a decomposition methods will be compared to this natural partitioning.*

Images of the random ball packing were generated at five different resolutions, ranging from 100^3 to 500^3 pixels. Figure 3.16 presents one element at different resolutions and level of dilation. These images were also used to test the robustness of our artefact removal method, with examples shown in figure 2.8, because the number of contacts, and thus the number of branches that should appear in the skeleton, is known.

3.3.2 Procedure

The testing procedure is straightforward. Each image is taken independently, and all separations between the different elements erased. The input is a binary image, in which we define the porous phase to be the set of balls. Pore positioning presented in chapter 2 locates the markers for the delimitation. The watershed delimitation is then performed, along with the throat construction methods to compare their behaviour (even though we have shown that they present some fundamental flaws).




Figure 3.16: Illustration of one element generated from the ball packing digitisation, extracted from the images at different resolutions (the columns, from 100^3 to 500^3 pixels by increments of 100), and at different stages of radius dilation (the lines, respectively 100%, 110% and 120%). The grey faces represent the contact with the solid phase, and each other colour corresponds to a distinct element.

The results take the form of the porous phase having each pixel assigned a value corresponding to the label of the element. We wish to compare these partitions with the reference partition from the random packing digitisation process. The problem is that element i in the reference does not necessarily correspond to element i in the resulting decomposition, as the pore(s) delimited in that area might have been assigned any other label. Before any comparison, a relabelling procedure must first be performed.

3.3.3 Relabelling for similar partitions

The problem can be summarised as followed: given two decompositions \mathcal{D}_r and \mathcal{D}_m of the same object X into pores \mathcal{P} such that $\mathcal{D}_x = \{\mathcal{P} | \mathcal{P} \subseteq X \wedge \mathcal{P} \text{ is simply connected}\}$ and $\bigcup_{\mathcal{P} \in \mathcal{D}_i} \mathcal{P} = X$, what is the optimal permutation array $P[\cdot]$ for \mathcal{D}_m such that the difference between the two is minimal, or $\sum_i \mathcal{V}(\mathcal{P}_i \cap \mathcal{P}_{P[i]})$ is maximal[†].

To find this array, we have developed a greedy algorithm, starting with the two elements $\mathcal{P}_i \in \mathcal{D}_r$ and $\mathcal{P}_j \in \mathcal{D}_m$ such that $\mathcal{P}_i \cap \mathcal{P}_j$ is maximal, and set $P[j] = i$, and then moving to the two next elements with maximum overlap, and so forth. Unless there exists a stage at which this maximum is not unique (and this pathological situation is rare), we conjecture that this relabelling is indeed optimal.

3.3.4 Results

Illustrations of a few partitionings are presented in figure 3.17. We observe that the watershed performs well compared to the reference partition, as differences vary by at most one pixel. When ball radius increases, the watershed still shows fairly small variations, except in the case of low resolution, where the initial difficulty resides in properly locating the markers (1267 markers for a packing of 1707 balls). Figure 3.18 shows the biggest pore obtained from our partitioning method on the 100^3 pixels image with 10% radius increase. This pore is 7019 pixels (when the average number is 247 for the reference), and consists of a cluster of balls with no cavity, which justifies their merging.

Figure 3.19 presents the number of markers found for the partitioning as a function of resolution (from 100^3 to 500^3), for three radius dilations up to 10%. Given one or two markers that define regions of one or two pixels, these three digitisations converge towards the correct number of elements. For a radius dilation of 15%, pore positioning completely breaks down, reaching 101 markers for the 500^3 pixel image. This is to be expected, as the object comprises 88% of the volume, and 2454 isolated connected components of the background can be found (which would correspond, in pore space decomposition, to solid particles floating in the porosity).

Figure 3.20 presents the amount of difference (computed as a percentage of pixels differing in label) between the relabelled decompositions.

[†]The problem is written supposing that \mathcal{D}_m and \mathcal{D}_r have an equal number of elements, but implemented for the general unequal case

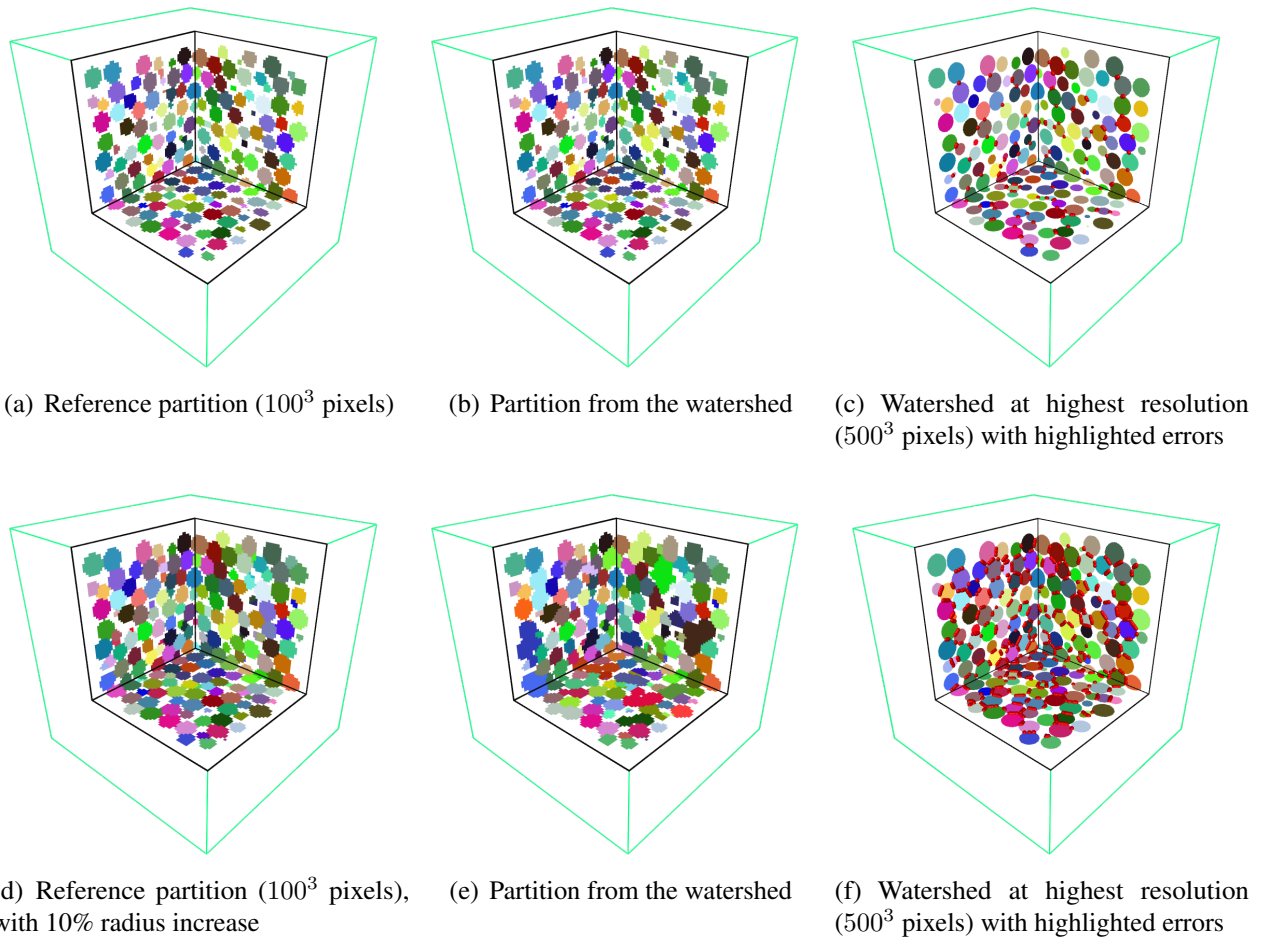


Figure 3.17: Illustration of partitioning results using the watershed method, for two different resolutions and two stages of ball dilation. At the first stage, with no dilation, the method performs well, with differences at the pixel scale. With a 10% radius dilation, especially at low resolution, the differences become more severe, mainly because the pore positioning is not able to locate as many markers. At higher resolutions, however, the watershed method still performs fairly efficiently (again with mostly pixel-wide variations).




Figure 3.18: A merging of a cluster of balls in the low-resolution packing. This pore, the biggest from the low resolution packing with a 10% radius increase, consists of 7019 pixels, corresponding to about 28 balls in the reference partitioning. This cluster contains no cavity, which can justify the merging.




Figure 3.19: Number of detected pores as a function of resolution. The reference packing consists of 1707 elements, and given sufficient resolution, our pore positioning method converges towards that value (1709, 1707 and 1708 for the respective radius dilations of 0%, 5% and 10%).




Figure 3.20: Relative error as functions of resolution (times 100), from left to right, and as radius coefficient, from 1 to 1.1 (i.e. 10% increase). The throat construction methods present strong discontinuities, because they either perfectly locate the constriction, or completely miss it. In all situations, the difficulty in delimiting the pores increase exponentially at low resolutions, but the overall results are strongly in favour of the watershed method.

3.3.4.1 Computation times

The objective of decomposition is first of all to quickly describe geometric characteristics of the material. Figure 3.21 shows the timings of the (single thread) computations on the test images performed on a Xeon E5320 processor at 1.86GHz, with 16% of the time used to remove digitisation artefacts, and 79% for the skeletonisation procedure. The remaining steps use 1% or less of the total computation time.




Figure 3.21: Average time (in minutes) taken to compute the decomposition of the ball packing as a function of resolution (in pixels). The test images range from 100^3 to 500^3 pixels, with computation times all under one hour.

Conclusion

This chapter presented two main approaches to the problem of pore delimitation. The first, which consists in identifying the volume of the pore, uses a watershed approach. We have described and implemented an algorithm for computing a watershed with good topological properties, mainly that of preserving the most significant features in the image. In order to improve the placement of the resulting throats, a new seed maximisation technique was used.

The second approach directly identifies the throat, and thus naturally delimits the volumes of the pores. We have shown that if a rigid application of Dullien's definition for a pore is used, then results inaccurately identifies the throats. A more practical solution was investigated, that of constructing triangulated non-planar separating surfaces for the throats, but have shown that it fails to identify curved throats, and proved that this type of approach cannot handle all types of situations.

Once all the tools have been put in place, we developed a methodology for testing the robustness of a decomposition method. The novelty does not lie in the use of a random close ball packing, but in the fact that we iteratively increase the ball radii in order to reduce the conspicuity of the constrictions. To compare the results, we had to propose a new relabeling algorithm to identify the same regions between two different decompositions with the same label.



4 Applications of pore-space decomposition

Contents

4.1	Evolution of permeability with geometry	85
4.1.1	Presentation	85
4.1.2	Material characterisation	86
4.2	Pore network models	93
4.2.1	Merging non-pairwise connected pores	94
4.2.2	Inserting interstitial pores	96
4.2.3	Inserting special branches in the network	97

The previous chapters have compared different methods for identifying and delimiting pores in a porous material, and the one that seems to provide the most representative decomposition was selected. Its purpose is twofold. Firstly, it allows to extract local geometric characteristics, considering each pore individually, which can help understand the relation between the geometry and macroscopic transport properties, namely intrinsic permeability. It is not known exactly what local factors influence transport phenomena in real porous media, due to their complexity. The prospect of extracting geometric parameters from a great number of porous structures with significant variations and correlating them with permeability does not seem pertinent. A wealth of parameters can be calculated from the decomposition, as well as from the entire pore space, therefore one combination of these parameters that correlates well with permeability is sure to be found. Nevertheless no generalisation can be proved from this correlation, and it would not help in understanding the influence of local geometry on permeability.

Instead, we prefer to examine how geometric parameters evolve compared to permeability on a given type of geometry. Section 4.1 presents this study on images of sintered glass powders and two types of evolution.

The second purpose of this decomposition is to define a network model for flow computation. The use of a network model is explained in section 4.2, along with a discussion on the fundamental condition of pairwise connection between pores.

4.1. Evolution of permeability with geometry

4.1.1 Presentation

4.1.1.1 Sintering experiment

A sintering experiment was performed on a commercial soda-lime glass powder (Sovitec, Glaverbel SA) sample. The powder was sieved at 100-120 μm , binded with 1wt% of PVA (Polyvinyl alcohol) and compacted uniaxially at 1.5 MPa to form a cylinder 5 mm high and 3mm in diameter. A first heat treatment at 650 $^{\circ}\text{C}$ for 1 hour burned out the PVA. The sample was then repetitively sintered at 720 $^{\circ}\text{C}$ for fixed durations and scanned on beamline ID19 at the ESRF at a resolution of 1.9 μm , with an energy of 17.5keV (with 900 projections per scan, excluding the reference and dark field images) [Bernard et al., 2005a]. Perpendicular cross-sections of the first scan is shown in figure 4.1. From each reconstruction, a 400 pixel³ subvolume far from the substrate is extracted, and shown in figure 4.2. This set of nine images represents the first evolution examined.

4.1.1.2 Precipitation simulation

The second type of evolution we consider is a precipitation simulation, performed on the first scan of the glass powder, shown in figure 4.2(a). The objective was initially to verify the accuracy of the Grain and




Figure 4.1: *Cross-sections of the complete reconstruction of the sample after 10 minutes at 720°C, (a) parallel and (b) perpendicular to the substrate, which can be seen on the bottom. The quality of the reconstruction and the dynamic of the intensities are such that binarisation is trivial.*

Pore Throat Model to predict hydrodynamic properties in consolidated granular media [Bernard, 2005; Chauveteau, 2002; Chauveteau et al., 1996]. The precipitation consists in performing morphological dilations using, 6-connectivity, of the solid phase. Nine dilations can be applied before the porosity becomes non-percolating, i.e. the permeability drops to zero. Three stages are illustrated in figure 4.3.

4.1.2 Material characterisation

4.1.2.1 Permeability

Figure 4.5 presents the permeability computed for these images as a function of the percolating porosity[†], ranging from 37% to 0.049%. The difference between porosity and percolating porosity, as plotted in figure 4.4, remains minimal when moving away from the limit of non-percolation.

The values of permeability, ranging from 7.424 to 0.007 mD, are the means of the diagonal values of the computed tensor (non-diagonal values are close to zero, due to the isotropic nature of the material). We first notice that some variation exists in the sintering experiment, this is because the subvolumes were extracted at different positions, due to the global shrinkage. Otherwise we can observe that the relation between porosity and permeability is similar in both evolutions and follows a power law.

The GPT model uses the notion of critical porosity, i.e. the porosity left in the sample once the evolution reaches the stage of non-percolation. It can be calculated only in the precipitation evolution, and equals 3.49%. If we plot the difference of porosity and critical porosity versus permeability, shown in

[†]i.e. the porosity connected to the outside and through which there exists a flow, taking into account conditions of periodicity in the same way as for the permeability calculation.




Figure 4.2: Binarised subvolumes of the glass powders taken at different sintering times. The structure begins to visibly change past 90 minutes.




Figure 4.3: Images of the precipitation simulation on the image shown in figure 4.2(a). At the tenth precipitation stage, the porosity becomes non-percolating.




Figure 4.4: Difference between porosity and percolating porosity as function of porosity. The value is on average below 0.1 % of the total porosity.




Figure 4.5: Permeability K (in milliDarcy) as a function of the percolating porosity Φ (as a fraction of the total volume) for (a) the sintering experiment, and (b) the precipitation simulation. Despite small variations at the beginning of the sintering evolution, we notice that both follow the same trend.

figure 4.6 on logarithmic scales, we see a clear power relation.




Figure 4.6: Permeability κ (in milliDarcy) as a function of percolating minus critical porosity. The dashed line represent a range of $\pm 20\%$ around the values of the model $\kappa = 210.61(\phi - \phi_c)^{3.0013}$, with ϕ_c the critical porosity determined from the precipitation evolution. We observe that all computed values are within this range, proving the accuracy of this model.

To ascertain the representativity of the volume for permeability, values were computed on subvolumes of increasing sizes, ranging from 100 to 400 pixels. The evolution of the permeability for the initial state is drawn in figure 4.7. We notice that past 350 pixels, the value stabilises around 7.4 mD.




Figure 4.7: Permeability κ (in milliDarcy) as a function of computational cell size in pixels. We observe that the measure converges around 7.4 mD, and can affirm that an elementary volume representative of the material has been reached.

4.1.2.2 Geometric characteristics using the decomposition

Figure 4.8 presents four quantities computed with the decomposition, for the two evolution processes. They are plotted as a function of the percolating porosity fraction (which decreases during the two evolution processes), and each point represents a measure computed for one volume. Several observations can be made:

- Apart from the mean pore connexity (i.e. the number of connections with adjacent pores), where the two image sets are distinct, both evolutions follow the same trends.
- The sintering experiment present small variations for the 5 first volumes, mostly because the subvolumes were extracted at different locations. Porosity remains about constant, but the decompositions provide stronger dissimilarities.

- Precipitation simulation contained a significant discontinuity at the first stage of dilation (the second point). Average pore volume is lower than expected, therefore pore number is greater. This is naturally linked to the fact connectivity is greater and thus the average throat surface is lower. It is not clear as to why the decomposition presents such a difference at this particular stage of the precipitation.
- Two stages can be identified in all the measures. It is especially clear for the specific number of pores, where the values follow a steep linear trend up to about 20% of porosity, then decreases slightly.
- Mean pore volume follows a polynomial relation of order 3 with porosity, with a correlation coefficient R^2 of 0.9851.

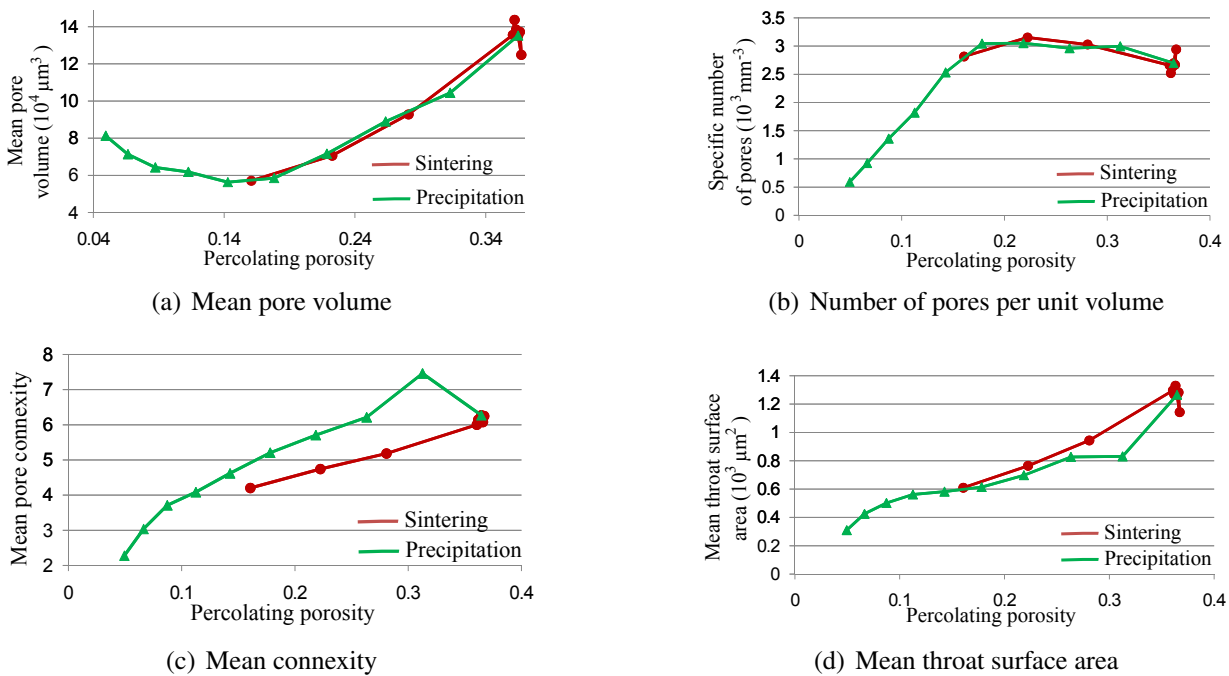


Figure 4.8: Illustration of several local quantities computed for the two evolutions.

In order to examine more closely the behaviour of the decomposition method, distributions of pore volume for three different stages of the precipitation simulation are plotted in figure 4.9. We notice that the initial state presents a clear majority of large pores especially between 5000 and 10000 μm^3 . After three stages of precipitation, a shifting of the number of pores to the left is natural, and the distribution becomes more uniform, resembling a gaussian. After 8 stages of percolation, two main classes of pores can be identified: the pores which were initially largest, situated in significant openings of the porosity, and that will represent most of the critical porosity, and the small pores below 2000 μm^3 , which probably represents the remainder of the percolating connections to the outside.

The representativity of values computed from the decomposition has been evaluated in the same manner as for permeability. Figure 4.10 presents the evolution of those values as a function of the size of the subvolume. We can also observe a stabilisation past 350 pixels.







Figure 4.9: Pore volume (in μm^3) distribution for the precipitation simulation at three different stages. The vertical axis represents the number of pores and horizontal axis, the volume.




(a) Mean pore volume (in $10^4 \mu\text{m}^3$)



(b) Number of pores per unit volume (in 10^3mm^{-1})



(c) Mean connectivity




(d) Mean throat surface area (in $10^3 \mu\text{m}^2$)

Figure 4.10: Illustration of local quantities computed for the glass beads image set.


4.1.2.3 Geometric characteristics without decomposition

Geometric characteristics can also be measured without using a decomposition, and are plotted in figure 4.11 as a function of porosity. The mean distance from the solid, shown in figure 4.11(a), was computed by averaging the pixel values in the Euclidean distance map. This type of measure is obviously very well correlated to porosity, with a value of 0.998, therefore not presenting any new information. The three other values are derived from the Minkowski functionals [Blasquez and Poiraudau, 2003], which provide four quantities: volume, surface, integral of mean curvature, and Euler number. The surface is rendered specific by dividing by the image volume. It shows that the two evolutions are of different nature, as material diffusion on the solid surface will tend to reduce specific surface more rapidly.


We also clearly observe on figures 4.11(c) and 4.11(d) two main stages, with an initial linear evolution for high porosities. The differences in the two types of evolution are also apparent in the Euler characteristic, once sintering significantly modifies the geometry. In the precipitation simulation, there is a clear cut between the two stages, as connections begin to be severed and the number of isolated pores increases.




(a) Mean distance from the solid



(b) Specific surface



(c) Integral of the mean curvature



(d) Euler characteristic, normalised by the image volume

Figure 4.11: Illustration of the global geometric quantities that can be computed.

4.1.2.4 Conclusion

Direct measurements from pore geometry can provide information on the organisation of the porous structure in great detail. The information extracted from the decomposition is complementary to the one that can be computed on the entire pore space, and together provide more information than that which is represented by a macroscopic property such as intrinsic permeability. We have verified that on this family of materials, a power law in relation to the porosity accurately describes permeability, therefore we

have examined geometric parameters in regards to porosity. The complex variations of these parameters leads us to ask whether they can be used independently to predict transport phenomena. We have noticed in many of these measures that two stages can be identified, which is representative of the nature of the topological modification in the porosity (grain neck growth, which has less impact on the geometry than pore elimination), but does not seem to play a role as such in the evolution of the permeability.

Geometric parameters provide a qualitative appreciation of the nature of the porous structure, but more work is needed to relate these directly to transport phenomena. The second objective of pore space decomposition is to produce a simplified version of the porosity geometry so as to compute flow and hydraulic resistance more efficiently.


4.2. Pore network models

Pore decomposition gives access to pore-level numerical modeling based on the sample geometry but with a greatly simplified data structure. The purpose is to make use of the graph created in chapter 2, along with the decomposition and extracted values from chapter 3, to compute transport properties such as permeability. In order to do this, conductivities are computed locally for each branch in the graph, and then assembling all the computed values to obtain a averaged measure over the entire porosity, thus greatly simplifying the computational costs. The underlying hypothesis is that the conductivity of one separation does not depend on any other.

The details of the method for computing the conductivities are presented elsewhere [Combaret, to appear]. One crucial factor in this approach is the fact that one branch represents the flow between two pores, and only two pores. Therefore, a fundamental and implicit property of a constriction (the separating surface) between two pores, is that it must be entirely bounded by the solid phase, or put in other terms, there must always be a pairwise connection between pores [Prodanović et al., 2006]. Were this not the case, intersections of pore contacts would exist. This is computationally difficult to justify as, from a numerical modelling standpoint, there is no proper limit condition for the intersections of constrictions.

Figure 4.12 shows two examples of pores delimited in the initial glass powder sample from the previous section. The left pore presents a typical configuration, having exclusively pairwise connections. The right pore, however, has four throats (each having an identifying colour) in contact with each other, thus contradicting the pairwise connection condition. In such a case, the statement from Lindquist et al. [2000]: "every pore channel (i.e. connection between pores defined by the nodes of the skeleton) has a corresponding path in the skeleton" simply is not true: not all contacts have a corresponding branch in the skeleton, otherwise artificial loops would be created during the skeletonisation process.

Figure 4.13 is a 2D illustration of this problem: it shows two similar pore configurations, along with the computed graphs and pore delimitations. In figure 4.13(a), each branch can be considered separately in order to compute local conductivity. In figure 4.13(b), the three central separations cannot be considered independently, and the flow through the pores is greatly modified by this fact. Figure 4.14 presents the



(a) Example of a well-defined pore

(b) Example of a problematic pore

Figure 4.12: Illustration of the type of situations encountered in real data. The examples are taken from the cropped regions of the decomposition of image 10mc3 shown in figure 4.2(a). Figure (a) shows one pore as a set of pixels, at the centre of the image, surrounded by four glass beads shown in transparent. The coloured faces are the contact with adjacent pores. Each colour signifies a contact with a distinct pore: the separations form simple and isolated surfaces around the pore. Each of these contacts accurately represents a link with another pore in the network, and thus hydraulic conductivity can be effectively computed for each independently. The pore depicted in figure (b), on the other hand, shows four contacts with adjacent pores that are not isolated, and therefore present a fundamental problem with the pore network concept.

flow fields computed for an example of both configurations, in the vertical direction and with periodic conditions. The latter geometry is indeed 16% more permeable.

This type of situation tends to appear fairly often, and has the inconvenient property of not being detectable until after the delimitation process. Therefore, handling this problem must be performed as a post-processing of the decomposition. Several approaches have been tested in order to remove this type of configuration.

4.2.1 Merging non-pairwise connected pores

The first attempt was to merge all adjacent pores that have a non-pairwise connection. Figure 4.15 presents the result on the previous illustration. Implementation of this solution is fairly straightforward, and the merging can be done sequentially (as long as three pore should be merged) or in parallel (find the complete set of pores that should be merged), the result will be identical.

Figure 4.15 shows this process for the minimum number of three pores, while figure 4.12(b) shows a simple situation involving five pores, but in practice a much greater number of pores can be involved. In fact, for the glass powder image from figure 4.2(a) this process merges 85% of the 1381 pores, resulting in one element occupying 97.5% of the pore space. It was initially hoped that this was partly caused by too few markers retained after the graph node merging presented in section 2.5.1, when in fact the choice




Figure 4.13: 2D illustration of the pairwise connection problem. The grey area represents the solid, the blue lines the skeleton of the porous phase, the black discs the pore markers and the red dotted lines the resulting separations. Figure (a) shows three spherical pores having three distinct pairwise connection between them (in 2D, the central solid portion appears to be floating inside the porosity, but in 3D it would be connected to the rest of the solid phase above or below the pores). Figure (b) shows almost the same configuration, but with the three pores slightly larger, such that the three separations come in contact. One conspicuous difference is the number of branches, as in the latter case one of the three central branches is missing (were it still present, the skeleton would not have the same topology as the object, as it would create an additional loop).

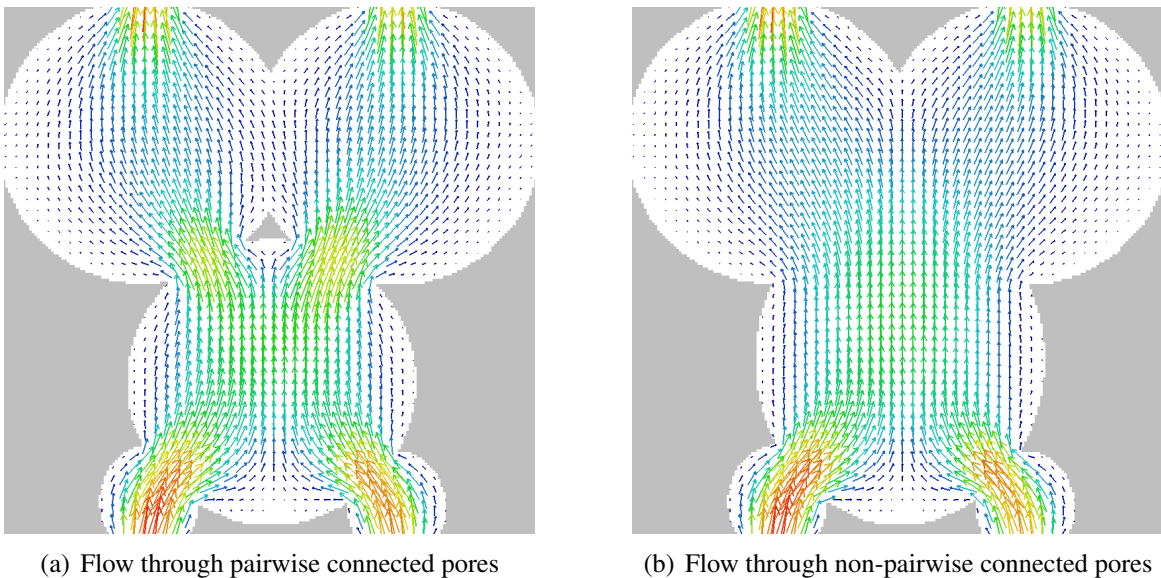


Figure 4.14: Vertical flow field computed inside three pores to illustrate the importance of pairwise connection. Figure (a) shows three pores with separate connections. The flow through the left connection does not depend on the right. Figure (b) shows a very similar situation, but the throats are connected. The flow field is then modified significantly, which is reflected in the ensuing permeability measure.




Figure 4.15: *Illustration of the merging of removing non-pairwise connected pores. This figure presents the subsequent pore obtained by post-processing the delimitation from figure 4.13(b).*

of the node merging condition hardly influences at all this occurrence. On the contrary, when no node merging is performed, the biggest pore then occupies 98.3% of the pore space !

Obviously this method is unusable in practice. Prodanović et al. [2006] reaches similar conclusions, and proposes an iterative suppression of throats (and therefore a merging of two pores), beginning with the smallest surface area throats, which effectively corresponds to a pore merging. This iterative approach is said to provide a minimal number of modifications (although it is not proved), but the fact it is iterative implies non-local conditions for pore merging (since the decision to suppress a throat depends on all the other throats to be eventually suppressed). Although defining a pore in real porous media is difficult, one basic assumption was always that this definition was based on local features. For this reason, this type of method was deemed unsuitable.

Since pore merging resulted in oversized pores, the opposite solution of pore insertion was also tested.

4.2.2 Inserting interstitial pores

The idea in this approach would be to insert a new pore in between those not pairwise connected, so that each would then be pairwise connected to the newly inserted pore. Returning to the example depicted in figure 4.13(b), the intended result is shown in figure 4.16.

The choice of the amount of volume occupied by this new pore is ad hoc. A simple method is to first assign the pixels containing the separations involved in non-pairwise connections, then artificially dilate the pore to give it more volume.

Unfortunately, this approach suffers from the same symptom as the merging post-process. There are actually very long branching chains of non-pairwise connections, resulting in newly inserted pores spanning through the entire volume. Figure 4.17 presents two of the 130 new pores inserted by this method in the glass powder image of figure 4.2(a). In most cases, the non-pairwise separations involve three pores, but they can also be chained along several pores, as in figure 4.17(a). The worst case is




Figure 4.16: *Insertion of new pores to force pairwise connections. The intended result from the example of figure 4.13(b) is shown here: the initial pores are no longer directly connected to each other, but are instead each in a pairwise connection with the new central pore.*

presented in figure 4.17(b), showing the biggest insertion: a pore twelve times bigger than any of the original partitioning, and containing 440 tunnels.

This problem again renders the solution unusable. The fact that the inserted pores are irreducible, i.e. cannot be subdivided to reduce their size (otherwise new non-pairwise connections would be created), compels the need for another solution.

4.2.3 Inserting special branches in the network

The implemented method that offers the best compromise consists in two stages. The first is to modify the pore network in order to incorporate the separations for which no branch existed in the skeleton, by inserting new branches. This effectively changes the topology of the network, but to a controlled extent. This alone cannot solve the difference between pairwise and non-pairwise connected systems (recall figure 4.14), which is why these newly created branches, along with the other ones involved in the non-pairwise connections, are tagged for special processing. Figure 4.18 illustrates this alternative. The thick green branch on the right has been inserted, and all three central branches are marked.

When the hydraulic resistance is computed between two of these pores, their volume will not be considered completely separate from the third pore. Instead, we suppose that a portion of the third pore must also be examined when computing the conductivity along one branch. Of course the furthest portion of the third pore will not influence at all the flow through the other two, but the nearer portion will. Therefore when computing the flow between two of these marked pores, an extra pore volume with a certain geodesic distance of these two pores is taken into account. Figure 4.19 illustrates the influence of this solution. With the use of this extra volume, the computed flow field (and thus the conductivity represented by the branch) is much more representative of the real flow in this region of the pore space.

Figure 4.20 illustrates this alternative on real data. The subvolume shows three pores and their mutual separations. The bottom separation is not associated to any branch in the graph, therefore a new branch




Figure 4.17: Application of pore insertion on the glass powder image: two newly created pores, prior to volume dilation, are shown (the pixels of the new pores are only those that contained the non-pairwise separations). Figure (a) shows a small example, where the newly inserted pore, although tortuous, remains of reasonable size. Figure (b), however, presents the worst case in this image: the pore spans the entire 400^3 pixel volume, and in its initial undilated state, is already more than twelve times bigger in volume than any other pore.




Figure 4.18: Insertion and marking of new branches in the pore network, illustrated for the example of figure 4.13(b). The thick green branch on the right has been inserted, and the other two central ones are marked for special processing.




Figure 4.19: Local flow computed through two pores in a non-pairwise connection. Figure (a) presents the initial pore space, and (b) its decomposition into three pores. If local flow is simulated using only the pore space of the two considered pores, as in figure (c), the fluid passing through the third pore is not taken into account, thus reducing the resulting hydraulic conductivity. Figure (c) shows the flow field computed if the volume of all pores adjacent to the non-pairwise connection are used. The difference in permeability between (c) and (d) is almost 40%.

is inserted. To define an associated path to the branch, in the same manner as the pixels of the skeleton define paths for their respective branches in the graph, two pixelated segments, each linking one node and the centre of mass of the separation, are defined.





Figure 4.20: Example of a branch insertion taken on the same glass powder image of figure 4.2(a). The bottom separation was initially not identified by a branch, therefore the red one was added. Its path is defined as two line segments passing through the centre of mass of the separation.

Conclusion

The presented post-processes stem from the fact that the provided decomposition does not have every pore pairwise-connected to their adjacent counterparts. In this situation, defining the location of a hydraulic minimum as in Dullien's definition is very difficult. The most reasonable solution to handling non-pairwise connections is to consider the volume of all adjacent pores in the computation of the hydraulic resistivity of one edge in the pore network model. This naturally leads to the question of whether there always exists a decomposition that satisfies the following conditions :

- Each element must be a simple 6-connected component of pixels (i.e. containing neither tunnels nor cavities).
- If two elements are adjacent, then they are pairwise-connected.

If the answer is yes, then how can a solution be determined ? This is a challenging problem in itself, and remains an open question at the end of this work. The continued effort in using the pore network model defined above, along with this new post-processing of non-pairwise connections, will hopefully confirm the accuracy of the approach for the prediction of flow properties.



Conclusion

The aim of this work was to better understand the influence the geometric structure of porosity can have on transport phenomena and associated macroscopic properties. If an efficient way to characterise the structure of a porous material is found in that regard, it would allow to predict measures such as permeability or conductivity, which is of particular importance in numerous fields, from earth science and metallurgy to the food industry.

When studying the pore space of a material, one examines, among global characterisations, any geometric features of the pores themselves. The greatest difficulty comes from the vagueness surrounding the notion of pore. Their definition, positioning and delimitation is not unique, and depends strongly on the application. We have described that they should obey certain rules and have certain properties, and attempted to develop a robust method for extracting them, i.e. decomposing the pore space.

Two separate stages were distinguished: marker positioning, and region delimitation. The first stage provides the number and position of the pores in the material. We have seen that using only geometric information was not sufficient, and that topology is necessary and complementary information. This implied using a skeletonisation algorithm, a tool which is unfortunately sensitive to small features, such as digitisation artefacts. We have created a new method to identify pixels of the pore space image that are responsible for these artefacts and removed them. We have also quantified the impact of this method on macroscopic measurements.

The skeleton is converted to a graph, and the vertices of the graph are used to define the markers (additionally, the edges can give an indication as to the connectivity between these markers). A common post-processing performed on the graph is marker merging. To this end, we have defined the notion of redundant information between markers.

The second stage is delimitation. Several methods were surveyed, classified in two categories: throat construction and watershed. The former, which directly constructs the separations between pores, shows serious limitations when handling complex structures that are commonly found in porous materials. The latter, which iteratively assigns pixels to a specific region (pore), naturally generates these separations. We have implemented a watershed algorithm to compute such separations.

To verify the robustness of our decomposition method, we developed a methodology designed to highlight its limitations. From a random close ball packing, we have developed a digitisation process in which a reference partition is readily available, so that the quality of our decomposition can be assessed. Of course, prior to this comparison, pore relabeling must be performed. A greedy but quasi-optimal algorithm has been implemented to this purpose. We observe that our results begin to significantly differ from the reference when dilation becomes too important to conserve the features of the individual balls that constitute the packing.

From a representative decomposition of the pore space into meaningful elements that we call pores, several geometric quantities become available, such as pore size or throat surface area. These can provide valuable information on the nature of the geometric structure of the porosity. In order to understand the relation with intrinsic permeability these variables can have, we have examined two sets of images, based on a sample of a sintered glass powder. The dissimilarities in the nature of the image evolutions (a sintering experiment and a precipitation simulation) are highlighted by geometric quantities such as mean

pore connectivity, or specific surface. We can also notice similar behaviours for both processes on quantities measured on the pores, such as mean pore volume, which clearly follows a polynomial trend of order 3 in regards to porosity, or the number of pores, which presents two evolution regimes. It is difficult to directly conclude to a relation between such geometric measures and macroscopic properties, as perhaps they represent geometric information not relevant to intrinsic permeability.

Nevertheless, decomposition also provides a pore network that greatly simplifies the pore structure, and on which numerical modelling can be performed. A strong assumption of pore network models is that each connection an edge represents is only between the two corresponding pores, but all decomposition methods require a post-processing to produce only such pairwise connections. We have presented an alternative method to handle this problem, which is to avoid considering a pair of pores independently when non-pairwise connections appear. We retain the local geometry for flow computation, but consider a more pertinent portion of the porosity for assigning a hydraulic resistance to a given edge.

The fundamental problem studied in this thesis was the manner in which to define a pore in a pore space, based on local and logical geometric concepts. The classical definition proposed by Dullien, "part of the pore space bounded by the solid and planes erected where the hydraulic radius is minimal", has been shown to be inappropriate in practice. Through the methodology developed in this work, we are able to propose an improved definition of a pore: "part of the pore space, homotopic to a ball, bounded by the solid, and connected to other pores by throats of minimal surface area". This improvement eliminates the strict condition of planar separations, and adds the constraint of simple topology to the decomposed elements. This new definition, although more restrictive, does not imply uniqueness of the solution, since the decomposition depends on the application (for instance for network models, the non-pairwise connection property is sought).

Future work should first focus on a algorithmic point of view of the decomposition process, such as parallelisation, block processing for very large data sets, and optimisation for novel algorithms (such as the digitisation artefact removal method, which takes more than 15% of the computation time). The use of the framework of cubical complexes [Bertrand and Couprie, 2006] for the skeletonisation process could be tremendously beneficial for a parallelisation strategy, accurate surface detection in the skeleton, easier digitisation artefact detection and graph conversion. The node merging post-processing of the graph would also require a detection of created cycles in individual pores.

Two interesting geometric problems were encountered. The first was finding the optimal pore relabelling for matching a reference partition. A greedy implementation was provided, but an optimal algorithm would require a more complex strategy. The second problem is whether there systematically exists a decomposition that satisfies both element simplicity and pairwise connections. Unless a counter-example can be found, it might prove very difficult to prove that this fact is true.

Statistical analysis of the influence of geometric parameters on permeability can be developed further. Mercury intrusion curves can be simulated on the decompositions we produce, and compare them with experimental results in order to confirm the accuracy of this decomposition. Quantity distributions as a function of distance from the image border might also provide new and potentially pertinent information.



Acknowledgements

This work has been funded by the Corning European Technology Center. I am grateful to the people at Corning, especially Daniela Lavric, Pierre Woehl and Pernelle Barlier, who have allowed me to work on this project and have been very encouraging of my work.

I want to thank the director Claude Delmas, for having welcomed me at ICMCB, and I am very grateful to the members of the jury, Peter Lee, Jean-Michel Missiaen, Christian Germain, and Mario Maglione (excluding those who are cited elsewhere), who have given me very positive reviews and very helpful criticism to help improve the quality of my work. I owe special thanks to my PhD supervisor, Dominique Bernard, without whom not only this PhD but all my other contracts at the laboratory would not have been possible. He has been an invaluable help and support. Thanks to Ali Chirazi for his advice and comments and to Nicolas Combaret for our numerous constructive discussions, as well as several algorithms he has implemented.

I would like to acknowledge the people at ESIEE, Michel Couprie, Hugues Talbot and John Chaussard, who have been of precious help during our many collaborative meetings.

A special thanks to the people of the Ferroelectrics Materials, Ceramics and Composites group with whom I have worked, and who have produced such a pleasant work environment. Thanks to the young ones, Guillaume Lacombe, Grégory Lalet, Grégory Hauss, Cécile Vincent, Emmanuel Arveux, Cindy Mauriac, Elias Castel, and including those who have left before me: Frédérique Frezouls, and Valérie Denis-Lutard. Thanks to the old ones, Jean-François Silvain, Jean-Marc Heintz, and Lionel Teule-Gay.

I am also grateful to all those at ICMCB who take care of the non-academic aspects of research, Carole Malburet, Gérard Lebreton, Virginie Cosseron-Dacosta, Frédérique Bonnet, Sylvie Létard, Jacques Domingie, Bernard Lestienne, Alain Larroche, and Sandrine Quillateau.

Finally I want to thank my family for always being there. And thanks to my beloved Amélie, who has made these years at Bordeaux the best years of my life.



Bibliography

- A.A. Adedeji and M.O. Ngadi. The use x-ray micro-ct for characterization of micro-structural properties of deep-fat fried breaded chicken nuggets. In *American Society of Agricultural and Biological Engineers Annual International Meeting 2008*, volume 9, pages 5781–5793, 2008.
- A.S. Al-Kharusi and M.J. Blunt. Network extraction from sandstone and carbonate pore space images. *Journal of Petroleum Science and Engineering*, 56(4):219–231, 2007.
- R. Al-Raoush and K.A. Alshibli. Distribution of local void ratio in porous media systems from 3d x-ray microtomography images. *Physica A: Statistical Mechanics and its Applications*, 361(1-4):441–456, 2006.
- R. Al-Raoush and C. Willson. Extraction of physically realistic pore network properties from three-dimensional synchrotron x-ray microtomography images of unconsolidated porous media systems. *Journal of Hydrology*, 300(1-4):44–64, 2005.
- R. Al-Raoush, K. Thompson, and C. Willson. Comparison of network generation techniques for unconsolidated porous media. *Soil Science Society of America Journal*, 67(6):1687–1700, 2003.
- M. Ando and S. Hosoya. An attempt at x-ray phase-contrast microscopy. In G. Shinoda, K. Kohra, and T. Ichinokawa, editors, *Proceedings of the 6th International Conference on X-ray Optics and Microanalysis*, pages 63–68. University of Tokyo Press, 1972.
- YY Anguy, D. Bernard, and R. Ehrlich. The local change of scale method for modelling flow in natural porous media (i): Numerical tools. *Advances in Water Resources*, 17(6):337–351, 1994.
- J.-Y. Arns, V. Robins, A.P. Sheppard, R.M. Sok, W.V. Pinczewski, and M.A. Knackstedt. Effect of network topology on relative permeability. *Transport in Porous Media*, 55(1):21–46, 2004.
- T. Aste, M. Saadatfar, and T.J. Senden. Geometrical structure of disordered sphere packings. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 71(6):1–15, 2005.
- F.H. Attix, W.C. Roesch, and E. Tochilin. *Radiation Dosimetry*. Academic Press, New York, 1968.

- R.C. Atwood, J.R. Jones, P.D. Lee, and L.L. Hench. Analysis of pore interconnectivity in bioactive glass foams using x-ray microtomography. *Scripta Materialia*, 51(11):1029–1033, 2004.
- S. Bakke and P. Øren. 3-D pore-scale modelling of sandstones and flow simulations in the pore networks. *Society of Petroleum Engineers Journal*, 2(2):136–149, 1997. ISSN 1086-055X.
- C. A. Baldwin, A. J. Sederman, M. D. Mantle, P. Alexander, and L. F. Gladden. Determination and characterization of the structure of a pore space from 3d volume images. *Journal of Colloid and Interface Science*, 181:79–92, July 1996. URL <http://www.ingentaconnect.com/content/ap/cs/1996/00000181/00000001/art00358>.
- J. Baruchel, J. Buffière, E. Maire, P. Merle, and G. Peix. *X-Ray Tomography in Material Science*. Hermes Sciences Publications, 2000. ISBN 2-7462-0115-1. URL <http://www.hermes-science.com>.
- R. Bayer. Symmetric binary b-trees: Data structure and maintenance algorithms. *Acta Informatica*, 1(4): 290–306, 1972.
- M.A. Bender and M. Farach-Colton. The LCA problem revisited. In *Latin American Theoretical Informatics*, pages 88–94, 2000.
- D. Bernard. Using the volume averaging technique to perform the first change of scale for natural random porous media. In G. Gambolati and G. Verri, editors, *Advanced methods for groundwater pollution control*, pages 9–24. Springer, 1995.
- D. Bernard. Expérience numérique de précipitation uniforme dans un lit de billes de verre ; conséquences sur le tenseur de perméabilité. unpublished, 2005.
- D. Bernard and A. Chirazi. Numerically enhanced microtomographic imaging method using a novel ring artefact filter. In Jacques Desrues, Gioacchino Viggiani, and Pierre Bésuelle, editors, *Advances in X-ray tomography for geomaterials*, pages 119–124. ISTE, 2006.
- D. Bernard and G. Vignoles. A numerical study of the coupled evolutions of micro-geometry and transport properties of simple 3d porous media. In J. Crolet, editor, *Computational methods for flow and transport in porous media*, pages 217–229. Springer, 2000.
- D. Bernard, D. Gendron, J.-M. Heintz, S. Bordère, and J. Etourneau. First direct 3d visualisation of microstructural evolutions during sintering through x-ray computed microtomography. *Acta Materialia*, 53(1):121–128, 2005a.
- D. Bernard, Ø. Nielsen, L. Salvo, and P. Cloetens. Permeability assessment by 3d interdendritic flow simulations on microtomography mappings of al-cu alloys. *Materials Science and Engineering: A*, 392(1-2):112–120, February 2005b. URL <http://www.sciencedirect.com/science/article/B6TXD-4DNB09M-1/2/8ea06908d7097fbc4e22fc5759c7001>.

- G. Bertrand. Simple points, topological numbers and geodesic neighborhoods in cubic grids. *Pattern Recognition Letters*, 15(10):1003–1011, 1994.
- G. Bertrand. On topological watersheds. *J. Math. Imaging Vis.*, 22(2-3):217–230, 2005. ISSN 0924-9907.
- G. Bertrand and M. Couprie. A new 3d parallel thinning scheme based on critical kernels. In *Discrete Geometry for Computer Imagery*, pages 580–591, 2006.
- G. Bertrand and M. Couprie. Transformations topologiques discrètes. In David Coeurjolly, Annick Montanvert, and Jean-Marc Chassery, editors, *Géométrie discrète et images numériques*, pages 187–209. Hermès, 2007. URL <http://www.lavoisier.fr/fr/livres/index.asp?texte=2746216430&select=isbn&from=Hermes>.
- G. Bertrand and G. Malandain. A new characterization of three-dimensional simple points. *Pattern Recognition Letters*, 15(2):169–175, 1994.
- G. Bertrand and G. Malandain. Note on 'building skeleton models via 3-d medial surface/axis thinning algorithms'. *Graphical Models and Image Processing*, 57(6):537–538, 1995.
- I. Blasquez and J.F. Poiraudau. Efficient processing of minkowski functional on a 3d binary image using binary decision diagrams. *Journal of WSCG*, 11:–, 2003.
- H. Blum. An associative machine for dealing with the visual field and some of its biological implications. *Biological Prototypes and Synthetic Systems*, 1:244–260, 1961.
- G.T. Bolton and K.M. Primrose. An overview of electrical tomographic measurements in pharmaceutical and related application areas. *AAPS PharmSciTech*, 6(2):–, 2005.
- E. Cendre, P. Duvauchelle, G. Peix, J.-Y. Buffière, and D. Babot. Conception of a high resolution x-ray computed tomography device: Application to damage initiation imaging inside materials. In *Proceedings of the 1st World Congress on Industrial Process Tomography*, pages 362–369, 1999.
- R. Cesareo, A. Brunetti, R. Tadeu Lopes, G. Galli, D.V. Rao, A. Castellano, G.E. Gigante, S. Mascarenhas, R. Robert, V. S. Filho, M. Gilardoni, H. Pereira Da Silva, and P. Quarta Colosso. X and γ -ray tomography for non destructive material testing. *Proceedings of SPIE - The International Society for Optical Engineering*, 3772:292–303, 1999.
- J. Chaussard, G. Bertrand, and M. Couprie. Characterizing and detecting loops in n-dimensional discrete toric spaces. In D. Coeurjolly et al., editor, *Lecture Notes in Computer Science*, volume 4992 of *Lecture Notes in Computer Science*, pages 129–140. Springer, 2008.
- G. Chauveteau. The grain and pore throat model: A tool to predict mass transport and formation damage. In *SPE International Symposium and Exhibition on Formation Damage Control*, 2002.
- G. Chauveteau, L. Nabzar, Y. El Attar, and C. Jacquin. Pore structure and hydrodynamic in sandstones. In *Society of Core Analysts*, pages –, 1996. Conference Paper Number 9607.

- N. Combaret. *Modèles de réseaux, pores, et calculs de perméabilité*. PhD thesis, Université de Bordeaux 1, to appear.
- G. N. Constantinides and A.C. Payatakes. A theoretical model of collision and coalescence of ganglia in porous media. *Journal of Colloid and Interface Science*, 141(2):486–504, 1991.
- J. H. Conway and N.J. A. Sloane. *Sphere Packings, Lattices and Groups*. Springer, second edition, 1992. ISBN 978-0387979120.
- M. Couprie and G. Bertrand. Topological grayscale watershed transformation. In *SPIE Vision Geometry VI*, pages 136–146, 1997. ISBN 0-8194-2590-7.
- M. Couprie, F.N. Bezerra, and G. Bertrand. Topological operators for grayscale image processing. *Journal of Electronic Imaging*, 10(4):1003–1015, 2001.
- M. Couprie, L. Najman, and G. Bertrand. Quasi-linear algorithms for the topological watershed. *Journal of Mathematical Imaging and Vision*, 22(2):231–249, 2005.
- M. Couprie, D. Coeurjolly, and R. Zrour. Discrete bisector function and euclidean skeleton in 2d and 3d. *Image and Vision Computing*, 25(10):1543–1556, 2007.
- O. Cuisenaire and B. Macq. Fast and exact signed euclidean distance transformation with linear complexity. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 6: 3293–3296, 1999.
- H. Darcy. *Les fontaines publiques de la ville de Dijon*. Victor Dalmont, 1856.
- S. Deans. *The Radon transform and some of its applications*. Wiley and sons, New York, N.Y., 1983.
- J.F. Delerue. *Segmentation 3D, application à l'extraction de réseaux de pores et à la caractérisation hydrodynamique des sols*. PhD thesis, Université Paris XI Orsay, 2001.
- F. A. L. Dullien. *Porous Media: Fluid Transport and Pore Structure*. Academic Press, 2 edition, 1991. ISBN 0122236513.
- F. A. L. Dullien and G. K. Dhawan. Characterization of pore structure by a combination of quantitative photomicrography and mercury porosimetry. *Journal of Colloid and Interface Science*, 47(2):337–349, 1974.
- I.A. Feldkamp, L.C. Davis, and J.W. Kress. Practical cone-beam algorithm. *Journal of the Optical Society of America A: Optics and Image Science, and Vision*, 1(6):612–619, 1984.
- J. D. Frost. *Studies on the monotonic and cyclic behavior of sands*. Civil engineering, Purdue University, 1989.
- M. Frucci and G. Sanniti di Baja. Oversegmentation reduction in watershed-based grey-level image segmentation. *International Journal of Signal and Imaging Systems Engineering*, 1(1):4–10, 2008.

- A. Ganssen, W. Loeffler, A. Oppelt, and F. Schmidt. Tomography by nuclear magnetic resonance. *Computertomographie Sonographie*, 1(1):2–10, 1981.
- E. J. Garboczi. Three-dimensional mathematical analysis of particle shape using x-ray tomography and spherical harmonics: Application to aggregates used in concrete. *Cement and concrete research*, 32(10):1621–1638, 2002.
- T. Hirata. A unified linear-time algorithm for computing distance maps. *Information Processing Letters*, 58(3):129–133, 1996.
- J.G. Hocking and G. S. Young. *Topology*. Dover Publications, new ed edition, 1988. ISBN 0-486-65676-4.
- D.S. Holder. *Electrical Impedance Tomography: Methods, History and Applications*. Taylor & Francis, 2004.
- M. A. Ioannidis, M. J. Kwiecien, I. Chatzis, I. F. MacDonald, and F. A. L. Dullien. Comprehensive pore structure characterization using 3d computer reconstruction and stochastic modeling. In *Proceedings of the 1997 SPE Annual Technical Conference and Exhibition. Part Omega (pt 1)*, 1997.
- R.J. Jaszczak. The early years of single photon emission computed tomography (spect): An anthology of selected reminiscences. *Physics in Medicine and Biology*, 51(13):–, 2006.
- A. Kak and M. Slaney. *Principles of computerized tomographic imaging*. IEEE Press, New York, N.Y., 1988.
- E. S. Kikkinides and V. N. Burganos. Permeation properties of three-dimensional self-affine reconstructions of porous materials. *Physical Review E*, 65(5):6906–6915, 2000.
- T. Kong and A. Rosenfeld. Digital topology: introduction and survey. *Comput. Vision Graph. Image Process.*, 48(3):357–393, 1989.
- A.G. Konstandopoulos. Flow resistance descriptors for diesel particulate filters: Definitions, measurements and testing. *SAE transactions*, 112(4):700–712, 2003. ISSN 0096-736X. SAE technical paper No. 2003-01-0846.
- C. Lantuéjoul. *La squelettisation et son application aux mesures topologiques de mosaïques polycristallines*. PhD thesis, Ecole des Mines, Paris, France, 1978.
- L. Latecki, U; Eckhardt, and A. Rosenfeld. Well-composed sets. *Computer Vision and Image Understanding*, 61(1):70–83, 1995.
- L.J. Latecki. 3d well-composed pictures. *Graphical Models and Image Processing*, 59(3):164–172, 1997.
- T.C. Lee, R.L. Kashyap, and C.N. Chu. Building skeleton models via 3-d medial surface/axis thinning algorithms. *CVGIP. Graphical models and image processing*, 56(6):462–478, 1994.

- Y. Li, W. G. Laidlaw, and N. Wardlaw. Sensitivity of drainage and imbibition to pore structures as revealed by computer simulation of displacement process. *Advances in Colloid and Interface Science*, 26:1–68, 1986.
- Z. Liang, M.A. Ioannidis, and I.I. Chatzis. Geometric and topological analysis of three-dimensional porous media: Pore space partitioning based on morphological skeletonization. *Journal of Colloid and Interface Science*, 221(1):13–24, January 2000.
- A. Lieutier. Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis. In *Proc. 8th ACM Symposium on Solid Modeling Applications*, pages 65–75. Academic Press, 2003.
- W.B. Lindquist and A. Venkatarangan. Investigating 3d geometry of porous media from high resolution images. *Physics and Chemistry of the Earth, Part A: Solid Earth and Geodesy*, 24(7):593–599, 1999.
- W.B. Lindquist, S.M. Lee, D.A. Coker, K.W. Jones, and P. Spanne. Medial axis analysis of three dimensional tomographic images of drill core samples. *J. Geophys. Res. B*, 101:8296–8310, 1996.
- W.B. Lindquist, A. Venkatarangan, J. Dunsmuir, and T.-F. Wong. Pore and throat size distributions measured from synchrotron x-ray tomographic images of fontainebleau sandstones. *Journal of Geophysical Research B: Solid Earth*, 105(B9):21509–21527, 2000.
- D.P. Lymberopoulos and A. C. Payatakes. Derivation of topological, geometrical, and correlational properties of porous media from pore-chart analysis of serial section data. *Journal of colloid and interface science*, 150(1):61–80, 1992.
- G. Malandain, G. Bertrand, and N. Ayache. Topological segmentation of discrete surfaces. *International journal of computer vision*, 10(2):183–197, 1993.
- P. Van Marcke. *Development of a pore network model to perform permeability computations on X-ray computed tomography images*. PhD thesis, Katholieke Univeriteit Leuven, 2008.
- A. Meijster, J.B.T.M. Roerdink, and W.H. Hesselink. A general algorithm for computing distance transforms in linear time. *Mathematical Morphology and Its Applications to Image and Signal Processing*, 1:331–340, 2000.
- F. Mendoza, P. Verboven, H.K. Mebatsion, G. Kerckhofs, M. Wevers, and B. Nicolai. Three-dimensional pore space quantification of apple tissue using x-ray computed microtomography. *Planta*, 226(3): 559–570, 2007.
- C. Mercet. *Modélisation tridimensionnelle de l'évolution des transferts à l'échelle du pore*. PhD thesis, Université de Bordeaux 1, 2000.
- F. Meyer. Algorithmes à base de files d'attente hiérarchique. Technical Report NT-46/90/MM, Centre de Morphologie Mathématique, Ecole des Mines de Paris, Fontainebleau, France, 1990.

- P.A. Midgley and R.E. Dunin-Borkowski. Electron tomography and holography in materials science. *Nature Materials*, 8(4):271–280, 2009.
- J.K. Mitchell and K. Soga. *Fundamentals of soil behaviour*. John Wiley & Sons Inc., New York, third edition, 1993.
- P. Mrázek and M. Navara. Consistent positive directional splitting of anisotropic diffusion. In Boštjan Likar, editor, *Proceedings of Computer Vision Winter Workshop*, pages 37–48, February 2001.
- G.G. Muntean, D. Rector, D. Herling, D. Lessor, and M. Khaleel. Lattice-boltzmann diesel particulate filter sub-grid modelling-a progress report. In *SAE Technical Papers*, 2003. No. 2003-01-0835.
- L. Najman and M. Couprie. Quasi-linear algorithm for the component tree. In *Proceedings of SPIE - The International Society for Optical Engineering*, volume 5300, pages 98–107, Laboratoire A2SI, Groupe ESIEE Cité Descartes, BP99, 93162 Noisy-le-Grand Cedex, France, 2004.
- F. Natterer. *The Mathematics of Computerized Tomography*. SIAM: Society for Industrial and Applied Mathematics, 2001.
- S. Neethirajan, D.S. Jayas, N.D.G. White, and H. Zhang. Investigation of 3d geometry of bulk wheat and pea pores using x-ray computed tomography images. *Computers and Electronics in Agriculture*, 63(2): 104–111, 2008.
- M. Oda. Deformation mechanism of sand in triaxial compression tests. *Soils and Foundations*, 12(4): 45–63, 1972.
- N. Passat, M. Couprie, and G. Bertrand. Minimal simple pairs in the 3-d cubic grid. *Journal of Mathematical Imaging and Vision*, 32(3):239–249, 2008.
- A.K. Patel and B.N. Thorat. Gamma ray tomography—an experimental analysis of fractional gas hold-up in bubble columns. *Chemical Engineering Journal*, 137(2):376–385, 2008.
- P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- S. Peth, R. Horn, F. Beckmann, T. Donath, J. Fischer, and A.J.M. Smucker. Three-dimensional quantification of intra-aggregate pore-space features using synchrotron-radiation-based microtomography. *Soil Science Society of America Journal*, 72(4):897–907, 2008.
- D.L. Pham, C. Xu, and J.L. Prince. A survey of current methods in medical image segmentation. *Annual Review of Biomedical Engineering*, 2:315–338, 2000.
- M.E. Phelps, E.J. Hoffman, N.A. Mullani, and M.M. Ter Pogossian. Application of annihilation coincidence detection to transaxial reconstruction tomography. *Journal of Nuclear Medicine*, 16(3):210–224, 1975.

- L. Pothuaud, P. Porion, E. Lespessailles, C.L. Benhamou, and P. Levitz. A new method for three-dimensional skeleton graph analysis of porous media: Application to trabecular bone microarchitecture. *Journal of Microscopy*, 199(2):149–161, 2000.
- W.K. Pratt. *Digital Image Processing*, chapter 17 - Image Segmentation, pages 551–587. John Wiley & Sons Inc., 2001.
- M. Prodanović, W.B. Lindquist, and R.S. Seright. Porous structure and fluid partitioning in polyethylene cores from 3d x-ray microtomographic imaging. *Journal of Colloid and Interface Science*, 298(1): 282–297, 2006.
- C. Pudney. Distance-ordered homotopic thinning: A skeletonization algorithm for 3d digital images. *Computer Vision and Image Understanding*, 72(3):404–413, 1998.
- J.A. Quiblier. A new three-dimensional modeling technique for studying porous media. *Journal of Colloid and Interface Science*, 98(1):84–102, 1984.
- J. Radon. Über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten. *Mathematische-Physische Klasse 69*, pages 262–267, 1917. Berichte Sächsische Akademie der Wissenschaften.
- G.N. Ramachandran and A.V. Lakshminarayanan. Three-dimensional reconstruction from radiographs and electron micrographs: application of convolutions instead of fourier transforms. *Proceedings of the National Academy of Sciences of the United States of America*, 68(9):2236–2240, 1971.
- M. Saadaftar, A. Shepard, and M. Knackstedt. Grain partitioning and its applications. In *Advances in X-ray tomography for geomaterials*, pages 269–276. iSTE, 2006.
- T. Saito and J.-I. Toriwaki. New algorithms for euclidean distance transformation of an n-dimensional digitized picture with applications. *Pattern Recognition*, 27(11):1551–1565, 1994.
- F. Sauli. Gas detectors: Recent developments and future perspectives. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 419 (2):189–201, 1998.
- A.J. Sederman, M.L. Johns, A.S. Bramley, P. Alexander, and L.F. Gladden. Magnetic resonance imaging of liquid flow and pore structure within packed beds. *Chemical Engineering Science*, 52(14):2239–2250, 1997.
- L.A. Shepp and Benjamin F. Logan Jr. The fourier reconstruction of a head section. *IEEE Transactions on Nuclear Science*, NS-21(3):21–43, 1974.
- A.P. Sheppard, R.M. Sok, and H. Averdunk. Improved pore network extraction methods. In *International Symposium of the Society of Core Analysts*, 2005.

- H. Shin. *A throat finding algorithm for medial axis analysis of 3D images of vesiculated basalts*. PhD thesis, State University of New York at Stony Brook, 2002.
- H. Shin, W.B. Lindquist, D.L. Sahagian, and S.-R. Song. Analysis of the vesicular structure of basalts. *Computers and Geosciences*, 31(4):473–487, 2005.
- D.B. Silin, G. Jin, and T.W. Patzek. Robust determination of the pore-space morphology in sedimentary rocks. *Journal of Petroleum Technology*, 56(5):69–70, 2004.
- P. Soille. *Morphological Image Analysis*. Springer-Verlag, 1999. ISBN 3-540-65671-5.
- R.M. Sok, M.A. Knackstedt, A.P. Sheppard, W.V. Pinczewski, W.B. Lindquist, A. Venkatarangan, and L. Paterson. Direct and stochastic generation of network models from tomographic images; effect of topology on residual saturations. *Transport in Porous Media*, 46(2-3):345–372, 2002.
- Edwin H. Spanier. *Algebraic Topology*. 1966.
- G. Stockman and L.G. Shapiro. *Computer Vision*. Prentice-Hall, 2001.
- V. Tariel. *Image analysis of cement paste: relation to diffusion transport*. PhD thesis, Ecole Polytechnique, 2008.
- K. Thompson, C. Willson, and W. Zhang. Quantitative computer reconstruction of particulate materials from microtomography images. *Powder Technology*, 163(3):169–182, 2006.
- J. F. Thovert, J. Salles, and P. M. Adler. Computerized characterization of the geometry of real porous media : their discretization, analysis and interpretation. *Journal of microscopy*, 170(1):65–79, 1993. ISSN 0022-2720.
- C.D. Tsakiroglou and A.C. Payatakes. A new simulator of mercury porosimetry for the characterization of porous materials. *Journal of Colloid and Interface Science*, 137(2):315–339, 1990.
- C.D. Tsakiroglou and A.C. Payatakes. Effects of pore-size correlations on mercury porosimetry curves. *Journal of Colloid and Interface Science*, 146(2):479–494, 1991.
- C. Vachier. *Extraction de Caractéristiques, Segmentation, d’Image en Morphologie Mathématique*. PhD thesis, Ecole des Mines, Paris, 1995.
- A.B. Venkatarangan. *Geometric and statistical analysis of porous media*. PhD thesis, State University of New York at Stony Brook, 2000.
- L. Vincent and P. Soille. Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):583–598, 1991.
- L.V. Wang. Prospects of photoacoustic tomography. *Medical Physics*, 35(12):5758–5767, 2008.

- N. Wardlaw, Y. Li, and D. Forbes. Pore-throat size correlation from capillary pressure curves. *Transport in Porous Media*, 2(6):597–614, 1987.
- S. Whitaker. *The Method of Volume Averaging*, volume 13 of *Theory and Applications of Transport in Porous Media*. Springer, 1999. ISBN 978-0-7923-5486-4.
- R. A. Williams and M. S. Beck. *Process tomography: principles, techniques and applications*. Butterworth-Heinemann, 1995.
- O. Wirjadi. Survey of 3d image segmentation methods. Technical report, Fraunhofer ITWM, 2007.
- M. Xu and L.V. Wang. Photoacoustic imaging in biomedicine. *Review of Scientific Instruments*, 77(4):–, 2006.
- E.C. Zeeman. On the dunce hat. *Topology*, 2(4):341–358, 1963.