

FEB 26 1999

MS 0619 Review and Approval Desk, 15102 (1)  
For DOE/OSTI

RECEIVED  
MAR 03 1999  
OSTI

# SANDIA REPORT

SAND99-0369  
Unlimited Release  
Printed February 1999

## Development, Implementation, and Experimental Validation of the Lattice Boltzmann Method for Modeling Three- Dimensional Complex Flows

J. T. Fredrich, D. R. Noble, R. M. O'Connor, and W. B. Lindquist

Prepared by  
Sandia National Laboratories  
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,  
a Lockheed Martin Company, for the United States Department of  
Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

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

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831

Prices available from (703) 605-6000  
Web site: <http://www.ntis.gov/ordering.htm>

Available to the public from  
National Technical Information Service  
U.S. Department of Commerce  
5285 Port Royal Rd  
Springfield, VA 22161

NTIS price codes  
Printed copy: A03  
Microfiche copy: A01



## **DISCLAIMER**

**Portions of this document may be illegible  
in electronic image products. Images are  
produced from the best available original  
document.**

# **Development, Implementation, and Experimental Validation of the Lattice Boltzmann Method for Modeling Three-Dimensional Complex Flows**

J.T. Fredrich  
Geomechanics Department

D.R. Noble  
Incompressible Fluid Dynamics Department

R.M. O'Connor  
Geomechanics Department

Sandia National Laboratories  
P.O. Box 5800  
Albuquerque, NM 87185-0751

W.B. Lindquist  
Department of Applied Mathematics and Statistics  
State University of New York  
Stony Brook, New York 11794-3600

## **Abstract**

Lattice Boltzmann (LB) methods are ideally suited for modeling flow and transport at intermediate scales where neither continuum models nor molecular dynamics methods are practical, and also for modeling very large scale continuum flows where conventional Computational Fluid Dynamics methods suffer severe difficulties. This report documents an LDRD project that developed an advanced computational framework based on the LB method for investigating the fundamental phenomena characterizing micro- and macro-scale complex flows. Specific accomplishments include: (1) Development of LB codes for modeling heat and mass transport in complex geometries, including implementation of scaleable, distributed algorithms. (2) Validation of the LB codes by comparison with analytic solutions, calculations based on other numerical methods, and laboratory experiments. (3) Development of a distributed computational platform based on commercial off-the-shelf computer components that is capable of production-scale LB simulations. (4) Development of methods and computational algorithms to extract quantitative statistical descriptions of the microgeometry of porous media from three-dimensional image data and a computational framework (numerical pipeline) for simulating flow and transport in experimentally determined porous microstructures.

This page intentionally left blank.

## **Acknowledgement and Preface**

This report documents work performed under the Laboratory Directed Research and Development (LDRD) Program under a project entitled "Development, implementation, and experimental validation of the lattice Boltzmann method for modeling three-dimensional complex flows." The project was funded by the Engineering Sciences Research Foundation in fiscal years 1997 and 1998, and we thank Chuck Hartwig for his enthusiasm for our work.

We thank Steve Kempka and Charles Hickox for their contributions to this project, and also Harlan Stockman for his initial participation. The work on cluster-based computing benefited from discussions with Pete Wyckoff, David Greenberg, and Robert Clay, and from our participation in the Scalable Cluster Workshop convened by David Greenberg and Robert Clay in November 1997. Jeremy Baca and Todd Kuzior both shared some hands-on knowledge for assembling computer components. We thank John Torczynski and Dale Preece for reviewing an earlier draft of this report.

We have sought to write this report as a comprehensive work, and in that regard, note that as originally-proposed, the LDRD project advantaged previous and ongoing research sponsored by the US DOE Office of Basic Energy Sciences (Geosciences Program), in the form of independent grants to J.T. Fredrich and W.B. Lindquist.

The work on the porous ceramic insulations was initiated originally with funding from NASA Ames Research Center in collaboration with Jochen Marschall of NASA Ames.

This page intentionally left blank.

## Table of Contents

Abstract .....	iii
Acknowledgement and Preface .....	v
Table of Contents .....	vii
Figures .....	ix
Tables .....	xiii
1    Introduction .....	1
1.1    Background to this project .....	1
1.2    Work Plan .....	2
1.3    Project accomplishments .....	3
1.4    Organization of this report .....	4
2    LB code development .....	5
2.1    Introduction .....	5
2.2    Physics .....	5
2.3    Algorithmic .....	8
2.3.1    Moving and non-conforming obstacles .....	9
2.3.2    Complex irregular geometries .....	9
2.3.3    Stencil representation .....	9
2.3.4    Volume visualization .....	11
2.3.5    Parallel I/O .....	12
2.4    Code architecture .....	12
3    Araneæ— A scalable system for distributed computing .....	15
3.1    Space: flow modeling in reconstructed porous media .....	15
3.2    Time: evolution to steady state .....	16
3.3    System architecture .....	17
3.3.1    Networking .....	17
3.3.2    Cost .....	18
3.4    Operating system and environment .....	19
3.4.1    Programming model .....	19
3.4.2    System definition .....	19
3.4.3    Compute node definition .....	22
3.4.4    Useful URLs .....	22
3.5    Discussion .....	23
4    Validation problems .....	24
4.1    Flow over a backward facing step .....	24
4.2    Three dimensional duct flow .....	25
5    Application to a coupled thermal-hydrodynamic flow on non-uniform grids .....	27
5.1    Lattice Boltzmann model for heat transfer on a non-uniform grid .....	28
5.2    Simulation of mixed convection over a backward-facing step .....	32
5.3    Accuracy of the LB model for mixed convection on a non-uniform grid .....	34

6	Application to flow about moving and non-conforming obstacles .....	40
6.1	Lattice Boltzmann method for partially saturated cells.....	40
6.2	Steady flow about a periodic line of cylinders .....	43
6.3	Oscillatory flow about a periodic line of cylinders .....	46
6.4	Discussion .....	49
7	Experimental characterization of porous media.....	51
7.1	Approaches for characterizing the microstructure of porous materials .....	52
7.2	Principles of laser scanning confocal microscopy .....	55
7.3	Technical and practical considerations.....	57
7.4	Applications .....	59
7.5	Image processing and segmentation.....	63
7.5.1	Anisotropic diffusion.....	63
7.5.2	Segmentation.....	63
7.6	n-point correlation functions .....	67
7.7	Medial axis analysis .....	67
7.8	Application to geometrically complex porous media.....	70
7.9	Summary .....	74
8	Application to microscale flows in complex porous media .....	75
8.1	Lattice Boltzmann method .....	76
8.2	Distributed Implementation.....	77
8.3	Flow boundary conditions .....	80
8.4	Distributed computing.....	81
8.5	Application .....	81
8.6	Discussion .....	84
9	Summary .....	85
10	References .....	86
	Appendix A– List of refereed publications and presentations .....	92

## Figures

<b>Figure 1.1</b> Flow chart illustrating the work plan for the LDRD project. ....	2
<b>Figure 2.1</b> Schematic illustration of the LB method: (upper left) velocity distribution, (upper right) collision step, showing current versus equilibrium particle distribution along each of the eight links in two dimensions (there are a total of 18 links in three dimensions), (lower left) translation step, showing the streaming of particles along each of the links, and (lower right) simple boundary condition.....	6
<b>Figure 2.2</b> Schematic illustration of the traditional mapping approach versus the stencil algorithm. ....	10
<b>Figure 3.1</b> Schematic illustrating a scalable unit (SU) in the Araneæ cluster.....	20
<b>Figure 3.2</b> Schematic illustrating the current Araneæ cluster. ....	21
<b>Figure 4.1</b> Illustration of the backward facing step problem. ....	24
<b>Figure 4.2</b> Convergence of LB simulation to analytical solution for 3D duct flow.....	26
<b>Figure 5.1</b> Sample non-uniform computational grid for lattice Boltzmann simulations. For clarity, the grid shown here is 61x21 whereas the actual grid used in the simulations is 241x81. Also note that the horizontal and vertical scales are not the same. Gravity, $g$ , acts in the negative $x$ direction as indicated by arrow. ....	32
<b>Figure 5.2</b> Profiles of $U$ vs. $Y$ at indicated $X$ positions for $Gr=0$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ). ....	34
<b>Figure 5.3</b> Profiles of $U$ vs. $Y$ at indicated $X$ positions for $Gr=1000$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ). ....	35
<b>Figure 5.4</b> Profiles of $V$ vs. $Y$ at indicated $X$ positions for $Gr=0$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ). ....	36
<b>Figure 5.5</b> Profiles of $V$ vs. $Y$ at indicated $X$ positions for $Gr=1000$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ). ....	36
<b>Figure 5.6</b> Profiles of dimensionless temperature, $\Theta$ , vs. $Y$ at indicated $X$ positions for $Gr=0$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ) ....	37
<b>Figure 5.7</b> Profiles of dimensionless temperature, $\Theta$ , vs. $Y$ at indicated $X$ positions for $Gr=1000$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ). ....	38
<b>Figure 5.8</b> Profiles of the dimensionless skin friction vs. $X$ along the walls of the channel for $Gr=0$ and $Gr=1000$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ). ....	38
<b>Figure 5.9</b> Profiles of the local Nusselt number vs. $X$ along the walls of the channel for $Gr=0$ and $Gr=1000$ . FIDAP simulation data taken from published tabulated data ( <i>Torczynski, 1993</i> ). ....	39

**Figure 6.1** Schematic diagram (to scale) of the validation problem in the (a) moving-cylinder reference frame and (b) fixed-cylinder reference frame. ....43

**Figure 6.2** Convergence with grid refinement for lattice Boltzmann simulations of steady flow using various boundary conditions. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution. ....44

**Figure 6.3** Components of velocity near the surface of the cylinder. The cylinder is moving with a constant velocity in the positive x direction, and the velocities are expressed relative to the velocity of the cylinder. Velocities are plotted as a function of angle for a fixed radial distance of  $D/32$  from the surface of the cylinder. The solid lines denote the FIDAP predictions. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution. ....45

**Figure 6.4** FIDAP streamlines during the fifth cycle for oscillatory flow at indicated normalized times and phases, with cylinder velocities in the fixed frame proportional to the arrows. ....46

**Figure 6.5** History of normalized horizontal velocity component  $u_x/U_0$  during the fifth cycle of oscillation for the point located at  $r = 0.75D$  and  $\theta = 90^\circ$ . The solid line denotes the FIDAP prediction. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution. ....47

**Figure 6.6** Spatial variation of the normalized horizontal velocity component  $u_x/U_0$  along the centerline ( $y = 0$ ) at indicated times during the fifth cycle. The solid lines denote the FIDAP predictions. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution. ....48

**Figure 7.1** Scanning electron micrographs of (top) Berea sandstone (backscattered imaging mode), and (bottom) a porous rigid fibrous refractory insulation. The sandstone has a grain size of  $\sim 150$  micron, and the pore space has been impregnated with epoxy (black). There is no continuous path through the pore phase in the imaging plane. The insulation is composed of ceramic fibers that occupy less than  $\sim 20\%$  of the bulk volume, and that have a diameter of several microns. The image reveals the complex geometry of both the solid and pore phases, but it is unfeasible to extract a quantitative description of the three-dimensional microstructure. (Bottom image courtesy of NASA Ames Research Center, Thermal Protection Materials and Systems Branch). ....53

**Figure 7.2** Schematic illustrating the principle of laser scanning confocal microscopy. ....55

<b>Figure 7.3</b> 3D volume renderings of pore space in Berea sandstone. The pore phase is opaque whereas solid grains are translucent. Each image is $512 \times 512 \times 71$ voxels, with an $(x \times y \times z)$ image volume equal to $0.526 \times 0.526 \times 0.071 \text{ mm}^3$ .....	60
<b>Figure 7.4</b> (Top) 3D volume rendering of crack damage in a triaxially deformed low-porosity rock. The fine crack network in the upper right corner would not be resolvable using conventional LM. The image is $768 \times 512 \times 71$ voxels, with an image volume of $430 \times 287 \times 14.4 \mu\text{m}^3$ . (Bottom) A 2D image slice from the 3D data set. The relationship of the cracks trending NE-SW to the grain boundary (running diagonally NW-SE) is indisputable in the 3D image, but less so in the 2D slice. ....	61
<b>Figure 7.5</b> 3D volume renderings of the solid phase in two rigid ceramic insulations. Each image is $768 \times 512 \times 85$ voxels, with cubic voxels with dimension $0.3 \mu\text{m}$ ( $230 \times 154 \times 25.2 \mu\text{m}^3$ ). The two insulations are composed of the same three fiber constituents, but have been pressed to different final densities. The insulation shown at top has a bulk density of $0.35 \text{ g/cc}$ and bulk porosity of $0.86$ , whereas that shown at bottom has a bulk density of $0.19 \text{ g/cc}$ and bulk porosity of $0.92$ .....	62
<b>Figure 7.6</b> Example illustrating application of diffusion algorithm to confocal image data. The top image shows a raw $768 \times 512$ image, and below that is the same image after application of the diffusion algorithm (400 iterations, length scale $k=6$ ). ....	64
<b>Figure 7.7</b> Example illustrating application of diffusion algorithm (400 iterations, length scale $k=6$ ) to confocal image data and resultant enhancement of void and solid phase discrimination. The two plots show histograms (dashed lines) of the inverted fluorescence intensity for the raw (top) and diffused (bottom) images shown in Figure 7.6. Also shown (solid lines) is the variation in the calculated porosity as a function of threshold. ....	65
<b>Figure 7.8</b> Schematic illustrating a single nodal pore body that due to its irregular geometrical shape, has a medial axis reduction containing two, rather than one, vertices. ....	69
<b>Figure 7.9</b> (Top) 3D volume rendering of pore space in Berea sandstone. The data set is $768 \times 512 \times 101$ voxels in size, with cubic voxels of dimension $1 \mu\text{m}$ . (Bottom) The corresponding medial axis calculated for the pore phase. The medial axis encodes the distance to the pore-solid interface, so that the geometry of the original structure is retained, and consists of a series of links (or paths) that connect and/or fork at branch points (or clusters). For example, consider a regular 2D cubic network consisting of nodes and bonds; in a medial axis representation, the nodes would correspond to branch clusters and the bonds to paths. In like manner, a 3D geometrically complex porous medium reduces to a topologically complex medial axis that consists of an irregular network of pore paths that interconnect at branch clusters. ....	71
<b>Figure 7.10</b> Two-point correlation function $S_2$ for the data set shown in Figure 7.9. The porosity (equal to the y-intercept) is $28\%$ , and the specific surface area (inversely proportional to the slope at the origin) is $23 \text{ mm}^2/\text{mm}^3$ . The upper and lower curves represent one standard deviation. The point at which $S_2$ reaches an asymptotic values defines a characteristic length for the microstructure ( $\sim 250 \mu\text{m}$ ). ....	72
<b>Figure 7.11</b> (Top) Cluster coordination for branch points in pore medial axis for Fontainebleau sandstone with porosities as indicated. (Middle) Density of pore medial axis clusters. (Bottom) Number of paths and branch clusters for disconnected medial axis. ....	73

<b>Figure 8.1</b> Schematic illustrating coupled experimental-computational pipeline for studying flow in complex porous media.....	78
<b>Figure 8.2</b> Example illustrating the decomposition of a two dimensional geometric domain into four subdomains, and showing the communication between neighboring subdomains that is accomplished using message passing. In the 3D case, each subdomain also communicates along two shared boundaries that would lie in the plane of the image. The implementation of periodic flow boundary conditions on the boundary of the entire geometric domain is illustrated by the large arrows outside of the image area.....	79
<b>Figure 8.3</b> Plot showing the evolution to a steady state permeability during the simulation. The simulation is for the reconstructed volume shown following in Figure 8.4 (top)...	82
<b>Figure 8.4</b> (Top) Volume rendering of the pore space in Berea sandstone used in the simulation with results shown in Figure 8.3. The image size is $768 \times 512 \times 128$ , with cubic voxels of dimension $1 \mu\text{m}$ . (Bottom) Three dimensional rendering of the steady state flow field. The red color corresponds to the highest velocities, and blue to the lowest velocities. Shown in green are intermediate velocities.....	83

## Tables

<b>Table 4.1</b> Comparison of LB solution with experimental data and conventional CFD methods including finite element (FE) and finite difference (FD) for the classic backward facing step problem. ....	25
<b>Table 7.1.</b> Lateral resolution $R$ (Eq. 7.1) and optical section thickness at $\lambda=514$ nm for various settings of the confocal aperture.....	56

This page intentionally left blank.

# 1 Introduction

The purpose of this project was to develop an advanced capability for modeling complex three-dimensional (3D) flows based on the Lattice Boltzmann (LB) method. This project was funded by the Laboratory Directed Research and Development (LDRD) program in FY97 and FY98 under the Phenomenological Modeling and Engineering Simulation initiative that is managed by the Engineering Sciences Research Foundation.

Our interest in LB is two-fold. First, many applications of interest to Sandia involve time and length scales that are too small to be addressed easily by continuum models while at the same time are too large to be addressed using molecular dynamics methods. LB has shown particular promise for modeling phenomena at these intermediate, or mesoscopic, scales, and is especially well suited for tracking the evolution of interfaces (such as the boundaries between immiscible liquids), and for modeling flow, transport, and reaction in porous media at low to moderate Reynolds numbers and Peclet numbers.

Our second major area of interest is determining the applicability of LB to what is generally referred to as conventional Computational Fluid Dynamics (CFD). Here, CFD is taken to address continuum flows, and larger Reynolds numbers are of primary interest. For very large problems, conventional CFD methods for incompressible flows become prohibitively expensive. Considering the relative ease with which LB methods are implemented on large parallel machines, it could be the method of choice for very large-scale flows.

LB is a kinetic theory based method that recovers the Navier-Stokes and advection-dispersion equations at the macroscale (*Chen, Chen and Matthaeus*; 1992). The principal advantages of LB over more traditional methods are its ability to handle arbitrarily complex geometric boundaries and the local, explicit solution algorithm which is readily implemented on massively parallel computers. This project examines a number of applications of LB, including classic engineering flows, coupled thermal-hydrodynamic flows, flow about moving obstacles, and flow in realistic porous media. Particular emphasis is placed on development of a coupled experimental-simulation framework for modeling slow or creeping flows in complex porous microstructures.

## 1.1 Background to this project

LB methods typically use fixed geometric grids and a limited set of discrete particle velocities to model transport. Macroscopic averages over many sites have been shown to recover the Navier-Stokes and advection-dispersion equations. At the time that this project was initiated, there were several limitations and uncertainties associated with existing LB codes that this project sought to address.

First, there were few quantitative comparisons of LB with other numerical methods, particularly for  $Re > 100$ , and even fewer comparisons with experiments. Therefore, the

ranges of Re and Pe numbers for which the technique was practical were largely unknown. Second, while there were LB codes to model both one-component systems as well as two immiscible fluids, we were aware of no LB codes that could model reaction and transport in multi-component systems. Third, existing LB codes were not designed to handle systems with a large (>20%) fraction of solids with arbitrary geometries; this latter problem is particularly severe for three-dimensional calculations. Fourth, the specification of boundary conditions is critical; it is was already widely accepted that the lattice-gas “bounce-back” condition is only first-order accurate when applied to LB, and that the error increases with the magnitude of the pressure gradient (e.g., *Noble*, 1996). Yet, for problems in which there is some degree of uncertainty regarding the specific location of the boundaries, it is not necessary to specify extremely accurate no-slip boundaries.

This project sought to address these general questions and deficiencies, and to develop general purpose three-dimensional LB codes for modeling coupled multi-physics flows.

## 1.2 Work Plan

The general work plan undertaken is summarized in the flow chart below. Both high Reynolds Number (Re) flows as well as slow, creeping flows were investigated. Also

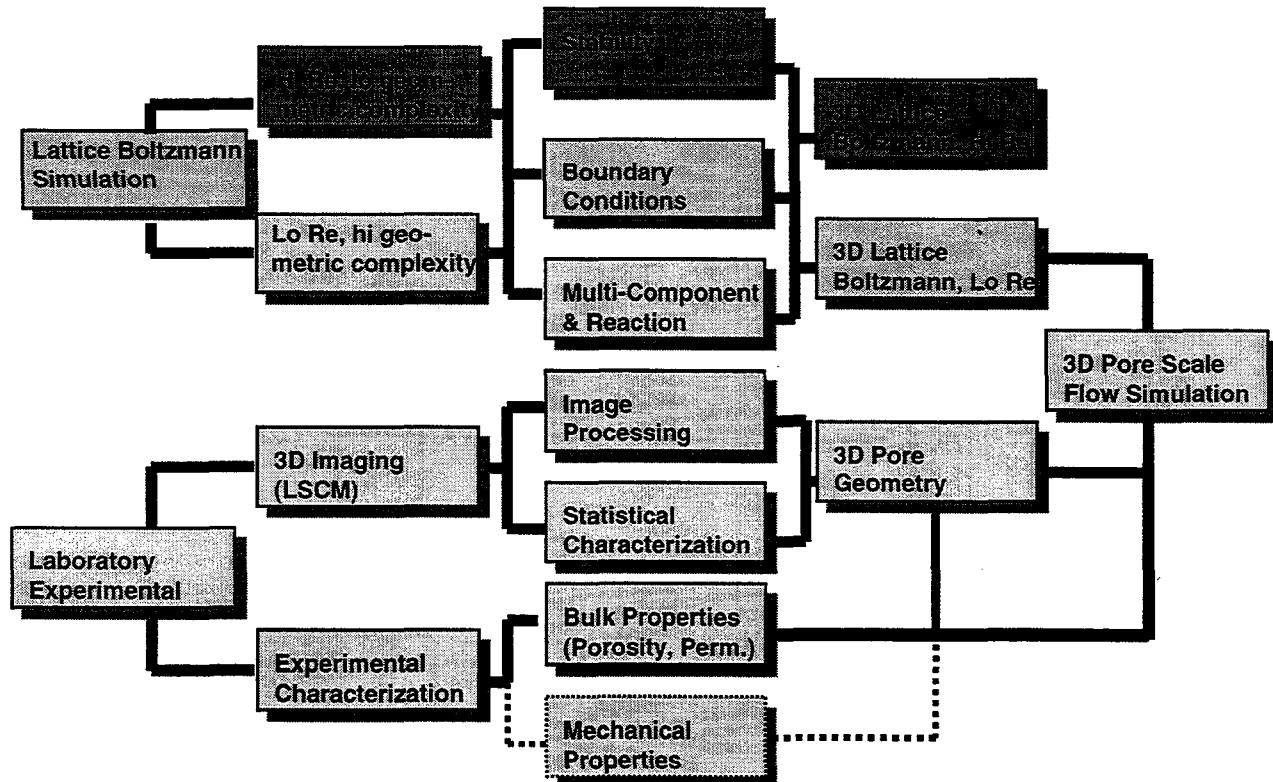


Figure 1.1 Flow chart illustrating the work plan for the LDRD project.

investigated were various boundary condition implementations, and application to multicomponent reactions. Some of the project's efforts in this latter areas are summarized by *Stockman* (1999) rather than in this report.

### **1.3 Project accomplishments**

Some specific accomplishments achieved during the course of this two-year project include:

- 1) Development of numerical LB codes for modeling non-reactive transport in complex geometries, including implementation of scalable, distributed algorithms.
- 2) Validation of the LB codes by comparison with analytic solutions, calculations based on other numerical methods, and laboratory experiments.
- 3) Design and development of a distributed computational platform using commercial off-the-shelf components that is capable of production-scale simulations (tens to hundreds of millions of lattice sites).
- 4) Development of methods and computational algorithms to extract quantitative statistical descriptions of the microgeometry of porous media from three-dimensional image data and a computational framework (numerical pipeline) for simulating flow and transport in experimentally determined porous microstructures.

Applications for this work abound in the geosciences, materials science, and chemical, thermal, and mechanical engineering, and include for example:

- 1) Flows in geologic media such as the displacement of oil from the pore space during waterflooding or enhanced oil recovery (EOR) processes, and the transport of contaminants in the subsurface.
- 2) Metals solidification and processing, in which the rate of growth of sub-millimeter dendrites controls the microstructural evolution and hence mechanical properties of the alloy.
- 3) Transport processes in porous-bed chemical reactors and filters, including fuel cells.
- 4) Heat and mass transfer processes in fibrous, ceramic insulations and ablative tiles such as those used on spacecraft.
- 5) Materials characterization and design, including evolution of physical properties during powder compaction, sintering, and hot isostatic pressing.
- 6) Fluid and thermal flows in MicroElectroMechanical Systems (MEMS).

In our future work we hope to incorporate capabilities for modeling reactive flows, as well as interfacial phenomena, including immiscible and multicomponent flows.

## 1.4 Organization of this report

This report is organized as follows.

Section 2 introduces LB methods, provides an overview of the basic physics, and summarizes the new algorithms that were developed in this project and that are described further in ensuing sections. Section 2 also describes the overall architecture of the parallel three-dimensional LB code that is a primary product of this LDRD project.

In Section 3 we describe the design and development of a distributed computational platform (named Araneæ) that was motivated by our need for a prototyping environment for parallel code development for application to very large scale simulations that require massively parallel supercomputers such as the ASCI Red supercomputer, and by our need to perform modest-sized simulations locally with rapid turnaround.

Section 4 documents validation of the numerical LB codes by comparison of LB solutions with numerical and analytic solutions for two classic engineering flows.

Sections 5-8 describe in detail fundamentally new LB models and their application to complex flows not amenable to solution with conventional CFD methods. Section 5 describes application to a coupled thermal-hydrodynamic flow, and also implementation of the LB method on an irregular grid. Section 6 details development of a LB model for simulating flow about moving and non-conforming obstacles, with applications to modeling flow in reconstructed porous media. Section 7 describes development of experimental methods for three-dimensional imaging and statistical characterization of complex porous media. Section 8 then describes development and application of LB models for simulating flow and transport in these experimentally reconstructed porous media.

## 2 LB code development

### 2.1 Introduction

The lattice Boltzmann (LB) method (*McNamara and Zanetti*, 1988; *Higuera and Jiménez*, 1989; *Chen and Doolen*, 1997) is a relatively new numerical technique for solving transport problems. The method is based on concepts from kinetic theory, but unlike particle-based methods such as Molecular Dynamics (MD) or Direct Simulation Monte Carlo (DSMC), LB does not simulate individual particle motion. Although detailed particle simulations recover Navier-Stokes behavior in the continuum limit, they are too expensive to use for simulating continuum-scale hydrodynamics. The LB method also recovers Navier-Stokes behavior (*Chen, Chen and Matthaeus*, 1992) but incorporates a simpler, probabilistic model of particle motion that is far less expensive to compute. Rather than resolving the detailed molecular scale transport, the method provides a kinetic theory-motivated solution technique for macroscopic transport equations.

### 2.2 Physics

In lattice Boltzmann, the Navier-Stokes equations are not solved directly. Instead, a discretization of the underlying Boltzmann equation is formulated such that the velocity and pressure fields satisfy the Navier-Stokes equations (and possibly an advection-diffusion equation for heat or mass transfer as described in Section 5) to within the discretization error of the method. The form of the Boltzmann equation that is solved includes discrete time, space, and velocity. The problem is formulated in terms of the particle distribution function,  $f_i(x, t)$ . Normalized by the local density, this quantity indicates the probability of finding a particle at location  $x$  and time  $t$  that is moving with velocity  $e_i$ . Unlike the continuous particle distribution function, this discrete equivalent is defined only for a fixed set of (lattice) velocities (Figure 2.1). The directions of the microscopic velocity vectors are denoted by the subscript  $i$ .

The primary variables are calculated from moments of the particle distribution, in direct analogy with the integral moments from kinetic theory. Thus, the density,  $\rho$ , and velocity,  $\mathbf{u}$ , are recovered by forming the first two moments

$$\sum_i f_i = \rho \quad (2.1)$$

$$\sum_i f_i e_i = \rho \mathbf{u} \quad (2.2)$$

where the sums are performed over all directions.

The particle distribution obeys the (discrete) velocity Boltzmann equation given by

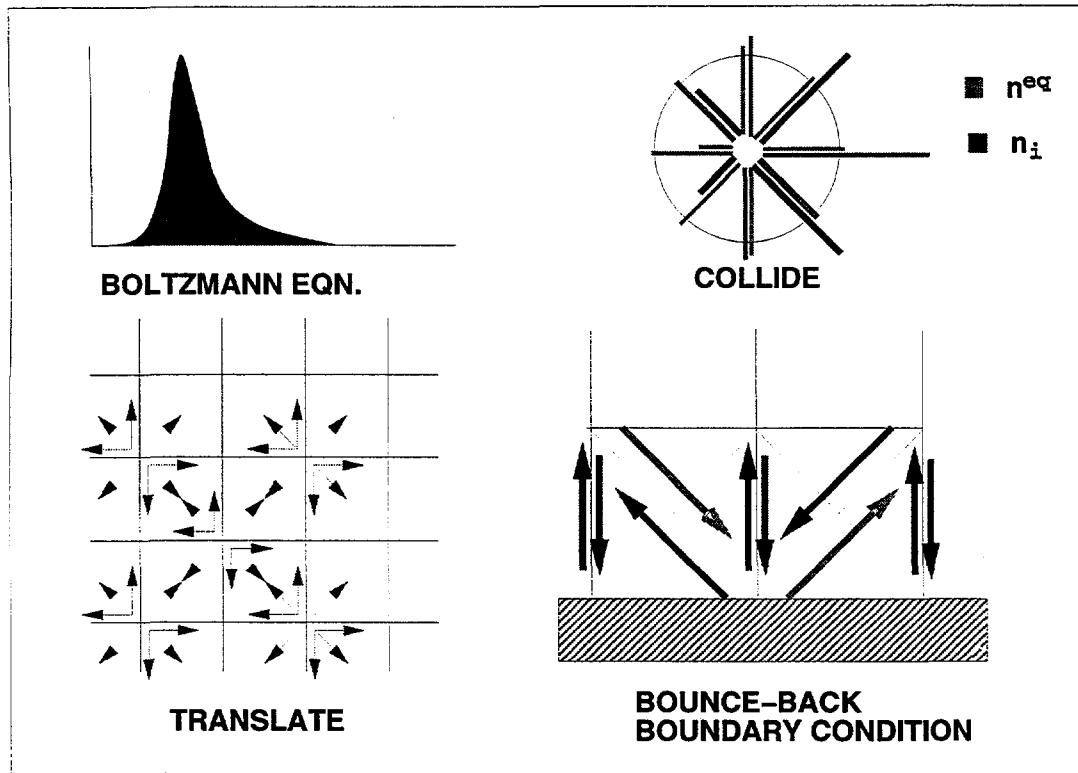


Figure 2.1 Schematic illustration of the LB method: (upper left) velocity distribution, (upper right) collision step, showing current versus equilibrium particle distribution along each of the eight links in two dimensions (there are a total of 18 links in three dimensions), (lower left) translation step, showing the streaming of particles along each of the links, and (lower right) simple boundary condition.

$$\frac{\partial f_i}{\partial t} + \mathbf{e}_i \cdot \nabla f_i = \Omega_i(f(\mathbf{x}, t)) \quad (2.3)$$

where  $\Omega_i(f(\mathbf{x}, t))$  is a collision term which accounts for the addition and removal of particles moving with velocity  $\mathbf{e}_i$  due to particle collisions.

The discrete velocity Boltzmann equation (2.3) is spatially and temporally discretized using a first-order Lagrangian discretization which yields

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \Omega_i(f(\mathbf{x}, t)) \Delta t \quad (2.4)$$

Utilizing the linearized, single time relaxation model (credited to *Bhatnagar et al.* 1954) applied to lattice Boltzmann (*Chen et al.* 1991), the collision operator is written as

$$\Omega_i(f) = -\frac{1}{\lambda} (f_i - f_i^{(0)}) \quad (2.5)$$

where  $f_i^{(0)}$  is an equilibrium distribution analogous to the Maxwellian distribution. Using this simplification, the lattice Boltzmann evolution equation is written as

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \frac{\Delta t}{\lambda} (f_i^{(0)}(\mathbf{x}, t) - f_i(\mathbf{x}, t)) \quad (2.6)$$

It is useful to define a dimensionless relaxation time,  $\tau = \lambda / \Delta t$ . Substituting gives the final form of the lattice Boltzmann evolution equation

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \frac{1}{\tau} (f_i^{(0)}(\mathbf{x}, t) - f_i(\mathbf{x}, t)) \quad (2.7)$$

Using the above formulation, the integration of the fluid momentum equations is thus reduced to two major steps. First, in a collision and forcing step, the distributions are modified according to right-hand-side of (2.7). Second, the particle distributions stream to their nearest neighbors. It is noted that the collision and forcing calculation is completely local in that all quantities are evaluated at location  $\mathbf{x}$ . All of the non-local interaction takes place during the streaming step and involves only the nearest neighbors.

Below is an outline of the recovery of the Navier-Stokes equations from the lattice Boltzmann discretization. Complete derivations are given elsewhere (*Chen et al.* 1991; also refer to the appendix of *Hou et al.* 1995). The starting point is a multi-dimensional Taylor series expansion of the particle distribution about the point  $(\mathbf{x}, t)$ :

$$f_i\left(\mathbf{x} + \frac{\mathbf{e}_i}{c} \Delta \mathbf{x}, t + \Delta t\right) = f_i(\mathbf{x}, t) + \sum_{n=1}^{\infty} \frac{1}{n!} \left[ (\Delta t) \frac{\partial}{\partial t} + \frac{\Delta \mathbf{x}}{c} \mathbf{e}_i \cdot \nabla \right]^n f_i(\mathbf{x}, t) \quad (2.8)$$

This expansion is substituted into the lattice Boltzmann evolution equation, Eq. (2.7) to give,

$$\sum_{n=1}^{\infty} \frac{1}{n!} \left[ (\Delta t) \frac{\partial}{\partial t} + \frac{(\Delta \mathbf{x})}{c} \mathbf{e}_i \cdot \nabla \right]^n f_i(\mathbf{x}, t) = \frac{1}{\tau^*} [f_i^{(0)}(\mathbf{x}, t) - f_i(\mathbf{x}, t)] \quad (2.9)$$

Retaining only the first two terms in the Taylor series expansion introduces a truncation error of  $O(\Delta x^2, \Delta t^2)$  relative to the leading order term. The Chapman-Enskog expansion, an asymptotic expansion method used in kinetic theory, is applied next. Here the expansion is in terms of a computational Knudsen number, which is defined as the time between collisions relative to the convective time scale of the flow:

$$\delta_{\Delta t} = \frac{\Delta t U}{L} \quad (2.10)$$

where  $L/U$  is the convective time scale with  $L$  and  $U$  as the characteristic length and velocity scales of the flow, respectively. This quantity characterizes the degree to which the system deviates from equilibrium. The expansion of the particle distribution and evolution time scales are described by,

$$f_i = f_i^{(0)} + \delta_{\Delta t} f_i^{(1)} + \delta_{\Delta t}^2 f_i^{(2)} + \dots \quad (2.11)$$

$$\partial_t = \partial_{t_0} + \delta_{\Delta t} \partial_{t_1} + \delta_{\Delta t}^2 \partial_{t_2} + \dots, \quad (2.12)$$

where  $\partial_t = \partial/\partial t$  is used to indicate the time derivative.

A hierarchy of evolution equations are derived by substituting these expansions into Eq. (2.9) and collecting like powers of the Knudsen number. By substituting the equilibrium distribution and taking the first two moments of the resulting equations, the macroscopic conservation equations for mass and momentum are obtained. The specifics of this last step depend somewhat on the spatial and velocity discretization. For the 2D square grid, that is later used and described further in Sections 5 and 6, the following expressions are obtained:

$$\partial_t \rho + \partial_\gamma (\rho u_\gamma) = O(\Delta x^2, \Delta t^2) \quad (2.13)$$

$$\partial_t (\rho u_\alpha) + \partial_\beta (\rho u_\alpha u_\beta) = -\partial_\alpha P + \partial_\beta [\mu (\partial_\alpha u_\beta + \partial_\beta u_\alpha)] + E_\alpha + O(\Delta x^2, \Delta t^2) \quad (2.14)$$

where the Einstein notation is used, and the subscripts  $\alpha, \beta, \gamma$  are used to denote components of the vectors. The notation  $\partial_\alpha = \partial/\partial x_\alpha$  is used to indicate the spatial derivatives. The quantity  $\mu$  is the viscosity and is given by

$$\mu = \rho v = \frac{\rho (\Delta x)^2}{(\Delta t)} \left( \frac{2\tau - 1}{6} \right) \quad (2.15)$$

Equation (2.13) implies that the correct form of the continuity equation is obtained to  $O(\Delta x^2, \Delta t^2)$ . The term  $E_\alpha$  in the momentum equation (2.14) is an error term indicating the deviation of the LBM momentum equation from the Navier-Stokes equation:

$$E_\alpha = -\frac{3v}{c^2} \partial_\beta \partial_\gamma (\rho u_\alpha u_\beta u_\gamma) \quad (2.16)$$

Examination of this term shows that the deviations are proportional to the square of the Mach number of the flow.

Summarizing the results of this subsection, LBM has been shown to simulate the Navier-Stokes equations with a truncation error proportional to the square of the discretization parameters,  $\Delta x$  and  $\Delta t$ , and the square of the Mach number.

## 2.3 Algorithmic

In this section we summarize the essential aspects of some of the algorithms developed in this project.

### 2.3.1 Moving and non-conforming obstacles

Many applications involve flow about complex geometries that are not readily fitted with a boundary conforming grid. In addition, the obstacles in the flow may move about, making it completely impossible to maintain such a grid. For these applications, which include suspension and porous media flows, a method is needed to accommodate moving and non-conforming obstacles in the flow. Section 6 focuses on this issue, and presents and validates a LB algorithm for these complex flows. In this approach, the collision term in LB is modified so that it shifts smoothly between regular hydrodynamics at nodes occupied solely by fluid to rigid body motion at nodes occupied solely by particles. Thus nodes that are only partially filled with fluid can be addressed. This extends LB capability to a wide class of problems that previously could not be modeled accurately.

### 2.3.2 Complex irregular geometries

One of the project goals was the development a simulation capability for modeling microscale flows in complex geometries, with a specific focus on complex porous microstructures as described in Section 7. The description of a complex 3D porous medium geometry is captured directly from the medial axis (3dma) analysis output stream. The basic format for the streamed data is a raster-encoded description of the voxels. In the simplest case this is a 3D array of boolean values indicating the presence of either solid or void at that site. Other more compact encodings such as the BioRad™ (BioRad, Inc.) image file format can also be used. This allows us to manipulate directly 3D image data recovered using laser scanning confocal microscopy as is described further in Section 7. The 3D geometry is mapped to a lattice in the form of either a binary mask or a stencil as summarized below and later in Section 8. Also, alternative approaches to the partial voxel problem, which is intrinsic to all imaging techniques, are described in Sections 6 and 7.

### 2.3.3 Stencil representation

The complex porous geometry can be mapped onto the lattice such that every voxel from the imaged domain is mapped onto a fully defined lattice site. As the lattice is fully populated (defined), the binary data thus act as a mask designating whether the fluid physics at that site is included in the overall system and needs to be updated. However, while this 1:1 mapping is straightforward and automatically creates a mesh of the geometry, it is exceedingly wasteful in terms of storage for materials with large solid fractions.

A fundamentally different mapping approach that we refer to as stenciling was developed to address this problem (Figure 2.2). The essence of our method is to use the binary data set as a stencil for memory allocation and layout, rather than as a simple mask. By this we mean that lattice memory is allocated only for sites in the volumetric data set that correspond to pore space. Additionally, the location of the memory associated with each void voxel is encoded and stored in a corresponding location in a three dimensional image map on a site-by-site basis. With this approach, the algorithm traverses the image, and only

2D Geometry  
(Binary Map)  
Porosity = 50%

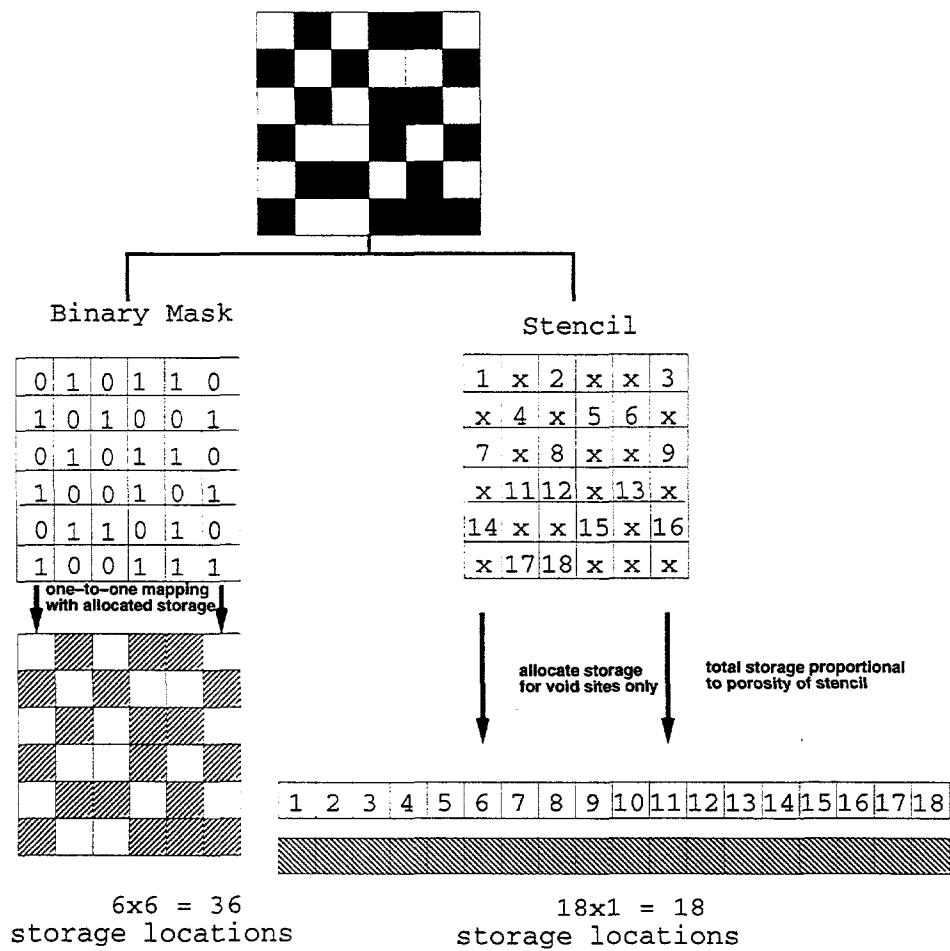


Figure 2.2 Schematic illustration of the traditional mapping approach versus the stencil algorithm.

when a site corresponding to a void space is encountered are the contents of the associated memory accessed to determine the state of the physics at that location. This scheme requires that we store a representation of the fluid state at pore sites only. Thus the storage requirement is identically proportional to the porosity of the medium being modeled.

### 2.3.4 Volume visualization

The Araneæ cluster system that is described in Section 3 is utilized directly in a number of ways to perform or facilitate visualization of simulation results. Each compute node has access to a number of 3D graphics libraries, including Mesa, an emulation library for OpenGL. Mesa forms the basis for all visualization performed on the cluster.

The simplest form of visualization is to direct render each subdomain on a separate processor, in a unique window or frame, and display the results on a single monitor accessible by all processes. The windows can be manually tiled and a reasonable depiction of the progress of a simulation can be monitored remotely. While visually disparate, the rendering does require a certain amount of global synchronization to apply consistent color levels. For example, a color-coded profile of phase density requires knowledge of the global maximum and minimum density and any normalization coefficients. When applied to data structures such as regular lattices we find it adequate to direct render onto the viewing plane of interest. For more complex visual analyses this can be inadequate as sufficient occlusion (hidden line removal) and consistent lighting calculations may not be possible for accurate rendering.

The second general method used is to display the sub-domains in a single window on the root process using a method referred to as compositing. The idea is to perform a full 3D rendering of each subdomain as before and store it in a data structure we call a depth image. The depth image is composed of a record of the 2D spatial extents of the image, a 2D RGB image containing the color of each pixel making up the sub-image (i.e. what would be displayed for the subdomain), and a depth buffer of equivalent extent that stores the depth of each pixel (with respect to a common camera location). Depth images from all sub-processes are combined logically on the root process to create a complete 3D image that is then stored or displayed as a single frame. This operation simply compares the  $xy$  position and depth of an incoming pixel with the one that is already stored. If the incoming pixel is closer to the viewer then it and the corresponding RGB pixel replace those in the primary depth image. When all depth images are combined the resulting image describes correctly the composite 3D image with occlusion.

To better investigate physical phenomena arising in the numerical simulations that Araneæ is used to perform we have implemented a simple but useful distributed isosurface algorithm and rendering functionality. Isosurface rendering is the mainstay of visualization in fluid dynamics and this allows us to apply common techniques to probe the physical systems being investigated. The basic concept is to tessellate discretizations of scalar fields on a regular lattice and image an interpolation surface through the sample points. It should be noted that such techniques are based on the assumption of continuity and may not be applicable to all examinations. Specifically, if we are interested in phenomena that arise at length scales smaller than that of the sampling mesh resolution, we can only make statements about average properties at this scale.

The forth and final visualization method that we use to generate 3D image data sets from the simulation results (velocity components, phase densities, etc) and visualize or

analyze them uses the hardware-supported volume visualization software called VoxelView™ (Vital Images, Inc., Fairfield, IO) that runs on Silicon Graphics workstations.

### 2.3.5 Parallel I/O

I/O on distributed memory platforms is a challenging issue. This problem has been acute on the ASCI-RED system where we have experienced I/O latencies that have consumed upwards of 90% of the time slot normally allocated to perform numerical simulation. This diminishes or excludes the ability to complete a large-scale simulation within a standard allocated time slot.

We have addressed this problem with a simple but effective approach applicable to dedicated cluster computers such as the Araneae system described in Section 3 of this report. The basic premise is to maintain a local disk on each node of the system with sufficient storage to allow useful amounts of intermediate and final results to be retained for subsequent analysis or visualisation without recourse to external media. This approach is possible in an environment where specific nodes can be allocated on demand, such as with small private clusters.

We have also implemented a suite of routines to facilitate the consolidation and decomposition of very large files within the parallel environment. This allows us to collect distributed data into a single file for archival purposes, or to decompose such large files over differing processor allocations and layout without recourse to manual migration.

## 2.4 Code architecture

The primary LB code developed in this work is named PoSTAL (**P**ore **S**cale **F**low and **T**ransport on **A** **L**attice). The code is written in C with strong typing (data encapsulation) to provide a level of abstraction similar to that of less efficient languages such as C++. The code architecture is based on defining data types that map to components of the mathematical model of the problem (i.e. LB) and a supporting software framework with which to interact with the model (I/O). Without loss of generality, the data structures are defined with a view to being manipulated in a distributed computing environment. That is, we define global and local descriptors for the problem in such a way that they become synonymous when handled in a single process environment.

For the LB model we define the following data types:

Links_t	- vector description of a lattice site
LBParams_t	- physical and numerical parameters used to compute the model system (viscosity, density, time step, etc)
IOP3D_t	- I/O parameters and descriptors

SDomain3d_t	- a description of the computational space associated with a single process. It forms a subset of a complete lattice.
Lattice_t	- a global description of the problem domain, including the extents of the geometry, the number and identity of sub-domains, boundary conditions, bulk or average descriptors (density, flow speed, permeability, etc), embedded data types describing the sub-domains, physical and I/O parameters, and information to support distributed computation

For each broad grouping of descriptive information we define a library of routines to perform actions on the data. The functionality implemented includes:

Lattice – operations on the lattice independent of the LB algorithm:

- memory allocation for geometry, sites, state
- classification of sites (solid, void, boundary)
- statistics (porosity, surface area, etc).
- stencil operations

I/O – operations to perform I/O:

- initialize description of model I/O
- access/create buffers, files
- access 3D geometry from input stream
- generate and capture state information
- generate visualization information

params – operations to handle model parameters:

- allocate parameter storage
- form/set/retrieve parameters

mpi – message passing operations:

- set up communications patterns
- manage message buffers
- communications

viz – operations to support 3D visualization:

- windowing operations
- lighting operations
- rendering operations

- isosurface (sampling and tessellation)

LB – operations to perform the LB algorithm

- apply boundary conditions

- collision operation

- stream particles from site to site

The application of these routines and the sequence of operations used in the LB algorithm are described in more detail in Section 8 of this report.

### 3 Araneæ— A scalable system for distributed computing

Araneæ is a prototype distributed memory parallel computer architecture based on interconnecting a number of dedicated processing units built from commodity components and recently-developed high bandwidth, low latency network fabrics. Our development was motivated by four factors:

- 1) A need to routinely perform large-scale numerical simulations
- 2) Limited access to large-scale computing resources
- 3) Availability of high performance computing components at commodity costs
- 4) Recent introduction of high bandwidth, low latency network switching technologies at or near commodity costs

For the purposes of this report, large-scale computation refers to computing the solution to a problem that requires resources beyond those available on a standard scientific workstation. That is, the amount of memory required to represent the problem and/or the amount of sustainable, uninterrupted CPU time required to complete the solution are not feasible.

The specific problem driving our development is constructing a representation of the geometric microstructure of a porous material using an advanced imaging technique known as laser scanning confocal microscopy (LSCM) and a statistical characterization technique known as Medial Axis (MA) analysis as described in Section 7. The reconstructed geometric model forms the basis for an automatically generated mesh that is used to simulate physical processes such as flow and transport using numerical techniques known as lattice Boltzmann methods (LB). Further details of this specific application are provided in Section 8.

From the representation of the geometry and the specific numerical technique applied (i.e., LB), we can isolate the two principal characteristics that constrain the computation of the system behavior. The first is the amount of storage required to represent numerically the system, i.e., the numerical representation of the geometry in the form of a mesh. The second characteristic is an estimate of the CPU time required to compute (simulate) the evolution of the system to a desired state.

#### 3.1 Space: flow modeling in reconstructed porous media

Imaging techniques are a means to recover directly a detailed description of the two-phase microstructure of a porous medium (i.e., a material consisting of solid and void phases). In the work reported here, we use a high-resolution three-dimensional imaging technique known as laser scanning confocal microscopy (Section 7). However, high-resolution data comes at a cost that is proportional to the amount of storage required to record the image data. As is described in detail in Section 7, we typically collect image data at 8 bit (1 byte) resolution (equivalent to 0-255 gray levels). In our case, the image data consists of the spatial fluorescence intensity, with data collected at regularly spaced

locations in  $xyz$  space. The volumetric data set thus consists of a three dimensional array of voxels (i.e., the three dimensional equivalent of a two dimensional pixel, or picture element).

Sampling at a resolution of 1 micron ( $\mu\text{m}$ ) is typical for granular porous materials where the characteristic length scale is  $\sim 100 \mu\text{m}$  (one grain dimension). Thus, several hundred microns of coverage is required to capture the geometry of several interconnected, cemented grains. As an example (and one that we use in practice), consider a volumetric image containing  $768 \times 512 \times 128$  voxels. Each voxel has a volume of  $1 \mu\text{m}^3$  and requires 1 byte of storage to record the image data for that site. Overall, the complete data set requires  $768 \times 512 \times 128 = 50 \times 10^6$  bytes of information. In terms of storage, this is equivalent to 50 Megabytes of disk space.

The next step is to binarize the voxels using one of several computational algorithms whereby each voxel is interpreted as either the solid or void phase (see Section 7). The binarized data set forms a discrete representation of the pore space that is used subsequently in numerical simulations of flow and transport. Each voxel is directly and automatically mapped to a site in the model domain. What then remains is to represent the physical state of the fluid phases residing in the pore phase. The conventional implementation of the lattice Boltzmann method requires the physical state of the fluid to be represented at each and every site of the mesh, i.e. every voxel. For the physics of interest in our application, approximately 100 bytes of information are required to represent the fluid state for each lattice site. Scaling this quantity by the size of the computational domain defines a storage requirement of approximately  $50 \times 10^6 \times 100$  bytes, or 5 Gigabytes. Currently, this requirement greatly exceeds the amount of storage available on a standard scientific workstation.

### 3.2 Time: evolution to steady state

The second constraint on computational resources is the time required to evolve the simulated behavior to reach a steady state condition. LB are numerically explicit as formulated, i.e., a finite propagation time is required to transfer information about the physical state of any one site to the neighboring sites in the model domain. And to reach a steady state, every site must receive at least one set of information from every other site in the model domain. As a general metric, the number of update steps required is proportional to the square of a characteristic dimension,  $L_c^2$ , that is typically equal to the dimension of the largest connected flow path.

While no specific measure can be given *a priori*, the specific implementation of LB used in our work can update  $10^6$  sites in 4 seconds of CPU time on a standard scientific workstation (e.g. a 200 MHz Pentium Pro running the LINUX operating system). If the characteristic dimension is estimated conservatively at 150 lattice sites, this leads to a requirement of at least  $2.25 \times 10^4$  site updates. To update the complete system ( $50 \times 10^6$  sites) requires  $50 \times 4 = 200$  CPU seconds, and to reach steady state requires  $200 \times (2.25 \times$

$10^4$ ) CPU seconds. This is equal to over 50 CPU days, a result that is unacceptable for practical application and usage.

### **3.3 System architecture**

As discussed above, the key requirement of our application is the need to represent large amounts of data (space) and to quickly process this data (time). And although the supercomputing resources available at Sandia are unparalleled, routine, regular, and uninterrupted access to such resources is generally not the case. We had a need for a local computational environment capable of modest-sized (but still production-scale) simulations in less than a day's time, and also in which we could perform code prototyping for very large-scale simulations that could only be performed using resources such as the ASCI Red supercomputer.

By 1996, commodity components such as those used in desktop personal computers (PCs) had reached a price-performance ratio comparable, and more often better, than those components used in high-end scientific workstations and name-brand parallel computers from companies such as SGI/Cray, IBM, and Intel. Simultaneously, high bandwidth low latency network switching technologies at or near commodity costs were being introduced. For these reasons we chose a distributed memory multiprocessing model for the system architecture.

For the purposes of this discussion, the term node is used generically to refer to a unit resource containing some amount of memory to store data and a processor capable of performing some amount of computation. No distinction is made between nodes containing single or multiple processors, dedicated or shared memory, or the activity the resource is intended to perform. The term cluster is used to refer to a collection of nodes that are configured so as to allow communication between nodes. The term scalable unit (SU) is meant to connote the basic building block that is used to construct the larger system while maintaining uniform network capacity.

The requirements of large-scale computation can be translated simply into having a large amount of memory distributed over an appropriate number of processing nodes. What remains is the issue of how to interconnect the processing nodes so that they work cooperatively, sharing the data stored on each node, in a manner such that the computation is not restricted by the requirement to share information between nodes.

#### **3.3.1 Networking**

Networking describes the general means used to connect groups of nodes together so that they can communicate and work cooperatively. The primary techniques used are bus and switched interconnects.

The simplest interconnect is the bus, which is a single communication channel (wire) with taps to connect two or more processes together. A bus with multiple taps is known collectively as a hub. When three or more processes communicate, they must time share use of the wire. Obviously, the more processes that share use of the hub, the higher the likelihood of simultaneous contention for use of a limited resource. A resolution strategy or protocol (that forms the basis for Ethernet) is to allow a single pair of nodes to communicate through the resource at any one time while the remaining nodes wait until the hub is clear.

A switch consists of dedicated wires to interconnect either statically or dynamically pairs of communicating processes. When there is a single wire available to connect simultaneously every unique pair of processes, the switch is called a cross-bar. While providing optimal connectivity, this form of switch does not scale well and is expensive to implement.

Between the hub and the cross-bar are a spectrum of packet and circuit switched networks. These are fabrics that provide a limited number of wires coupled with sophisticated routing algorithms implemented in software and hardware to direct efficiently messages between communicating processes. Hub and packet switched variants of Ethernet are commercially available, ranging from 10baseT – Ethernet (10Mbit/sec) to 100baseT – Fast Ethernet (100Mbit/sec) and up to 1000 baseT – Gigabit Ethernet (1000Mbit/sec). The calibration of these networks refers to the peak bandwidth of wires connecting a pair of processing nodes.

We chose to utilize a dual network system. A 100BaseT – Ethernet connectivity to provide services, file system access, diagnostics and a portion of the I/O. The second network is a Gigabit bandwidth, low latency, switched fabric called Myrinet™, developed by Myricom, Inc. (Arcadia, CA). Myrinet is used exclusively for message passing.

Each SU is internally connected using a 3Com Superstack II baseline switch and a 16 port Myricom M2M-OCT-SW8 switch. Ethernet connection between SUs is provided by routing through the service nodes. Myrinet connectivity between SUs is provided through inter-switch links. In this way, the Myrinet is scalable in a manner that preserves the concept of minimum bisection bandwidth. On the other hand, the 100 baseT ethernet is restricted to a flat communications space for groups of 16 nodes. Thus, to scale the ethernet would require a means to uplink to a higher bandwidth fabric (for example, through a 1000 baseT packet switch).

### 3.3.2 Cost

The two most common commodity based processors used in clustering are Intel architecture processors and clones, and Compaq Digital Alpha processors. At the time of this writing, the cost-performance ratio for the two components is approximately equal. However, market factors tend to drive the Intel based components to better price-performance ratios more rapidly than Compaq Digital products. In the past year, price reductions in the components market have been approximately 50% when compared to

equivalent costs at the start of this development. The effective cost of a cluster capable of 1 Gigaflop in August 1997 was approximately \$25K. The same cluster can now be purchased and constructed for approximately \$12.5K. However, it should be noted that smaller numbers of components are needed to achieve this example capability (1 Gigaflop) because current processors are  $\sim 2\times$  as fast as one year ago, and thus scaling is not necessarily linear.

### 3.4 Operating system and environment

The base operating system used to control the Araneæ nodes is LINUX. LINUX is a multi-processing operating system based on UNIX that schedules processes, controls peripheral devices, and provides a low-level software platform on which to develop and run applications.

LINUX implements a UNIX kernel with no proprietary code; that is, the source code for the operating system is freely available. Red Hat Software, Inc. (Research Triangle Park, NC) bundles LINUX with a suite of utilities that allow applications to be developed on top of the operating system. For example, compilers (gcc, g77, g++) and editors (emacs, sed) developed by the Free Software Foundation under the GNU project are standard issue with the Red Hat distribution.

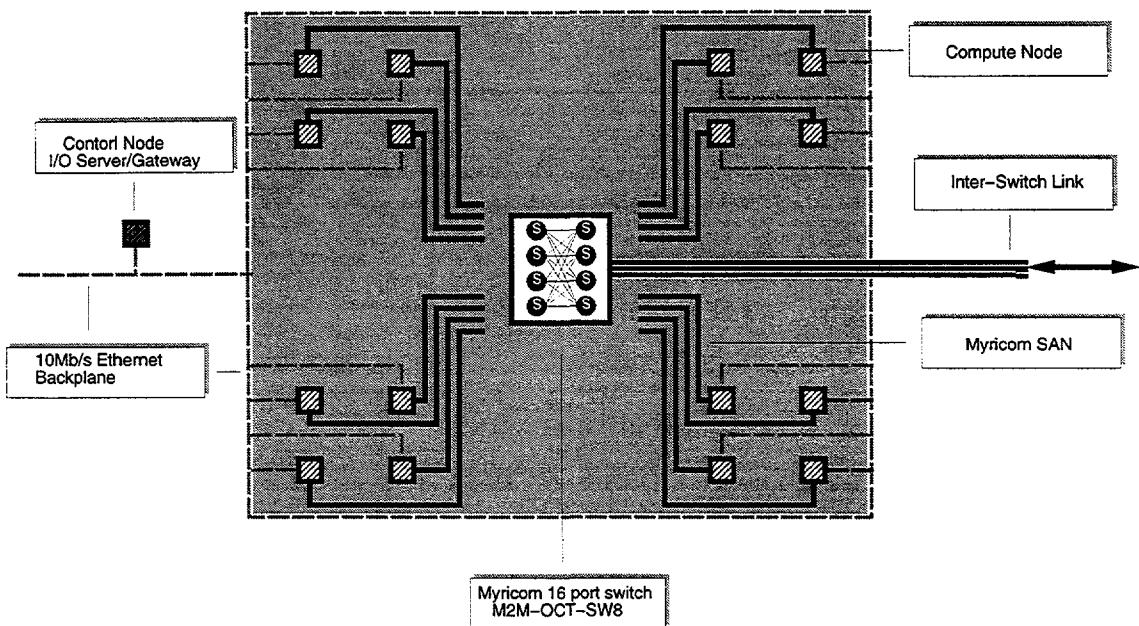
To use LINUX in a networked environment requires several layers of additional operating system software, most of which are also freely available from the LINUX community and/or bundled with Red Hat. Specifically TCP/IP networking libraries, network interface drivers, networked file system (NFS) support and high level application libraries to support message passing are custom installed and configured to create the necessary environment to perform distributed computing.

#### 3.4.1 Programming model

The application level programming interface to perform message passing is MPI (*MPI Forum*, 1995). Specifically, we use the MPICH library implementation (*Gropp and Lusk*, 1997) of the *de facto* standard MPI to ensure portability to multiple hardware platforms and vendors. Message passing is a programming method that allows multiple processes to communicate with one another (whether they are running on the same processor, or on another processor connected over a network), in a way that makes the computation appear contiguous at the shared boundaries of the subdomains.

#### 3.4.2 System definition

As of November 1998, the Araneæ system consists of 32 computational nodes that are configured as 2 sub-clusters, or scalable units (SU) of 16 nodes each (Figure 3.1).



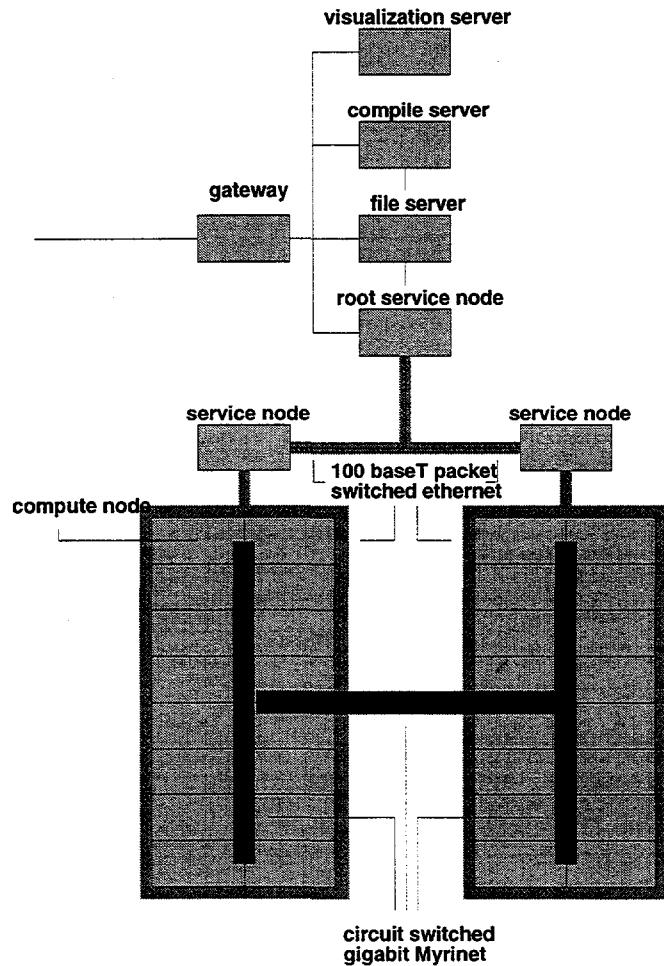
**Figure 3.1** Schematic illustrating a scalable unit (SU) in the Araneæ cluster.

As illustrated in Figure 3.2, the system is supported through a:

Root service node	– to administer the service nodes
Hierarchy of service nodes	– used to mediate I/O transactions between the file server and a SU, administer user and application resources, etc.
File server	– to maintain system and user files, binaries, libraries etc.
Compile node	– to develop and link applications
Visualization node	– for post-processing
Gateway node	– to access the system from the internet

Each computational node (CN) is configured with:

CPU	– 333 MHz Intel Pentium-II
Motherboard	– Intel SE440BX with 100MHz memory bus
RAM	– 256M PC100 compliant memory
Storage	– 4.3 Gigabyte local hard drive
Network interface card	– 3Com 3c905bTX 10/100 auto-negotiation
Myrinet interface card	– Myricom M2MPCI32B SAN interface card
Case	– ATX form factor



**Figure 3.2** Schematic illustrating the current Araneæ cluster.

And as discussed above, each group of 16 nodes is configured with 2 networks, a 100 baseT ethernet and a Gigabit Myrinet.

Each service node (SN) including the top level servers are configured at a minimum as follows:

CPU	– 200 MHz Intel Pentium-Pro
Motherboard	– Intel 440LX
RAM	– 128M, 60 ns EDO
Storage	– 4.3 Gigabyte local hard drive
Network interface card	– 3Com 3c905bTX 10/100 auto-negotiation
Myricom interface card	– Myricom M2MPCI32B SAN interface card
Case	– ATX form factor

Cumulatively, Araneæ is capable of ~ 6 Gflops, with 8G of RAM, 160 GB of combined storage, and 1 Gbit networked communications bandwidth. The system consists of

two scalable units (SU) each of which contains 16 computational nodes that are controlled by intermediate servers and I/O workstations.

### 3.4.3 Compute node definition

All compute nodes are stateless until provided with an operating system image and access to a common file system. Compute nodes may maintain local disk storage exclusively for intermediate computational results. While these results can be recovered as a distributed file system, this usage is not guaranteed. A common file system is mounted from the file system server using the network file system (NFS) protocol. This allows users to read/write files on any compute node and provides a level, albeit minimal, of transaction management. Users have common access to this file system on all nodes that they utilize. That is a user's identity (password, user id, group id, etc.) is the same on every node. This facility is not provided through the standard networking information service (NIS) for security reasons.

To enable this, each compute node boots from a single image supplied on the network. This requires a custom boot floppy containing a kernel with preset instructions to broadcast a request for standard boot information from a file server using the BOOTP protocol in the form of a client program. This information consists of the node's designated IP address, the location of the kernel image it should use to perform a full installation, and the location of the file system that it should mount from the server. The kernel on the boot floppy is configured to download the complete boot image using the trivial file transfer protocol (TFTP).

The files mounted include a common set of system binaries /sbin, /bin, and a slightly modified /etc. The directory /var is defined locally, and to adhere to the stateless model for the nodes, we implement it in the form of a ramdisk, i.e. we emulate a mountable file system in RAM. Each node then mounts the local disk to be used exclusively as scratch space.

The nodes are configured as a private network, i.e. they cannot be accessed directly from the internet. This is implemented using a predefined series of IP addresses. Access to the nodes is provided through a single gateway server that connects the private networks with the external network (internet). This requires a kernel with support for multiple NIC drivers on the gateway node. Drivers supporting the Myrinet interface are accessed through the mounted file system and are installed dynamically in the kernel using the LINUX modules functionality.

### 3.4.4 Useful URLs

<a href="http://www.redhat.com">www.redhat.com</a>	– Red Hat homepage
<a href="http://www.sunsite.unc.edu/pub/Linux">www.sunsite.unc.edu/pub/Linux</a>	– LINUX archives
<a href="http://www.kernel.org">www.kernel.org</a>	– Home of the LINUX kernel source
<a href="http://www.myri.com">www.myri.com</a>	– Myricom homepage

### **3.5 Discussion**

We have demonstrated the design and development of a distributed computational platform for the performance of numerically intensive simulations such as the pore scale flow modeling that is described later in Sections 7 and 8. While not a foreseen goal of this project, this result has greatly extended our ability to routinely perform large-scale simulations. Using this system, we can achieve single-phase flow simulations for  $\sim 0.25$  billion sites for porous media with representative porosities of 25%.

Furthermore, the computational platform as developed is scalable and we envision continued expansion. Cluster systems provide a vital and necessary bridge between the current resource discontinuity facing the typical researcher, i.e. individual scientific workstations versus massively parallel supercomputers such as the ASCI Red supercomputer.

This is evidenced uniquely by the duality of the cluster. That is, the cluster also provides the necessary post-processing platform to analyze very large simulations performed using the ASCI Red supercomputer as described previously in Section 2.

## 4 Validation problems

The accuracy of the LB method was determined quantitatively by comparing LB simulation results with analytical solutions and published numerical results obtained using conventional CFD methods. The two classical flows considered are flow over a backward facing step and flow through a three-dimensional duct.

### 4.1 Flow over a backward facing step

Flow over a backward facing step is a classical 2D flow that involves recirculation (Figure 4.1). Due to the presence of the step, the flow separates, creating a primary recirculation region behind the step. At moderate Reynolds number  $Re$  ( $Re > 400$ ), the downward motion behind the step causes the flow to separate along the upper wall forming a secondary re-circulation region. Comparing LB predictions for these lengths to previously published experimental and numerical results provides a good benchmark of LB for moderate Reynolds number flows.

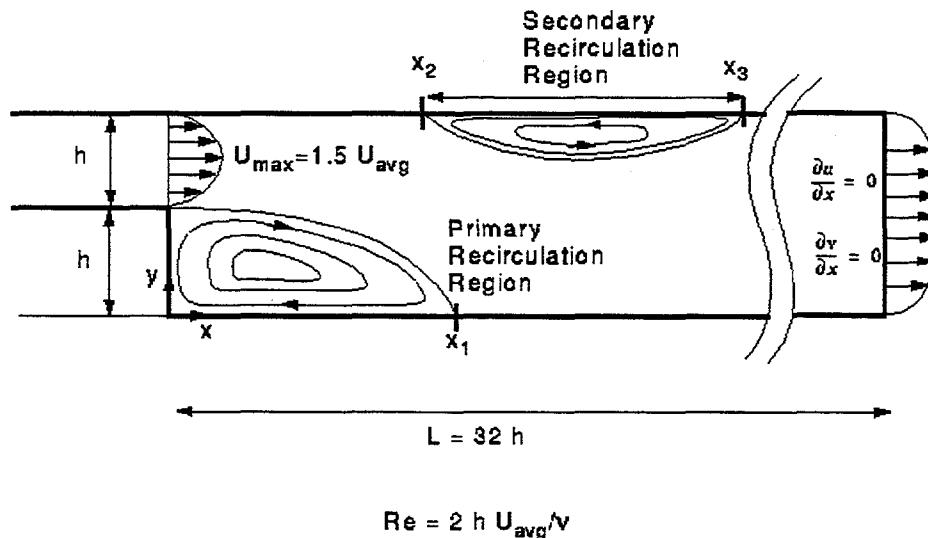


Figure 4.1 Illustration of the backward facing step problem.

LB simulations of flow over a backward-facing step were performed for  $Re$  equal to 200 and 450. Due to the considerable variation found among published results, quantitative evaluation of the accuracy is difficult. Nonetheless, for these Reynolds numbers, the LB results agree very well with previous results (Table 4.1). Simulations at  $Re=800$  were also

attempted, but the velocity field predicted by LB continued to oscillate in time and did not reach a steady solution.

A number of researchers have examined this flow previously to determine if the Navier-Stokes equations have a steady solution at this Reynolds number (*Gresho et al.* 1993). As the Reynolds number is increased, the flow is expected to become unstable, eventually becoming turbulent. For a Reynolds number of 800, however, their conclusion was that the flow does possess a steady solution. However, the flow may be termed “convectively unstable.” At this Reynolds number, a perturbation introduced continually at the inlet will not be damped out. Thus, a slight oscillation in the inlet boundary condition can cause an oscillatory solution throughout the domain. This underscores the importance of accurate boundary conditions for quantitative predictions of moderate to high Reynolds number flows.

These simulations establish LB as a quantitative method for predicting low and moderate Reynolds number flows and point to areas of further investigation. The application of LB to flow over a backward-facing step with heat transfer is examined in Section 5. In that investigation, a non-uniform grid is used to refine the region near the step. Further boundary condition development and the application of non-uniform and adaptive grids will enable LB simulations to resolve the sharp gradients that occur at higher Reynolds numbers.

**Table 4.1** Comparison of LB solution with experimental data and conventional CFD methods including finite element (FE) and finite difference (FD) for the classic backward facing step problem.

<i>Re</i>	LB	<i>Armaly et al.</i> (experimental)	<i>Gartling</i> (FE)	<i>Kim &amp; Moin</i> (FD)	FLOW-3D (FE)	FLUENT (FE)
200	$x_1=5.28$	5.0		5.1	4.87	4.98
450	$x_1=9.16$ $x_3-x_2=3.82$	9.4 3.7		9.0	8.54 3.03	8.50 4.89
800	$x_1=?$ $x_3-x_2=?$	14.4 8.2	12.20 11.26	12.0 11.5		

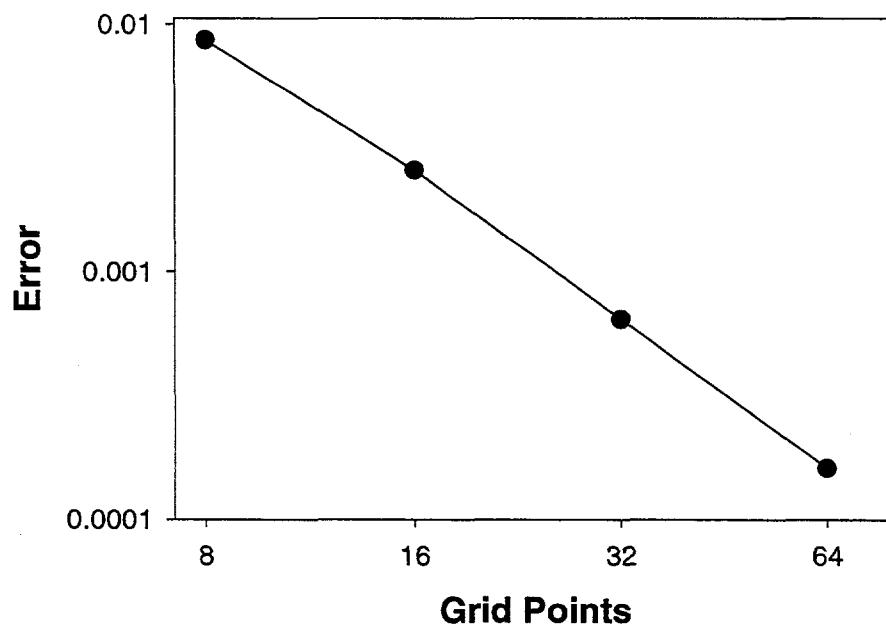
## 4.2 Three dimensional duct flow

Quantitative assessment of the 3D LB code was obtained by performing simulations of flow through a duct which is square in cross-section. The accuracy was directly evaluated by comparing LB simulation results with the analytical solution, and the root mean square of the error was calculated for simulations with varying grid sizes. The velocity field for this flow is given by,

$$u(y, z) = \frac{16a^2}{\mu\pi^3} \left( -\frac{\partial p}{\partial x} \right) \sum (-1)^{(i-1)/2} \left[ 1 - \frac{\cosh(i\pi z/2a)}{\cosh(i\pi/2)} \right] \frac{\cos(i\pi y/2a)}{i^3} \quad (4.1)$$

where  $u$  is the velocity component in the  $x$  direction,  $\partial p/\partial x$  is the imposed pressure gradient,  $\mu$  is the viscosity and  $-a \leq y \leq a, -a \leq z \leq a$ .

The results given in Figure 4.2 show that for this flow, the LB method incurs less than 1% error with 8 grid points spanning the duct in each direction. In addition, for simulations with  $\geq 8$  grid points, the error decreases with the square of the grid size and exhibits second order accuracy. The boundary conditions used for this benchmark were proposed by *Chen, Martinez and Mei* (1996). This demonstrates the quantitative accuracy of the 3D LB implementation.



**Figure 4.2** Convergence of LB simulation to analytical solution for 3D duct flow

## 5 Application to a coupled thermal-hydrodynamic flow on non-uniform grids

In many applications in materials science and engineering, a coupled heat transfer and fluid dynamics model is needed. In this section, a coupled thermal-hydrodynamic LB model is developed and assessed. Recently, multiple LB models have been proposed for solving the Navier-Stokes equations coupled with the energy equation. A LB model for natural convection was developed by *Shan* (1997) using a multicomponent LB model previously developed for simulations of partially immiscible fluids. The temperature was approximated by the density field of the second component of the mixture, which was shown to obey the advection-diffusion equation.

*Noble, Georgiadis and Buckius* (1997) developed a LB model for simulating inert solute transport in randomly packed beds using a slightly simpler model for the advection-diffusion equation. By analogy, these results directly apply to heat transfer as well. Using the Chapman-Enskog expansion technique from kinetic theory, this work showed that the scalar transport equation can be accurately simulated using a four-speed LB model as compared to the nine-speed model used previously (*Shan*, 1997). More importantly, an expression for the LB particle distribution is derived in terms of the solute concentration and its gradients. This allows the prescription of accurate boundary and initial conditions for the thermal particle distribution function in terms of known macroscopic quantities, such as the temperature or heat flux. This model is extended here to the problem of mixed convection using a Boussinesq approximation.

The majority of LB models have utilized a uniform grid. Recently, a number of approaches have been proposed for implementing the LB method on irregular grids (*Nannelli and Succi et al.*, 1992; *He, Luo and Dembo*, 1996; *Chen, Martinez and Mei*, 1996). In the interpolation-supplemented LB equation (ISLBE) method proposed by *He, Luo and Dembo* (1996), the normal LB algorithm is followed at all of the grid points, which do not need to be uniformly distributed. This algorithm includes a linear advection step in which the components of the particle distribution are uniformly shifted according to the microscopic velocity. Then in an additional step, the updated values at each grid point are calculated by interpolating between the shifted quantities. The method has successfully been applied to flow about a circular cylinder using a cylindrical coordinate system. This method is applied here to the coupled mixed convection problem. This approach is appropriate for investigating flow past a backward facing step in light of the aspect ratio of the physical domain.

Another topic that has drawn significant attention is the specification of accurate hydrodynamic boundary conditions in the LB method. Several methods have been proposed based on heuristic rules, such as the ‘bounce-back’ condition (*Cornubert, d’Humières and Levermore*, 1991), finite difference concepts (*Chen, Martinez and Mei*, 1996), and analysis of the discrete velocity Boltzmann equation (*Skordos*, 1993; *Noble et al.*, 1995; *Zou et al.*, 1995; *Ginzburg and d’Humières*, 1996). This list of references is not exhaustive; *Chen, Martinez and Mei* (1996) give a more complete review. The goal of each method is the

same: to formulate a boundary condition for the particle distribution function in terms of the known velocity or pressure at the boundaries. For thermal boundary conditions, analysis of the discrete velocity Boltzmann equation gives explicit constraints for the thermal particle distribution in terms of macroscopic quantities (Noble, Georgiadis and Buckius, 1997).

*Ginzbourg and d'Humières* (1996) recently have performed a similar analysis for the hydrodynamic LB model. The finite difference method of applying boundary conditions (Chen, Martinez and Mei, 1996) is used here, however, for the hydrodynamics for two reasons. First, the method is readily applied to various types of boundary conditions and boundary orientations. Also, the method is relatively simple to incorporate into the LB algorithm.

In this section, the accuracy of a coupled thermal-hydrodynamic LB model applied to a benchmark mixed convection problem is quantified. In previous work LB methods have been shown to accurately model isothermal flows (Martinez *et al.*, 1994; Reider and Sterling, 1995; Hou *et al.*, 1995; Noble, Georgiadis and Buckius, 1996) and inert solute transport (Noble, Georgiadis and Buckius, 1997). The model used here is an extension of the previous inert solute LB model (Noble, Georgiadis and Buckius, 1997) to coupled flows accurately described by a Boussinesq approximation. In addition, the use of a non-uniform grid (He, Luo and Dembo, 1996) and finite difference-based boundary conditions (Chen, Martinez and Mei, 1996) are addressed for this problem.

## 5.1 Lattice Boltzmann model for heat transfer on a non-uniform grid

Using the LB method, the Navier-Stokes and advection-diffusion equations are not solved directly. Instead, a discretization of the underlying Boltzmann equation is formulated such that the hydrodynamics and scalar fields (i.e. temperature) satisfy the correct transport equations to within the discretization error of the method. The problem is solved in terms of particle distribution functions,  $f_i$  and  $g_i$ . The quantity  $f_i$  relates to the probability of finding a particle in the vicinity of  $x$  at time  $t$  that is moving with velocity  $e_i$ . In mass transfer problems,  $g_i$  relates to the probability of finding a solute particle in the vicinity of  $x$  at time  $t$  that is moving with velocity  $e_i$ . By using a heat/mass transfer analogy, the same method is used here to compute the heat transfer in the channel. Unlike the continuous particle distribution function in kinetic theory,  $f_i$  and  $g_i$  are defined only for a fixed set of velocities denoted by the subscript  $i$ . This discretization of the microscopic velocity space is similar to the discrete ordinate method used for radiative transport. The LB model in this work uses an underlying orthogonal, square lattice superimposed on the non-uniform computational grid. The hydrodynamics are solved on a lattice in which communication occurs over both the Cartesian directions and the diagonal directions. For two-dimensional problems the velocity space is, therefore, discretized into the four Cartesian directions,

$$e_i = \frac{\Delta x}{\Delta t} \left( \cos \frac{\pi(i-1)}{2}, \sin \frac{\pi(i-1)}{2} \right), i = 1, 2, 3, 4 \quad (5.1)$$

and the four diagonal directions,

$$\mathbf{e}_i = \sqrt{2} \frac{\Delta x}{\Delta t} \left( \cos \frac{\pi(2i-9)}{4}, \sin \frac{\pi(2i-9)}{4} \right), i = 5, 6, 7, 8 \quad (5.2)$$

where  $\Delta t$  is the time step, and  $\Delta x$  is the grid spacing of the underlying LB grid. Also included is the probability that particles are at rest ( $\mathbf{e}_0 = \mathbf{0}$ ). This is called a nine-speed model since the microscopic velocity field is discretized into nine different velocity states. For the heat transfer, a four-speed model is sufficient for recovering the linear advection-diffusion equation with  $g_i$  being defined for the four Cartesian given in Eq. 5.1 (*Noble, Georgiadis and Buckius, 1997*). The primary variables are calculated from moments of the particle distribution, in direct analogy with the integral moments from kinetic theory. For the incompressible LB model (*Zou et al., 1995*) the hydrodynamics quantities, velocity and pressure, are found according to

$$p = \sum_{i=0}^8 f_i, \quad p_0 \mathbf{u} = \sum_{i=0}^8 f_i \mathbf{e}_i \quad (5.3)$$

where  $p_0$  is a reference pressure. The temperature is found from the first moment of  $g_i$  given by

$$T = \sum_{i=1}^4 g_i \quad (5.4)$$

The particle distributions  $f_i$  and  $g_i$  are governed by the discrete velocity Boltzmann equation expressed as

$$\frac{\partial f_i}{\partial t} + \mathbf{e}_i \cdot \nabla f_i = \Omega_i [f_i(\mathbf{x}, t)] + F_i \quad (5.5)$$

and

$$\frac{\partial g_i}{\partial t} + \mathbf{e}_i \cdot \nabla g_i = \Omega_i [g_i(\mathbf{x}, t)] \quad (5.6)$$

where  $\Omega_i$  is a collision operator that accounts for the increase or decrease of particles moving with velocity  $\mathbf{e}_i$  due to interparticle collisions. The quantity  $F_i$  is a forcing term used to implement a body force. This quantity is related to the imposed body force,  $\mathbf{F}$ , according to

$$\mathbf{F} = \sum_{i=0}^8 F_i \mathbf{e}_i \quad (5.7)$$

In this work, the body force term is given by

$$F_i = \frac{1}{6c^2} \mathbf{F} \cdot \mathbf{e}_i \quad (5.8)$$

where  $c = \Delta x / \Delta t$  is the lattice Boltzmann computational speed of sound.

Using the linearized, single time relaxation model (*Bhatnagar, Gross, and Krook*, 1954) from kinetic theory applied to lattice Boltzmann (*Qian*, 1990; *Chen et al.*, 1991), the collision term is written as

$$\Omega_i [f] = -\frac{1}{\tau_f} (f_i - f_i^{(0)}) \quad (5.9)$$

where  $f_i^{(0)}$  is an equilibrium distribution, and  $\tau_f$  is the relaxation time, which characterizes the rate of decay toward equilibrium. The collision term in the scalar transport equation is handled identically. Analogous to the Maxwellian distribution, the equilibrium distribution is the distribution to which the system will evolve in the absence of forcing gradients.

Using a first-order Lagrangian discretization of the discrete velocity Boltzmann equation, the LB equations for hydrodynamics and scalar transport are expressed as

$$f_i(x + e_i \Delta t, t + \Delta t) = f_i(x, t) - \frac{\Delta t}{\tau_f} [f_i(x, t) - f_i^{(0)}(x, t)] + F_i \Delta t \quad (5.10)$$

$$g_i(x + e_i \Delta t, t + \Delta t) = g_i(x, t) - \frac{\Delta t}{\tau_g} [g_i(x, t) - g_i^{(0)}(x, t)] \quad (5.11)$$

Examination of these equations reveals that the solution of the transport problem is reduced to two major steps. First, in a collision and forcing step, the right hand sides of Eqs. (5.10-5.11) are computed, modifying the distributions at location  $x$ . Then, a streaming step occurs in which the particle distributions stream to their nearest neighbors. In the case of a uniform spatial grid, this produces the particle distribution at each location for the new time step. It is noted that the collision and forcing step is entirely local since all quantities are evaluated at location  $x$ . The non-local communication occurs during the streaming step and involves only the nearest neighbors.

For the case of the non-uniform grid, the information is shifted according to the displacement vector  $e_i \Delta t$ , but this does not generally correspond to the location of the nearest neighbor. The values of the particle distribution at the neighbors can be computed however by interpolating between these post-streaming values. *He, Luo and Dembo* (1996) first proposed this method of implementing lattice Boltzmann on a non-uniform grid. The method has successfully been applied to flow about a circular cylinder using a grid that conformed to the cylindrical geometry (*He and Doolen*, 1997). The interpolation scheme proposed by *He, Luo and Dembo* (1996) is used here for the coupled hydrodynamics and scalar transport on the non-uniform grid.

The remaining step in the LB algorithm is the calculation of boundary conditions for the particle distributions. For the hydrodynamics, the finite difference-based method proposed by *Chen, Martinez, and Mei* (1996) is used here. The particle distribution,  $f_i$ , is extrapolated from the values at the interior using second order extrapolation. For the thermal particle distribution,  $g_i$ , local boundary conditions are used. For isothermal conditions, Eq. (5.4) is used to calculate an unknown component on the boundary from the known components and the temperature on the boundary. For flux conditions, the constraint developed by *Noble, Georgiadis and Buckius* (1997) is used:

$$\sum_{i=1}^4 g_i \mathbf{e}_i = T \mathbf{u} - \frac{(\Delta x)^2 (\tau_g / \Delta t)}{2 \Delta t} \nabla T \quad (5.12)$$

This allows an unknown component of the distribution to be calculated from the known components and the temperature gradient at the boundary.

The selection of the equilibrium distribution determines the macroscopic partial differential equation solved by the LB method. For the two-dimensional, square grid, the equilibrium distribution for the incompressible LB hydrodynamics (*Zou et al.*, 1995) is given by

$$f_i^{(0)} = w_i \left\{ p + p_0 \left[ \frac{1}{c^2} (\mathbf{e}_i \cdot \mathbf{u}) + \frac{3}{2c^4} (\mathbf{e}_i \cdot \mathbf{u})^2 - \frac{1}{2c^2} (\mathbf{u} \cdot \mathbf{u}) \right] \right\} \quad (5.13)$$

where  $w_i$  is given by

$$w_i = \begin{cases} \frac{4}{9} & i = 0 \\ \frac{1}{9} & i = 1, 2, 3, 4 \\ \frac{1}{36} & i = 5, 6, 7, 8 \end{cases} \quad (5.14)$$

For the coupled scalar transport, the equilibrium distribution is given by

$$g_i^{(0)} = \frac{T}{4} + \frac{T}{2c^2} (\mathbf{e}_i \cdot \mathbf{u}) \quad (5.15)$$

These equilibrium distributions are selected such that the incompressible Navier-Stokes equations and advection-diffusion equations are recovered to within the second-order accuracy of the method. The derivations of these distributions are more thoroughly described elsewhere (Appendix of *Hou et al.*, 1995; *Noble*, 1997), but a short summary is given here. In order to determine the equilibrium distributions, a form of the distribution must be assumed. Typically, a power series in velocity is assumed (which approximates the exponential found in the Maxwell-Boltzmann distribution). The coefficients of the series are determined by comparing the desired transport equation to the near equilibrium behavior of the LB equation. This is accomplished by performing a Taylor series expansion of the particle distribution about the location  $\mathbf{x}$  and time  $t$ . This is followed by a near equilibrium

expansion, called the Chapman-Enskog expansion, in which the particle distribution is assumed to be equal to the equilibrium distribution plus higher order terms. The expansion is in terms of a Knudsen number that compares the mean free collision time to the time scale of the flow. Substituting the assumed form of the equilibrium distribution (which is defined in terms of macroscopic quantities) into the expanded LB equation gives the macroscopic equations. Comparison to the desired transport equations yields the coefficients for the equilibrium distribution along with the definition of the transport coefficients.

The transport coefficients, kinematic viscosity and thermal diffusivity, are controlled by the relaxation times. The Taylor series and Chapman-Enskog analysis of the LB models given above yield

$$\nu = \frac{2(\tau_f/\Delta t) - 1}{6} \frac{(\Delta x)^2}{\Delta t} \quad (5.16)$$

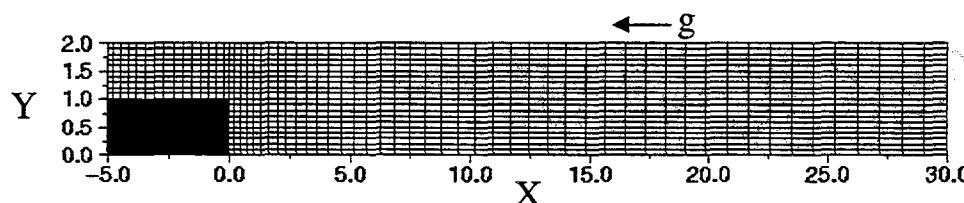
and

$$\alpha = \frac{2(\tau_g/\Delta t) - 1}{4} \frac{(\Delta x)^2}{\Delta t} \quad (5.17)$$

for the kinematic viscosity and thermal diffusivity respectively.

## 5.2 Simulation of mixed convection over a backward-facing step

To evaluate the accuracy of the non-uniform grid implementation of the mixed convection LB algorithm, the steady, two-dimensional flow of a laminar Newtonian fluid in a heated vertical channel with a backward facing step is simulated. A coarse version of the non-uniform grid used in the LB simulations is shown in Fig. 5.1. Note that the horizontal and vertical scales of the figure are not the same. This problem was the focus of a recent benchmarking study (*Blackwell and Armaly, 1993*) and is defined as follows. The walls forming the step are adiabatic. The wall downstream of the step is maintained at temperature,  $T_w$ , while the wall opposite the step is held at  $T_0$ . The total channel height is twice the step height. A parabolic velocity profile of temperature  $T_0$  is assumed at the inlet, which is located five step heights upstream of the step.



**Figure 5.1** Sample non-uniform computational grid for lattice Boltzmann simulations. For clarity, the grid shown here is 61x21 whereas the actual grid used in the simulations is 241x81. Also note that the horizontal and vertical scales are not the same. Gravity,  $g$ , acts in the negative  $X$  direction as indicated by arrow.

The dimensionless governing equations for the heat transfer and hydrodynamics are given by

$$\frac{\partial U}{\partial X} + \frac{\partial V}{\partial Y} = 0 \quad (5.18)$$

$$U \frac{\partial U}{\partial X} + V \frac{\partial U}{\partial Y} = - \frac{\partial P}{\partial X} + \frac{1}{Re} \left( \frac{\partial^2 U}{\partial X^2} + \frac{\partial^2 U}{\partial Y^2} \right) + \frac{Gr}{Re^2} \Theta \quad (5.19)$$

$$U \frac{\partial V}{\partial X} + V \frac{\partial V}{\partial Y} = - \frac{\partial P}{\partial Y} + \frac{1}{Re} \left( \frac{\partial^2 V}{\partial X^2} + \frac{\partial^2 V}{\partial Y^2} \right) \quad (5.20)$$

$$U \frac{\partial \Theta}{\partial X} + V \frac{\partial \Theta}{\partial Y} = \frac{1}{Pe} \left( \frac{\partial^2 \Theta}{\partial X^2} + \frac{\partial^2 \Theta}{\partial Y^2} \right) \quad (5.21)$$

Here the velocity is made dimensionless using the average inlet velocity,  $\bar{u}$ :

$$U = u/\bar{u}, V = v/\bar{u} \quad (5.22)$$

The spatial quantities are made dimensionless with the step height,  $s$ :

$$X = x/s, Y = y/s \quad (5.23)$$

A dimensionless pressure and temperature are defined as

$$P = (p + \rho g x)/\rho \bar{u}^2, \Theta = (T - T_0)/(T_w - T_0) \quad (5.24)$$

Dimensionless parameters of interest are given as

$$Re = \bar{u}s/v, Pe = \bar{u}s/\alpha, Gr = g\beta(T_w - T_0)s^3/v^2 \quad (5.25)$$

Of interest are the local Nusselt number and skin friction coefficient given by

$$Nu = \frac{s}{k(T_w - T_0)} \left( -k \frac{\partial T}{\partial Y} \right) = - \frac{\partial \Theta}{\partial Y} \quad (5.26)$$

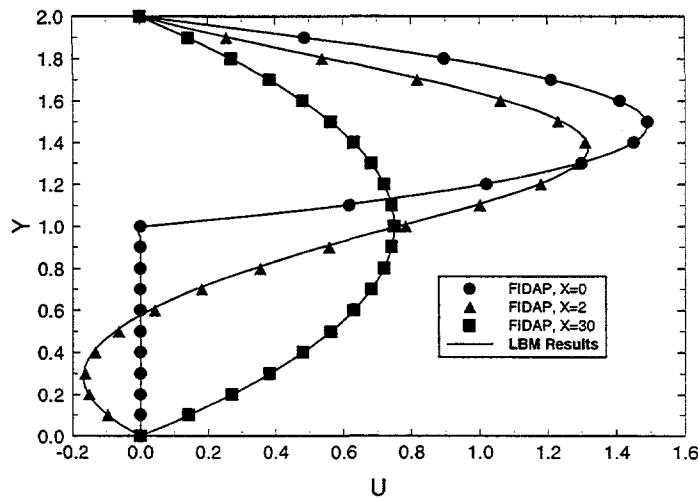
and

$$C_{f_w} Re = \frac{\tau}{(1/2)\rho \bar{u}^2} Re = \pm \frac{\partial U}{\partial Y} \quad (5.27)$$

respectively, where  $\tau$  is the wall shear stress and the  $\pm$  sign is chosen to be  $+$  along  $Y = 0$  and  $-$  along  $Y = 2$ .

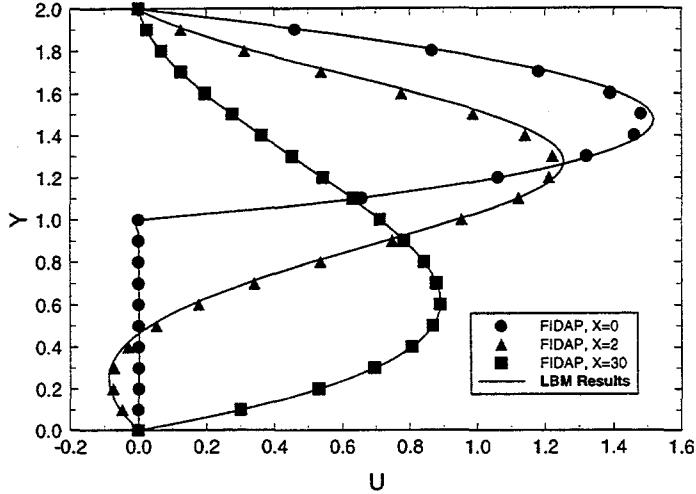
### 5.3 Accuracy of the LB model for mixed convection on a non-uniform grid

LB simulation results are compared with predictions from the commercial code FIDAP v6.0 (Fluid Dynamics International) obtained by *Torczynski* (1993) as part of a benchmarking study (*Blackwell and Armaly*, 1993). The results by *Torczynski* were tested for grid independence and are in good agreement with predictions by other participants in the benchmarking study. A grid refinement study showed the results to be accurate to three significant digits. The tabulated data included by *Torczynski* (1993) therefore provide a good method for evaluating the accuracy of the LB simulations performed here.



**Figure 5.2** Profiles of  $U$  vs.  $Y$  at indicated  $X$  positions for  $Gr=0$ . FIDAP simulation data taken from published tabulated data (*Torczynski*, 1993).

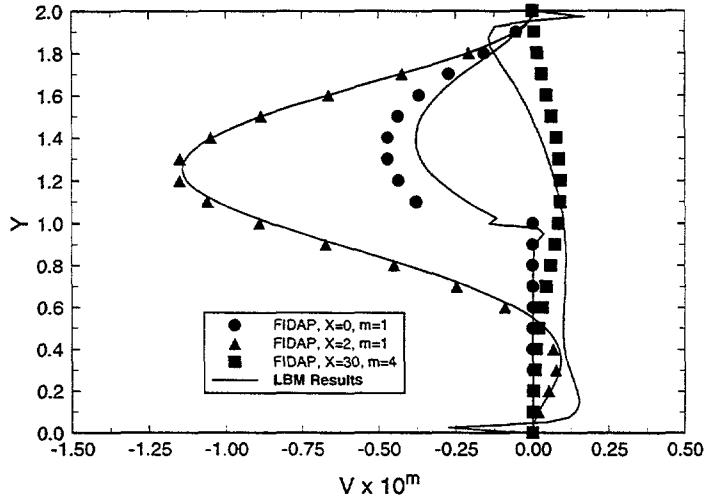
Pointwise velocity profile comparisons are made for  $X = 0$ ,  $X = 2$ , and  $X = 30$ . In addition, the local Nusselt number and skin friction coefficient along the boundaries of the channel are compared. Figures 5.2-5.3 show profiles of  $U$  at the three streamwise locations for  $Gr = 0$  and  $Gr = 1000$ , respectively. The Grashof number is controlled by varying the coefficient of thermal expansion,  $\beta$ , while keeping the temperature difference across the channel and the viscosity fixed. For  $Gr = 0$  the agreement is very good with the LB curves passing through the FIDAP data points. At  $Gr = 1000$  the agreement is still good, but some disparity is apparent in the maximum  $U$  velocity predicted at  $X = 0$  and  $X = 2$ . This is due to the fact that the boundary condition used at the inlet does not succeed in setting the prescribed condition.



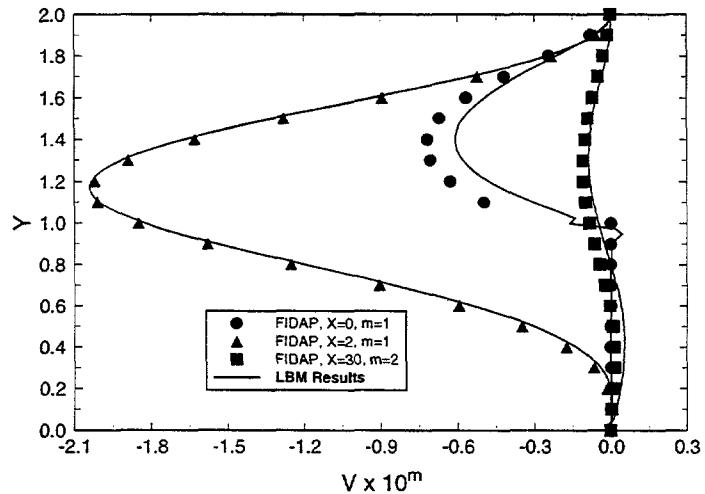
**Figure 5.3** Profiles of  $U$  vs.  $Y$  at indicated  $X$  positions for  $Gr=1000$ . FIDAP simulation data taken from published tabulated data (Torczynski, 1993).

In the finite-difference method for hydrodynamic boundary conditions, the particle distribution at the boundary is obtained by extrapolating from the interior of the computational domain. This extrapolation may cause wiggles in the solution, however, that may remain in the final solution or lead to the simulation becoming unstable. The desired macroscopic conditions (velocity or pressure) are imposed by using them to calculate the equilibrium distribution. It should be noted that during the collision step, the system decays toward the equilibrium distribution, but in the presence of forcing gradients, the equilibrium condition is never fully attained. Consequently, the desired conditions are not prescribed exactly. There may be a small steady-state difference between the desired boundary conditions and the resulting velocity or pressure at the boundary. For the case of  $Gr = 1000$ , the effect is apparent in the approximately 1.5% error in the maximum velocity. Although the effect is small, it pervades far into the domain, with an approximately 1% error in the maximum  $U$  velocity at  $X = 2$ . This shows the importance of the development of accurate hydrodynamic boundary conditions for LB models.

The plots of the  $V$  velocity component given in Figs. 5.4-5.5 also show that LB accurately models the fluid dynamics in the bulk of the fluid but that small wiggles and errors are present near the corners of the computational domain. Although the agreement is good for both  $Gr = 0$  and  $Gr = 1000$  at  $X = 2$ , significant disagreement is found at  $X = 0$  and  $X = 30$ . It should be noted, however, that the magnitudes of the errors are less than 1.5% of the average velocity  $\bar{u}$ . The relative error appears large because of the small magnitude of the  $V$  velocity component. The errors at  $X = 0$  and  $X = 30$  are also significant because both of these locations involve corners of the computational domain. The corner of the step (at  $X = 0$ ,  $Y = 1$ ) produces wiggles in the solution where the velocity gradient is discontinuous. Corners provide particular challenges in LB simulations because a large number of the components of the particle distribution are unknown and must be extrapolated from the fluid interior. At  $X = 30$  the exit boundary condition and bounding walls also meet



**Figure 5.4** Profiles of  $V$  vs.  $Y$  at indicated  $X$  positions for  $Gr=0$ . FIDAP simulation data taken from published tabulated data (Torczynski, 1993).



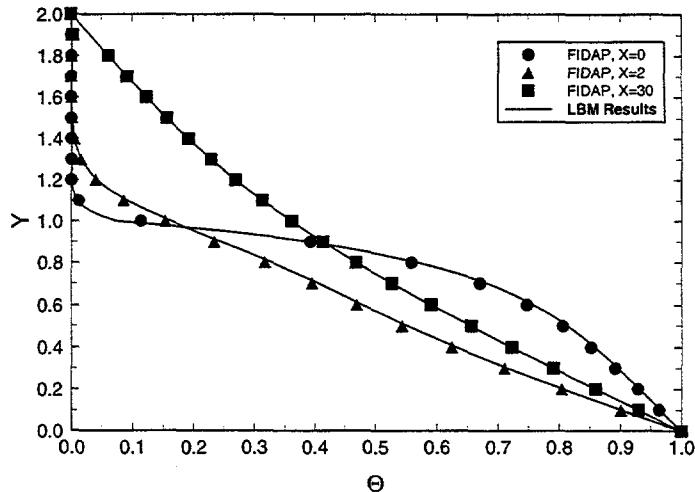
**Figure 5.5** Profiles of  $V$  vs.  $Y$  at indicated  $X$  positions for  $Gr=1000$ . FIDAP simulation data taken from published tabulated data (Torczynski, 1993).

These results show that further work is needed to develop boundary conditions for LB methods that avoid generating unphysical wiggles at corners of the computational domain. One possible improvement might be to use a method similar to that proposed by *Ginzbourg and d'Humières* (1996), in which the velocity gradient information is extracted from the local particle distribution rather than from a finite difference approximation. This may eliminate the error associated with extrapolating the particle distribution from the interior of the computational domain.

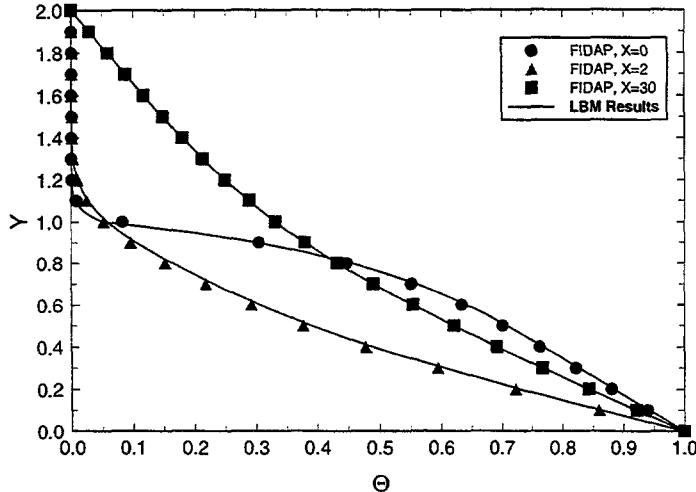
The combination of the irregular grid and finite difference-based boundary conditions increases the boundary condition error. When extrapolation is necessary, it is

desirable to have the extrapolation distance be small compared to the data spacing on the interior. However, the grid spacing near the boundary is typically decreased in order to resolve the steep gradients. The extrapolation distance is equal to the LB grid spacing,  $\Delta x$ , (based on the magnitude of the microscopic velocities), and this quantity is constant. Thus, the relative distance of the extrapolation is increased near the boundary. As shown in Figs. 5.6-5.7, the temperature field is accurately modeled using LB methods. For both  $Gr = 0$  and  $Gr = 1000$ , the LB and FIDAP predictions for the dimensionless temperature are in close agreement over the entire width of the channel at each of the streamwise positions examined.

Comparisons of the skin friction and local Nusselt number shown in Figs. 5.8-5.9 show very good agreement between the LB and FIDAP predictions by *Torczynski* (1993). From the skin friction along the heated side of the channel, the flow reattachment locations are calculated. For  $Gr = 0$  and  $Gr = 1000$ , the LB simulations predict reattachment at  $X = 5.05$  and  $X = 3.08$  respectively. These predictions differ from the FIDAP predictions of  $X = 4.984$  and  $X = 2.977$  by 1.3% and 3.5%, respectively. The larger error for  $Gr = 1000$  is likely due to the inlet boundary condition error for this case which gives a larger inlet velocity than prescribed.

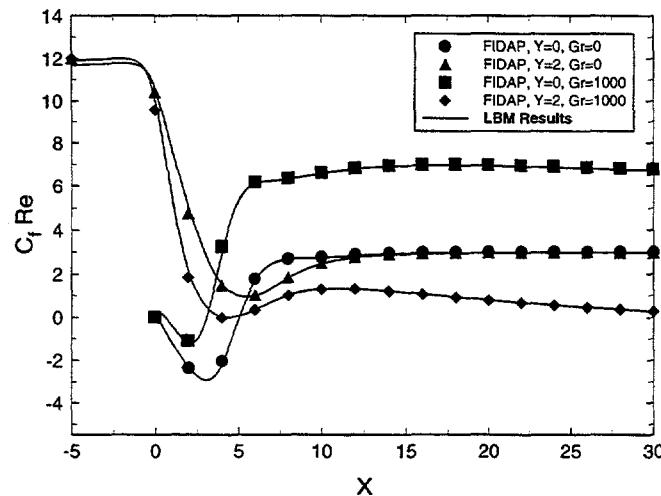


**Figure 5.6** Profiles of dimensionless temperature,  $\Theta$ , vs.  $Y$  at indicated  $X$  positions for  $Gr=0$ . FIDAP simulation data taken from published tabulated data (*Torczynski*, 1993).

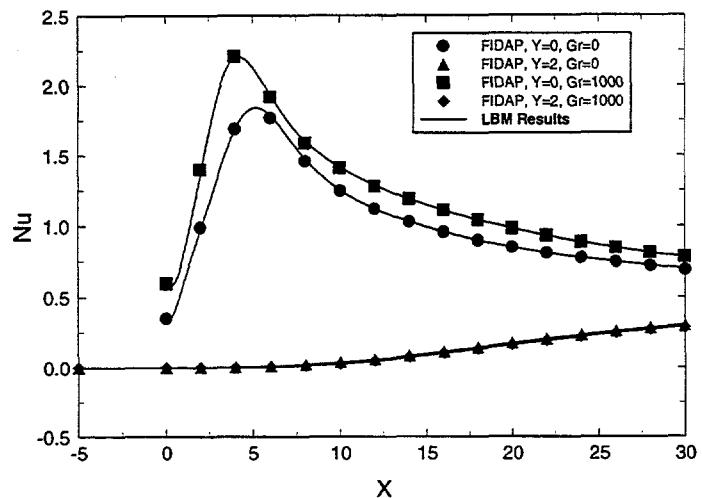


**Figure 5.7** Profiles of dimensionless temperature,  $\Theta$ , vs.  $Y$  at indicated  $X$  positions for  $Gr=1000$ . FIDAP simulation data taken from published tabulated data (*Torczynski, 1993*).

This work shows how a lattice Boltzmann (LB) model for mixed convection is developed and applied to flow in a vertical heated channel with a backward facing step. The comparisons with FIDAP predictions by *Torczynski (1993)* validate LB as a quantitative method for thermal problems including mixed convection. Velocity profile comparisons show agreement between LB and FIDAP solutions to within 1.5% of the average streamwise velocity. LB models using finite difference-based boundary conditions (*Chen, Martinez and Mei, 1996*) have been shown to exhibit some difficulties near the boundaries, however. Wiggles in the solution are formed near the corners of the computational domain, and the maximum velocity in the channel is over-predicted for the case of mixed convection due to an error in the inlet boundary condition. This demonstrates that further boundary condition work is warranted. LB methods are shown to provide alternate techniques for coupled heat transfer and viscous fluid dynamics.



**Figure 5.8** Profiles of the dimensionless skin friction vs.  $X$  along the walls of the channel for  $Gr=0$  and  $Gr=1000$ . FIDAP simulation data taken from published tabulated data (*Torczynski, 1993*).



**Figure 5.9** Profiles of the local Nusselt number vs. X along the walls of the channel for  $Gr=0$  and  $Gr=1000$ . FIDAP simulation data taken from published tabulated data (*Torczynski, 1993*).

## 6 Application to flow about moving and non-conforming obstacles

One concern when modeling boundary-dominated problems like porous media and suspension flows is the accuracy with which the geometry is represented in LB simulations. This concern is even stronger when the boundaries between the solid and fluid phases do not conform to the computational lattice. In this section, a LB model pertaining to non-conforming, moving obstacles is developed and assessed. Several methods have been proposed for applying boundary conditions including the “bounce-back” condition (*Cornubert, d’Humières and Levermore*, 1991), finite difference-based methods (*Chen, Martinez and Mei*, 1996), and techniques resulting from analysis of the discrete velocity Boltzmann equation (*Skordos*, 1993; *Noble, Chen, Georgiadis and Buckius*, 1995; *Ginzbbourg and d’Humières*, 1996). This list of references is not exhaustive; *Chen, Martinez and Mei* (1996) give a more complete review.

Below, a method is developed for simulating two-dimensional and three-dimensional fluid flow in the presence of large numbers of arbitrary-shaped particles from dilute to dense packing. In brief, the conventional lattice Boltzmann method is extended by modifying the collision term so that it shifts smoothly between hydrodynamics at nodes occupied solely by fluid and rigid body motion at nodes occupied solely by particles. The fluid and solid volume fractions are used to weight the corresponding portions of the collision term for nodes with both phases present. In a recently published paper, a similar method is applied to simulations of porous media in the dilute limit. In their work, *Dardis and McCloskey* (1998) modified the regular lattice Boltzmann scheme allowing for the presence of solid scatterers that resist fluid motion. In contrast, the focus here is on the partially saturated computational cells that appear at the interface between clear fluid and solid obstacles. The goal of this work is to model flows about obstacles that do not conform to the computational grid and to allow the objects to move relative to the grid. For validation purposes, simulations are performed of the flow produced by a periodic line of cylinders moving in a straight channel. For both steady and periodic motion of the cylinders, detailed comparisons are made with finite element simulations in order to quantify the accuracy of the method. The capability developed here will be applied in future work to coupled fluid-flow/particle-motion simulations.

### 6.1 Lattice Boltzmann method for partially saturated cells

The LB model described here for partially saturated cells uses a two-dimensional, square lattice. As described in previous sections, the velocity space is discretized into the four Cartesian directions,

$$\mathbf{e}_i = \frac{\Delta x}{\Delta t} \left( \cos \frac{\pi(i-1)}{2}, \sin \frac{\pi(i-1)}{2} \right), i = 1, 2, 3, 4 \quad (6.1)$$

and the four diagonal directions,

$$\mathbf{e}_i = \sqrt{2} \frac{\Delta x}{\Delta t} \left( \cos \frac{\pi(2i-9)}{4}, \sin \frac{\pi(2i-9)}{4} \right) \quad i = 5, 6, 7, 8 \quad (6.2)$$

where  $\Delta t$  is the time step, and  $\Delta x$  is the grid spacing of the LB grid. Also included is the probability that particles are at rest ( $\mathbf{e}_0 = \mathbf{0}$ ). The primary variables are calculated from moments of the particle distribution, in direct analogy with the integral moments from kinetic theory. The hydrodynamic quantities, density  $\rho$  and velocity  $\mathbf{u}$ , are found according to

$$\rho = \sum_{i=0}^8 f_i, \quad \rho \mathbf{u} = \sum_{i=0}^8 f_i \mathbf{e}_i \quad (6.3)$$

The particle distribution  $f_i$  is governed by the lattice Boltzmann equation expressed as

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta t, t + \Delta t) - f_i(\mathbf{x}, t) = \Omega_i [f_i(\mathbf{x}, t)] \quad (6.4)$$

where  $\Omega_i$  is a collision operator that accounts for the increase or decrease of particles moving with velocity  $\mathbf{e}_i$  due to interparticle collisions.

Using the linearized, single time relaxation model (*Bhatnagar, Gross, and Krook*, 1954) from kinetic theory applied to lattice Boltzmann (*Qian*, 1990; *Chen et al.*, 1991), the collision term for regions completely saturated with fluid is written as

$$\Omega_i^{\text{BGK}} [f] = -\frac{\Delta t}{\tau} (f_i - f_i^{(0)}) \quad (6.5)$$

where  $f_i^{(0)}$  is an equilibrium distribution, and  $\tau$  is the relaxation time, which characterizes the rate of decay toward equilibrium. Analogous to the Maxwellian distribution, the equilibrium distribution is the distribution to which the system will evolve in the absence of forcing gradients. Substituting the BGK collision term into the lattice Boltzmann equation gives the standard lattice Boltzmann evolution equation:

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) - \frac{\Delta t}{\tau} [f_i(\mathbf{x}, t) - f_i^{(0)}(\mathbf{x}, t)] \quad (6.6)$$

The goal of this study is to introduce a modification to this equation that allows for large numbers of nonconforming, evolving boundaries but retains the advantages of the conventional lattice Boltzmann method (i.e. a collision step that is local to nodes and a molecular-motion step that is identical to the above). This is accomplished by introducing additional collision terms,  $\Omega_i^s$ , that account for the interactions with each solid obstacle that lies within the computational cell. The evolution equation then becomes,

$$f_i(x + e_i \Delta t, t + \Delta t) = f_i(x, t) + \left[ 1 - \sum_s B(\varepsilon_s, \tau) \right] \Omega_i^{\text{BGK}} + \sum_s B(\varepsilon_s, \tau) \Omega_i^s \quad (6.7)$$

where  $\varepsilon_s$  is the volume fraction of each solid particle that intersects the volume surrounding the node, and  $B(\varepsilon_s, \tau)$  is a weighting function. At its limits,  $B(\varepsilon_s = 0, \tau) = 0$  and  $B(\varepsilon_s = 1, \tau) = 1$ , which correspond to pure fluid and pure solid, respectively. For generality the nature of the transition is allowed to depend on  $\tau$ , however. Two possible forms for  $\Omega_i^s$  are investigated in this study. The first is based on the concept of “bounce-back” of the non-equilibrium part (Zou and He, 1997) of the particle distribution and is given by,

$$\Omega_i^s = f_{-i}(x, t) - f_i(x, t) + f_i^{\text{eq}}(\rho, \mathbf{U}_s) - f_{-i}^{\text{eq}}(\rho, \mathbf{u}) \quad (6.8)$$

where  $\mathbf{U}_s$  is the velocity of the solid particle at time  $t + \Delta t$  at the node and the notation  $f_{-i}$  is used to denote the “bounce-back” state from  $f_i$  (obtained by reversing all molecular velocities, i.e.,  $\mathbf{e}_i \rightarrow -\mathbf{e}_i$ ). In cases where the particle velocity is unknown, it must be calculated by a force analysis. For this investigation, however, the particle velocity is prescribed. In the limit of pure solid, this collision term causes the non-equilibrium part of the distribution to be “bounced-back”. A second type of solid boundary collision term is investigated which has the form,

$$\Omega_i^s = f_i^{\text{eq}}(\rho, \mathbf{U}_s) - f_i(x, t) + \left( 1 - \frac{\Delta t}{\tau} \right) [f_i(x, t) - f_i^{\text{eq}}(\rho, \mathbf{u})] \quad (6.9)$$

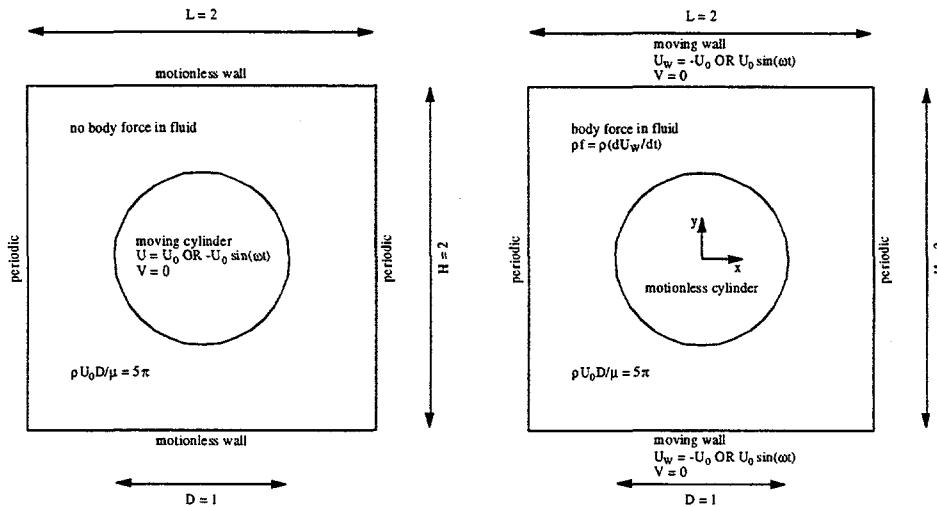
This form is designed such that the resulting particle distribution is equal to the equilibrium distribution evaluated at the obstacle velocity plus a term that depends on the deviation of the current distribution from its equilibrium value. This superposition method might be expected to increase in accuracy as the distribution approaches equilibrium. Another feature of this collision term is that it causes the distribution to match the equilibrium distribution when the relaxation time equals the time step in the same way that this occurs in the pure fluid case. There are many other collision terms that might be successful, but both of these proved to give good representations of moving circular cylinders. Possible implementations of the boundary enhanced evolution equation are distinguished by the expression used for  $B(\varepsilon_s, \tau)$ . While the near equilibrium, superposition approach given in Eq. (6.9) is found to perform well with the simple weighting function,  $B(\varepsilon_s, \tau) = \varepsilon_s$ , the non-equilibrium reflection model required a more involved function. Knudsen layer effects appear when applying reflection type boundary conditions, but it was found that very good approximation of the cylinder surface was obtained for  $\tau/\Delta t$  between 0.6 and 0.9 (it was not tested outside this range) using the expression,

$$B(\varepsilon_s, \tau) = \frac{\varepsilon_s (\tau/\Delta t - 1/2)}{(1 - \varepsilon_s) + (\tau/\Delta t - 1/2)} \quad (6.10)$$

Based on these concepts four different boundary conditions were tested:

- i. Non-equilibrium reflection with relaxation time dependent weighting [Eqs. (6.8, 6.10)]
- ii. Non-equilibrium reflection with simple solid fraction weighting [Eq. (6.8) and  $B(\varepsilon_s, \tau) = \varepsilon_s$ ]
- iii. Near equilibrium superposition with simple solid fraction weighting [Eq. (6.9) and  $B(\varepsilon_s, \tau) = \varepsilon_s$ ]
- iv. Simple, binary non-equilibrium reflection (“bounce-back” of the non-equilibrium part of the distribution) [Eq. (6.8) and  $B(\varepsilon_s, \tau) = \text{ROUND}(\varepsilon_s)$ ]

The last case is used to compare the methods proposed here to an approach in which computational cells that are more than half full of solid simply replace the regular BGK collision term with that given in Eq. (6.8). This helps to determine the importance of a smooth transition from fluid to solid rather than an abrupt ‘stair-step’ representation.

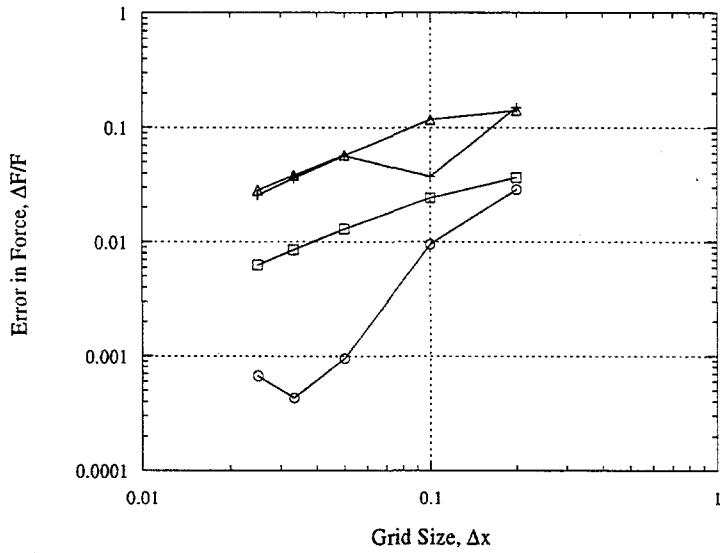


**Figure 6.1** Schematic diagram (to scale) of the validation problem in the (a) moving-cylinder reference frame and (b) fixed-cylinder reference frame.

## 6.2 Steady flow about a periodic line of cylinders

The accuracy of the lattice Boltzmann method presented in the previous section is assessed by comparing to fully resolved finite-element simulations obtained with FIDAP, a commercially available computational fluid dynamics code produced by Fluid Dynamics International. Figure 6.1 shows a schematic diagram (to scale) of the validation problem. A two-dimensional situation was considered in which a linear periodic array of circular cylinders moves with a constant or oscillating velocity in the horizontal direction without rotation. In the lattice Boltzmann simulations, one full unit cell of the periodic array is simulated (as in Figure 6.1), whereas only the upper half of one unit cell is simulated in the FIDAP simulation since the centerline ( $y = 0$ ) is a line of symmetry. In the FIDAP

simulation, the significant deformation of the domain would require extensive remeshing during each oscillation if this problem were treated in the same fashion as in the lattice Boltzmann simulation. Instead, a coordinate transformation to the fixed-cylinder frame of reference was employed. In this reference frame, the cylinder is motionless, and the walls travel with a velocity opposite to the original cylinder. All results herein are presented in the fixed-cylinder reference frame. To ensure that the FIDAP solutions were grid-independent, simulations of the oscillatory flow described below in Sect. 6.4 were performed with two different meshes and two different time steps. The elements of the second mesh had about half the side length of the elements of the first mesh (about four times as many elements in the second mesh as in the first mesh). Simulations were performed with each mesh using a fixed time step of 0.001, and an additional simulation was performed with the coarse mesh using a time step of 0.0005. The results of all simulations were found to be virtually indistinguishable at all times and all locations. The same fine grid was then used for the steady state simulations.

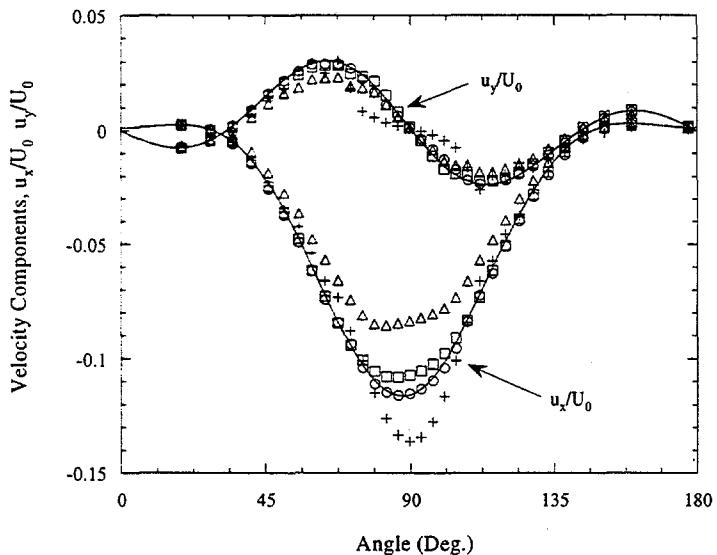


**Figure 6.2** Convergence with grid refinement for lattice Boltzmann simulations of steady flow using various boundary conditions. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution.

The LB and FIDAP simulation results for the drag force as well as the velocity components near the surface of a cylinder are compared. Figure 6.2 shows the relative error in the drag force predictions as the computational grid is refined for each of the four LB boundary conditions. For LB, the drag force per unit length on a solid  $s$  is given by,

$$\frac{\mathbf{F}_s}{L} = \frac{1}{\Delta t} \sum_i \Omega_i^s \mathbf{e}_i \Delta x \Delta y \quad (6.11)$$

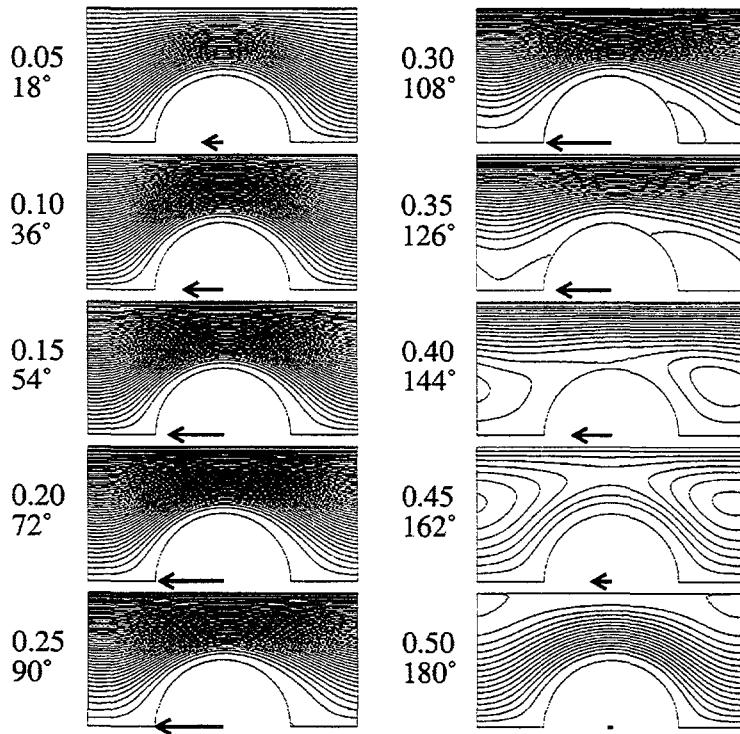
where  $\Delta x = \Delta y$  is the grid spacing. From the slope of the curves on this log-log plot, it is seen that linear convergence is obtained except for boundary condition (i) which uses the relaxation time dependent weighting function. This boundary condition results in an order-of-magnitude smaller error than the other methods and appears to produce quadratic convergence. Further investigations would be useful in confirming this. The apparent saturation of the error at 0.0005 is likely due to the fact that the FIDAP solution has an error of this same magnitude. Thus the LB solution continues to converge toward the true solution as the grid is refined but begins to differ from the FIDAP prediction.



**Figure 6.3** Components of velocity near the surface of the cylinder. The cylinder is moving with a constant velocity in the positive x direction, and the velocities are expressed relative to the velocity of the cylinder. Velocities are plotted as a function of angle for a fixed radial distance of  $D/32$  from the surface of the cylinder. The solid lines denote the FIDAP predictions. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution.

Figure 6.3 shows the horizontal and vertical velocity components very near the surface of a cylinder as the line of cylinders is steadily moved. The solution at a radial distance of  $D/32$  from the surface of a cylinder is examined and compared to the FIDAP predictions. The FIDAP grid has several nodal points along this radius, but the cylinder is moving relative to the lattice Boltzmann mesh. The lattice Boltzmann solutions at this radius are found by interpolation, evaluating cubic splines of the solutions above and below a cylinder along lines of constant  $x$  at the desired radius. This figure shows that boundary

condition (i) using the non-equilibrium reflection with viscosity-dependent weighting is particularly successful in modeling the velocity field near a moving, smooth cylinder. This explains the accuracy of the drag force calculated using this method as well. The simple solid fraction weighting (ii) instead of the viscosity-dependent one, however, cripples the method, causing it to smooth out the gradients near the surface of a cylinder. The near equilibrium superposition method (iii) yields reasonable accuracy even with the simple solid fraction weighting. The simple “bounce-back” model (iv), which allows cells to be either pure fluid or pure solid, shows sharp jumps in the solution due to the ‘stair-step’ representation of the surface.



**Figure 6.4** FIDAP streamlines during the fifth cycle for oscillatory flow at indicated normalized times and phases, with cylinder velocities in the fixed frame proportional to the arrows.

### 6.3 Oscillatory flow about a periodic line of cylinders

To test the behavior and accuracy of the proposed boundary conditions for unsteady flows, the case of oscillatory motion of the cylinders is examined. The horizontal velocity component of the cylinder was taken to be

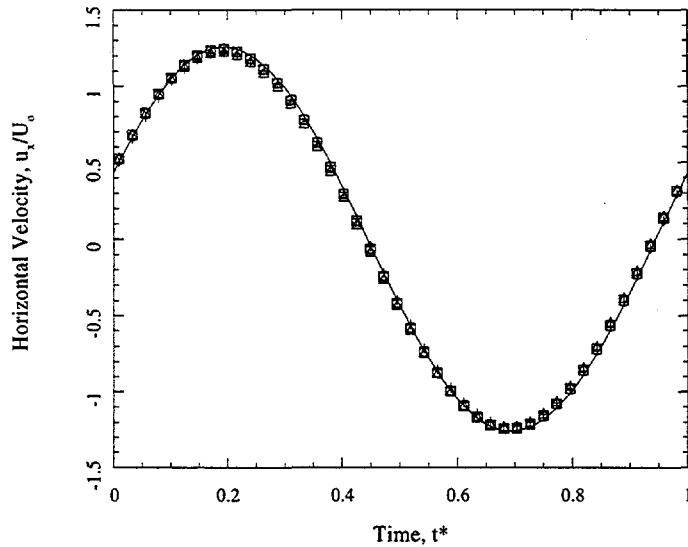
$$U = -U_0 \sin(\omega t) \quad (6.12)$$

The Reynolds and Strouhal numbers were given by,

$$Re = \rho U_0 D / \mu = 5\pi \approx 15.7 \text{ and } Sr = \omega D / 2\pi U_0 = 1/\pi \approx 0.32 \quad (6.13)$$

suggesting that eddies will form.

In the lattice Boltzmann simulation, the cylinder center undergoes a maximum peak-to-peak displacement of one diameter. To prevent the cylinder from crossing the domain boundaries, the cylinder and domain right edges are tangent when  $\omega t = 0, 2\pi, 4\pi, \dots$ , and the cylinder and domain left edges are tangent when  $\omega t = \pi, 3\pi, 5\pi, \dots$ . Thus, the cylinder crosses a large fraction of the nodes during each oscillation: nodes are “covered” and “uncovered” repeatedly during the motion, which represents a stringent test of the accuracy of the treatment of moving boundaries.



**Figure 6.5** History of normalized horizontal velocity component  $u_x/U_0$  during the fifth cycle of oscillation for the point located at  $r = 0.75D$  and  $\theta = 90^\circ$ . The solid line denotes the FIDAP prediction. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution.

As in the steady case a coordinate transformation to the fixed-cylinder frame of reference is employed and the walls travel with a velocity opposite to the original cylinder.

$$U_w = U_0 \sin(\omega t) \quad (6.14)$$

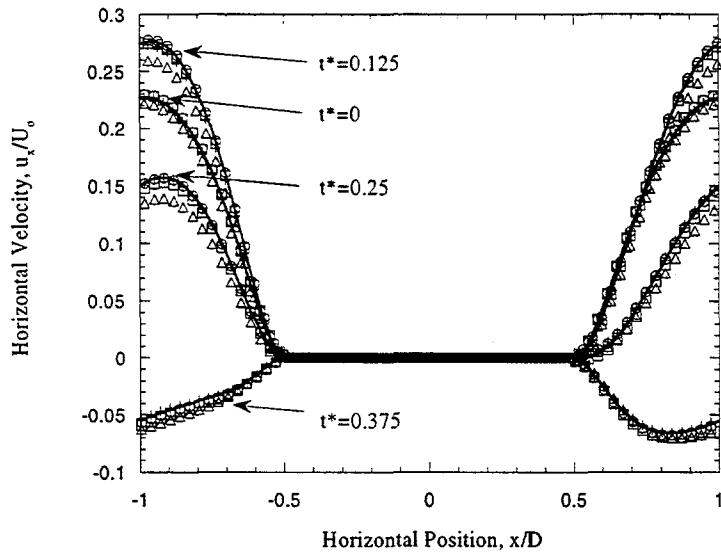
Since the new frame of reference is accelerated with respect to the original frame, the fluid in the new reference frame experiences a body force in the horizontal direction:

$$\rho f = \rho \frac{dU_w}{dt} \quad (6.15)$$

Since the simulations were initialized with a zero-velocity condition throughout the domain, five cycles of the oscillation were simulated to ensure that the time-periodic solution was obtained. This was verified by quantitative comparison of the fourth and fifth cycles, which were found to differ by well less than 0.1%.

Figure 6.4 shows the streamlines during the first half of the fifth cycle based on the FIDAP simulations at various normalized times  $t^* = (t - 4T)/T$  and phases

$\phi = 2\pi(t - 4T)/T$ , where  $T = 2\pi/\omega$  is the period of the oscillation. As the particle begins to move leftward at the beginning of the cycle, the streamlines exhibit approximate fore-and-aft symmetry. As the particle velocity increases, an eddy appears to the right of the cylinder at a normalized time of 0.25. This eddy grows in extent and strength and fills the domain at 0.40, detaches from the cylinder at 0.45, and has almost vanished by 0.50, when the particle has come to rest and is ready to begin moving rightward.



**Figure 6.6** Spatial variation of the normalized horizontal velocity component  $u_x/U_0$  along the centerline ( $y=0$ ) at indicated times during the fifth cycle. The solid lines denote the FIDAP predictions. The circles denote method (i): non-equilibrium reflection with viscosity-dependent weighting. The squares denote method (ii): non-equilibrium reflection with simple solid fraction weighting. The triangles denote method (iii): near-equilibrium superposition with simple solid fraction weighting. The pluses denote method (iv): “bounce-back” of the non-equilibrium part of the distribution.

Figures 6.5 and 6.6 show detailed comparisons of the lattice Boltzmann and FIDAP solutions. Figure 6.5 shows histories of the horizontal velocity component  $u_x$  normalized by the velocity amplitude  $U_0$  at the location  $r = 0.75D$  and  $\theta = 90^\circ$ . This point is midway between the cylinder boundary and the domain boundary. This point is a nodal location of the FIDAP mesh but is continuously moving relative to the lattice Boltzmann mesh. The

lattice Boltzmann solution at this location is found by interpolation, evaluating cubic splines of the solutions along lines of constant  $y$  at the desired  $x$  values. The agreement between the FIDAP results and the lattice Boltzmann results is excellent for all of the employed boundary conditions for this resolution of  $60 \times 60$ . It is seen that at this location, which lies outside the wake of the oscillating cylinders, the boundary condition specifics have little effect.

Figure 6.6 shows the spatial variation of horizontal velocity  $u_x/U_o$  along the centerline ( $y = 0$ ) at the indicated times throughout the cycle for LB simulations using a  $60 \times 60$  computational grid. The effect of the viscosity-dependent weighting is apparent along the centerline due to the larger effect of the boundaries in this region. However, all of the other boundary conditions are reasonably successful. This shows that any of the proposed methods can be used providing reasonable accuracy even in this wake region. Only in the vicinity very near the obstacle, shown in Fig. (6.3), and in the force calculation, shown in Fig. (6.2), can it be seen that boundary condition (i) provides superior modeling of the moving, non-conforming cylinder.

#### 6.4 Discussion

In this section, a lattice Boltzmann method is developed and implemented for simulating coupled fluid flow and particle motion. The method is successfully applied to the flow about a periodic line of cylinders that move relative to the computational grid. To model moving obstacles that do not conform to the grid, the collision term in the conventional lattice Boltzmann method is modified so that it shifts smoothly between hydrodynamics at nodes occupied solely by fluid and rigid body motion at nodes occupied solely by particles. The modified collision operator uses the fluid and solid volume fractions to weight the corresponding portions of the collision term for nodes with both phases present. Various possible forms for the modified collision operator are examined and possible weighting functions are explored. It is found that a collision operator that employs reflection of the non-equilibrium part of the distribution is particularly successful in modeling smooth surfaces when using the suggested weighting function. The method appears to exhibit quadratic convergence with grid refinement suggesting that the overall second-order convergence of the lattice Boltzmann method is maintained. This method promises to be an economical and accurate method for modeling moving obstacles that do not conform to the computational grid.

There are at least three types of problems in which a formulation in terms of a smooth transition from pure liquid to pure solid is useful. One is the type of application examined here where an obstacle is moving relative to the grid or does not conform to the grid. The ability to accommodate partially saturated cells allows for curved and other complex surfaces to be modeled with more integrity than a 'stair-step' representation. A second application is for porous media simulations (*Dardis and McCloskey, 1998*). The partial saturation concept can be used in this context to model non-uniform porous media with prescribed porosity variation.

A third related, but distinct, application is modeling flow in reconstructed porous media (Sections 7 and 8). The direct investigation of pore-scale flow in porous media typically starts with reconstruction of the geometry, followed by pore-scale simulation of the flow. One potentially ambiguous step in this process involves segmenting the geometry into areas of void and solid. Often, the experimental reconstruction method captures intrinsically the smooth transition from the solid to void phase, i.e. some voxels encompass both the void and solid phase. But, in order to use a conventional lattice Boltzmann algorithm, researchers design and apply algorithms to designate the transition zone as either solid or void so that a binary description of the geometry is arrived at. The formulation developed in this work provides an alternative to this approach since the experimentally measured smooth transition can be directly employed in the fluid simulation. This allows higher fidelity modeling since some ambiguity in the binary reconstruction is inevitable due to the intrinsic resolution of the data. The method for partially saturated computational cells promises an accurate and economic method for modeling each of these problems, which are difficult to handle using a conventional lattice Boltzmann scheme.

## 7 Experimental characterization of porous media

Natural as well as many engineering materials exhibit complex pore structures which greatly influence their bulk physical and mechanical properties, including elastic moduli, compressibility, seismic velocity, permeability, electrical conductivity, thermal conductivity, poroelastic parameters, strength, and failure behavior. The extent to which the macroscopic properties of an arbitrarily complex porous medium can be predicted using a microscale modeling approach follows directly from the sophistication of the model, which in turn relies on the completeness of the microgeometric description available as input.

The transport properties in particular depend critically on the geometry of the void space (e.g., *Adler*, 1992; *Dullien*, 1992; *Sahimi*, 1995). The size, shape, connectivity, and tortuosity of the pore space can in fact affect the permeability of a porous medium more strongly than the total void fraction available to transmit that flow. As an example, consider a typical quartz sandstone such as Berea with a bulk porosity of ~20-25% and a permeability of ~200-800 millidarcy. Contrast that with a sedimentary diatomaceous rock with more than twice the porosity (~60%), but a permeability that is lower by well over two orders of magnitude (~0.1-1 millidarcy). This dramatic inversion of the usual trend between porosity and permeability is a direct consequence of differences in the geometry of the pore space for the two rocks.

A variety of geoscience and engineering applications stand to benefit significantly from an improved understanding of flow and transport processes in geometrically complex porous media. Geoscience applications include understanding the displacement of oil from the pore space during waterflooding or enhanced oil recovery processes, and predicting the transport of contaminants in an underground aquifer. Understanding the physics of microscale flow processes is likewise central to many chemical, materials, and thermal engineering applications, including porous bed chemical reactors, powder compaction and sintering, gas flow in thermal protection systems, heat dissipation in microelectronic components, and fluid and thermal flows in MicroElectroMechanical Systems (MEMS).

Fundamental understanding requires both statistical characterization of the geometry of the porous microstructure, as well as study of the geometry of the flow itself. The former aspect can in principle be addressed experimentally; however, the geometry of the flow through a complex porous medium is difficult to ascertain in the laboratory, but amenable to numerical investigation. For example, using idealized two-dimensional network models, *David* (1993) demonstrated the complexity of hydraulic flow paths in heterogeneous porous media, and that the critical paths for hydraulic flow are not necessarily identical to those for electric current.

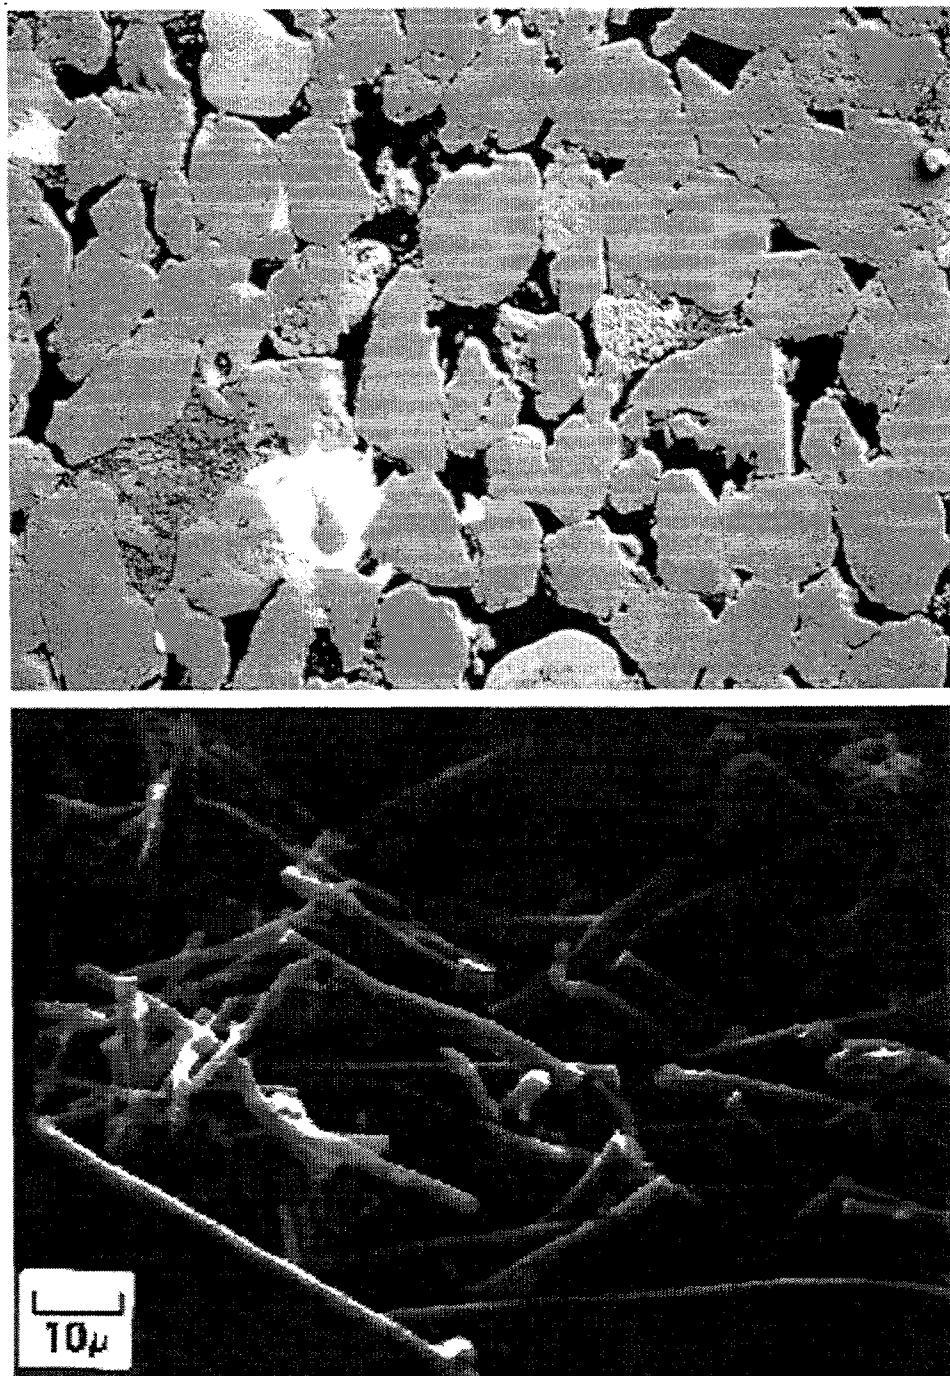
In this section we describe our work towards developing new experimental techniques and approaches for characterizing the three-dimensional microgeometry of complex porous media. Following in Section 8, we describe the application of lattice Boltzmann methods for modeling and simulating mass and heat transfer processes in complex porous media as imaged using confocal microscopy.

## 7.1 Approaches for characterizing the microstructure of porous materials

The simplest geometric attribute of a porous material is its porosity, or void fraction,  $\phi$ . This property can be measured on bulk samples to high accuracy using gas porosimetry techniques based on Boyle's Law, although less precise imbibition (saturation) techniques are also commonly applied. Specific surface area SA, defined as the grain-void interface area normalized by the total volume, is the next most elementary property characterizing a porous medium. SA as well can be readily measured on a bulk sample using gas adsorption techniques such as the BET method (Brunauer *et al.*, 1938). Limited additional geometric information can be obtained from bulk samples using mercury porosimetry (e.g., Van Brakel *et al.*, 1981), small-angle scattering (Wong *et al.*, 1986; Hall *et al.*, 1986), adsorption-desorption or other extensions of the basic gas adsorption technique (e.g., Avnir *et al.*, 1984), or nuclear magnetic resonance (Cohen and Mendelson, 1982; Mendelson, 1982). For geologic materials, mercury porosimetry is probably the most commonly applied indirect technique; information concerning "effective" pore sizes can be derived from the capillary pressure record. However, the geometric complexity of the pore space (i.e. its interconnectivity, geometric irregularity, etc.) leads to ambiguities in the physical interpretation of mercury capillary data, and the other indirect methods likewise require certain geometric assumptions that complicate their interpretation.

Imaging methods offer the possibility to characterize unambiguously, and completely, the microstructure of a porous material. Conventional techniques include reflected or transmitted light microscopy (LM) and scanning electron microscopy (SEM) of planar sections. Quantitative stereological methods (Underwood, 1970; Russ, 1986) can be applied to determine geometric parameters such as porosity and specific surface area directly from measurements made on planar sections (e.g., Wong, 1985). Similarly, stereological measurements of chord length can be used to infer pore and crack size distributions (Krohn, 1988; Wong *et al.*, 1989; Fredrich *et al.*, 1993). Alternatively, both porosity and specific surface area can be calculated from the 2-point probability function (defined later) that can be measured from micrographs of polished sections that have been processed to yield binary data (Berryman, 1985; Berryman and Blair, 1986). Methods have also been devised to characterize the fractal properties of the pore space using measurements made on images of polished sections (Thompson, 1991), as well as from images of rough fracture surfaces (Krohn and Thompson, 1986).

An important handicap of these traditional imaging approaches, however, is that one is restricted to examining a two-dimensional (2D) representation of a three-dimensional (3D) object, which greatly complicates, if not precludes, a determination of the true three-dimensional structure. As was pointed out by Bernabe (1991), certain pore types (such as tubular pores) are difficult to identify in 2D images. Likewise, it is impossible to determine the connectivity of a pore network from a 2D image. A quick examination of a 2D micrograph (Figure 7.1, Top) from a typical sandstone shows that the pore phase is discontinuous in the two-dimensional plane of the image; however, we know with certainty by the presence of the epoxy that the pore phase is in actuality connected in three dimensions.



**Figure 7.1** Scanning electron micrographs of (top) Berea sandstone (backscattered imaging mode), and (bottom) a porous rigid fibrous refractory insulation. The sandstone has a grain size of  $\sim 150$  micron, and the pore space has been impregnated with epoxy (black). There is no continuous path through the pore phase in the imaging plane. The insulation is composed of ceramic fibers that occupy less than  $\sim 20\%$  of the bulk volume, and that have a diameter of several microns. The image reveals the complex geometry of both the solid and pore phases, but it is unfeasible to extract a quantitative description of the three-dimensional microstructure. (Bottom image courtesy of NASA Ames Research Center, Thermal Protection Materials and Systems Branch).

Several innovative techniques have been developed by researchers to circumvent this limitation. Pore casts have been formed by impregnating a rock with an acid-resistant epoxy and then dissolving away the mineral matrix with acid (Pittman and Duschatko, 1970). The pore casts can then be imaged using SEM, and while they provide useful insight (e.g., Bourbie and Zinszner, 1985; Myer *et al.*, 1992; Zinszner *et al.*, 1997), the stereoscopic images are difficult to quantify. The same problem arises during a SEM study of the rough or irregular surface of a porous material itself (Figure 7.1, Bottom). Difficulties may also be encountered in impregnating and preserving the delicate pore structure of low porosity samples (Lin *et al.*, 1986).

*Koplik et al.* (1984) made sequential optical micrographs of a planar section after incrementally removing as little as 1  $\mu\text{m}$  of material by polishing. *Lin et al.* (1986) coupled such serial sectioning with image processing techniques to generate 3D images of porosity in a tight gas sand using conventional optical microscopy. Although effective, this technique is time consuming and tedious.

In the past ten or so years several new approaches have emerged for imaging the 3D structure of porous media. X-ray computed tomography (CT), originally developed by the medical community, can be used to measure porosity and to image porosity distribution in bulk porous samples (e.g., Wellington and Vinegar, 1987). However, quantitative measurements are complicated greatly by the need for beam hardening corrections that arise due to the polychromatic nature of X-ray sources, and that are particularly challenging for heterogeneous materials (Brown *et al.*, 1993; Fredrich *et al.*, 1994). Moreover, the typical resolution of 0.1-1 mm is not sufficient for detailed characterization of most porous media of interest in the geo- and materials- sciences. *Brown et al.* (1993) described a gamma ray CT system that relied on a monochromatic source, thereby obviating the need for complicated beam hardening relationships. However, the system resolution of 1 mm again indicates that while conventional X-ray and gamma ray CT are useful techniques for studying transport in fractured rock, they are not generally effective for characterizing porous media.

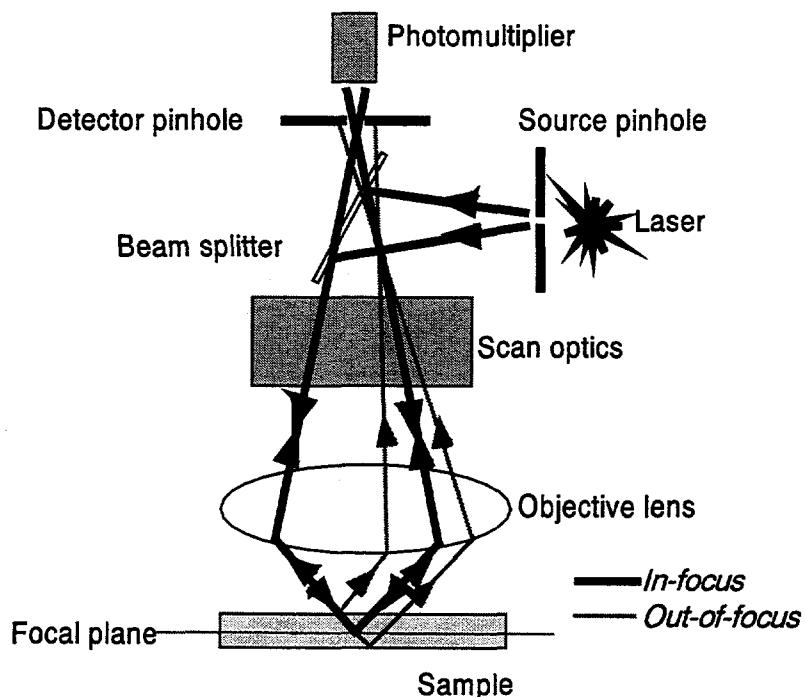
Nuclear magnetic resonance imaging (MRI), also developed primarily for medical uses, has seen increasing application for characterizing the structure of porous materials (Baldwin *et al.*, 1996; Doughty and Tomutsa, 1997). However, while significantly higher resolution than X-ray CT, the current maximum resolution of ~25  $\mu\text{m}$  is not sufficient to capture the narrow pore throats characteristic of most geologic media. For example, even in porous sandstone such as Berea or Fontainebleau, pore throats as narrow as a few microns are statistically significant (Doyen, 1988; Fredrich *et al.*, 1993). Besides being thought to control bulk transport properties such as permeability (e.g., Yale, 1985; Doyen, 1988; Bernabe, 1991), the compliant pore throats also impact critically the evolution of permeability with changes in pore or confining pressure (e.g., Walsh and Brace, 1984), an application of considerable importance in the geosciences.

3D imaging of porous media has recently been accomplished using synchrotron computed microtomography, with data at a resolution of 10  $\mu\text{m}$  reported by Spanne *et al.* (1994), and data at a resolution of 7.5  $\mu\text{m}$  reported by both Auzerais *et al.* (1996) and Coker *et al.* (1996). This resolution is only slightly greater than the resolution of 1-3  $\mu\text{m}$  that is

probably optimal for characterizing rocks such as sandstone. The tomographic reconstruction process favors the use of cylindrical samples and to obtain the quoted resolutions, the above workers used epoxy-impregnated samples ranging from 3 to 3.5 mm in diameter. While increased resolution is desirable and potentially feasible (e.g., *Coker et al.*, 1996), probably the most important factor limiting routine application of this technique is facility access and the availability of beam-line time.

## 7.2 Principles of laser scanning confocal microscopy

The unique optical sectioning capability of LSCM results from fundamental differences in image formation as compared to conventional light microscopy (LM). The essential feature of LSCM (Figure 7.2) is that both illumination and detection are confined to a single location on the specimen at any one time. This is achieved by the use of an hourglass-shaped beam to illuminate only a small volume at any single time, and by the insertion of a pinhole, or confocal, aperture in the returning optical path that blocks light emanating from planes above and below the focal plane. Because the confocal aperture admits light from only a specific plane of interest, thin optical sections can be resolved by varying the depth of the focal plane in the sample. (Conversely, light from both above and below the focal plane is admitted for image formation in conventional LM; this precludes the possibility for depth discrimination and also results in significant image degradation.) Because each point in the focal (image) plane is examined individually in LSCM (in the absence of light scattered from neighboring points), the lateral (in-plane, or x-y) resolution is greatly enhanced. (In conventional LM, an object is viewed under uniform illumination and the point to point resolution is degraded by scattering). In LSCM, imaging of a plane (or



**Figure 7.2** Schematic illustrating the principle of laser scanning confocal microscopy.

volume) is accomplished by scanning sequentially in a point to point manner. The scanning of the laser in the horizontal (x-y) plane is typically accomplished with a set of mirrors, whereas the vertical position of the stage (and therefore focal plane) is controlled with a precision stepper motor. Finally, the use of a laser light source in confocal systems further enhances resolution because lasers offer a high degree of monochromaticity, small divergence, high brightness, high degree of spatial and temporal coherence and plane polarized emission (Gratton and vandeVen, 1990).

The resultant lateral resolution  $R$  is considerably improved over conventional LM, and has been demonstrated (Brakenhoff *et al.*, 1979, 1989; Gard, 1993) to essentially reach the theoretical limit set by diffraction as defined by the Rayleigh criterion:

$$R = 0.61 \frac{\lambda}{N_A} \quad (7.1)$$

where  $\lambda$  is the wavelength and  $N_A$  the numerical aperture of the objective.

The numerical aperture also determines the "thickness" of the optical section, with higher  $N_A$  yielding "thinner" optical sections. In reality, the optical section does not have a discrete thickness, and is instead typically defined as the distance  $z_{1/2}$  between the focus positions at which the collected intensity decays to one-half the peak value. As discussed above, the function of the confocal aperture is to exclude light emitted from above and below the focus plane, and the opening of the confocal aperture thus influences directly the vertical resolution of the optical section. The axial resolution is always less than the lateral resolution, and in practice (Brakenhoff *et al.*, 1979; Gard, 1993) the maximum axial resolution (i.e., smallest  $z_{1/2}$  or narrowest optical section) achieved in confocal systems varies from ~3-10 times the theoretical lateral resolution.

With an objective of sufficiently high  $N_A$ , sub-micron resolution in both the lateral and axial planes can be achieved (Wilson, 1989). Table 7.1 reproduces the optical section thickness experimentally measured by Gard (1993) on a confocal imaging system manufactured by Bio-Rad (MRC-600) and fitted to a Nikon Optiphot for various openings of the confocal aperture at  $\lambda=514$  nm. In the MRC-600, the "closed" position corresponds to a minimum physical opening of 0.6 mm, whereas the "open" setting corresponds to the maximum opening of 8 mm.

**Table 7.1.** Lateral resolution  $R$  (Eq. 7.1) and optical section thickness at  $\lambda=514$  nm for various settings of the confocal aperture

Objective		R ( $\mu\text{m}$ )	Section thickness ( $\mu\text{m}$ )			
M	$N_A$		Open	1/3	2/3	Closed
$\times 10$	0.45	0.71	38	25	13	7.0
$\times 20$	0.75	0.42	14	10	6.7	5.0
$\times 40$	1.0	0.31	6.1	4.2	2.6	1.4
$\times 60$	1.4	0.22	3.7	2.0	1.0	0.7

### 7.3 Technical and practical considerations

*Fredrich et al.* (1993, 1995) and *Montoto et al.* (1995) described independently a technique for applying LSCM to image the void space of porous media. Here we provide additional technical details and discuss practical aspects important to consider for imaging porous materials.

Our technique consists of saturating the pore space of a bulk sample with a very low viscosity, slow-curing epoxy that is doped with a suitable fluorochrome. The wetting characteristics of the sample are generally improved if the sample is first flushed with a solvent such as acetone. Following vacuum drying of the sample, either vacuum impregnation or pressure impregnation can be used to force the epoxy into the void space.

We are partial to a four-component epoxy attributed to *Spurr* (1969) that has a viscosity of ~60 cps and an extended pot-life that approaches 24 hours (available from Ted Pella, Inc., Redding CA). Thorough mixing of the four components and dissolution of the fluorochrome (added as a powder) may take up to an hour, and air is typically incorporated during the process. The extended pot-life thus allows sufficient time both for complete degassing of the epoxy under vacuum prior to sample impregnation, and for the epoxy to penetrate completely into the void space during vacuum- or pressure- impregnation of the sample. Typically the sample is left overnight before curing at 60°C. Finally, the sample is slabbed using a diamond wafering saw, mounted on a glass slide, and a polished planar surface is prepared using standard abrasive polishing techniques, typically to a final finish with 0.05 µm alumina. The thickness is irrelevant. We have verified by direct observation that this procedure is sufficient to ensure epoxy penetration into cracks and pores at the resolution limit of the microscope (~200 nm) for one-inch diameter cores of geologic samples ranging from sandstone to granite.

The choice of fluorochrome is significant (*Wilson*, 1990), and the absorption peak of the fluorochrome should be well matched to the imaging wavelength. Confocal systems most commonly come equipped with either an argon ion laser with lines at 488 nm and 514 nm, or a krypton-argon mixed gas laser with lines at 488 nm, 568 nm, and 647 nm. In our work, we have used Rhodamine B (available from Polysciences Inc., Warrington PA), a popular fluorochrome that dissolves well in epoxy, usually at a concentration of 1:200 (rhodamine: total epoxy components) by mass. Rhodamine B is relatively resistant to photobleaching, and although it can be excited with the 514 nm line of an argon ion laser, it is excited more effectively with the 568 nm line of a krypton-argon mixed gas laser. The wider separation of lines possible with the krypton-argon mixed gas laser also facilitates dual excitation imaging. For example, one might impregnate the pore space with two different epoxies or resins that have been doped with different fluorochromes that are then both excited and distinguished individually using the 488 nm and 568 nm lines simultaneously.

During image collection, the fluorescence intensity as the laser is scanned point-to-point across the field of view is recorded. We find that data acquired at 8-bit resolution (0-255 gray levels) is sufficient for our interests. Besides the obvious influence of

fluorochrome concentration, there are three primary factors affecting the gray level, or brightness, of an image: the opening of the confocal aperture, the laser intensity, and the gain of the photomultiplier. The user must balance each of these simultaneously to arrive at an optimal solution. For example, a larger confocal aperture yields a brighter image, because a thicker optical section is being imaged. Increasing the laser intensity likewise increases image brightness; however, fluorochromes generally saturate at low laser intensity, and a high laser intensity also increases the risk of photobleaching (fading) of the fluorochrome. Finally, image brightness is improved by increasing the gain of the photomultiplier, but this also results in increased noise.

For the applications considered here there are two additional considerations. First, a major interest is to exploit the capability for optical sectioning, and clearly the laser is attenuated with depth due to both absorption and scattering by the material lying above the focal plane, so that greater laser intensity leads to greater penetration depth. Second, we are generally interested in minimizing the optical section thickness (Table 7.1) so that it is comparable to the lateral resolution. (As we discuss later on, subsequent image processing is facilitated if the data are collected to yield cubic voxels, a voxel being a three-dimensional volume element, i.e., the volumetric equivalent to the two-dimensional pixel, or picture element.)

For the fluorochrome concentrations used in our work, filtering the laser to only 10% transmission results in observable photobleaching (using a krypton-argon mixed gas laser with about 5 mW emitted power on the 568 nm line). However, 3% transmission results in no noticeable bleaching, and we generally operate with the laser filtered to either 3% or 1%. To compensate for the reduced intensity that results from the minimized opening of the confocal aperture, we increase the gain of the photomultiplier, and to compensate for the resultant increase in noise, we perform Kalman filtering during image collection so that each image is integrated over multiple scans (usually three).

For optimal resolution and image accuracy, data in the image plane should ideally be collected with a pixel size equal to one-half or less the theoretical lateral resolution imposed by the objective *NA* (*Inoue*, 1986; *Gard*, 1993). In confocal imaging systems, the lateral pixel dimension can be adjusted effortlessly using the microscope's computerized control software. (Physically, the angle at which the laser hits the sample is changed, so that the illuminated area is reduced. For example, with a  $\times 60$  objective, one can achieve a final magnification equivalent to that achieved with a  $\times 100$  objective using conventional light microscopy). But as noted above, the vertical resolution of confocal imaging systems is always less than the theoretical lateral resolution, and yet we would like to gather data such that the axial voxel dimension is equal to the lateral dimension. In our work we have generally sought to maximize the vertical resolution and accuracy.

In our experience, the vertical depths to which optical sectioning can be performed are variable and a function of the imaged material. For example, penetration depths of hundreds of microns ( $\sim 250$   $\mu\text{m}$ ) can be achieved with sandstone before significant image degradation is apparent. However, we observed penetration depths of only tens of microns ( $\sim 50$   $\mu\text{m}$ ) during our imaging of very fine fibrous insulations where the typical fiber

dimension is only slightly larger than the excitation wavelength of the light source (see next section).

In comparison to the alternative porous media characterization methods discussed earlier, the confocal imaging technique has several advantages. First, confocal imaging currently offers the highest resolution of any of the 3D imaging methods available. Furthermore, confocal systems are relatively inexpensive and therefore more widely available. Sample preparation is simple, and the rapid data acquisition time and absence of facility time restrictions expedites analysis of multiple samples and/or sample sub-volumes. Finally, post-processing is relatively straightforward since image data are acquired on a voxel-by-voxel basis, and because the data are intrinsically binary even in multiphase rocks. However, a significant limitation of LSCM compared to other direct imaging techniques such as MRI and synchrotron microtomography is the limit on the  $z$ -extent of the imaged volume due to absorption and scattering of the laser by the material overlying the optical plane.

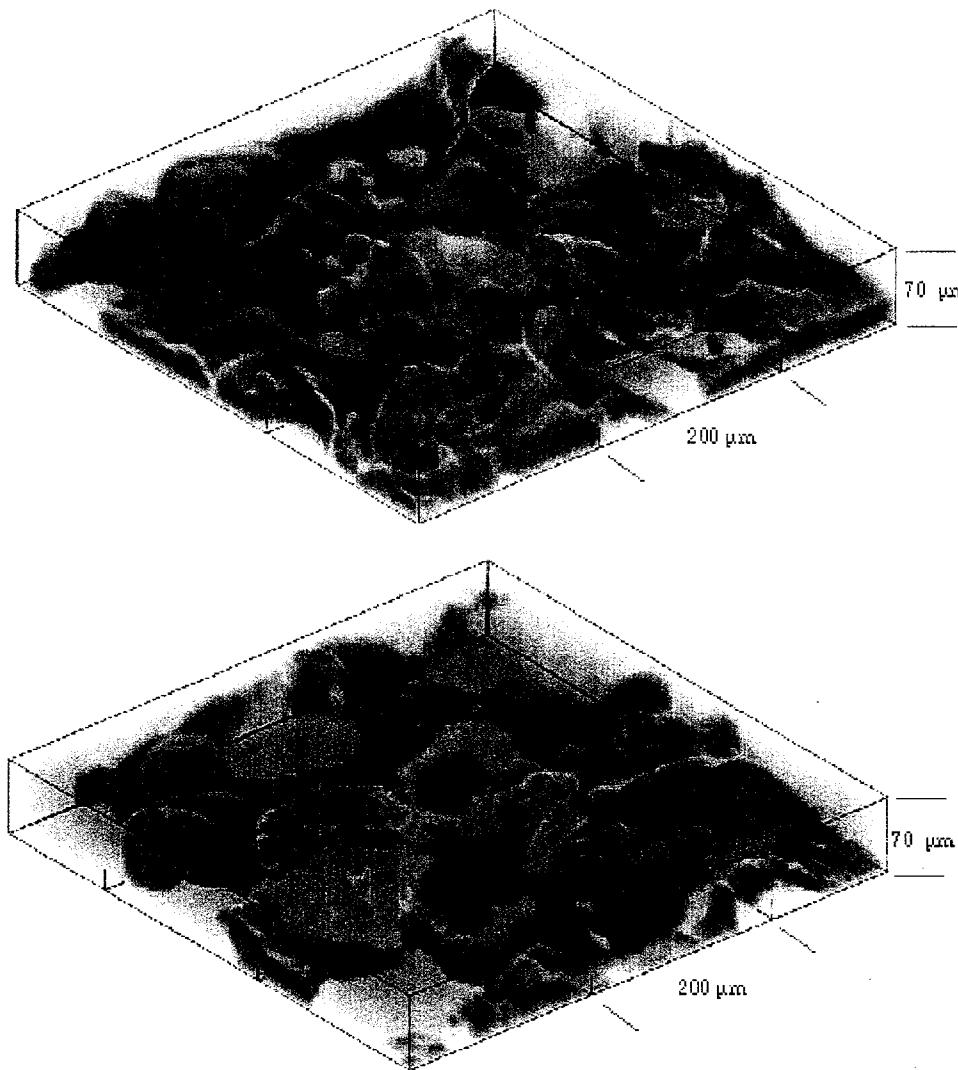
## 7.4 Applications

To illustrate some of the unique capabilities of confocal imaging, here we show several examples of its application. All of the data sets shown were collected with a Bio-Rad MRC-600 confocal imaging system fitted to a Nikon Diaphot, with a 568 nm excitor filter, YHS filter block, and with the opening of the confocal aperture at or very close to its minimal setting (0.6 mm). Collection times for the volumetric data arrays were in all cases  $<1$  hour. The three-dimensional reconstructions were performed using VoxelView® 2.5 (Vital Images, Inc.) on a Silicon Graphics High Impact workstation.

Figure 7.3 shows two data sets acquired from a single Berea sandstone sample with cubic voxels with dimension 1  $\mu\text{m}$ . The data clearly illustrate the fine structure and complicated topology and geometry of the pore space. Particularly conspicuous are the thin pore necks and throats connecting the larger (nodal) pore bodies.

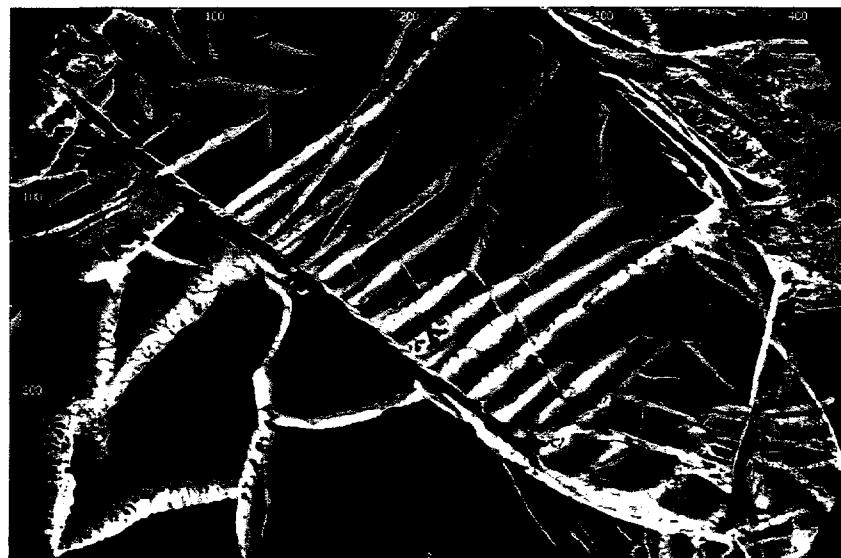
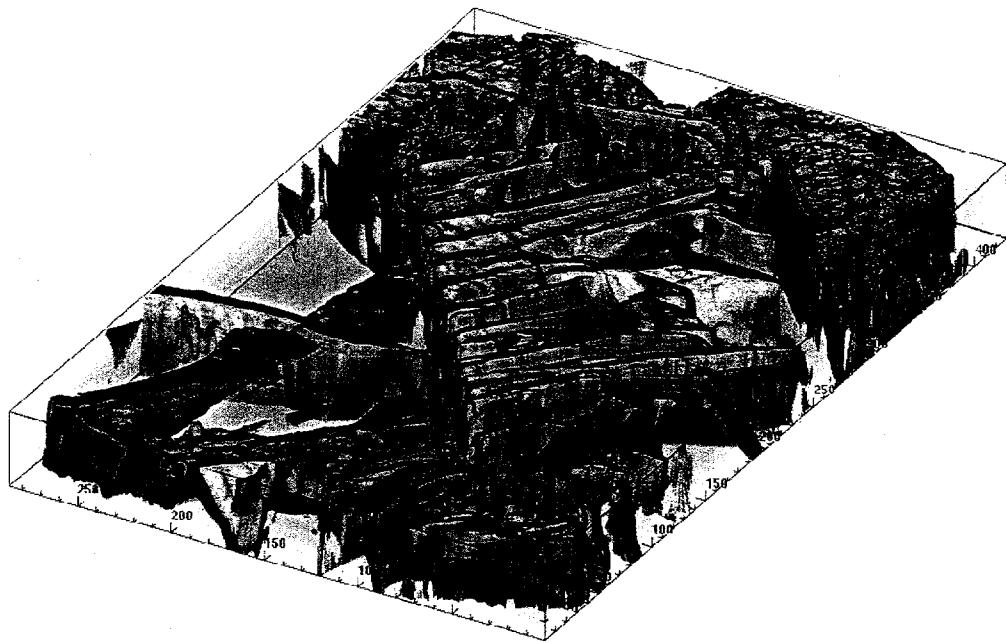
In Figure 7.4 we show a confocal volumetric image (top) from a low-porosity crystalline rock stressed under deviatoric all-compressive loading conditions in a triaxial deformation apparatus at very high confining pressure. The lateral pixel dimension is 0.56  $\mu\text{m}$ , and the spacing between sections is 0.2  $\mu\text{m}$ . With the addition of the third dimension it is clearly revealed that the micromechanical response is dominated by the growth of tensile cracks, with a remarkably fine crack spacing that is highly variable locally. The true crack geometry is difficult to infer from a single 2D image (bottom) because of the fine crack apertures and extremely high crack density. Moreover, using conventional LM, the sub-micron microcrack structure in the upper right would likely yield an undulatory texture suggestive of crystalline plasticity as opposed to brittle crack growth (e.g., *Tullis and Yund*, 1992).

In Figure 7.5 we demonstrate application to a porous engineering material that forms part of the Thermal Protection System (TPS) of hypersonic vehicles and spacecraft. Note

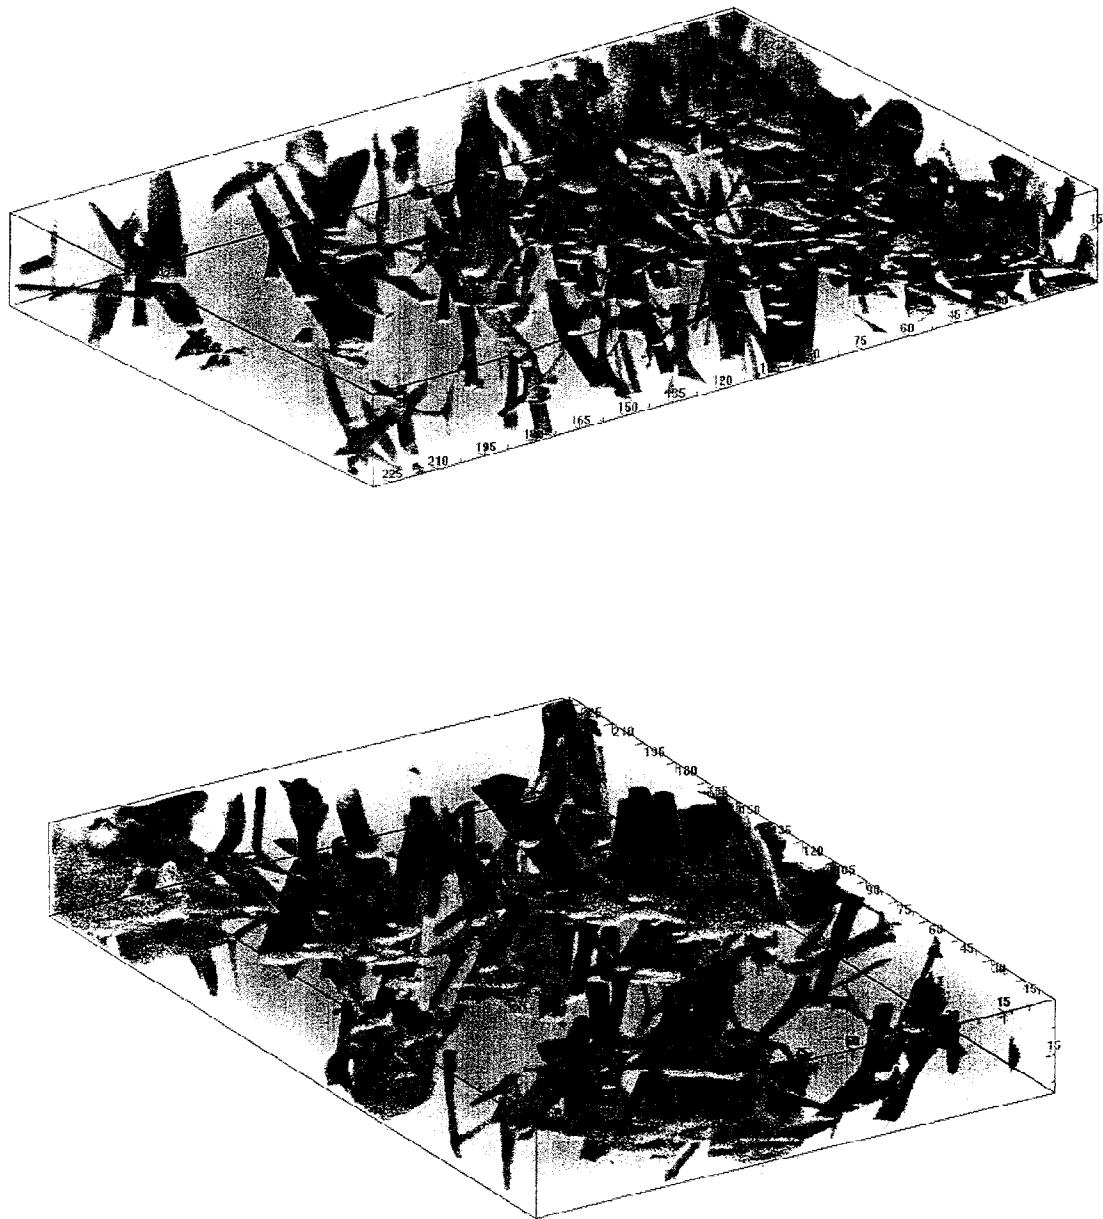


**Figure 7.3** 3D volume renderings of pore space in Berea sandstone. The pore phase is opaque whereas solid grains are translucent. Each image is  $512 \times 512 \times 71$  voxels, with an ( $x \times y \times z$ ) image volume equal to  $0.526 \times 0.526 \times 0.071 \text{ mm}^3$ .

that in these reconstructions, the solid structure is shown opaque, whereas the pore space is translucent. The insulation is made from a slurry of ceramic fibers that are pressed and then fired to yield a rigid material (Figure 7.1, bottom) with a bulk porosity of  $>80\%$  and a complex, anisotropic microstructure (Marschall and Milos, 1997). Central to the performance of TPS insulations is the maximization of their insulative properties with the minimization of their density, and quantitative characterization of the microstructure is thus critical for TPS modeling, development, and optimization. Because of the fine microstructure (fiber diameter of a few microns with length of tens of microns), LSCM is the only feasible technique by which to obtain 3D data that are useful for microscale modeling of bulk thermal and physical properties, and for materials design. The 3D image confocal data clearly reveal the complexity of the fibrous microstructure, and in this case, the complex topology of the solid phase.



**Figure 7.4** (Top) 3D volume rendering of crack damage in a triaxially deformed low-porosity rock. The fine crack network in the upper right corner would not be resolvable using conventional LM. The image is  $768 \times 512 \times 71$  voxels, with an image volume of  $430 \times 287 \times 14.4 \mu\text{m}^3$ . (Bottom) A 2D image slice from the 3D data set. The relationship of the cracks trending NE-SW to the grain boundary (running diagonally NW-SE) is indisputable in the 3D image, but less so in the 2D slice.



**Figure 7.5** 3D volume renderings of the solid phase in two rigid ceramic insulations. Each image is  $768 \times 512 \times 85$  voxels, with cubic voxels with dimension  $0.3 \mu\text{m}$  ( $230 \times 154 \times 25.2 \mu\text{m}^3$ ). The two insulations are composed of the same three fiber constituents, but have been pressed to different final densities. The insulation shown at top has a bulk density of  $0.35 \text{ g/cc}$  and bulk porosity of  $0.86$ , whereas that shown at bottom has a bulk density of  $0.19 \text{ g/cc}$  and bulk porosity of  $0.92$ .

## 7.5 Image processing and segmentation

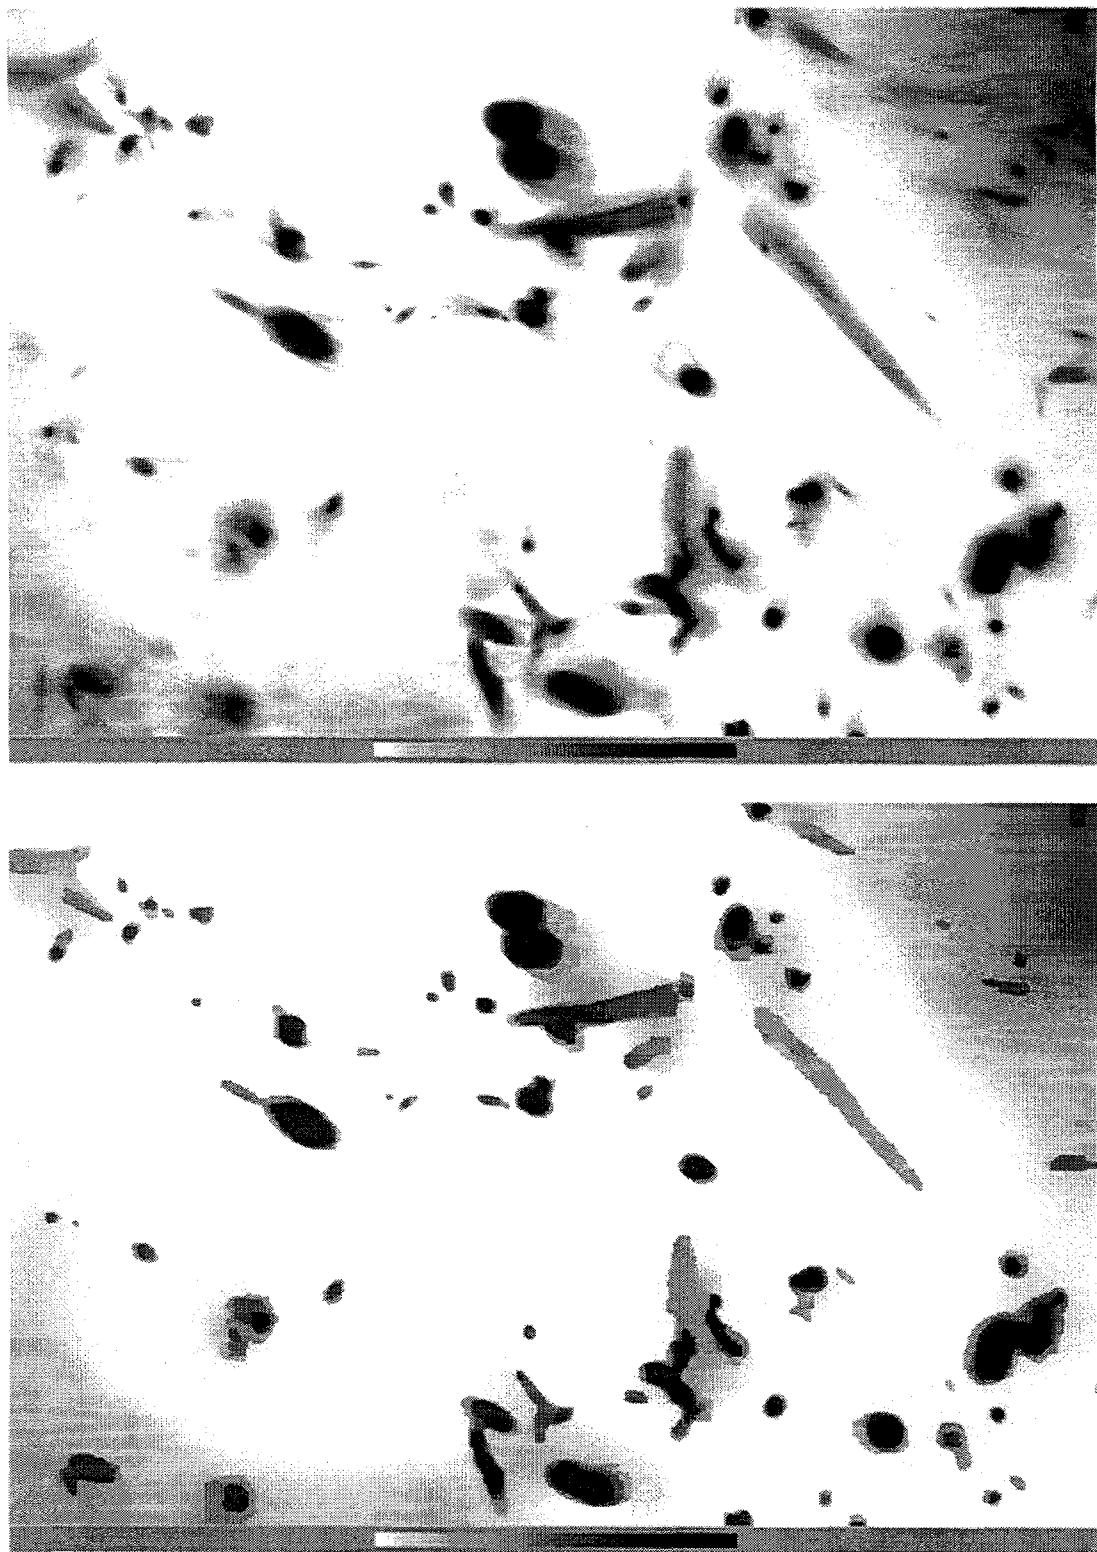
Extracting statistical descriptions of the microgeometry from the volumetric image data requires segmenting the data on a voxel-by-voxel basis into solid or pore space. Quantitative analysis is complicated mainly by the finite voxel resolution which results in some voxels encompassing both solid and void space. This leads to a nonbinary data distribution so that the fluorescence data are distributed across the available (8-bit) range of 0-255. This “partial voxel” problem is intrinsic to all imaging and tomographic techniques. A related difficulty caused by the finite voxel resolution is blurring of edges. To address these issues, we apply two independent algorithms sequentially. The first algorithm, anisotropic diffusion, “sharpens” the image, and the second algorithm, segmentation, determines the cut-off between the void and solid phases. We have investigated several different approaches for segmenting the image data.

### 7.5.1 Anisotropic diffusion

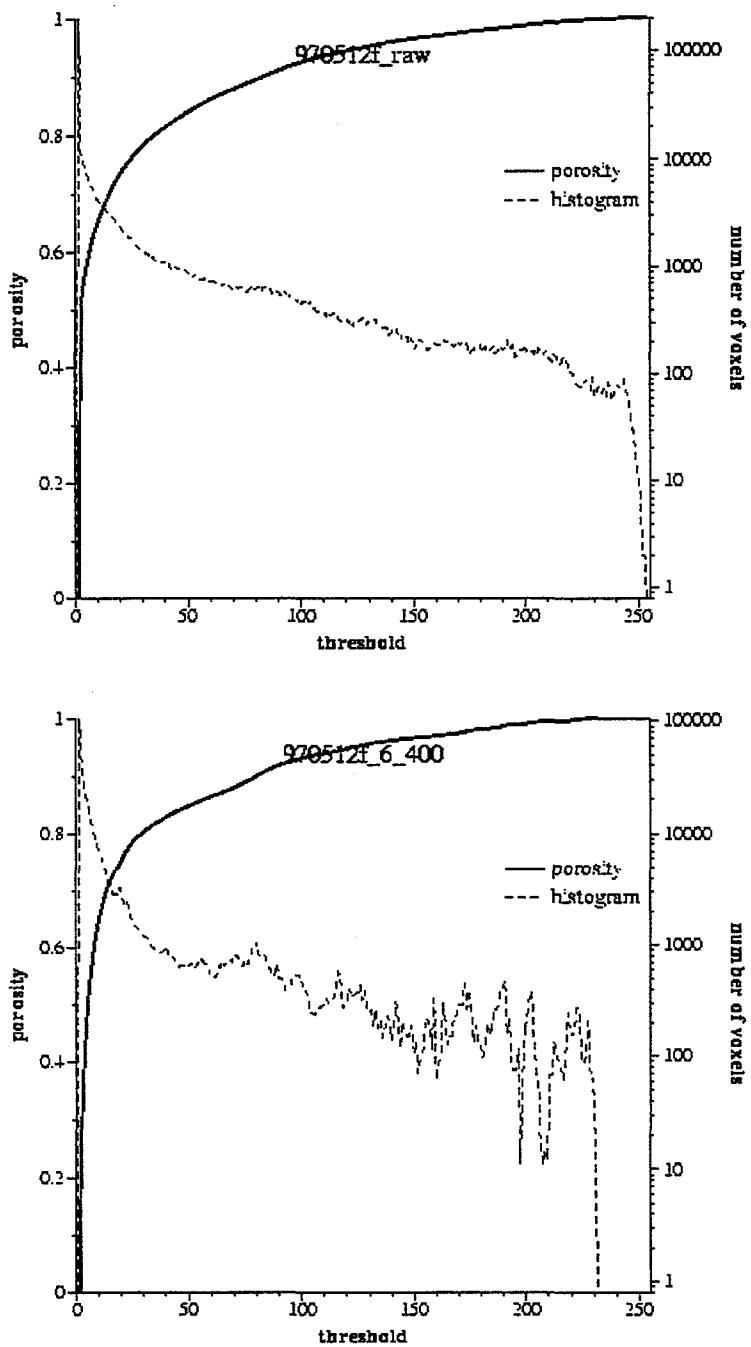
To account for the “blurring” of edges caused by the finite voxel resolution, the image (Figure 7.6, top) is smoothed (Figure 7.6, bottom) by a finite difference implementation of the anisotropic diffusion equation (*Perona et al.*, 1994). The algorithm can be implemented in two-dimensions, that is sequentially for each 2D image, or alternatively, in three-dimensions for the full volumetric data set. Gaussian blurring is assumed and the isotropic diffusion equation is generalized such that the diffusion coefficient is allowed to vary in  $x$ - $y$ - $z$  space, hence its ability to distinguish edges marked by large gradients in grey level from regions within object boundaries which contain small gradients in grey level (i.e. pore or grain interiors). Algorithmically, smoothing within a region in preference to smoothing across boundaries is achieved by setting the spatially variable diffusion coefficient proportional to the local gradient in intensity. The algorithm, as currently implemented, has two free parameters: the number of iterations for which each single image is successively smoothed, and a second parameter that sets a scale for measuring the strength of spatial changes in the diffusion coefficient. An example illustrating the application of this algorithm to image data a porous ceramic insulation (see Figure 7.5) is shown in Figure 7.6.

### 7.5.2 Segmentation

As noted above, an advantage of the confocal imaging technique compared to x-ray microtomography is that the acquired data are intrinsically binary in nature. That is, the pore space fluoresces whereas the solid matrix does not, and thus the distribution of fluorescence intensity in principle contains only two peaks, one corresponding to the void phase and a second corresponding to the solid phase. (In contrast, because each phase is characterized by a unique attenuation coefficient, volumetric image data acquired using x-ray microtomography contain multiple peaks, one each for every mineral phase plus the void phase.)



**Figure 7.6** Example illustrating application of diffusion algorithm to confocal image data. The top image shows a raw  $768 \times 512$  image, and below that is the same image after application of the diffusion algorithm (400 iterations, length scale  $k=6$ ).



**Figure 7.7** Example illustrating application of diffusion algorithm (400 iterations, length scale  $k=6$ ) to confocal image data and resultant enhancement of void and solid phase discrimination. The two plots show histograms (dashed lines) of the inverted fluorescence intensity for the raw (top) and diffused (bottom) images shown in Figure 7.6. Also shown (solid lines) is the variation in the calculated porosity as a function of threshold.

To facilitate efficient processing of multiple volumetric image data sets, it is advantageous to automate the thresholding procedure used to determine the cut-off (i.e., 0-255 in fluorescence intensity) separating the grain from the void phase. Although a simple thresholding procedure may be defined such that voxels above (or below) a certain cut-off are defined as void (or solid), this approach is generally considered inadequate.

Instead, it is desirable to determine the cut-off based on some physical characteristic of the fluorescence intensity distribution. One approach used in our work is based on a slice-wise analysis of the histogram of fluorescence intensity (Figure 7.8) and identification of a local minimum between the two peaks corresponding to the solid and void phase. Because segmentation of each two-dimensional slice is performed individually, intrinsic changes in the fluorescence intensity distribution caused by attenuation of the laser with depth can be taken into account. An example illustrating how the calculated porosity varies with the threshold is shown in Figure 7.8.

We have also implemented a segmentation algorithm based upon indicator kriging. The method has similarities to the general class of spatial thresholding algorithms developed by *Mardia and Hainsworth* (1988). However the Mardia-Hainsworth (MH) method relies on the maximization of a score function which is parameterized by the unknown means, standard deviations, and correlation functions of the populations in the image. Additionally, the score function is based upon an underlying assumption of Gaussian statistics. The indicator kriging algorithm is a non-parametric formulation, requiring only the estimation of the spatial covariance function for an indicator variable. In practice, information of the covariance function only over a limited range of (short) length scales is required.

Implementation of the kriging method requires a-priori population identification of some percentage of the image. In practice this is not difficult to achieve as, based upon the fluorescence intensity histogram, it is (usually) possible to identify the population type of those voxels lying interior to objects (i.e. non-edge voxels) in the image with negligible identification error. Thus a first pass over the image produces population assignments for that fraction of the (in general, non-edge) voxels whose population type is, with negligible error, determinable solely from fluorescence intensity value. Minimum variance estimation (kriging) is then used to complete the segmentation for the remaining voxels in the image.

The kriging-based segmentation algorithm has several advantages over the local minima-based algorithm. First the kriging algorithm requires a fixed number (seven) of passes through the data set, whereas the iterative diffusion algorithm requires  $n$  passes, where  $n$  is set by the user. Typically  $n$  is on the order of 100. Secondly, CPU efficient implementation of the kriging algorithm can be achieved using total storage of  $N$ , where  $N$  is the amount of storage necessary to store a single copy of the image. CPU efficient implementation of the diffusion algorithm requires storage of at least  $2N$ . This is especially significant for 3D segmentation.

As slicewise segmentation does not take into account local variation in the third direction, it is intrinsically poorer than volume segmentation. Consequently our efforts have been directed towards three-dimensional segmentation, which is more memory intensive

than segmentation of a 2D image. For images that are too large to segment as a single volume, the image can be divided into overlapping sub-volumes. If the overlap region is large enough, the segmented image obtained by taking the union of the sub-volumes after discarding the overlapping regions is equivalent to that obtained by performing a single volume segmentation.

## 7.6 n-point correlation functions

Once the data have been segmented (i.e., binarized), first-order descriptions of the microgeometry, that is, the porosity  $\phi$  and specific surface area  $S_A$ , can be calculated using the  $n$ -point probability functions (e.g., *Berryman*, 1985). The one-point probability function  $S_1$  is the probability that a voxel is located in the void phase. It is thus equal to the volume fraction of the void phase ( $\phi$ ), and can be determined by a straight voxel count. The two-point correlation function  $S_2$  is the probability that any two points separated by a distance  $r$  lie in the same phase:

$$S_2(r) = \langle f(x)f(x+r) \rangle \quad (7.2)$$

where the brackets denote volume averaging over the position vector  $x$  and where the function  $f$  is equal to 0 if  $x$  lies in the void phase and 1 if  $x$  lies in the solid phase.  $S_2$  has the important attributes:

$$S_2(0) = S_1 = \phi \quad (7.3)$$

$$\left. \frac{d}{dr} S_2(r) \right|_{r=0} = \frac{1}{4} S_A \quad (7.4)$$

and can be calculated in 1D, 2D, or 3D. Thus, variation in  $S_2$  along the primary  $x$ ,  $y$ , or  $z$  axes can be used to test for isotropy, or conversely, to quantify anisotropy. Finally, the shape of the function  $S_2$  defines a characteristic length of the microstructure (the correlation length).

## 7.7 Medial axis analysis

Geometric analysis of a three-dimensional, irregularly shaped object such as the void phase of a porous medium is difficult. Skeletonization (or medial axis) algorithms (*Lee et al.*, 1994) provide a lower dimensional representation of the object which is easier to analyze than the original object (*Thovert et al.*, 1993, *Lindquist et al.*, 1996). Unless the

original object has embedded cavities, its skeleton is a one-dimensional object (union of curves) in  $R^3$ . The skeleton has a strict geometrical relationship to the object's surface and preserves important geometric properties of the object. Thus various quantitative measures of the geometric structure can be obtained directly from the object's skeleton. Additionally, the medial axis can be utilized as an embedded search structure to find specific sites in the object. An example of the analysis for a confocal data set is shown in Figure 7.9.

With two minor changes, we employ the skeletonization algorithm of *Lee et al.* (1994). The first is to implement conditions at the boundary of the imaged region that reflect the fact that the shape of the object is unknown beyond the boundaries of the imaged region. The second is a reorganization of the order in which the (current) set of surface voxels of the eroded object are treated in order to avoid the systematic north-south, east-west, up-down biasing present in the original algorithm.

As it is defined relative to the object's surface, construction of the medial axis is sensitive to surface noise; different algorithms have different levels of sensitivity. Consider the problem of identifying the void-grain surface in a digitized image of porous (two-phase) medium. Surface noise occurs in two forms, irregularities in an otherwise smooth digitized void-grain surface, and disconnected clusters of void or grain voxels that occur due to voxel misidentification. In constructing the medial axis of the void phase, irregularities in the void-grain surface can have the same effect as 'dead-end pores', producing extraneous 'branches' on the void medial axis. These extraneous branches are generally identified and removed before the medial axis is used to characterize the object. We employ a user set cutoff on the length of such branches to distinguish between true dead end pore paths and those produced by surface irregularities. If only that portion of the void space that supports fluid flow is to be investigated, all dead-end branches can be trimmed from the void space axis.

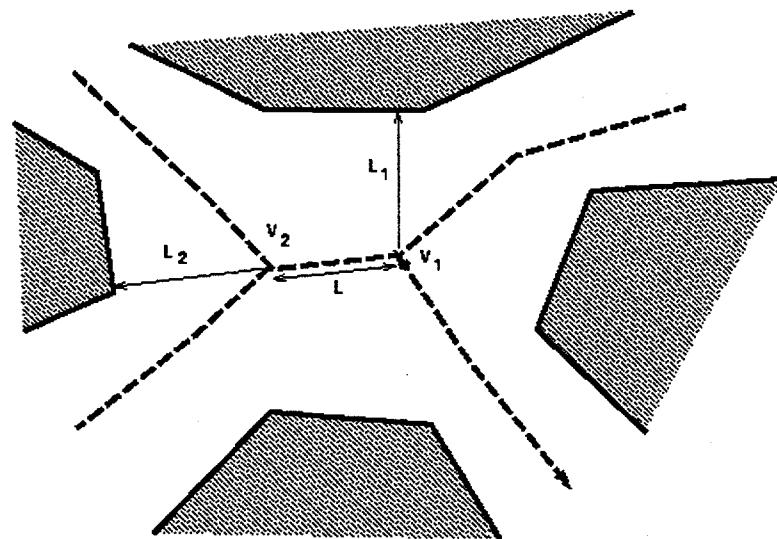
Isolated clusters of void voxels appear as disconnected 'objects' each with its own skeleton piece. Since isolated regions of void phase may indeed be physical (in the case of laser scanning confocal microscopy such physical isolation is due to boundary effects), it is more difficult to distinguish between small isolated void space clusters which are physical and those which are artifacts. Since the probability of misidentifying a cluster of voxels falls rapidly with the increase in number of voxels in the cluster, we employ a size filter; all isolated void clusters less than a user defined size limit are re-identified as grain phase.

Isolated clusters of grain voxels appear as 'cavities' within the void space. Construction of the medial axis for the void space will result in a skeleton that is not reduced to a union of one dimensional curves but contains segments of surfaces that surround each such embedded cavity. Except at boundaries of the imaged domain such isolated grain voxels are clearly unphysical; we identify all such cases and re-identify the affected voxels as void space. If an edge of the imaged region cuts through a grain, the imaged portion of the grain may appear to be isolated from the rest of the imaged grain phase.

Since the information on how this boundary grain segment connects is lost, the medial axis constructed for the void phase will contain a segment of a surface around the inward projecting segment. Since such surface segments in the medial axis are a nuisance to

deal with algorithmically, we identify all such projections and 'remove' them by re-identifying their voxels as void voxels. This produces error in the medial axis at the edges of the imaged volume in proportion to the ratio of boundary area to volume of the imaged region, but has the advantage of producing a medial axis for the void space that is only a union of 1 dimensional curves.

As discussed above, the medial axis requires trimming as dangling ends (dead end paths) can correspond either to spurious elements or to non-percolating components of the pore space. Distinguishing between the two types of dead end paths is very difficult. Consequently we generally trim all dead end paths from the medial axis structure; paths that exit through the boundary of the imaged area are not considered dead ends and are not deleted. What remains is the percolating backbone of the pore space.



**Figure 7.8** Schematic illustrating a single nodal pore body that due to its irregular geometrical shape, has a medial axis reduction containing two, rather than one, vertices.

The medial axis can be viewed as a graph of edges and vertices. Unfortunately, there is not quite a 1-to-1 correspondence between the vertices and edges on the skeleton and the nodal pores and pore pathways in the image. While every pore pathway has a corresponding edge in the skeleton, not every edge has a corresponding pathway. Similarly, while some vertex in the skeleton represents each nodal pore, several vertices may occur in the same nodal pore. This problem occurs within the nodal pores and is illustrated in Figure 7.8 that shows schematically in two dimensions a nodal pore that has a coordination number of four. Rather than containing a single vertex with four edges, the medial axis skeleton of the nodal pore consists of two coordination number three vertices joined by a common edge of length  $L$ . In the usual view of the pore space this edge does not correspond to a pore pathway but is internal to a nodal pore.

We account for this problem in the following manner. In constructing the medial axis of an object the shortest distance between any medial axis voxel and the object's surface (in

this case the pore-grain surface) is computed and stored. Consider an edge, length  $L$ , which joins two vertices,  $V_1$  and  $V_2$ . Let  $L_i$  denote the shortest distance from  $V_i$  to the pore-grain surface.

*Vertex-pair merge condition:*

If  $L < \max(L_1, L_2)$  than the edge and the two vertices it joins are considered to be a single unit, namely a vertex  $V$ .

*Iterative vertex merge:*

The vertex-pair merge condition is iterated in the following manner:

if the vertex pairs  $V_1$  and  $V_2$  are to be merged and  $V_2$  and  $V_3$  are also to be merged, then the three vertices and the two edges involved are all considered to belong to a single vertex unit.

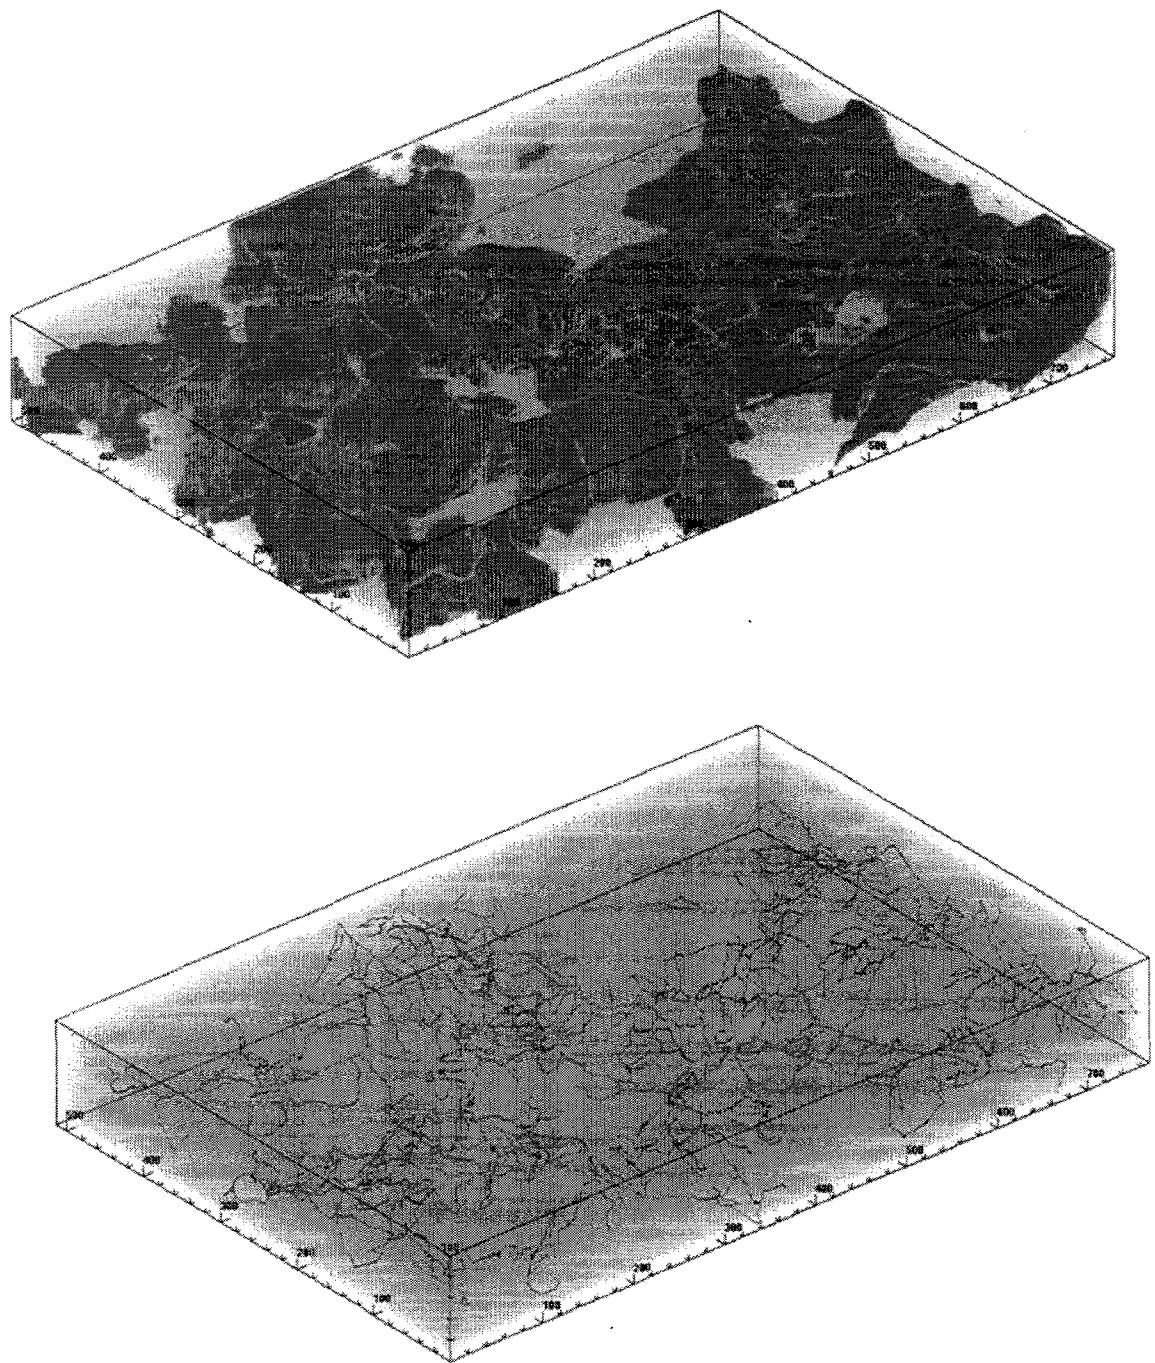
The vertex-pair merge condition has physical interpretation. The distance  $L_i$  is the radius of the largest sphere, centered at vertex  $V_i$ , that fits completely inside the pore space surrounding  $L_i$ . It is thus a measure of the smallest size of the nodal pore associated with vertex  $V_i$ . The vertex-pair merge condition thus states that if another vertex lies within this sphere, it and the edge joining the two must be considered part of the same nodal pore. Iterating the condition allows for the possibility that a nodal pore may, in fact, have more than two vertices lying within it.

When statistics are gathered from the medial axis, each such merged vertex unit identified is treated as a single object. Thus in computing a coordination number, this section of the medial axis is treated as a vertex of coordination number 4. Further, in computing the distribution of edge lengths of the medial axis, the edge joining  $V_1$  and  $V_2$  is ignored. After use of the vertex merge conditions, the edges and (possibly merged) vertices of the medial axis are now in 1-1 correspondence with the pore pathways and nodal pores of the image.

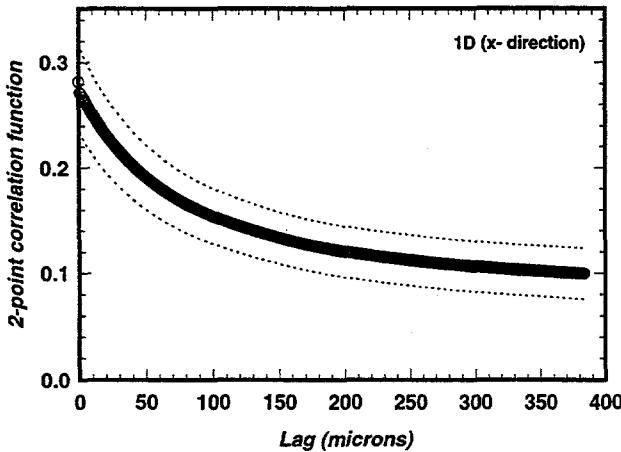
Note that the vertex merge conditions provide protection against finite resolution effects. It is entirely possible that, imaged under coarser resolution (larger voxel size), the medial axis for the nodal pore in Fig. 7.8 would appear as a single vertex with coordination number 4. In both cases (either the coarser resolution or the finer resolution sketched in Fig. 7.8, the statistics gathered from the medial axis after the vertex merge condition identify correctly a nodal pore of coordination number 4.

## 7.8 Application to geometrically complex porous media

As an illustration of the reduction of the complex pore space of geologic media to a medial axis, we show in Figure 7.9 a three dimensional rendering of the pore volume for Berea sandstone, and the corresponding medial axis calculated for this pore volume. Shown in Figure 7.10 is the corresponding two-point correlation function for this volumetric data set that defines the porosity, specific surface area, and characteristic length. The sets of surface points removed during each iteration of the thinning procedure to derive the medial



**Figure 7.9** (Top) 3D volume rendering of pore space in Berea sandstone. The data set is  $768 \times 512 \times 101$  voxels in size, with cubic voxels of dimension  $1 \mu\text{m}$ . (Bottom) The corresponding medial axis calculated for the pore phase. The medial axis encodes the distance to the pore-solid interface, so that the geometry of the original structure is retained, and consists of a series of links (or paths) that connect and/or fork at branch points (or clusters). For example, consider a regular 2D cubic network consisting of nodes and bonds; in a medial axis representation, the nodes would correspond to branch clusters and the bonds to paths. In like manner, a 3D geometrically complex porous medium reduces to a topologically complex medial axis that consists of an irregular network of pore paths that interconnect at branch clusters.

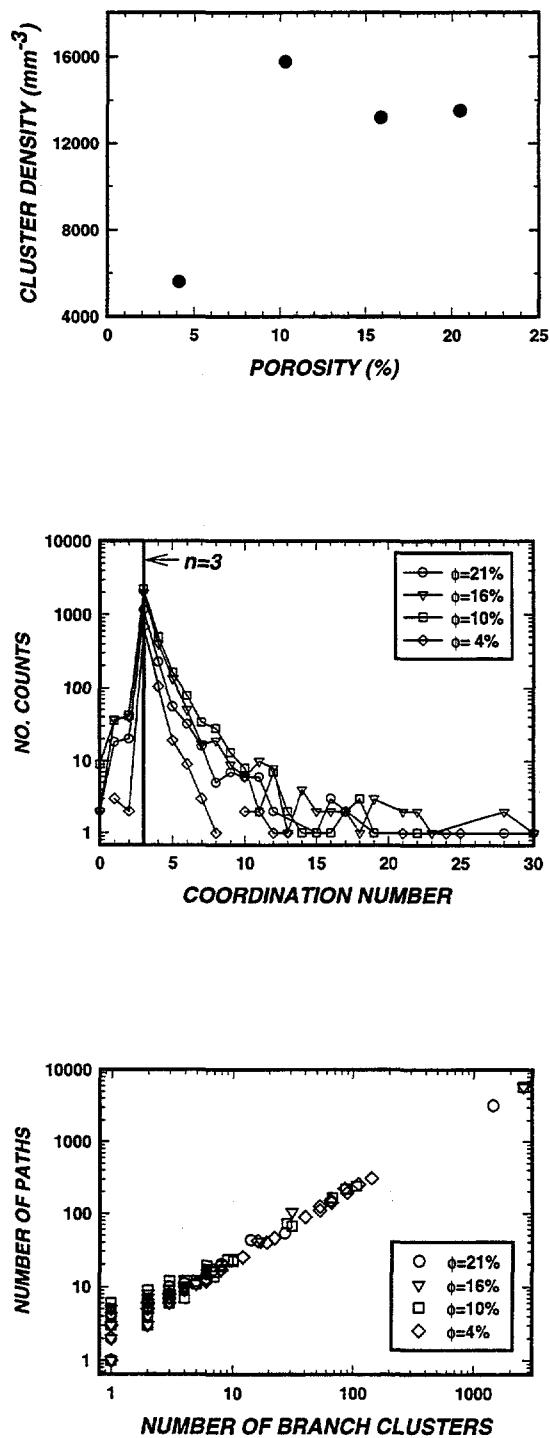


**Figure 7.10** Two-point correlation function  $S_2$  for the data set shown in Figure 7.9. The porosity (equal to the y-intercept) is 28%, and the specific surface area (inversely proportional to the slope at the origin) is 23  $\text{mm}^2/\text{mm}^3$ . The upper and lower curves represent one standard deviation. The point at which  $S_2$  reaches an asymptotic value defines a characteristic length for the microstructure ( $\sim 250 \mu\text{m}$ ).

axis contain information regarding the pore size (or grain size if the object being thinned is the solid phase of the medium rather than the void phase) distribution. Likewise, the medial axis itself can be probed for a variety of statistical metrics, including, coordination, pore path lengths, minimum throat size, tortuosity, etc.

As an example of our results, *Fredrich and Lindquist (1997)* applied medial axis analysis to 3D confocal image data for Fontainebleau sandstone, and showed quantitatively that the internal connectivity of the pore space undergoes drastic evolution with decreasing porosity. Image data were acquired for four samples of Fontainebleau sandstone with porosities of 21%, 16%, 10%, and 4%, and following segmentation, medial axes were calculated for the pore phase for each of the image volumes. Figure 7.11 (top) shows the distribution of the number of discrete branches of the medial axis emanating from clusters, which are branch points in the medial axis. The pore space of each of the four samples is dominated by a coordination number of 3. However, when the density of clusters, which is the number of intersections in the pore network per unit volume, is considered, it is seen that the lowest porosity sample F1 contains a low density of pore space interconnections in comparison to the three higher porosity samples (Figure 7.11, middle).

Finally, the medial axis was probed to examine statistically the distribution of medial axis segments that are internally disconnected within the imaged volume (Figure 7.11, bottom). The latter shows that the three higher porosity samples each contain a single medial axis segment which in size dwarfs all other segments within the imaged volume (this single largest segment corresponds to more than 99% of the pore space in the imaged volume). However, the largest segment of the medial axis in the low porosity sample corresponds to only  $\sim 25\%$  of the pore space. Furthermore, this largest medial axis segment contains about



**Figure 7.11** (Top) Cluster coordination for branch points in pore medial axis for Fontainebleau sandstone with porosities as indicated. (Middle) Density of pore medial axis clusters. (Bottom) Number of paths and branch clusters for disconnected medial axis.

an order of magnitude fewer branch clusters and pore paths than the largest medial axes segments in the three higher porosity samples.

As discussed by *Zhu et al.* (1995), the sample with 4% porosity is in the regime where the power law relationship (with an exponent  $\sim 3$ ) between porosity and permeability is no longer obeyed, and instead permeability undergoes an accelerating decline with porosity reduction (*Fredrich et al.*, 1995). Furthermore, *Fredrich et al.* (1993) showed that the Kozeny-Carman equivalent channel model (*Walsh and Brace*, 1984) over-predicts the bulk permeability of this sample by nearly two orders of magnitude, implying that the pore phase does not uniformly accommodate the bulk fluid flow. In turn, the medial axis analysis reveals quantitatively the decreased connectivity of the pore space, and suggests that the accelerated permeability reduction is due to the resultant increased tortuosity of the fluid flow paths.

## 7.9 Summary

Understanding the physics of microscale flow processes is central to many geoscience, chemical, materials, and thermal engineering applications. Recent developments in three-dimensional imaging offer an unprecedented opportunity to determine the microgeometry of porous media at high resolution. In this chapter we have summarized the technical and practical considerations for applying laser scanning confocal microscopy to porous media, and demonstrated its application to several geologic and engineering materials. Although confocal microscopy has limited ability to penetrate solid materials, it is nevertheless the highest resolution technique currently available for three-dimensional imaging. Other benefits include the wide availability of the technology, and the relatively simple sample preparation requirements. The size of the volumetric data sets attainable using confocal microscopy are in any event of sufficient size ( $\sim 50$  million voxels) that they challenge significantly the current capabilities of high-end workstations.

We believe that besides statistical characterization of the three-dimensional geometry of the microstructure, fundamental understanding of flow processes in complex porous media also requires study of the geometry of the flow itself. Although MRI has been applied successfully to image flow (local fluid velocity) in packed beds (*Lebon et al.*, 1996a,b), the current resolution of MRI is typically not sufficient for imaging transport processes in complex geologic media.

Numerical simulation is the most promising technique at this time, and to this end, we describe in the next section the application of Lattice Boltzmann Methods for modeling heat and mass transport processes in these data sets. The simulations are used to predict macroscopic properties such as permeability, and volume visualization software is used to study the three-dimensional flow geometry in the complex geometric domains characteristic of porous geomaterials. Laboratory measurements of bulk flow properties are also performed for direct comparison with the numerically computed flow properties.

## 8 Application to microscale flows in complex porous media

In this section we describe how three dimensional imaging techniques (Section 7), numerical methods for simulating flow and transport (Sections 1 and 2), and emergent computational architectures (Section 3) are combined to enable fundamental studies of microscale fluid flow in complex porous media.

The microscale flow characteristics of most porous geologic and engineering are not spatially uniform. In particular, the heterogeneous nature of the pore space may cause the flow to be non-uniformly carried. That is, the flow is affected critically by the local geometry and topology as typically there is a distribution of nodal pore and throat sizes, with variable connectivity and coordination (David, 1993; Fredrich *et al.*, 1993). For example, using two-dimensional network models David (1993) found that for heterogeneous pore systems, only a small number of preferential paths carry the majority of the fluid flow. Earlier, Agrawal (1991) found indirectly a similar result by artificially obstructing the first percolating path in Berea sandstone, and then measuring the resultant effect on the bulk permeability. For coupled multi-phase or multi-physics processes, the geometry of the flow likely becomes even more complex and potentially significant to the macroscopic behavior.

To capture and understand how the microstructural geometry affects the creation of preferential paths and their topology, high-resolution reconstructions of pore geometry are obtained using laser scanning confocal microscopy as described in Section 7. The high-resolution volumetric image data of geologic and engineering materials obtained using laser scanning confocal microscopy reduce sampling artifacts to sub-micron features, and simultaneously capture multiple grain length scales. However, the volumetric image data sets are extremely large, and as a result, there are significant computational challenges to utilize this information effectively. The principal problem lies in the complexity of the geometry and the retention of this detailed structure in subsequent numerical analyses.

The geometric complexity (quantity of detail) oftentimes precludes the use of most numerical techniques, such as finite element (FEM), finite difference (FD), molecular dynamics (MD), or network models (NM). FEM is constrained by the quantity of storage required to discretise consistently the pore space (in the mathematical sense) and retain geometric detail. Smoothing (averaging) to achieve representative continuum elements (and numerical tractability) may erase the very features that dictate the evolution of the processes of interest. This is particularly true for multi-physics systems. While similar numerically to LB methods both in the manner in which the domain is discretised and the explicit numerical formulation, FD methods do not provide the mathematical framework in which to develop the multi-physics that are sought ultimately. Discretisation using MD requires meshing the pore geometry in the same manner as FEM, and then discretizing the fluid phase from a molecular length scale up to multiple pore lengths. The storage requirement alone is on the order of a terabyte of RAM per cubic micron of pore space. Using NM, it is possible to discretise the domain in a manner that retains the topology, but idealizes the geometry. Thus, while NM is a powerful tool for the qualitative study of flow paths in complex networks, it is not well suited for quantitative analysis of real systems.

Lattice Boltzmann (LB) methods are ideally suited for modeling phenomena at time and length scales intermediate to the continuum (macro) and molecular dynamics (micro) scales, that is, at the mesoscopic scale. LB methods do not suffer as severely from the discretisation constraints of FEM, but still require some ingenuity to be applied to the scale of real problems such as those considered here. To appreciate the computational obstacles involved, consider a typical data set obtained using confocal microscopy that consists of  $768 \times 512 \times 128$  voxels ( $\approx 50$  million sites). In the standard LB formulation a voxel maps to a single lattice site, and to model the physics of interest, each site requires  $\approx 100$  bytes of state information. So while the meshing can be automated as a 1:1 mapping, a total of  $\approx 5$  gigabytes of memory is required. Because of the numerically explicit formulation of LB, an additional constraint is introduced in that state information from each site must be transmitted to every other site before a steady state flow can be reached. For a data set of the size mentioned, it is not unusual to require of order  $10^4$ - $10^5$  iterations. On a modest scientific workstation  $\sim 0.25 \times 10^6$  sites per second might be processed, thus a data set of the above size would require weeks to months of CPU time. Of course this assumes that the data set can be represented in memory in the first place; in reality these spatio-temporal requirements render this technique unfeasible for application on a single workstation.

We overcome these constraints through two distinct but complementary innovations in the treatment of large data sets. The first is the derivation of a compression algorithm (a stencil of the pore space) to represent the geometry and simultaneously reduce the compute time. Further, because LB methods are numerically explicit as formulated, this characteristic can be exploited through a mapping of the numerical domain to distributed computing architectures. To further advantage this aspect, we designed a cost effective hardware platform to store and compute tractably systems of this size. Combined, the two techniques permit the simulation of transport processes in large data sets over representative length and time scales.

Here we describe the application of these techniques to perform single-phase flow simulations in volumetric image data sets obtained from cores of Berea sandstone using confocal microscopy. Simulations are performed using both the purpose-built distributed processor system described in Section 3 and the Sandia/Intel ASCI Red massively parallel supercomputer.

## 8.1 Lattice Boltzmann method

As described in Sections 1 and 2, LB methods are kinetic theory based techniques for simulating fluid flow that, for appropriate choices of energy distribution at the microscale, recover the Navier-Stokes and advection-diffusion equations at the macroscale. The microscale (in this model) consists of populations of fluid particles,  $n_i$ , associated with each site of a regular grid or lattice. The sites are connected with their neighbors through a set of prescribed numerical links that describe the paths for particle migration through the system. The number and spatial orientation of these links are defined by constraints on the symmetry and isotropy imposed by the underlying physics being modeled. In this work we use the D3Q19 lattice (*Qian, 1992*).

The state of each site, i.e. energy distribution, evolves according to the Boltzmann equation:

$$n_i(\mathbf{x} + \mathbf{e}_i, t + 1) = \Omega(n_i(\mathbf{x}, t)) + n_i(\mathbf{x}, t) \quad (8.1)$$

where  $\mathbf{x}$  denotes a position in the lattice,  $\mathbf{e}_i$  is the direction of link  $i$ ,  $t$  is time, and  $\Omega$  is a Boltzmann collision operator. Eq. 8.1 can be decomposed into two parts, a collision term and a streaming (translation) term (Fig. 2.1). The collision term seeks to relax the local population of particles to equilibrium, and the streaming term migrates any net imbalance in the population to neighbouring sites along the set of prescribed links. Repeating these two steps over time allows the global system to evolve to an equilibrium or steady state condition.

In practice we use the linearised Bhatnagar, Gross and Krook (LBGK) form of the Boltzmann equation (*Bhatnagar et al.*, 1954):

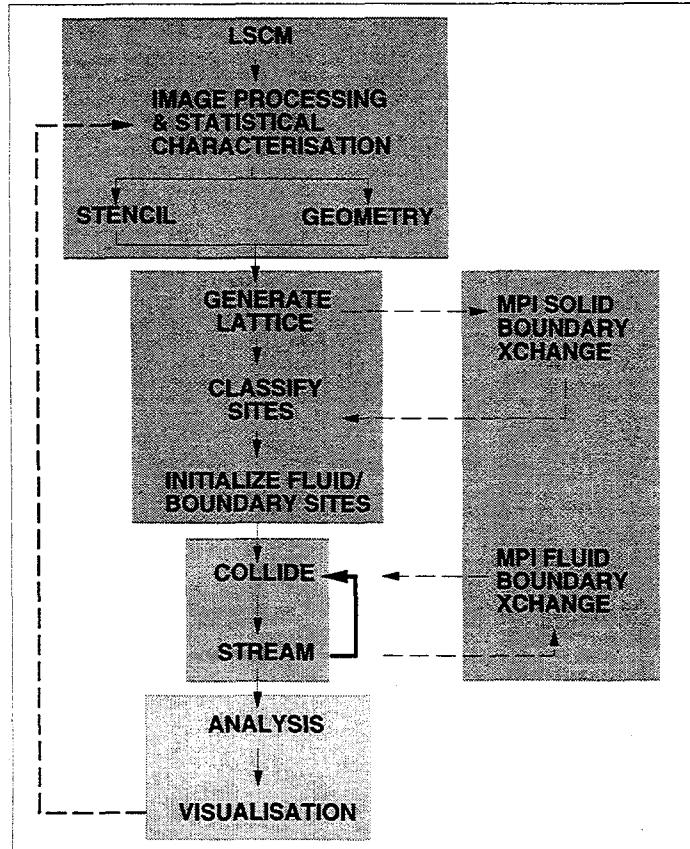
$$n_i(\mathbf{x}, t + 1) = n_i(\mathbf{x}, t) + \frac{1}{\tau} [n_i^{eq}(\mathbf{x}, t) - n_i(\mathbf{x}, t)] + \mathbf{F}_i(\mathbf{x}, t) \quad (8.2)$$

As before  $n_i$  describes the particle populations at each site  $i$ , but the collision operator  $\Omega$  has been reduced to a linear expression in terms of the current distribution  $n_i$  and a prescribed equilibrium distribution  $n_i^{eq}$ . This linearization is described as a single time-step relaxation where  $\tau$  is the critical time step. In this expression we also show the term  $\mathbf{F}_i$ , that describes a forcing component used to drive the flow, e.g. gravity. Choosing a suitable equilibrium distribution,  $n^{eq}$ , allows one to recover Navier-Stokes flow at the macroscale, with stability over a range of Reynolds Number,  $Re$ , that bounds the flow regimes of interest in this work.

## 8.2 Distributed Implementation

The porous media simulation framework is shown schematically in the form of a pipeline in Figure 8.1. The pipeline starts with three-dimensional volumetric image data acquired using confocal microscopy (LSCM), statistical analysis and characterization, and the geometric reconstruction of the porous medium as a segmented (binary) three-dimensional image. The reconstructed volume forms the input to the flow simulation proper, that is, the sites describing the solid and void phases that together define the geometry of the connected pore space. These data are mapped onto the lattice in one of two ways.

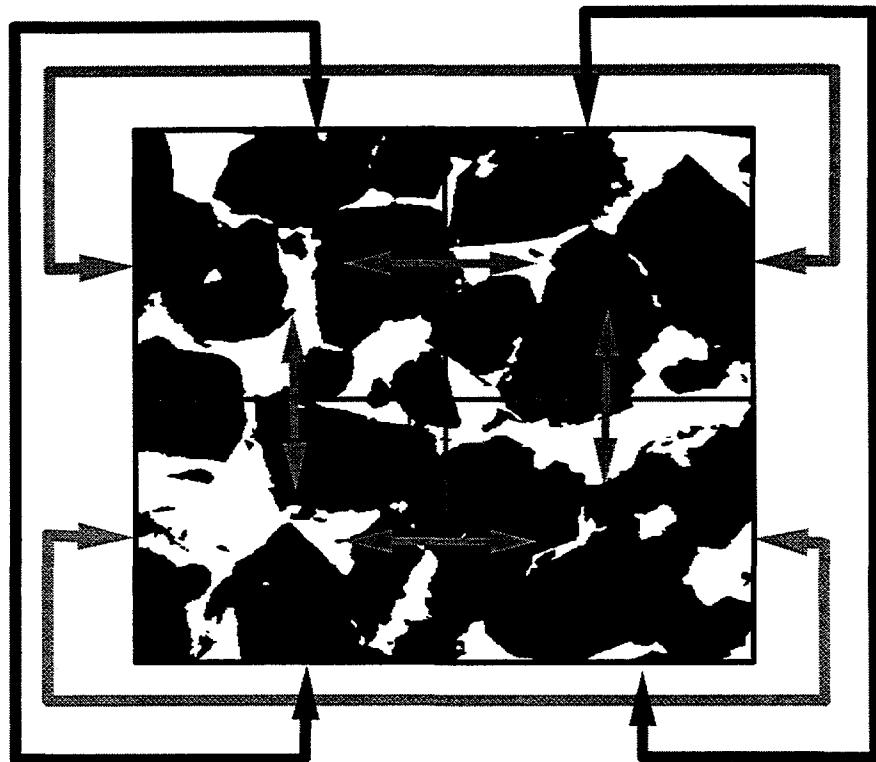
In the traditional method, each binary voxel from the imaged domain is mapped onto a fully defined lattice site. Each voxel designates whether a site is treated as solid or void (containing fluid), but regardless of the phase, each site reserves sufficient memory to represent the fluid state. As the lattice is fully populated (defined), the binary data act as a mask designating whether the fluid physics at that site is included in the overall system and needs to be updated. So while this 1:1 mapping is straightforward and creates automatically a mesh of the geometry, it is very wasteful in terms of storage. Consider that the porosity of engineering materials is rarely  $>70\%$  and in the case of hydrocarbon-bearing sandstones,



**Figure 8.1** Schematic illustrating coupled experimental-computational pipeline for studying flow in complex porous media.

more typically  $\sim 5\text{-}35\%$ . That is, in many cases the majority of the domain may consist of solid material that does not take part in the flow simulation. Yet the traditional implementation reserves memory to represent the state of the fluid at every site. Thus, for sandstone with porosity of say 20%, this approach wastes 80% of allocated storage.

A second mapping approach that we refer to as stenciling was developed specifically to address this problem (Figure 2.2). Although more sophisticated, it nonetheless generates automatically a mesh consistent with the traditional approach. The essence of our method is to use the binary data set as a stencil for memory allocation and layout, rather than as a simple mask. That is, lattice memory is allocated only for sites that correspond to the void phase. Additionally, the location of the memory associated with each void voxel is encoded and stored in a corresponding location in a 3D image map on a site-by-site basis. With this approach, the algorithm traverses the image, and only when a site corresponding to void is encountered are the contents of the associated lattice memory accessed to determine the physics at that location. Thus, for Berea sandstone, this method enables the storage of models up to 5 times larger with the same memory resource as the traditional approach. In the general case, the increase is equal to  $1/\phi$  where  $\phi$  is porosity. Because the stencil representation retains the surface description of the solid phase, it does not preclude



**Figure 8.2** Example illustrating the decomposition of a two dimensional geometric domain into four subdomains, and showing the communication between neighboring subdomains that is accomplished using message passing. In the 3D case, each subdomain also communicates along two shared boundaries that would lie in the plane of the image. The implementation of periodic flow boundary conditions on the boundary of the entire geometric domain is illustrated by the large arrows outside of the image area.

application to coupled processes such as dispersion or multiphase flow. Multi-physics flows involving the solid phase (e.g., reactive flows such as precipitation/dissolution) clearly require the solid geometry to be mapped explicitly.

For the distributed implementation we also developed algorithms to decompose the geometric domain into subdomains such that each subdomain is computed as a separate process. The decomposition was facilitated through use of a technique known as message passing (*MPI Forum*, 1995). Specifically, we use the MPICH library implementation (*Gropp and Lusk*, 1997) of the *de facto* standard MPI to ensure portability to multiple hardware platforms and vendors. Message passing is a programming method that allows multiple processes to communicate with one another (whether they are running on the same processor, or on another processor connected over a network), in a way that makes the computation appear contiguous at the shared boundaries of the subdomains. A decomposition and communication pattern is shown schematically for a two dimensional domain in Figure 8.2. The arrows connecting the common internal boundaries of the decomposed domain indicate how neighbouring processes communicate. The larger arrows connecting the outer boundaries of the domain indicate the communication pattern for periodic flow conditions as described in the next section.

In Figure 8.1 we show the message passing component of the pipeline as a parallel task to indicate that the simulation framework is independent of the underlying computer architecture. Message-passing specific operations are performed only if the system is being run on a distributed or parallel computer.

To initialize the system, copies of the shared boundary geometry and stencil information are sent to neighboring processes (subdomains). The state of the fluid sites on the shared boundaries are also initialized by way of message passing so that consistent information is used to perform the computation of the flow physics.

The flow field is calculated iteratively, with the update loop treated as two distinct phases. First the local physics is calculated for each fluid site using the collision term of Eq. 8.2. This establishes a new distribution of particles on each of the links. The second phase is to translate (or stream) the contents of the links to the nearest-neighbor lattice sites. This numerical communication extends to the logical boundaries of the subdomain, at which point an additional message passing step is performed to transmit this information to the other subdomains.

The system is computed in this way until a steady state condition is reached. This is typically recognized as a stabilization of the local velocity and the decay of long range fluctuations over the domain. At this point the state of the system is analyzed for the physical phenomena of interest, e.g. intrinsic permeability. An advantage of the coupled experimental-simulation technique is manifest in the form of the rich data set that is derived. For example, it is possible to resolve the flow characteristics at the sub-pore scale. Finally, insight can be gained into the evolution and detailed behavior of the flow fields through use of three dimensional visualization.

### 8.3 Flow boundary conditions

Two general forms of flow boundary conditions may be applied. The first defines sources and sinks at the boundaries to maintain a constant flow rate through the system. Realistic inflow conditions require a steady-state but spatially non-uniform velocity profile to match the effects of flow emanating from and being absorbed by the neighboring pore space. In general it is not possible to define consistently this condition without detailed *a priori* knowledge of the flow conditions at the boundaries.

In the second form, the flow is defined to be periodic at the boundaries. That is, as shown schematically in Figure. 8.2, the flow exiting the system is wrapped around to the ingress; however, this approach is not locally consistent at the boundaries because of the inevitable geometrical mismatch. To address this issue, the geometry is explicitly mirrored in the flow direction (doubling the  $x$  extent of the domain). Mirroring can also be introduced in the other two directions, either explicitly as in the flow direction, or, alternatively, by imposing symmetry boundary conditions. In the absence of a means to define consistent inlet-outlet or

source-sink profiles, we have implemented the latter approach. (Note that symmetry and mirroring are not shown in Figs. 8.1.)

The remaining component of the LB implementation is the definition of a no-slip boundary condition to describe how fluid particles behave in the presence of solid obstacles (e.g., at pore walls). For simplicity the “bounce-back” boundary condition is used. This rule was derived originally for Lattice Gas Automata (Cornubert *et al.*, 1991) and is applied in a continuum form for LB. The premise is to reflect fluid particles impinging upon a solid wall (Fig. 2.1). While this is sufficient for modeling single-phase flow, a more complex treatment is necessary for modeling dispersion or multiphase flow (e.g., Noble *et al.*, 1995, Chen *et al.*, 1996). (For example, for dispersion, the carrier fluid is not characterized accurately in the boundary layer between the solid and void phase, leading to physical anomalies in solute transport.)

## 8.4 Distributed computing

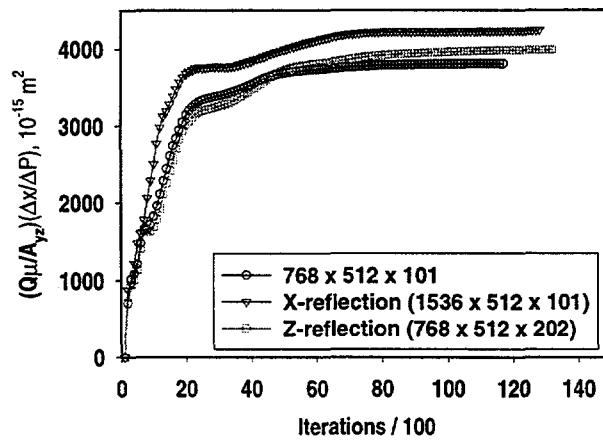
As discussed above, the computation is intense and it is not possible to implement problems of the size considered here,  $768 \times 512 \times 128$  ( $\approx 50$  million sites), on a standard scientific workstation. While supercomputing resources are available, regular and uninterrupted access to such resources is not typical.

To address the latter aspect, we used commercial off-the-shelf components to construct a 32-node distributed computer that is capable of 6 Gflops, with 8G of RAM (see Section 3). Using the cluster, we can achieve single-phase flow simulations for  $\approx 100$  million sites for representative porosities. However, extending the size of the data set further, or examining scaling properties still requires use of supercomputing resources. We have used 1024 nodes of the ASCI Red (Teraflops) supercomputer to perform single phase flow calculations in a data set sized  $(768 \times 2) \times (512 \times 2) \times (128 \times 2) \approx 0.25$  billion sites.

## 8.5 Application

We have applied our experimental-computational framework to model single phase fluid flow in three dimensional reconstructed volumes of Berea sandstone. Simulations have been performed for several data sets; however, here we show results only for the data set shown in Figure 7.9.

The intrinsic permeability  $k$  can be calculated based on the standard Darcy formulation as indicated in Eqs. 8.3 and 8.4, where  $\mathbf{U}_{Darcy}$  is the average fluid velocity over a cross-sectional area  $A_{yz}$  orthogonal to the pressure difference  $\Delta P$  driving the flow,  $\mu$  is the prescribed fluid viscosity, and  $\mathbf{Q}$  is the volumetric flow rate through the medium. To drive the flow under gravity conditions one can define an equivalence between an applied body force  $\mathbf{F}_x$  and the pressure difference experienced using Eq. 8.5. Substituting in Eq. 8.3 yields Eq. 8.6 that is used directly to derive  $k$ :



**Figure 8.3** Plot showing the evolution to a steady state permeability during the simulation. The simulation is for the reconstructed volume shown in the following figure (Figure 8.4, Top).

$$\mathbf{U}_{\text{Darcy}} = -\left(\frac{k}{\mu}\right) \left(\frac{\Delta P}{\Delta x}\right) \quad (8.3)$$

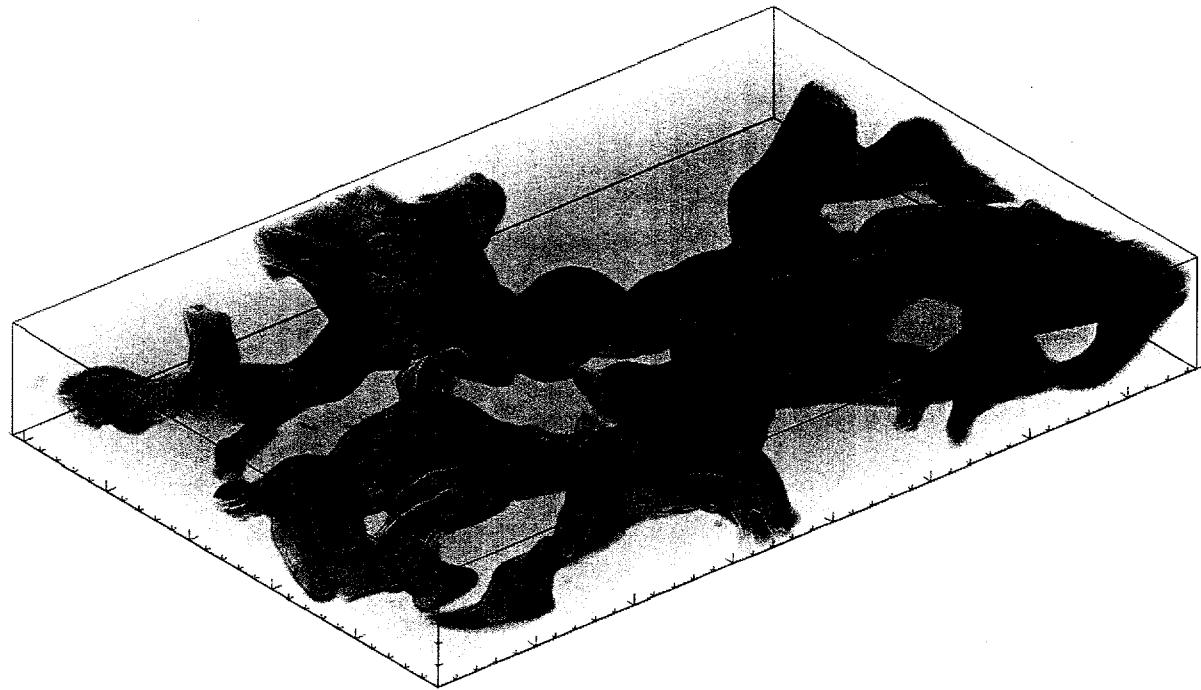
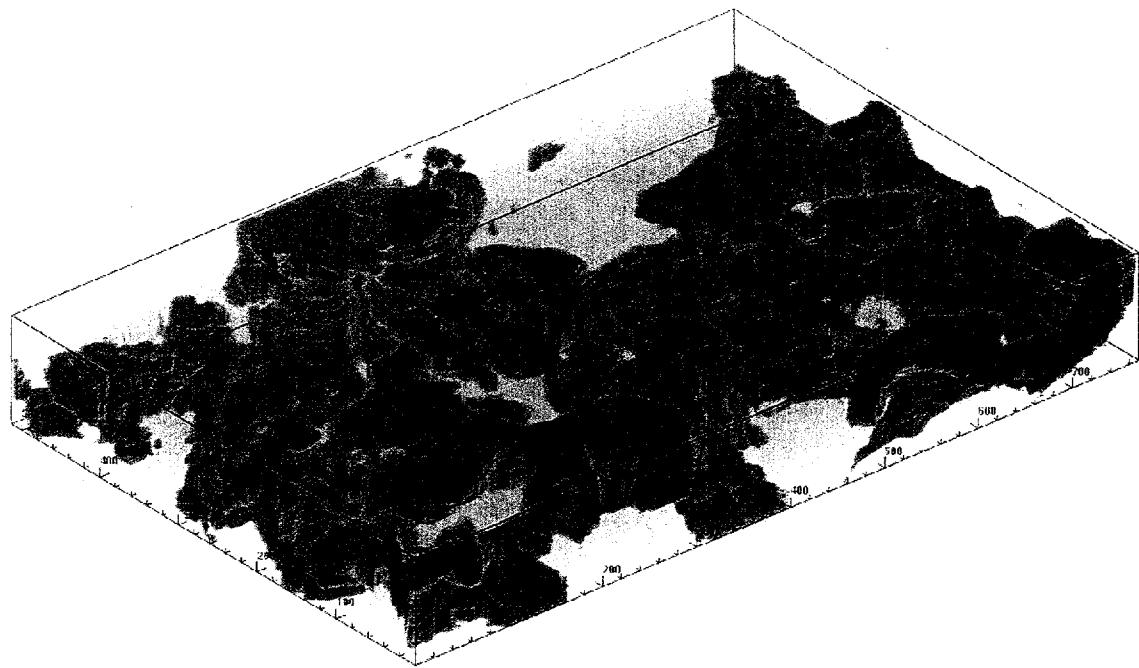
$$\mathbf{Q} = \mathbf{U}_{\text{Darcy}} A_{yz} \quad (8.4)$$

$$\Delta P = -F_x \Delta x \quad (8.5)$$

$$k = \frac{\mu \mathbf{U}_{\text{Darcy}}}{F_x} = \frac{\mu \mathbf{Q}}{F_x A_{yz}} \quad (8.6)$$

Note that the flow is being driven in the  $x$ -direction that corresponds to the longest dimension in our model (i.e., 768 voxels as shown in Figure 7.8).

The results from three simulations are shown in Figure 8.3 and 8.4. The lowermost curve in Figure 8.3 corresponds to the simulation on the basic reconstructed volume shown in Figure 8.4 (top). The uppermost curve in Figure 8.3 corresponds to a simulation where the volumetric domain has been reflected about the  $y$ - $z$  plane in the direction of the flow (i.e., the  $x$ -axis). The middle curve in Figure 8.3 is for a simulation where the volumetric domain has been reflected about the  $x$ - $y$  plane, that is, doubling the minimum dimension of the model in the  $z$ -direction.



**Figure 8.4** (Top) Volume rendering of the pore space in Berea sandstone used in the simulation with results shown in Figure 8.3. The image size is  $768 \times 512 \times 128$ , with cubic voxels of dimension  $1 \mu\text{m}$ . (Bottom) Three dimensional rendering of the steady state flow field. The red color corresponds to the highest velocities, and blue to the lowest velocities. Shown in green are intermediate velocities.

Figure 8.4 (bottom) shows a three-dimensional rendering of the magnitude of the velocity field for the basic reconstructed volume (corresponding to the lowermost curve in Figure 8.3). The velocities have been thresholding at three levels. The complexity of flow patterns that arise in 3D “real” geometries is captured clearly: the simulation reveals accelerated flow in localized areas within pore bodies and through some pore throats, as well as the presence of low-velocity or “dead” regions within pore bodies where the flow is more than 10× slower than that observed in high velocity zones.

The range of  $k$  exhibited is approximately a factor of 3-5× the experimentally determined intrinsic permeability of Berea sandstone. Because the data sets used corresponds to a original imaged volume of  $0.8 \text{ mm} \times 0.5 \text{ mm} \times 0.1 \text{ mm}$ , one may expect the numerically derived permeability to be higher than that measured in laboratory tests conducted on core samples (typical cylindrical dimensions of  $25.4 \text{ mm} \times 50.8 \text{ mm}$ ). In other words, there may be artificially connected flow paths that are either unconnected or else may not contribute to bulk flow in a larger sample, i.e. the image volume size may not capture all of the representative length scales.

These issues are being addressed by gathering multiple sample volumes from a single sample, by statistical analyses of the image data, and by numerical simulations on the imaged volumes. In addition, we are performing simulations on imaged volumes obtained from multiple samples of Fontainebleau sandstone with porosities that vary from 5% to 25% to compare more directly the computed versus simulated permeabilities.

## 8.6 Discussion

We have created a coupled experimental-computational framework to integrate 3D imaging with numerical flow simulation techniques using approaches conducive to distributed computing. Our preliminary simulations examine single-phase flow phenomena in Berea sandstone, and our results, including numerical estimates of bulk permeability, are encouraging. We achieve the tractable computation of large 3D data sets through development of new storage algorithms to represent complex geometries, in addition to development of a distributed hardware platform. In combination, these advances have dramatically increased the time and length scales over which transport processes can be simulated using commodity computer components readily available at the single-investigator level.

The complexity of flow patterns that we observe for single-phase flow simulations in “real” porous materials demonstrates clearly the power of the LB simulation method. The simulations reveal the portions of the void space that are most readily accessed during macroscopic fluid flow, and can be used to investigate the relationship between the microscale geometry and the macroscopic flow properties. Further, ability to visualize the flow field has great utility for studying dispersion in porous media. Likewise, we anticipate great utility in addressing multiphase immiscible flows, where using the methods developed here can be used to directly determine the efficiency of an invading fluid in displacing an original pore fluid, for example, during simulated enhanced oil recovery (EOR) processes.

## 9 Summary

The purpose of the work described here was the development of a lattice Boltzmann simulation capability. This LDRD project has demonstrated and validated its application to classic engineering flows such as three-dimensional duct flow, and higher Reynolds number flows such as the backward facing step as described in Section 4.

Our work has also included application of the lattice Boltzmann method to coupled multiphysics flows such as the thermal hydrodynamic flow examined in Section 5. We have also developed a method (Section 6) to consider partially saturated computational cells, with applications towards problems with moving solid boundaries, or to modeling microscale flows using reconstructed porous media. Both of these problems demonstrate the intrinsic capabilities and utilities of the lattice Boltzmann method.

Significant emphasis was placed on developing a coupled experimental-computational framework to simulate microscale transport processes in porous geologic and engineering materials. Lattice Boltzmann methods have a unique capability in this area because of the inherent geometric complexity of the model domain that precludes use of conventional Computational Fluid Dynamics methods. Fundamental to this work was the development of high-resolution three-dimensional imaging methods and analysis techniques to characterize complex microstructures as described in Section 7.

Our intention is to use the coupled experimental-computational framework that is described in Section 8 to gain fundamental understanding of the nature of flow and transport in complex porous media to complement the laboratory scale experiments that we currently rely on. While our present focus is on single phase flows, ultimately we aim to extend this framework to simulate the multiphase flows and coupled processes that are even less amenable to experimental investigation. Our goal is to derive new understanding of transport and flow in complex porous media by complementary analyses of the 3D microgeometry, laboratory transport experiments, and pore-scale numerical flow simulations. The simulations and visualizations shown here clearly demonstrate the unique capability of LB methods to this class of applications.

Finally, as was described in Section 3, we also demonstrated the design and development of a distributed computational platform for the performance of numerically intensive simulations such as the pore scale flow modeling described above. While not an original goal of this project, this result has greatly extended our ability to routinely perform very large-scale simulations on the order of hundreds of millions of lattice sites. Furthermore, the computational platform as developed is scalable. As demonstrated here, such systems can provide a vital bridge between the current resource discontinuity facing the typical researcher, i.e. individual scientific workstations versus massively parallel supercomputers such as the ASCI-Red supercomputer.

## 10 References

Adler, P.M., *Porous Media: Geometry and Transports*, 544 pp., Butterworth-Heinemann Series in Chemical Engineering, Stoneham, Mass., 1992.

Agrawal, D.L., Cook, N.G.W., Myer, L.R., The effect of percolating structures on the petrophysical properties of Berea sandstone, *Proc. 32nd U.S. Rock Mech. Symp.*, 345-354, 1991.

Armal, B.F., Durst, F., Pereira, and Schonung, B., Experimental and theoretical investigation of backward-facing step flow, *J. Fluid Mech.*, 172, 473-496, 1983.

Auzerais, F.M., Dunsmuir, J., Ferreol, B.B., Martys, N., Olson, J., Ramakrishnan, T.S., Rothman, D.H., and Schwartz, L.M., Transport in sandstone: A study based on three dimensional microtomography, *Geophys. Res. Lett.*, 23, 705-708, 1996.

Avnir, D., Farin, D., and Pfeifer, P., Surface irregularity of particulate materials: A fractal approach, *J. Colloid Interface Sci.*, 103, 112-123, 1984.

Bhatnagar, P., Gross, E., Krook, M., A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems, *Phys. Rev.*, 94, 511, 1954.

Baldwin, C.A., Sederman, A.J., Matntle, M.D., Alexander, P., and Gladden, L.F., Determination and characterization of the structure of a pore space from 3D volume images, *J. Colloid Interface Sci.*, 181, 79-92, 1996.

Bernabe, Y., Pore geometry and pressure dependence of the tranpsort properties in sandstones, *Geophys.*, 56, 436-446, 1991.

Berryman, J.G., Measurement of spatial correlation functions using image processing techniques, *J. Appl. Phys.*, 57, 2374-2384, 1985.

Berryman, J.G. and Blair, S.C., Use of digital image analysis to estimate fluid permeability of porous materials: Application of two-point correlation functions, *J. Appl. Phys.* 60, 1930-1938, 1986.

Blackwell, B.F., and Armaly, B.F., Benchmark Problem definition and summary of computational results for mixed convection over a backward facing step, *ASME HTD-258*, 1-10, 1993.

Bourbie, T. and Zinszner, B., Hydraulic and acoustic properties as a function of porosity in Fontainebeau sandstone, *J. Geophys. Res.*, 90, 11524-11532, 1985.

Brakenhoff, G.J., Blom, P., and Barends, P., Confocal scanning light microscopy with high aperture immersion lenses, *J. Microsc.*, 117, 219-232, 1979.

Brakenhoff, G.J., van der Voort, H.T.M., van Spronsen, E.A., and Nanninga, N., Three-dimensional imaging in fluorescence by confocal scanning microscopy, *J. Microsc.*, 153, 151-159, 1989.

Brown, G.O., Stone, M.L., and Gazin, J.E., Accuracy of gamma ray computerized topography in porous media, *Water Res. Res.*, 29, 479-486, 1993.

Brunauer, S., Emmett, P.H., and Teller, T., The adsorption of gases in multi-molecular layers, *J. Am. Chem. Soc.*, 69, 309-316, 1938.

Chen, H., Chen, S., and Matthaeus, W.H., Recovery of the Navier-Stokes equations using a lattice Boltzmann method, *Phys. Rev. A*, 45, R5339-R5342, 1992.

Chen, S., Chen, H.D., Martinez, D., and Matthaeus, W., Lattice Boltzmann model for simulation of magnetohydrodynamics, *Phys. Rev. Lett.*, 67, 3776-3779, 1991.

Chen, S., and Doolen, G.D., Lattice Boltzmann method for fluid flows, *Ann. Rev. Fluid Mech.*, 30, 329-364, 1998.

Chen, S.Y., Martinez, D., and Mei, R., On boundary conditions in lattice Boltzmann methods, *Phys. Fluids*, 8, 2527-2536, 1996.

Cohen, M.H., and Mendleson, K.S., Nuclear magnetic relaxation and the internal geometry of sedimentary rocks, *J. Appl. Phys.*, 53, 1127, 1982.

Coker, D. A., Torquato, S. and Dunsmuir, J. H., Morphology and physical properties of Fontainebleau sandstone via tomographic analysis, *J. Geophys. Res.*, 101, 17497-17506, 1996.

Cornubert, R., d'Humières, D., and Levermore, D., A Knudsen layer theory for lattice gases, *Physica D*, 47, 241-259, 1991.

Dardis, O., and McCloskey, J., Lattice Boltzmann scheme with real numbered solid density for the simulation of flow in porous media, *Phys. Rev. E*, 57, 4834-4837, 1998.

David, C., Geometry of flow paths for fluid transport in rocks, *J. Geophys. Res.*, 98, 12267-12278, 1993.

Doughty, D.A., and Tomutsa, L., Imaging pore structure and connectivity in high resolution NMR microscopy, *Int. J. Rock Mech. & Min. Sci.*, 34:3-4, 1997.

Doyen, P.M., Permeability, conductivity, and pore geometry of sandstone, *J. Geophys. Res.*, 93, 7729-7740, 1988.

Dullien, F.A.L., *Porous Media – Fluid Transport and Pore Structure*, 2nd ed., Academic Press, New York, New York, 1992.

Fredrich, J.T., Greaves, K.H., and Martin, J.W., Pore geometry and transport properties of Fontainebleau sandstone. *Int. J. Rock Mech. Min. Sci.*, 30, 691-697, 1993.

Fredrich, J.T., Martin, J.W., and Owen, L.B., Application of X-ray computerized tomography for non-destructive measurement of physical properties of NTS tuff, Defense Nuclear Agency, DNA-TR-93-189, 121 pp., Alexandria, Virginia, 1994.

Fredrich, J.T., Menendez, B., and Wong, T.-f., Imaging the pore structure of geomaterials, *Science*, 268, 276-279, 1995.

Fredrich, J.T. and Lindquist, W.B., Statistical characterization of the three-dimensional microgeometry of porous media and correlation with macroscopic transport properties, *Int. J. Rock Mech. & Min. Sci.*, 34(3-4), 1997.

Gard, D.L., Confocal immunofluorescence microscopy of microtubules in amphibian oocytes and eggs, in *Methods in Cell Biology*, 38, 241-264, Academic Press, New York, NY, 1993.

Gartling, D.K., A test problem for outflow boundary conditions – flow over a backward-facing step, *Int. J. Num. Meth. Fluids*, 11, 953-967, 1990.

Ginzbourg, I., and d'Humières, D., Local second-order boundary methods for lattice boltzmann models, *J. Stat. Phys.*, 84, 927-971, 1996.

Gratton, E. and vandeVen, J.J., Laser sources for confocal microscopy, in *Handbook of Biological Confocal Microscopy*, Pawley, J.B. (ed.), 53-67, Plenum Press, New York, NY, 1990.

Gresho, P.M., Gartling, D.K., Torczynski, J.R., Cliffe, K.A., Winters, K.H., Garratt, T.J., Spence, A., and Goodrich, J.W., Is the steady viscous incompressible two-dimensional flow over a backward-facing step at  $Re=800$  stable?, *Int. J. Num. Meth. Fluids*, 17, 501-541, 1993.

Gropp, W.D., and Lusk, E., MPICH 1.1.1: A Portable Implementation of the Message-Passing Interface Standard, [www.mcs.anl.gov/mpi/mpich](http://www.mcs.anl.gov/mpi/mpich), 1997.

Grunau, D., Chen, S., and Eggert, K., A Lattice Boltzmann model for multiphase fluid flows, *Phys. Fluids A*, 5, 2557-2562, 1993.

Gunstensen, A.K., Rothman, D.H., Zaleski, S., and Zanetti, G., Lattice Boltzmann model of immiscible fluids, *Phys. Rev. A*, 43, 4320-4327, 1991.

Hall, P.L., Mildner, D.F.R., Borst, R.L., Small-angle scattering studies of the pore space of shaly rocks, *J. Geophys. Res.*, 91, 2183-2192, 1986.

He, X., and Doolen, G., Lattice Boltzmann method on curvilinear coordinates system: flow around a circular cylinder, *J. Comput. Phys.*, 134, 306-315.

He, X., Luo, L.-S., and Dembo, M., Some progress in lattice Boltzmann method. part i. nonuniform mesh grids, *J. Comput. Phys.*, 129, 357-363, 1996.

Higuera, F., and Jiménez, J., Boltzmann approach to lattice gas simulations, *Europhys. Lett.*, 9, 663-668, 1989.

Hou, S., Zou, Q., Chen, S., Doolen, G.D., and Cogley, A.C., Simulation of cavity flow by the lattice Boltzmann method, *J. Comput. Phys.*, 118, 329-347, 1995.

Inoue, S., *Video Microscopy*, Plenum Press, New York, NY 1986.

Inoue, S., Foundations of confocal scanned imaging in light microscopy, in *Handbook of Biological Confocal Microscopy*, Pawley, J.B. (ed.), 1-14, Plenum Press, New York, NY, 1990.

Kim, J., and Moin, P., Application of a fractional-step method to incompressible Navier-Stokes equations, *J. Comput. Phys.*, 59, 308-323, 1985.

Koplik, J., Lin, C., and Vermette, M., Conductivity and permeability from microgeometry, *J. Appl. Phys.*, 56, 3127-3131, 1984.

Krohn, C.E. and Thompson, A.H., Fractal sandstone pores: Automated measurements using scanning-electron-microscope images, *Phys. Rev. B*, 33, 6366-6374, 1986.

Krohn, C.E., Sandstone fractal and Euclidean pore volume distributions, *J. Geophys. Res.*, 93, 3286-3296, 1988.

Lebon, L., Oger, L., Leblond, J., Hulin, J.P. Martys, N.S., Schwartz, L.M., Pulsed gradient NMR measurements and numerical simulation of flow velocity distribution in sphere packings, *Phys. Fluids*, 6, 293-301, 1996.

Lebon, L., Leblond, J., Hulin, J.-P., Martys, N.S., and Schwartz, L.M., Pulsed field gradient NMR measurements of probability distribution of displacement under flow in sphere packings, *Mag. Reson. Imag.*, 14, 989-991, 1996.

Lee, T.-C., Kasyap, R. L. and Chu, C. -N. , Building skeleton models via 3-D medial surface/axis thinning algorithms, *CVGIP: Graphical Models and Image Processing*, 56, 462-478, 1994.

Lin, C., Pirie, G., and Trimmer, D.A., Low permeability rocks: Laboratory measurements and three-dimensional microstructural analysis, *J. Geophys. Res.*, 91, 2173-2181, 1986.

Lindquist, W.B., Lee, S.-M., Coker, D.A., Jones, K.W., and Spanne, P., Medial axis analysis of void structure in three-dimensional tomographic images of porous media. *J. Geophys. Res.*, 101, 8297-8310, 1996.

Lindquist, W.B., and Venkatarangan, A., Investigating 3d geometry of porous media from high resolution images, *Phys. Chem. Earth*, (this issue), 1999.

Marschall, J. and Milos, F.S., Gas permeability of rigid fibrous refractory insulations, *Proc. 32nd Thermophysics Conf.*, Amer. Inst. Aero. and Astro., Reston, VA, 1997.

Mardia, K.V., Hainsworth T.J., A spatial thresholding method for image segmentation, *IEEE Trans. Pattern Anal. Machine Intell.*, 10, 919-927, 1988.

Martinez, D.O., Matthaeus, W.H., Chen, S., Montgomery, D.C., Comparison of spectral method and lattice Boltzmann simulations of two-dimensional hydrodynamics, *Phys. Fluids*, 6, 1285-1298, 1994.

McNamara, G., and Zanetti, G., Use of the Boltzmann equation to simulate lattice-gas automata, *Phys. Rev. Lett.*, 61, 2332-2335, 1988.

Mendelson, K.S., Nuclear magnetic resonance in sedimentary rocks: Effect of proton desorption rate, *J. Appl. Phys.*, 53, 6465-6466, 1982.

Montoto, M., Martinez-Nistal, A., Rodriguez-Rey, A., Fernandez-Merayo, N., and Soriano, P., Microfractography of granitic rocks under confocal scanning laser microscopy, *J. Microsc.* 177, 138-149, 1995.

MPI Forum, MPI: A Message-Passing Interface Standard, Version 1.1, MPI Forum Technical Report, [www mpi-forum.org](http://www mpi-forum.org), 1995.

Myer, L. R., Kemeny, J.M., Zheng, Z., Suarez, R., Ewy, R.T., and Cook, N.G.W., Extensile cracking in porous rock under differential compressive stress, in *Micromechanical Modelling of Quasi-Brittle Materials Behavior*, V. C. Li (ed.), *Appl. Mech. Rev.*, 45, 263-280, 1992.

Nannelli, F., and Succi, S., The lattice Boltzmann equation on irregular lattices, *J. Stat. Phys.*, 68, 401-407, 1992.

Noble, D.R., Chen, S., Georgiadis, J.G., and Buckius, R.O., A consistent hydrodynamic boundary condition for the lattice Boltzmann method, *Phys. Fluids*, 7, 203-209, 1995.

Noble, D.R., Georgiadis, J.G., and Buckius, R.O., Comparison of accuracy and performance for lattice Boltzmann and finite difference simulations of steady viscous flow, *Int. J. Num. Meth. Fluids*, 23, 1-18, 1996.

Noble, D.R. Lattice Boltzmann study of the interstitial hydrodynamics and dispersion in steady inertial flows in large randomly packed beds, Ph. D. Dissertation, University of Illinois at Urbana-Champaign, 1997.

Noble, D.R., Georgiadis, J.G., and Buckius, R.O., Effect of disorder on dispersion in steady inertial flows in large randomly packed beds, *in preparation*, 1998.

Oh, W. and Lindquist, W.B., Image thresholding by indicator kriging, *IEEE Trans. Pattern Anal. Machine Intell.*, submitted, 1998.

Pawley, J.B., *Handbook of Biological Confocal Microscopy*, Plenum Press, New York, 1990.

Pittman, E.D. and Duschatko, R.W., Use of pore casts and scanning electron microscopy to study pore geometry, *J. Sed. Petr.*, 40, 1153-1157, 1970.

Qian, Y. and d'Humieres, D., and Lallemand, P., Lattice BGK models for the Navier-Stokes equation, *Europhys. Lett.*, 17, 479, 1992.

Qian, Y.H., Lattice gas and lattice kinetic theory applied to the Navier-Stokes equations, Ph. D. Dissertation, Université Pierre et Marie Curie, 1990.

Reider, M.B. and Sterling, J.D., Accuracy of discrete-velocity BGK models for the simulation of the incompressible Navier-Stokes equations, *Computers & Fluids*, 24, 459-467, 1995.

Russ, J.C., *Practical Stereology*, 186 pp., Plenum Press, New York, New York, 1986.

Sahimi, M., *Flow and Transport in Porous Media and Fractured Rock*, 482 pp., VCH, Weinheim, FRG, 1995.

Shan, X., Simulation of Rayleigh-benard convection using a lattice Boltzmann method, *Phys. Rev. E*, 55, 2780-2788, 1997.

Skordos, P.A., 1993, Initial and boundary conditions for the lattice Boltzmann method, *Phys. Rev. E*, 48, 4823-4842, 1993.

Spanne, P., Thovert, J.F., C.J. Jacquin, Lindquist, W.B., Jones, K.W., and Adler, P.M., Sychrotron computed microtomography of porous media: Topology and transports, *Phys. Rev. Lett.*, 73, 2001-2004, 1994.

Spurr, A.R., A low viscosity epoxy resin embedding medium for electron microscopy, *J. Ultrastructure Res.*, 26, 31-43, 1969.26, 31, 1969.

Thompson, A.H., Fractals in rock physics, *Ann. Rev. Earth Planet. Sci.*, 19, 237-262, 1991.

Thovert, J.F., Salles, J., and Adler, P.M., Computerized characterization of the geometry of real porous media: their discretization, analysis and interpretation, *J. Microsc.*, 170, 65-79, 1993.

Torczynski, J. R., Numerical solutions for a flow with mixed convection in a vertical geometry, *ASME HTD-258*, 105-113, 1993.

Tullis, J. and Yund, R., The brittle-ductile transition in feldspar aggregates: An experimental study, in *Fault Mechanics and Transport Properties of Rocks*, B. Evans and T.-F. Wong (eds.), 89-117, Academic Press, San Diego, CA, 1992.

Underwood, E.E., *Quantitative Stereology*, 274 pp., Addison Wesley, Reading, Massachusetts, 1970.

Van Brakel, J., Modry, S., and Svata, M., Mercury porosimetry: State of the art, *Powder Technol.*, 29, 1-12, 1981.

Walsh, J. B. and Brace, W. F., The effect of pressure on porosity and the transport properties of rock, *J. Geophys. Res.*, 89, 9425-943, 1984.

Wellington, S.L., and Vinegar, J.J., X-ray computed tomography, *J. Pet. Tech.*, 39, 885-898, 1987.

Wilson, T., Optical sectioning in confocal fluorescent microscopes, *J. Microsc.*, 154, 143-156, 1989.

Wilson, T., The role of the pinhole in confocal imaging systems, in *Handbook of Biological Confocal Microscopy*, Pawley, J.B. (ed.), 113-126, Plenum Press, New York, NY, 1990.

Wong, P.-Z., Howard, J., and Lin, J.-S., Surface roughening and the fracture nature of rocks, *Phys. Rev. Lett.*, 57, 637-640, 1986.

Wong, T.-F., Geometric probability approach to the characterization and analysis of microcracking in rocks, *Mech. Mat.*, 4, 261-276, 1985.

Wong, T.-F., Fredrich, J.T., and Gwanmesia, G.D., Crack aperture statistics and pore space fractal geometry of Westerly granite and Rutland quartzite: Implications for an elastic contact model of rock compressibility, *J. Geophys. Res.*, 94, 10,267-10,278, 1989.

Yale, D.P., Recent advances in rock physics, *Geophys.*, 50, 2480-2491, 1985.

Zinszner, B., Johnson, P.A., Rasolofosaon, P.N.J., Influence of change in physical state on elastic nonlinear response in rock: Significance of effective pressure and water saturation, *J. Geophys. Res.*, 102, 8105-8120, 1997.

Zhu, W., David, C., and Wong, T.-F., Network modeling of permeability evolution during cementation and hot isostatic pressing, *J. Geophys. Res.*, 100, 15451-15464, 1995.

Zou, Q., Hou, S., Chen, S., Doolen, G.D., An improved incompressible lattice Boltzmann model for time-independent flows, *J. Stat. Phys.*, 81, 35-48, 1995.

## **APPENDIX A**

**List of refereed publications and presentations**

## Refereed Publications

Fredrich, J.T., and Lindquist, W.B., Statistical characterization of the three-dimensional microgeometry of porous media and correlation with macroscopic transport properties, *Int. J. Rock. Mech. Min. Sci.*, v. 34:3-4, 1997.

Fredrich, J.T., 3D imaging of porous media using laser scanning confocal microscopy with application to microscale transport processes, *Phys. Chem. Earth*, 1999 (in press).

Lindquist, W.B., and Venkatarangan, A., Investigating 3d geometry of porous media from high resolution images, *Phys. Chem. Earth*, 1999 (in press).

Oh, W. and Lindquist, W.B., Image thresholding by indicator kriging, *IEEE Trans. Pattern Anal. Machine Intell.*, submitted, 1998.

Noble, D.R., and Torczynski, J.R. A lattice Boltzmann method for partially saturated computational cells, in press, *Int. J. Modern Phys. C*, 1999.

Noble, D.R., Mixed convection over a backward-facing step in a vertical channel by the lattice Boltzmann Method, *ASME Proc.*, HTD 357-2, p. 167, 1998.

O'Connor, R.M., and Fredrich, J.T., Microscale flow modeling in geologic materials, *Phys. Chem. Earth*, 1999 (in press).

O'Connor, R.M., and Fredrich, J.T., New approaches to distributed computing for the geophysical sciences, to be submitted to *Eos, Trans. Amer. Geophys. Un.*, 1999.

Stockman et al., Practical application of lattice-gas and lattice Boltzmann methods to dispersion problems, *InterJournal Complex Systems*, 1998.

Stockman, H.W., Accuracy and computational efficiency in 3d dispersion via lattice Boltzmann: models for double-diffusive fingering and dispersion in rough fractures, submitted, *Int. J. Modern Phys. C*, 1998.

## Presentations

Fredrich, J.T., and Lindquist, W.B., Statistical characterization of the 3D microgeometry of porous media, Amer. Geophys. Un., San Francisco CA, December 15-19 1996.

Fredrich, J.T., Three-dimensional imaging of porous media, Porous Media Processes Workshop, Los Alamos National Laboratory, April 2-4 1997.

Fredrich, J.T., and Lindquist, W.B., Statistical characterization of the 3D microgeometry of porous media and correlation with macroscopic transport properties, 36th U.S. Rock Mechanics Symp., New York NY, June 29- July 2 1997.

Fredrich, J.T., 3-D imaging of porous media & application to microscale modeling of transport properties, Invited lecture, Fluid Flow and Seismology Lecture Series, Los Alamos National Laboratory, September 18 1997.

Fredrich, J.T., Geometry and fluid flow processes in porous media, Invited lecture, Dept. of Geosciences, State University of New York, Stony Brook, February 9 1998.

Fredrich, J.T., 3D imaging of porous media and application to microscale modeling of transport, Europ.Geophys. Soc. XXIII General Assembly, Invited keynote talk, Symp. on Imaging, Analyzing, and Modeling Pore Structure in Geomaterials, Nice France, April 20-24 1998. (Abstract published in *Annales Geophysicae*, vol. 16, p. c254)

Fredrich, J.T., Fluid flow and microgeometry of porous media, Invited lecture at the GeoForschungsZentrum, Berlin Germany, April 30 1998.

Fredrich, J.T., Microgeometry of porous media and fluid flow, Invited lecture at the University of Edinburgh, Scotland, July 3 1998.

Fredrich, J.T., and O'Connor, R.M., Microscale Flow Modeling in Reconstructed Porous Media, Amer. Geophys. Un., San Francisco CA, December 6-10 1998.

Fredrich, J.T., Microscale Flow Modeling in Reconstructed Porous Media, Invited lecture at NM Tech, Socorro NM, February 4 1999.

Holdych, D., and Noble, D.R., Lattice Boltzmann methods for second-order-accurate simulation of initial value problems, 1998 Am. Phys. Soc., November 1998.

Lindquist, W.B., Medial axis analysis of porous media, Porous Media Processes Workshop, Los Alamos National Laboratory, April 2-4 1997.

Lindquist, W.B., Investigating three dimensional geometry of porous media from high resolution images, Europ. Geophys. Soc. XXIII General Assembly, Symp. on Imaging, Analyzing, and Modeling Pore Structure in Geomaterials, Nice France, April 20-24 1998. (Abstract published in *Annales Geophysicae*, vol. 16, p. c255).

Marschall, J., F. Milos, and J.T. Fredrich, Application of laser scanning confocal microscopy to heat and mass transport modeling in porous microstructures, National Heat Transfer Conf., DOE/NSF Technical Workshop on Thermophysical Phenomena in Microscale Sensors, Devices, and Structures, August 1997, Baltimore MD.

Noble, D.R., Evaluation of the Brinkman extension for randomly packed beds confined in a channel, Am. Phys. Soc., November 1996.

Noble, Mixed convection over a backward-facing step in a vertical channel by the lattice Boltzmann method, Joint ASME/AIAA Thermophysics Conf., Albuquerque NM, June 1998.

Noble, D.R. and Torczynski, J.R. A lattice Boltzmann method for partially saturated computational cells, 7th International Conference on the Discrete Simulation of Fluids, Oxford U.K., July 1998.

O'Connor, R.M., and Fredrich, J.T., Araneæ- A scalable system for distributed computing, Scalable Cluster Workshop, Monterey CA, November 21-22 1997.

O'Connor, R.M., Fredrich, J.T., and Stockman, H.W., Microscale flow modeling in geologic materials, Europ. Geophys. Soc., XXIII General Assembly, Symp. on Imaging, Analyzing, and Modeling Pore Structure in Geomaterials, Nice France, April 20-24 1998. (Abstract published in *Annales Geophysicae*, vol. 16, p. c255)

Stockman, H.W., Accuracy and computational efficiency in 3d dispersion via lattice Boltzmann: models for double-diffusive fingering and dispersion in rough fractures (invited), 7th Int. Conf. on the Discrete Simulation of Fluids, July 1998.

Stockman, Practical application of lattice Boltzmann and LGA to dispersion, Porous Media Processes Workshop, Los Alamos National Laboratory, April 2-4 1997.

Distribution:

MS0149 C.E. Meyers, 4000 (2)  
MS0701 J.W. Woodard, 6000  
MS0750 L.E. Shephard, 6100  
MS0735 P.B. Davies, 6115  
MS0735 S.A. McKenna, 6115  
MS0735 R.J. Glass, 6115  
MS0735 V.C. Tidwell, 6115  
MS0751 J.T. Fredrich, 6117 (15)  
MS0751 R.M. O'Connor, 6117 (5)  
MS0751 D.J. Holcomb, 6117  
MS0751 D.S. Preece, 6117  
MS0751 L.S. Costin, 6117  
MS0750 H.S. Westrich, 6118  
MS0750 H.W. Stockman  
MS0751 6117 File (5)  
MS9405 M.T. Dyer, 8700  
MS9402 C.M. Hartwig, 8701 (2)  
MS9405 J.M. Hruby, 8230  
MS9054 W.J. McLean, 8300  
MS9041 J.S. Binkley, 8345  
MS9042 S.K. Griffith, 8345  
MS9042 R.H. Nilson, 8345  
MS9056 L.A. Rahn, 8351  
MS9051 W.T. Ashurst, 8351  
MS0151 G. Yonas, 9000  
MS0337 R.D. Skocypec, 9002  
MS0841 P.J. Hommert, 9100  
MS0828 T.C. Bickel, 9101  
MS0828 R.K. Thomas, 9104  
MS0826 D.R. Noble, 9111 (15)  
MS0826 C.E. Hickox, 9111  
MS0826 W.L. Hermina, 9111  
MS0834 A.C. Ratzel, 9112  
MS0834 J.R. Torczynski, 9112  
MS0835 S.N. Kempka, 9113  
MS0836 C.W. Peterson, 9116  
MS0443 A.F. Fossum, 9117  
MS0443 H.S. Morgan, 9117  
MS0828 J.H. Biffle 9121  
MS0321 W.J. Camp, 9200  
MS1110 D.E. Womble, 9222  
MS1111 G.S. Heffelfinger, 9225  
MS1435 K.A. Hays, 1800  
MS1407 J.H. Aubert, 1815

MS0367 B.K. Damkroger, 1833  
MS0367 J.E. Glass, 1833  
MS1411 J.M. Phillips, 1834  
MS1349 W.F. Hammetter, 1846  
MS9018 Central Technical Files, 8940 (1)  
MS0899 Technical Library, 4916 (2)  
MS0619 Review and Approval Desk, 15102 (1)  
for DOE/OSTI