J. Fluid Mech. (1997), *vol.* 341, *pp.* 343–370. Printed in the United Kingdom

Two-fluid flow in sedimentary rock: simulation, transport and complexity

By JOHN F. OLSON AND DANIEL H. ROTHMAN

Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

(Received 25 March 1996 and in revised form 11 February 1997)

The macroscopic properties and structure of the flow of two immiscible fluids through Fontainebleau sandstone are studied by numerical simulation. The pore space geometry was obtained by X-ray microtomography (Kinney *et al.* 1993) and the numerical simulations were performed by a new lattice-gas cellular automaton method (Olson & Rothman 1995). We first validate the numerical method by showing that the drag on a cubic array of spherical drops matches theoretical predictions. As a further test, we present a comparison between computed relative permeability and experimental measurements on the same rock. We then present a study of fluid–fluid coupling; we find that it is significant, and that it appears to be reciprocal: the flux of one fluid due to forcing on the other is the same, regardless of which fluid is forced. Lastly, we characterize the complexity and organization of the flow by means of a statistical parameter, the *skewness* of the distribution of local velocities.

1. Introduction

The flow of two-phase mixtures in porous rock is a subject of considerable economic and scientific importance. For example, an oil reservoir typically contains at least two and usually three distinct fluid phases, oil, brine, and gas, and the transport of these species through the rock is of obvious interest. Of equal importance is two-phase flow in aquifers, where the economically important species is fresh water and any other phase which is present is likely to be a harmful pollutant. In the broader context of porous media, industrial processes may require two-phase mixtures to flow through heterogeneous catalysts, mixing drums, boilers or condensers. In yet broader terms, blood is a multiphase fluid mixture that percolates through a complex porous medium called flesh and bone. A better understanding of two-phase flow through sedimentary rock is thus of considerable practical value in its own right, and as an example of the larger class of multiphase flows through porous media.

Although flow through porous rock has been studied experimentally at least since the nineteenth century (Darcy 1956), numerical simulations can provide important information that is not easily accessible to experiments (van Genabeek & Rothman 1996). For example, one can measure the average flux through a sample of rock in the laboratory, but it is difficult to measure the local velocity in some pore in the rock. In a calculation, on the other hand, the velocity at every point is readily accessible. Such detailed information, appropriately analysed, may provide insight which can be tested in experiments. Simulation also affords more control than can be had in experiment: for example, in the present research it is straightforward to apply a body force to one fluid in a mixture and thus to probe the phenomenon of fluid–fluid coupling in porous

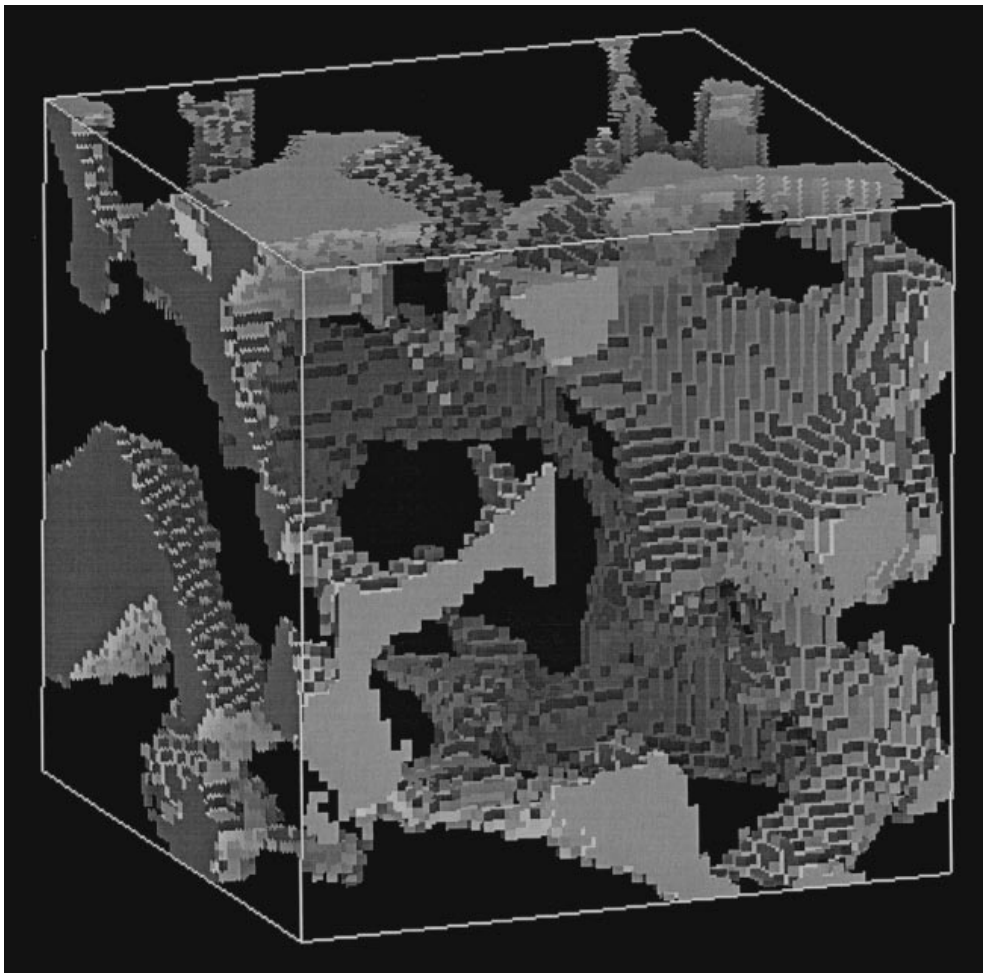

FIGURE 1. The pore space in Fontainebleau sandstone as determined by X-ray microtomography (Kinney *et al.* 1993). The data are discretized to cubes of side length 7.5 μm; this figure shows the surface of the void space in a cube 0.78 mm on a side.

rock, while the same study in the laboratory requires careful and elaborate preparations (Zarcone & Lenormand 1994; Zarcone 1994; Avram & Payatakes 1995). We thus seek a means to compute two-phase flow in a porous medium.

The problem would appear to be simple: we know the microscopic (pore scale) motion will be governed by Navier–Stokes equations with pressure-jump boundary conditions at the interfaces. But the pore space in sedimentary rock is enormously complex, as shown in figure 1. Not only is the geometry complicated, but for two-phase flows the interfaces are likely to break and merge, and certain to move, all of which would make a conventional numerical solution difficult. We therefore developed a new numerical method, a *lattice gas cellular automaton* (Olson & Rothman 1995; Frisch, Hasslacher & Pomeau 1986; Rothman & Zaleski 1994, 1997), for this work. The principle advantage of a lattice gas for this work is that interfaces, complete with pressure-jump boundary conditions, emerge from simple, microscopic rules that are applied uniformly throughout the simulation volume. This is much simpler than having to follow the interfaces and match boundary conditions explicitly.

The new method has been characterized in some detail, including demonstrations that phase separation is approximately isotropic (Appert *et al.* 1995) and that sheared interfaces transfer momentum correctly (Olson & Rothman 1995). In this paper, we present a further validation by computing the drag on fluid spheres and showing that our results are in good agreement with theory. We also present a preliminary comparison of relative permeability computed with the method to that measured on the same rock in the laboratory, and find good agreement. We then measure the macroscopic transport properties of the flow, the relative permeability of each species and the coupling between the fluids under a variety of conditions, and examine the details of the flow field, noting and characterizing its structure.

In studies of two-phase flow through porous rock, it is often assumed that the fluids do not couple to one another, that is, that forcing one fluid does not cause the other fluid to move. This assumption has been addressed by experiments (Kalaydjian 1994; Mannseth 1991; Bentsen & Manai 1993; Zarcone & Lenormand 1994; Zarcone 1994; Bentsen 1994; Avram & Payatakes 1995), theory (de Gennes 1983, Bacri, Chaouche & Salin 1990; S. Pride & E. Flekkøy 1996, personal communication) and simulation (Gunstensen & Rothman 1993), but remains controversial. In this paper, we will test the hypothesis that two-phase flow can be described by a simple extension of Darcy's law,

$$\begin{pmatrix} u_w \\ u_n \end{pmatrix} = k \begin{pmatrix} \kappa_{ww} & \kappa_{wn} \\ \kappa_{nw} & \kappa_{nn} \end{pmatrix} \begin{pmatrix} F_w \\ F_n \end{pmatrix} \tag{1.1}$$

where u_i is the flux of species i, the subscripts w and n corresponding to wetting and non-wetting fluid, respectively, F_i is the force on each species, k is the permeability of the medium, and the transport coefficients κ_{ij} are the relative permeabilities (diagonal terms) and coupling coefficients (off-diagonal terms). In general these κ_{ij} will depend on the viscosity of each fluid and the composition of the mixture. This is simply the most general linear force–flux relation, and it would not be surprising to find variation from it; but it serves as a basis for discussion. Earlier simulations (Gunstensen & Rothman 1993) in an artificial medium found significant values for the coupling coefficients, a forcing threshold below which no flow occurred, and some nonlinear dependence of the fluxes on the forcing. In the present work, we will again find significant coupling and a threshold forcing, but we do not study the nonlinear response above this threshold.

The present work is the first attempt either to measure or compute both coupling coefficients κ_{wn} and κ_{nw} in a real rock geometry. In addition, of those studies which have measured fluid–fluid coupling, only Bentsen (1994) makes any attempt to measure both coupling coefficients. A summary of the published research on fluid–fluid coupling in porous media is given in table 1. In our work, we computed both coefficients independently and found them to be equal within the uncertainty of the calculation.

Because transport properties are a macroscopic, average description of the flow, we would like to ask why they take the values they do. The answer must lie in the detailed velocity field in the rock pore space, but how can we characterize this velocity field in a comprehensible way? If we construe the velocity field as a distribution of local velocities, we can compute statistical properties – moments – of this distribution. We show that the third moment or *skewness* of the velocity field quantifies the degree of organization of the flow.

In the sections to follow, we first briefly introduce the lattice gas method and add to the already extensive literature of validations of such methods (for example,

κ	R?	Medium and fluids	k (μm^2)	m	Φ	pCa	Ref.
0.15	No	Sand pack, water and oil	20	30	0.35	7	Bentsen & Manai 1993
<0.01	A	Vosges sandstone, water and oil	0.12	11	0.233	8–7	Kalaydjian 1990
<0.003	N/A	Sand pack, water and mercury	34	1.5	0.36	7	Zarcone & Lenormand 1994
0 (< 10^{-4})	N/A	Sand pack, water and dodecane	34	1.7	0.36	7–6	Zarcone 1994
<0.02	A	Simulated tube network	N/A	1	N/A	N/A	Goode & Ramakrishnan 1993
<0.1	Yes	'Penetrable sphere' medium	large	1	0.65	0	Gunstensen & Rothman 1993
0.015	Yes	Sim. Fontainebleau sandstone	1.1	1	0.15	5	⋆

TABLE 1. Coupling coefficients in two-fluid flow through porous media. Experimental measures (top of table) and simulations (bottom) have been performed to deduce the importance of fluid–fluid coupling in two-fluid flow through porous media. The coupling coefficient κ (corresponding to the off-diagonal terms in equation (1.1) is dimensionless. The column labelled "R?" indicates the reciprocity of coupling terms, as found in each study; some researchers *assumed* the coupling coefficients would be equal, so these studies are labelled 'A.' k is the permeability of the medium, m is the ratio of dynamic viscosities (μ_n/μ_w) (the *mobility coefficient*) and Φ is the porosity of the medium. pCa is the negative, base 10 logarithm of the capillary number (analogous to pH), with the capillary number defined as $\mu_w u/\gamma$, where γ is the surface tension between the two fluids and u is the volume flux of the forced fluid. The row marked with '⋆' refers to this work. 'N/A' means 'not applicable' or 'not available.'

Rothman & Zaleski 1994; Olson & Rothman 1995; Appert *et al.* 1995) by showing that it correctly computes the drag on fluid spheres. We then validate the computation of flow through a digitized pore space by simulating the non-wetting invasion experiment by which relative permeability is measured in the laboratory and comparing the computed relative permeability to laboratory measurements on the same rock. We will find that a modified simulation is more effective than direct simulation for computing transport properties at the limited spatial resolution practical with current computers. We then compute macroscopic transport properties for a range of saturations and of forcings, by which means we can compute directly the fluid–fluid coupling. We then examine the flow field in detail, and show pictorially that the flow of pure fluid is more highly organized than the flow of a mixture, and that the flow of a mixture when the non-wetting fluid is forced is more organized than the flow of a mixture when the wetting fluid is forced. Inspired by the extensive literature on statistical measures of turbulence (see, for example, Tennekes & Lumley 1972), we compute the moments of the velocity fields, and see that the skewness indicates the degree of organization of the flow. Finally, we state our conclusions and indicate directions for future work, both with the computer and in the laboratory.

2. Simulated flow through Fontainebleau sandstone

Recent advances in X-ray microtomography have made it possible to image the three-dimensional pore space in sandstone at high resolution over macroscopic volumes. The complex geometry of the pore space can now be used to compute quantities

which could formerly be determined only by experiments. We have sought to simulate the flow of two-fluid mixtures through this pore space by means of a lattice-gas cellular automaton. Since the model used for this research (Olson & Rothman 1995) and a review (Rothman & Zaleski 1994) of lattice gases have been recently published elsewhere, we give only a brief outline here.

A lattice gas consists of particles (represented as bits in a computer memory) to which are assigned mass and momentum, and which are restricted to exist only at the nodes of a regular lattice in space and only at regular intervals in time. Particles move from node to node in the directions of their momenta; when particles meet at a node, they collide, conserving the total momentum and mass at the node. If a lattice with appropriate symmetry is chosen, the coarse-grained average density ρ and momentum $\rho\boldsymbol{u}$ satisfy equations like the Navier–Stokes equations:

$$\nabla\cdot\boldsymbol{u}=0 \quad \text{and} \quad \rho\frac{\partial\boldsymbol{u}}{\partial t}+\rho\, g(\rho)\,\boldsymbol{u}\cdot\nabla\boldsymbol{u}=-\nabla P+\eta\nabla^2\boldsymbol{u}, \tag{2.1}$$

where P is the pressure and η the dynamic viscosity. Ideally, the factor $g(\rho)$ should be unity, but for the present research fluid velocities will be small so it is of little relevance; in fact, we will choose a density so that $g(\rho)$ is identically zero and we model pure Stokes flow.

2.1. *The two-fluid model*

Each particle is labelled with a *colour,* either red or blue. The particles carry their colour as they move about the lattice, and the collisions preserve the number of particles of each colour at each site, as well as the total momentum. Since the collision scrambles particle directions, it is not clear how to assign a colour to each particle after a collision. We must thus choose a rule for recolouring the particles, but we are free to choose any rule that generates interesting dynamics and preserves the numbers of each species.

We compute a colour field ϕ which is the difference between the numbers of red and blue particles at each node,

$$\phi(\boldsymbol{x},t)=n_{\mathrm{red}}(\boldsymbol{x},t)-n_{\mathrm{blue}}(\boldsymbol{x},t). \tag{2.2}$$

We then compute the discrete gradient of this field,

$$\boldsymbol{f}(\boldsymbol{x},t)=\sum_i \boldsymbol{c}_i\phi(\boldsymbol{x}+\boldsymbol{c}_i,t) \tag{2.3}$$

where $\boldsymbol{c}_i$ is the vector which points to the ith neighbouring node. We then use this gradient both to guide the recolouring and to generate surface tension in a plane normal to the gradient. Surface tension is generated by reorienting particles at each node so that as many as possible align with the gradient, without changing the total momentum at the node. This is accomplished by reorienting oppositely directed pairs of particles, or *dumbbells*, which have no net momentum and so can be rearranged at will. After reorienting the dumbbells, the recolouring rule assigns red particles to directions mostly parallel to the gradient and blue particles antiparallel, so that in the following propagation step the red particles move toward nodes with higher concentrations of red, and blue towards higher concentration of blue. This rule causes a mixture of species to separate, and to remain separate. The model thus has fluid dynamic behaviour from the momentum- and particle-conserving collision rules as well as surface tension from the dumbbell rules and phase separation from the recolouring.

The surface tension has been computed from first principles in related *two*-dimensional two-fluid models (Adler, d'Humières & Rothman 1994), but no such theory has been carried out for the newer three-dimensional model yet; instead, the effective surface tension has been measured from simulations (Olson & Rothman 1995), and the non-equilibrium growth due to phase separation has been investigated in detail (Appert *et al.* 1995). In addition, the influence of interfaces on momentum transport was found to match theoretical predictions qualitatively (Olson & Rothman 1995) in the complex flow of a phase-separating mixture undergoing shear flow. We will show below that the model correctly computes the drag on spherical drops, which demonstrates that momentum is correctly transferred through the interface in all directions.

2.2. *Flow through digital rock*

Once we have shown that the lattice gas has sensible fluid and interfacial behaviour on its own, we simulate flow through the complex digitized pore space of the sandstone. A recently developed tomographic technique (Kinney *et al.* 1993) resolves rock structure in three dimensions at a scale of 7.5 μm over a parallelepiped roughly 2 mm on a side, containing nearly 19.5 million volume cells, or *voxels*. A cubic subset of these data, 64 voxels on a side, was used for most of the simulations reported here. This volume was large enough to contain several grains and a few channels, but small enough for practical calculations on workstations. Each voxel was mapped to a single node in the lattice gas, so the lattice spacing can be identified with the voxel scale, 7.5 μm. Solid voxels became *wall sites*, and empty voxels became void sites. The pore space was arranged to be periodic in the vertical direction by doubling the size of the simulation in the vertical direction and filling the upper half of the simulation volume with a vertically reflected copy of the original pore volume; the pore space was also jacketed by vertical walls. This periodic medium was employed in order to avoid end effects. The forcing is a body force, created by adding momentum to randomly selected sites in the lattice at each time step (Olson 1995); when one species is forced, the momentum is added to sites which contain only that species. No-slip conditions at the walls are simulated by a simple bounce-back rule: the collision rule at wall sites does not conserve momentum, but rather simply replaces each particle with a particle of the same colour travelling in the opposite direction. The effect of this is to create a layer between the wall node and its void neighbors where the velocity is zero on average. In fact, the proper way to model solid walls is a question of some subtlety (see, for example, Cornubert, d'Humières & Levermore 1991; Skordos 1993; Ginsbourg & Adler 1994; Ginsbourg & d'Humières 1996), but this simple rule suffices for the present work. In addition to the no-slip rule at the walls, the colour field ϕ used in equation (2.3) is fixed at a particular value for the entire simulation. We can then use the colour field at wall sites to compute the gradient and apply the surface tension rule at void sites adjacent to the walls; in this way, the fluid–wall interfaces also have a surface tension. If the wall colour field were set to zero, then both species would have the same surface tension at the wall, and neither species would be preferentially attracted to the wall. If the colour field is negative, then the blue fluid has a weaker surface tension on the wall, and the red fluid stronger; in this case, the red fluid avoids the wall and is called non-wetting, while the blue fluid spreads along the wall, and is called wetting.

The wetting layer may become very thin in flow through a porous medium, particularly when the concentration of the wetting species is low. Precisely this

situation arises in laboratory measurements of relative permeability, so we would like to be able to simulate it. However, because we model fluids by lattice-gas particles, at some low concentration the wetting layer must necessarily break up into disconnected particles adsorbed on the wall. It turns out that we can make even this case physically relevant by judicious selection of our model parameters. In the next section, we compare our model with both theoretical and experimental results, and briefly describe the method we use to select our model parameters.

3. Validation

We have performed a variety of simulations to demonstrate that the lattice gas is a valid model of two-fluid flows, some of which are reported elsewhere (Olson & Rothman 1995; Appert *et al.* 1995); below, we add to this literature a demonstration that the lattice-gas method correctly computes the drag on fluid spheres. After that, we compare the calculated transport properties of the simulated fluid through the porous rock geometry with experimental measurements, and show that the 'digital permeameter' developed here can compute experimental quantities with some accuracy. Thus, in this section we compare the new lattice gas simulation method with both theory and experiment.

3.1. *Drag on an array of spheres*

A fluid sphere moving through another fluid has complex interior and exterior motions, controlled by the transfer of momentum through and along the fluid–fluid interface. Computing the drag on a fluid drop is thus a sensitive test of the interface dynamics.

When a spherical drop falls through another fluid under the influence of gravity g, and if both fluids have the same viscosity μ, then the terminal velocity of the sphere – the speed at which the drag force balances the buoyancy force – is given by Stokes' law for liquid drops (Lamb 1932)

$$U_{\text{Stokes}} = \frac{4}{15}\frac{g\Delta\rho}{\mu}a^2, \tag{3.1}$$

where $\Delta\rho$ is the density difference between the two fluids and a is the radius of the drop. This equation is valid for a single, spherical drop falling slowly through an infinite extent of a second fluid, or more practically, in a container much larger than the drop. In order to test this theory, we would need to simulate a small drop – subject to large Brownian fluctuations, so requiring a great deal of averaging – in a large volume of fluid, which would require a long time to come to a steady state. This would be a costly calculation. Instead, we could place the sphere in a small volume of fluid with periodic boundary conditions, and thus simulate an infinite cubic array of identical spheres. Sangani (1987) deduced the drag per sphere for such a system. According to this theory, the actual terminal velocity of each sphere will be the Stokes velocity given in equation (3.1) multiplied by a correction factor ζ which depends on the volume fraction Φ occupied by the moving spheres, and which must be computed numerically:

$$U_{\text{actual}} = \zeta(\Phi)U_{\text{Stokes}}. \tag{3.2}$$

3.1.1. *Simulations*

We fill a cubic lattice with blue particles, and label a spherical region in the centre of the lattice red. Meanwhile we apply body forces to both fluids. We take care to

Φ	Sangani $\zeta(\Phi)$	Volumetric radius	Hydrodynamic radius
9.77%	0.370	1.78	1.17
		3.72	3.36
		5.82	5.32
		7.72	7.10
15.5%	0.282	1.78	1.26
		3.74	3.14
		5.81	5.49
		7.73	7.34
		9.76	9.09
		14.77	12.84
26.8%	0.172	7.76	7.93
		9.78	9.98
		11.75	11.83
		13.81	13.31
		15.78	14.55

TABLE 2. Summary of results for Stokes drag on a cubic array of spheres. The volume fraction Φ is the fraction of the particles which are red, which is the parameter needed to compute the Sangani ζ factor for equation (3.2). The values of ζ given in this table were read off Graph 1 in Sangani (1987). The volumetric radius is the radius of a sphere that would contain all the red particles, taking into account the fact that surface tension compresses the sphere. The hydrodynamic radius is the radius of a sphere of fluid that would have exactly the observed terminal velocity if the Sangani theory were followed exactly. All radii are in lattice units. Drops cannot get smaller than two lattice units in radius, and the drop radius must be a few lattice units smaller than half the box size to avoid the formation of cylinders instead of drops (see text). Drops could not be made arbitrarily large because large drops move faster, and both the theory and the model are valid only for slowly moving fluids.

add momentum to both species in such a way as to keep the total momentum of the system always zero: i.e. we may force the red sphere downward by adding a certain amount of momentum to it at each time step; then we will add the same amount of upward momentum to the blue fluid at each time step.

With this forcing scheme in place, we measure the velocity of the sphere, taken to be the average time derivative of the location of the centre of mass of all the red particles. After a few to a few thousand time steps (depending on the size of the suspended sphere) the velocity comes to a steady state, and an average velocity is computed over several thousand time steps. In addition, the experiment is repeated for several macroscopically identical but microscopically distinct initial states, and the average velocities are then averaged over this ensemble of states and histories.

The hydrodynamically meaningful velocity U is the average velocity of the background fluid in the frame fixed to the drop, but the simulation frame is not fixed to the drop; rather, it is the frame in which the total momentum is always zero. Denoting quantities in the simulation frame with tildes and quantities in the drop frame without,

$$\tilde{p}_{\text{drop}} + \tilde{p}_{\text{background}} = 0, \tag{3.3}$$

where $\tilde{p}_{\text{drop}}$ is the total vertical momentum of all the particles in the drop, i.e. all the red particles, and $\tilde{p}_{\text{background}}$ is the total vertical momentum of all the other (i.e. blue) particles. Denoting the total mass of the drop as m_{drop}, the average velocity of the

drop in the simulation frame will be

$$\langle \tilde{u}_{\text{drop}} \rangle = \tilde{p}_{\text{drop}}/m_{\text{drop}}. \tag{3.4}$$

Then velocity in the drop frame, $u(\boldsymbol{x})$, is simply the difference between velocity in the simulation frame $\tilde{u}(\boldsymbol{x})$ and the average velocity of the drop in the simulation frame:

$$u(\boldsymbol{x}) = \tilde{u}(\boldsymbol{x}) - \langle \tilde{u}_{\text{drop}} \rangle . \tag{3.5}$$

The average velocity of the background fluid relative to the drop, U, is easily computed by considering the volume flux through the simulation volume. Taking the forcing to be in the vertical (z) direction, the total flux through every horizontal plane should be the same since the fluid is incompressible; hence, we can define

$$U = \frac{\int_0^L \int_0^L u(\boldsymbol{x}) \, \mathrm{d}x \, \mathrm{d}y}{L^2} \tag{3.6}$$

which will be the same regardless of z. In fact, let us integrate over the vertical direction as well:

$$U = \frac{\int_0^L \int_0^L \int_0^L u(\boldsymbol{x}) \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}z}{L^3}. \tag{3.7}$$

Substituting for $u(\boldsymbol{x})$ from equation (3.5),

$$U = \frac{\int_0^L \int_0^L \int_0^L \tilde{u}(\boldsymbol{x}) \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}z}{L^3} - \langle \tilde{u}_{\text{drop}} \rangle . \tag{3.8}$$

If we multiplied the remaining integral in equation (3.8) by the density ρ, we would recognize it as the total momentum of the system in the simulation frame; but this is zero because of the forcing used in the simulation. Therefore,

$$U = - \langle \tilde{u}_{\text{drop}} \rangle . \tag{3.9}$$

From the hydrodynamic velocity U computed according to equation (3.9), the known Sangani $\zeta(\Phi)$ correction and equations (3.1) and (3.2), we can deduce the radius of a drop that should have that velocity. This is the hydrodynamic radius reported in table 2. In addition, we know the actual number of red and blue particles and the average density of particles in each phase (the pressure, and hence the density inside the drop is slightly higher due to the surface tension); with these data, we can calculate the radius of a sphere which would contain all the red particles. This is the volumetric radius reported in table 2.

3.1.2. *Comparison with theory*

Table 2 and figure 2 compare the volumetric and hydrodynamic radii. The theory is supposed to be valid up to the maximum possible volume fraction of 52.36%, at which value the spheres touch one another. Such a high volume fraction could not be simulated because, if the drop were to touch its image due to a fluctuation, it would merge with the image resulting in an array of cylinders instead of spheres. Simulations were performed at volume fractions of 9.7%, 15% and 27%, and the hydrodynamic radius is in good agreement with the volumetric radius in each case. The volumetric radius is larger than the hydrodynamic radius in almost every case, and the difference is on the order of half a lattice unit in almost every case. Such small discrepancies

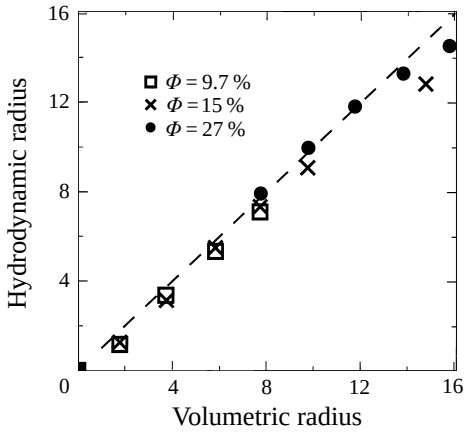

FIGURE 2. Comparing the hydrodynamic and volumetric radii of fluid spheres. If a point lay on the dashed line, the hydrodynamic and volumetric radii would be identical. The agreement is good for all cases. The hydrodynamic radius is deduced from the theory of Sangani (1987), and there are no free parameters in this fit. The radii are in lattice units, and the error bars are smaller than the symbols.

between the hydrodynamic and volumetric radii are not inconsistent with the Sangani theory: just as the bounce-back rule generates an effective no-slip surface at some small but hard-to-quantify distance from the wall, the surface tension rules generate an effective interface at a location that is hard to quantify precisely, but which is close to the apparent interface between species. This drag simulation can be taken as a measurement of the location of the fluid–fluid interface. The discrepancy is larger for the largest two drops, corresponding to drops which move more slowly than would be expected. These correspond to the largest simulations, which take much longer to come to steady state; perhaps the simulations were stopped before the steady state had been reached in these cases.

Ladd (1994) measured the hydrodynamic radii of *solid* spheres in a cubic array using a related lattice-Boltzmann technique, and found comparable results. He reported good agreement between theory (Hasimoto 1959) and his simulations for volume fractions around 10%, and further found good agreement between his simulations and more conventional numerical solutions of the Stokes equations for volume fractions up to the maximum possible. S. Zaleski & D. Gueyffier (1996, personal communication) have found good agreement between their calculations of drag on fluid spheres and the Sangani theory over a similar range of parameters.

3.1.3. *Parameter selection*

Having established that fluid–fluid interfaces behave correctly in the lattice gas model, we now choose parameters so as to compute useful results for very thin wetting layers. Simulations of concentric two-phase flow through a circular cylindrical pipe using related lattice-Boltzmann methods (Gunstensen & Rothman 1993) have been shown to reproduce theoretical predictions (Bacri *et al.* 1990) accurately when the wetting layer is at least one node thick. We verified that our new discrete model yields the same results. The present model can be extended to wetting layers less than one node thick by placing in the medium fewer wetting particles than are needed to cover the wall; then the random motion of the wetting particles along the surface simulates on average a continuous, thin wetting layer. In order to obtain correct

	Porosity(%)	V_p/S (μm)	k (μm)^2	κ^0_{nn}	$S_{r,w}$
Measured	15.2	9.6	1.1	0.92	0.03
Calculated	16.8	10.4	1.0 (FD) 1.2 (LB)	0.84 (I) 0.90 (II)	0.53 (0.03)

TABLE 3. Comparison of experiment with calculation. The porosity and the ratio of pore volume to surface area, V_p/S are geometric parameters and are well preserved by the digitization. The third column is the permeability of the medium, as computed by a finite difference (FD) and by a lattice-Boltzmann (LB) method. The last two columns are the relative permeability of the medium to non-wetting fluid κ^0_{nn} at the residual wetting saturation $S_{r,w}$ when the wetting fluid is displaced by non-wetting fluid. A direct simulation of the laboratory experiment (I) leaves 53% wetting fluid in the medium and computes a relative permeability of 0.84, while a simulation (II) with the correct residual saturation imposed yielded a better estimate (0.90) of the relative permeability.

results for such a thin layer, it is necessary to set the colour field ϕ at the walls to a moderate value (± 3), and to reduce the surface tension by applying the surface tension rule only some of the time. For the simulations reported here, we applied the surface tension rule 40% of the time, and the surface tension was measured at 1.38 ± 0.14(particle mass)(time step)$^{-2}$. We then set about comparing the experimental relative permeability of the rock with that computed by our simulations.

3.2. *Comparison with experiment*

The porosity and specific surface area (the ratio of pore volume to surface area) can be determined both from the digitized rock geometry and from laboratory experiments. We find that the calculation and the measurement give the same result, which gives us some confidence that the geometry of the rock has been adequately resolved by the tomography. This comparison has been published elsewhere (Auzerais *et al.* 1996). In addition, the permeability to single-phase flow in the rock has been computed by lattice-Boltzmann simulations and by a conventional finite-difference method, as well as measured in the laboratory. All three methods give consistent results, as reported also in Auzerias *et al.* (1996). In this section, we present in detail comparisons between simulations and measurements of two-fluid flow in porous rock, which were described briefly in Auzerias *et al.* (1996). Table 3 summarizes all these comparisons.

When non-wetting fluid displaces wetting fluid in a real rock, the non-wetting fluid first fills the widest channels and then only gradually finds its way into narrower pores, as the wetting fluid in these pores is slowly drained out of the rock. Eventually the wetting fluid is reduced to a thin film which does not flow, and the pore space is almost entirely filled with the non-wetting fluid. Since the non-wetting fluid is no longer flowing, its concentration cannot decrease further and the composition of the fluid in the pore space is at a steady state.

Precisely the same sequence can take place in a direct simulation of the laboratory experiment, as shown in figure 3. We will refer to such direct simulations of displacement experiments as simulations of type I. Because the lattice has a resolution of 7.5 μm, the thinnest wetting layer that can arise is macroscopically thick: in fact, the set of all void nodes which neighbour solid nodes contains about half of all the void nodes. This is part of the reason that the endpoint saturation – the concentration of wetting fluid remaining in the 'thin film' at steady state – is 53% in the simulation, while in the real rock, it is only 3%. It is hardly surprising that such a large discrepancy in the saturation should result in a discrepancy in the computed relative

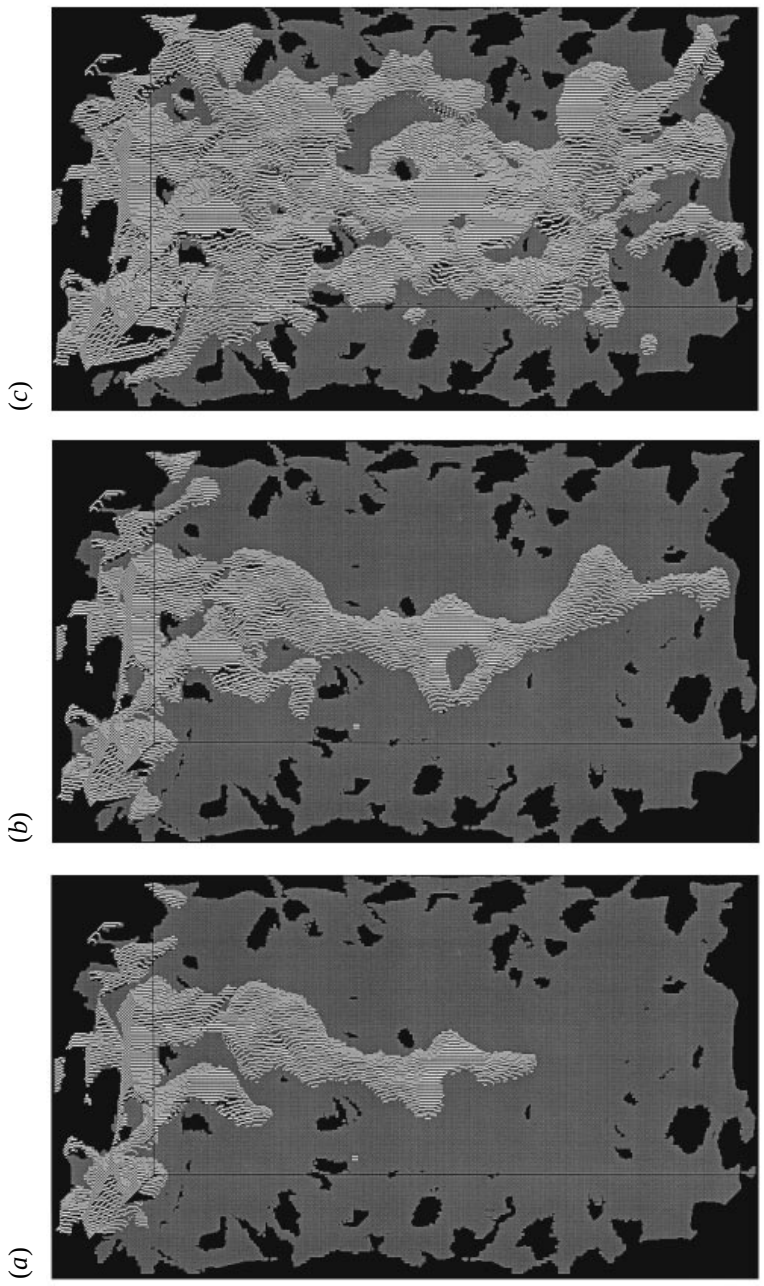

FIGURE 3. A simulation of a non-wetting displacement experiment. The rock (invisible) is initially filled with wetting fluid (dark grey, transparent). After 4000 time steps (*a*), the non-wetting fluid (light grey, opaque) has begun to travel through the medium; by 7000 time steps (*b*), the non-wetting fluid has found a channel through the medium; after this, the non-wetting fluid gradually fills up more of the pore space; (*c*) is at 52000 time steps.

permeability; it is nonetheless reassuring that the computed relative permeability, 0.84, is not that far from the experimental value of 0.92.

Rather than trying to duplicate the experimental flow, we can try to compute the relative permeability in a more convenient type of simulation, which we will call type II. We fill the pore space with a mixture containing some concentration of wetting fluid and force the mixture while keeping the vertical boundaries periodic so that the overall composition does not change. Now if there are only a few wetting particles, they cannot form a continuous layer of wetting fluid, but they can diffuse over the walls and *on average* create the effect of a thin wetting layer, as in the cylindrical pipe mentioned earlier. When we filled the void space with 3% wetting fluid to simulate the experimentally observed endpoint saturation, the computed relative permeability in the simulation was 0.90, in good agreement with the experiment. Thus, although the limited resolution of the lattice prevents us from computing some parameters accurately, others can be determined: specifically, those which can be averaged over a macroscopic volume, like the flux of each species at some particular concentration. It should be emphasized that the limited resolution is not an intrinsic problem with the numerical method; rather, it is a consequence of the computing technology which is currently available.

Both simulations used a uniform body forcing on the fluid to simulate the pressure forcing used in the experiments. It is not clear to what extent this is a valid substitution; body forcing was chosen for reasons of numerical convenience, but it is to be expected that the distribution of pressure in the medium will be different under body or pressure forcing. This remains a topic for further study. In the following sections, pressure forcing would obscure the coupling between fluids, while body forcing does not; indeed, experimentalists have proposed using body forcing (either buoyant, Kalaydjian 1990, or electromagnetic, Zarcone 1994) for this purpose.

4. Transport properties in two-fluid flow

While a 'simulation permeameter' (see also Soll *et al.* 1994; Buckles *et al.* 1994) would no doubt see valuable use, it is both simpler and of greater scientific interest to take the tomography data as representative of the kind of pore spaces found in rock, and to ask what are the generic properties of flow through such complex and confined geometries. While it is clear from experiments that flow through porous rock is qualitatively different from flow through a medium such as sand or a glass bead pack (A. Thompson, personal communication; Bear 1972; Scheidegger 1974), it is not clear what are the geometric properties of the rock which make it different. A number of techniques have been developed to characterize the geometry of porous rock, including both statistical measures (Yao *et al.* 1993) and topological properties (Lindquist *et al.* 1996; Spanne *et al.* 1994). It is possible to generate pore geometries which have the same properties as were measured from the original rock, but which are nonetheless distinct from the authentic rock; one may simulate flow through these synthetic geometries as well as through the original geometry (Adler, Jacquin & Quiblier 1990; Sallès, Thovert & Adler 1993, 1994). Research along these lines may elucidate the geometric properties of rock that control the flow. Here, we have chosen to use only the authentic pore geometry for our simulations.

In this section, we present the results of relative permeability measurements, and show that the fluid–fluid coupling in equation (1.1) is not only significant, but also that the coupling coefficients κ_{wn} and κ_{nw} are equal. This is the first time that both coupling coefficients have been independently measured for a natural rock geometry.

We first survey the variation of fluid response with concentration and with forcing, and then consider the fluid–fluid coupling in detail.

4.1. *Simulation details*

All the simulations reported in this section and the next were simulations of type II as described above. Two kinds of fluid flux are reported, *total flux* and *local flux*. The first is a sum of the momenta of all the particles throughout the lattice at some moment in time, and it may be calculated for each species of particle separately:

$$\text{Total Flux}_i(t) = \sum_{\boldsymbol{x}} \boldsymbol{p}_i(\boldsymbol{x}, t) \tag{4.1}$$

where the subscript i could refer to either species or both, and $\boldsymbol{p}_i(\boldsymbol{x})$ is the momentum of particles of species i at location $\boldsymbol{x}$. The local flux is the momentum at a particular node in the lattice, averaged over many (Ω) time steps:

$$\text{Local Flux}(\boldsymbol{x}) = \frac{1}{\Omega}\left(\sum_{t}^{\Omega} \boldsymbol{p}(\boldsymbol{x}, t)\right). \tag{4.2}$$

Note that the local flux is not defined for each species, but only for the mixture. We found that a local flux of particles of a particular species was strongly skewed at any interface that does not move, so much so that the quantity was not useful. It is, however, possible to define a local, time-averaged *concentration* of each species in analogy to equation (4.2).

4.2. *Concentration effect*

We calculated the total flux of each species through the medium when either species was forced, for a range of concentrations. For this study, each simulation was performed with the same amount of forcing per particle, 2.5×10^{-4} units of momentum per particle of the forced species per time step. The total flux was then divided by the total flux which is observed in the same medium when a single fluid is forced at this rate; this is the *normalized momentum* which is plotted in figure 4. We take care not to call this normalized momentum the relative permeability, because this calculation neglects the capillary threshold which will be measured in the next section, but the qualitative features of the plot reproduce the kinds of features observed in real experiments: the relative permeability of each species decreases with its concentration, and the viscous coupling terms are small. In fact, the coupling is not distinguishable from zero except at concentration 50%, though this too depends on the forcing.

4.3. *Capillary threshold*

For small forcings, non-wetting fluid may be excluded from the smallest pores in the medium. In the simulations performed here, that means that both species will form separate phases locally throughout the medium. Initially, drops of the non-wetting fluid will travel through the medium, but after a short time, the non-wetting fluid gathers into larger blobs that can no longer pass through any available opening. At this point the net flux of the non-wetting fluid will vanish; and under some conditions the non-wetting blobs may clog the medium sufficiently that the wetting fluid is also unable to flow. But if the forcing is strong enough, these blobs may be squeezed through the narrow channels, resulting in a steady flow. So when the forcing is below the capillary threshold, no flow occurs; and above this threshold, the flux increases with increasing force. Earlier simulations using artificial media

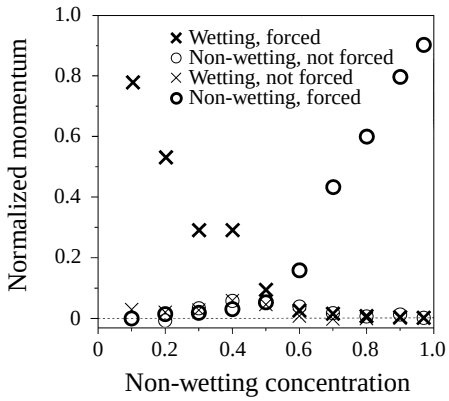

FIGURE 4. The normalized momentum is the total vertical momentum divided by the total vertical momentum of pure fluid at the same forcing in the same medium; it is roughly the same as the relative permeability. Error bars are about the size of the symbols. Each datum reported here represents an average of two to five simulations. Note that the viscous coupling is largest near concentration 0.5.

(Gunstensen & Rothman 1993) found evidence for nonlinear dependence of flux on force in this domain, but the present simulations cannot distinguish between linear and nonlinear dependence. We have assumed a simple linear dependence above the capillary threshold in what follows.

We take as our example the case with concentration 50%. When the non-wetting fluid is forced, it does not move at all until the force reaches 2.5×10^{-4} units of momentum per particle per time step, and at that forcing, some simulations show a significant steady-state non-wetting flux, while others show no detectable flux. For stronger forcings, the flux increases rapidly. On the other hand, when wetting fluid is forced, some wetting flux is observed even for small forcing, but the flux increases more slowly with increasing force. These results are shown in figure 5. This is in qualitative agreement with experimental measurements (for example, Bear 1972, pp. 449–452, 459–466). The different slopes can be understood in broad terms as a consequence of lubrication: when the wetting fluid moves, it has to creep along the walls, while the non-wetting fluid is lubricated by the wetting fluid.

Since the wetting fluid moves at all forcings and its response is linear with forcing, the 'normalized momentum' plotted in figure 4 for the wetting fluid can be interpreted directly as the relative permeability. In the non-wetting case however, one must carry out a series of simulations for each concentration, because both the capillary threshold as well as the slope of the linear response are likely to depend on the concentration. We expect that the actual relative permeabilities for the non-wetting fluid at each concentration will not be smaller than the normalized momenta plotted in figure 4.

4.4. *Fluid–fluid coupling*

The total flux varied from one simulation to another at the same forcing because the interfaces between species can take different locations in each simulation. The variation was not large compared to the flux of the forced species, but the variation was large compared to the much smaller fluxes of the unforced species in each simulation. We therefore found it necessary to perform many simulations – 87 in all – to be certain that the coupling coefficients κ_{nw} and κ_{wn} had been accurately determined (Olson 1995). The average fluxes and the best-fit lines through the data are shown

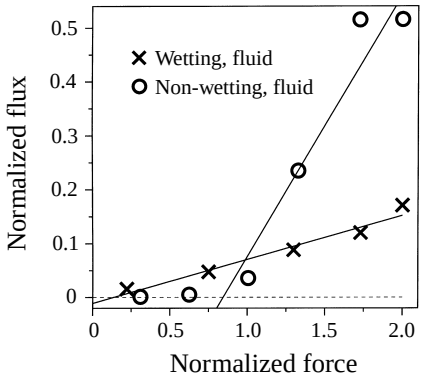

FIGURE 5. Each species can be forced independently in the simulations, and both species respond to forcing of either species. This plot shows the flux of each species when it is forced at 50% concentration. Each data point represents an average over five to fifteen simulations, and each simulation, an average over 10 000 time steps. Error bars are about the size of the symbols. The 'normalized flux' is the total flux of the species (defined in equation (4.1)) divided by the total flux of a pure fluid at a forcing of 0.0004 in lattice units (i.e. (particle mass)(lattice spacing)/(time step)^2). The 'normalized force' is the forcing per particle per time step divided by 0.0004; the forcing used in figure 4 corresponds to 0.625 on this scale. The lines are least-squares fits, and for the non-wetting fluid, the leftmost point is omitted from the fit. The slopes of the lines, which are the relative permeability coefficients, are $\kappa_{nn} = 0.48 \pm 0.03$ for the non-wetting fluid and $\kappa_{ww} = 0.081 \pm 0.006$ for the wetting fluid.

in figure 6. It is interesting to contrast this figure with figure 5, because while the unforced fluids respond so similarly to the forcing, the responses of the forced fluids are very different. This graph does still show the effect of the capillary nonlinearity, in that the weakest two wetting fluxes (corresponding to the weakest two non-wetting forcings) are essentially zero; but the non-wetting fluxes are also essentially zero at these forcings. The best-fit lines to the two sets of data have the same slope within the uncertainty of the fit, and the uncertainty is not large. The viscous coupling thus appears to be reciprocal, that is, $\kappa_{nw} = 0.015 \pm 0.003$ and $\kappa_{wn} = 0.014 \pm 0.004$ at 50% concentration in this porous medium.

An additional test of reciprocity, independent of regression, was performed on the data. The Kolmogorov–Smirnov test (Press *et al.* 1994) calculates the probability that two sets of data could be as different as they are if they came from the same underlying distribution. It is only approximate for two-dimensional distributions like this one, but the result of the test was that two sets of data shown in figure 6 are significantly alike.

4.5. *Reasons for reciprocity*

Why should the viscous coupling be the same (or nearly the same) in these two cases? On casual consideration, it might seem to be a consequence of the Onsager reciprocity relation from non-equilibrium thermodynamics (Onsager 1931*a*, *b*). In this theory, the rate of entropy production can be computed as the sum of products of forces F_i with their *conjugate* fluxes J_i,

$$\text{Entropy production} = \sum_i F_i J_i. \tag{4.3}$$

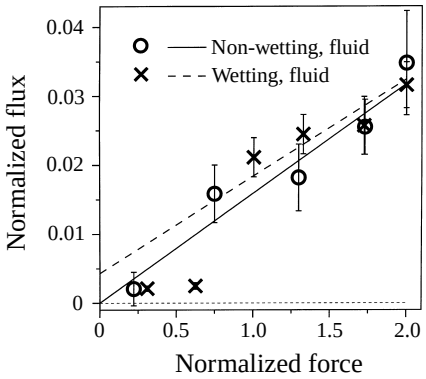

FIGURE 6. Fluid–fluid coupling at 50% concentration. The data points here are from the same simulations as for figure 5 and are averaged in the same way, but now it is the unforced species whose total flux is plotted. Note that the vertical scale is much smaller, so the error bars are larger than the symbols. Least-squares lines are fit to the data, and the leftmost wetting fluid point is omitted, corresponding to the point omitted from the fit in figure 5. Whereas the slopes of the forced fluid fluxes were quite different, here the slopes of the unforced fluid fluxes are the same, within the uncertainty of the fit: $\kappa_{nw} = 0.015 \pm 0.003$ (non-wetting fluid), and $\kappa_{wn} = 0.014 \pm 0.004$ (wetting fluid). The uncertainty of the y-intercept is large enough that both best-fit lines can be can be said to pass through the origin.

In addition, it postulates a general linear force–flux relation,

$$J_i = \sum_j L_{ij} F_j, \tag{4.4}$$

where the L_{ij} are the coupling coefficients; this equation is analogous to equation (1.1), with the coupling coefficients L_{ij} corresponding to $k\kappa_{ij}$. If one has a set of such conjugate forces and fluxes, and a fluctuation–dissipation theorem holds (i.e. the total correlation of rates of decay of fluctuations of the system in thermal equilibrium is the rate of dissipation of comparable macroscopic non-equilibrium perturbations, Reif 1965) then the theory states that the cross terms, i.e. the coupling coefficients between forces and their *non*-conjugate fluxes, will be reciprocal: $L_{ij} = L_{ji}$.

While our assumed linear force–flux relation and the reciprocity found in our simulations suggest a parallel to the Onsager theory, the fact that bubbles break and merge during the simulation indicates that the flow is maintained too far from equilibrium for the theory to apply. This complex, non-equilibrium, interfacial dynamics should have a significant influence on the macroscopic dissipation (Doi & Ohta 1991). In equilibrium, on the other hand, the interface topology does not change. Any fluctuation large enough to change the interface topology would be so improbable as to be negligible. Thus, though our macroscopic responses are approximately linear, there is no reason to believe that this linearity is related to the average decay of equilibrium fluctuations.

Thus we are left with a conundrum: we appear to have demonstrated a kind of Onsager reciprocity, but at the same time it does not appear to have a theoretical justification. It is interesting, then, to examine some cases that are understood. First of all, if the two fluids had the same wetting properties, then their coupling would be reciprocal by symmetry. If two fluids were confined to a straight channel with a flat interface between them, again symmetry would require reciprocity. It is not so obvious, but nonetheless true, that the coupling is reciprocal when two

fluids flow in a pipe with circular cross section, with wetting fluid on the outside, non-wetting fluid inside, and the interface between them a cylinder concentric with the pipe (Bacri *et al.* 1990). Pride & Flekkøy (1996, personal communication) have demonstrated a kind of differential reciprocity in a more general case: the change in fluxes due to small changes in forcing, such that the interface moves only a little bit due to the change in forcing, is reciprocal. None of these arguments can apply in the case of our simulations where interfaces move, break and merge during the flow, but they lead us to a speculation: when the interface takes the same shape regardless which fluid is forced, the coupling will be reciprocal. Perhaps the interfaces between the two fluids in the porous rock are mostly held in place by the complex rock geometry, and the flow takes place in essentially one simple channel regardless how the fluid is forced. The reciprocal coupling would then be a consequence of the constrained pore geometry. In what follows, we will examine the fluid–fluid interface and the velocity field in more detail, and find that this simple explanation is not sufficient. Along the way, we will learn that the flow is indeed organized, that this organization explains many of our observations, and that the organization of the flow can be characterized by considering statistical properties of the velocity field.

4.6. *Details of the flow*

Figure 7 shows the lattice nodes where the concentration (averaged over 20 000 time steps) is between 40% and 60% in both wetting and non-wetting forced flow. We can think of this as the average location of interfaces. While the interfaces are largely identical, most of them are in regions of little or no flow, as we shall see below. There is an important difference on the left side of the simulation volume. There, in the non-wetting forced case (figure 7*a*), the interface is clearly more connected than in the wetting forced case, figure 7(*b*). As shown in figure 8, this turns out to be a region of significant flow, so a different topology at this location is likely to be important. The vertical velocity fields are shown in figure 8, along with a plot for the velocity of a pure fluid. Although the pure fluid (figure 8*b*) and the non-wetting-forced mixture (figure 8*a*) are quite similar and similarly organized, the wetting-forced fluid is entirely different. The pure fluid flows in a single, well-organized channel which is connected through the medium. The plot does not show all of the channel, because some of the flow occurs more slowly than the top 1% of the velocity and is not plotted here. The non-wetting-forced fluid flows in the same channel, with a little bit of flow scattered around; but the wetting-forced mixture has no channel at all, only velocity scattered throughout the medium. In fact, the only place where the wetting fluid does not have any flow is in the channel where the pure fluid moves. Note that the channel through which the non-wetting-forced fluid moves is at the same location as the connected interface pointed out in figure 7.

In summary, the interfaces are largely at the same locations regardless which fluid is forced. However, the interface is connected in a particular region when the non-wetting fluid is forced, but disconnected at the same location when the wetting fluid is forced. Further, the flow is rapid at this location when the non-wetting fluid is forced, but slow when the wetting fluid is forced. These distinctions seem important, and yet the fluid–fluid coupling is the same in the two cases. One would expect different couplings since the interface geometries are different, so the observed reciprocity of the coupling remains unexplained. Perhaps it is an accidental consequence of our particular pore geometry, though it seems improbable that the first case chosen should happen to have such a convenient property.

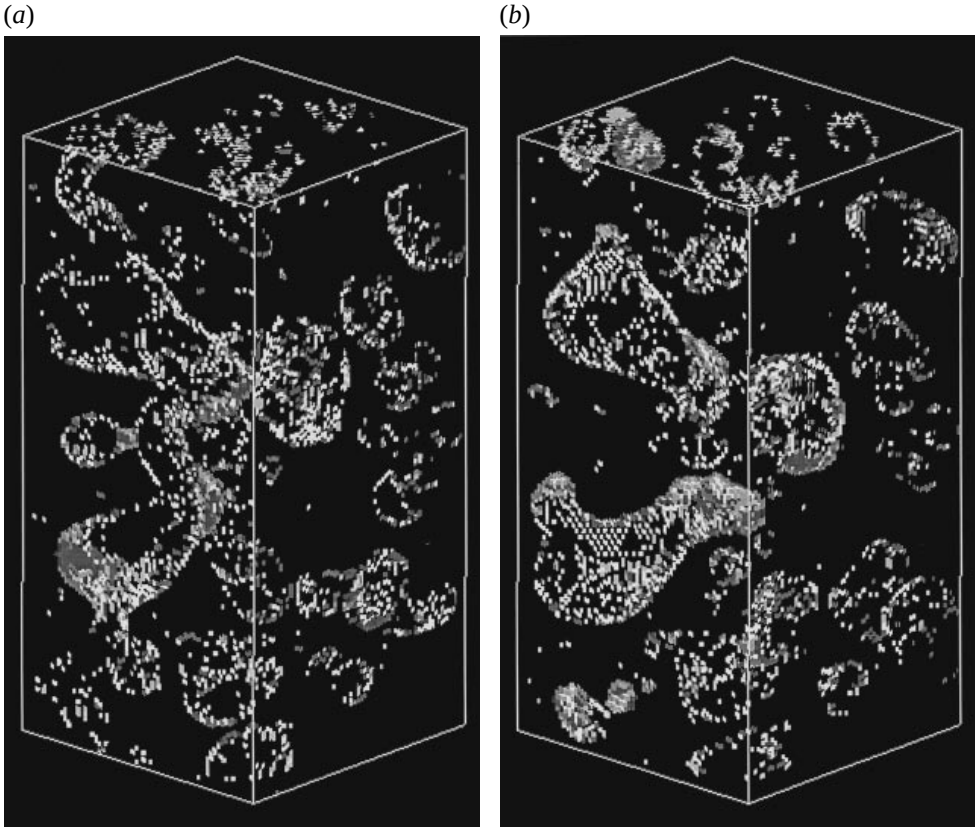

FIGURE 7. The locations of sites where the concentration, averaged over 20 000 time steps, is between 40% and 60%; such sites are deemed to be nearly stationary interfaces. (*a*) The location of interfaces when the non-wetting fluid is forced, and (*b*), when the wetting fluid is forced. Although much of the interface is the same in the two cases, note that on the left side of the volume, the non-wetting-forced case (*a*) is more connected than in the wetting-forced case (*b*). This is in a region where the vertical velocity is large (see figure 8), so the difference is salient.

5. Characterizing the flow

We now show that a simple statistical parameter, the *skewness* of the velocity distribution, characterizes the organization of flow evident in figure 8. One of the consequences of a complex flow is a range of velocities at different locations in the rock: a histogram of velocity would reveal something of this complexity. We have computed histograms for six reference cases: unforced pure and mixed fluids as controls; forced pure fluid and uniformly forced mixture; and mixtures wherein only the wetting or non-wetting fluid is forced.

5.1. *Unforced fluids*

Figure 9 show the distributions of velocity components in unforced fluids in the porous rock. As one would expect, the distributions are centred at and symmetric about zero, since there is no reason for a bias in any direction. But curiously, the distribution for the mixture is much wider than that for the pure fluid. This wide distribution does not grow narrower as more data are added, and a nearly identical distribution can be

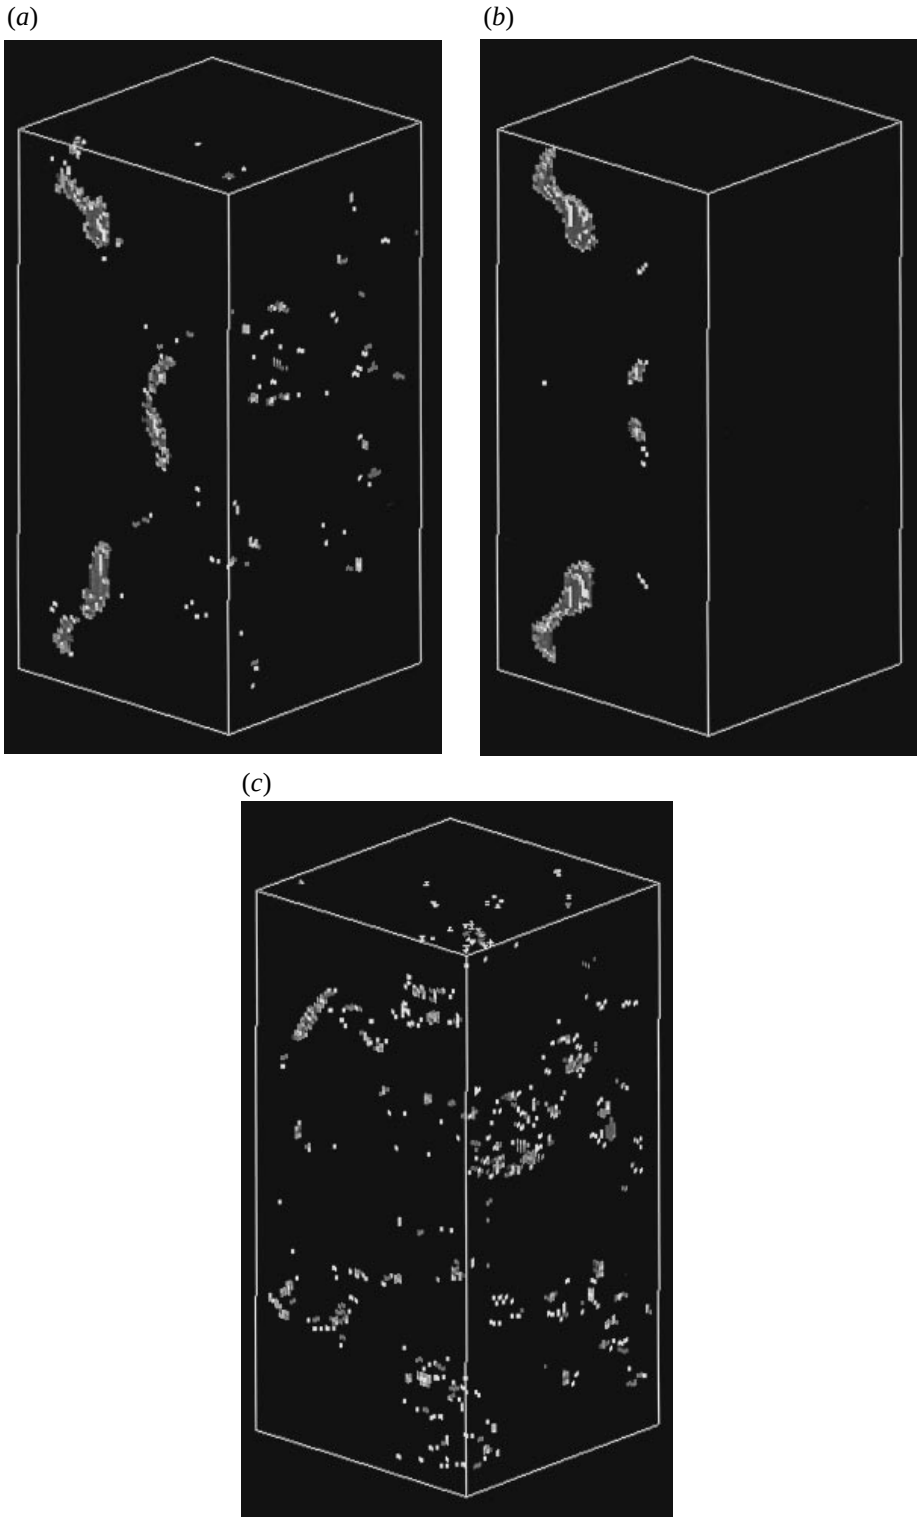

FIGURE 8. For caption see facing page.

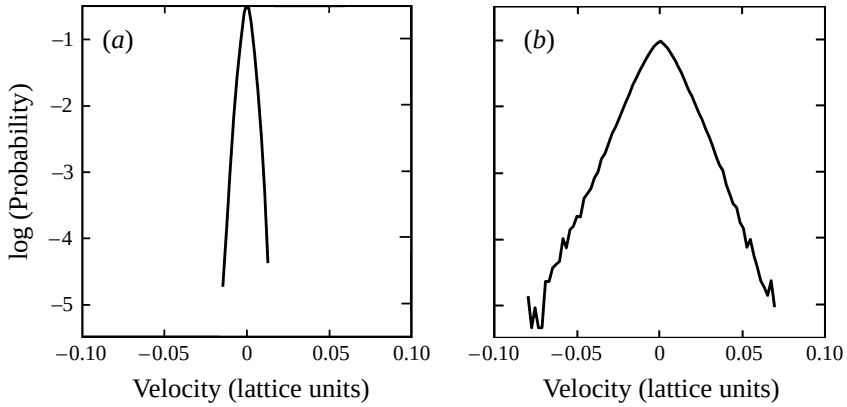

FIGURE 9. When there is no forcing, the average velocity is zero, but nonetheless there is a distribution of velocities in the porous rock: (*a*) the distribution for pure fluid, and (*b*) for a 50% mixture. The velocities were averaged for 20 000 time steps, and here we plot the probability that a location in the pore space chosen at random will have the velocity component indicated. The distributions for *x*-, *y*- and *z*-components of velocity are the same in each case.

found for a spherical drop in a spherical cavity. This distribution would seem, then, to be characteristic of stationary interfaces with arbitrary orientation.

Why should stationary fluid have a distribution of velocities at all? The simple explanation, and a sufficient one for the pure fluid, is that the lattice gas has fluctuations and that, even after long averaging, there is still some random variation from point to point in the medium. Such variation would be observed in real fluids as well, if they were sampled with fine enough time and space resolution. But the two-phase histogram is qualitatively different, and demands a different explanation. Here the cause appears to be *spurious currents* at the fluid–fluid interface. These currents are a well-known defect in two-phase lattice methods (see, for example, Gunstensen 1992), and are a result of the discrete underlying lattice. Figure 10 shows these currents on a slice through a spherical drop suspended in a spherical cavity. Experience shows that sensitive tests such as those reported in this paper (drag on fluid spheres and flow in cylindrical pipes) are not affected by these currents. We see, however, that they play an important rôle in our statistical analysis.

5.2. *Forced fluids*

Figures 11 and 12 show histograms for forced pure fluid and forced 50% mixtures, respectively. In every case the fluid or fluids were forced upward, in the positive *z*-direction. Consequently the distributions of horizontal (*x* and *y*) components are always found to be the same, while the vertical (*z*) component has a different distribution. The horizontal distributions are always symmetric about zero, while the vertical component is always skewed towards positive velocities.

FIGURE 8. The vertical velocity in the porous rock, when the non-wetting fluid (*a*) or wetting fluid (*c*) in a 50% mixture is forced; or when pure fluid (*b*) is forced. The plots show the locations in space where the fastest hundredth of the vertical velocities are found. Note that in (*a*) and (*b*), the fastest velocities are located in a coherent channel, the same one in both cases; while in (*c*), no such coherent channel appears, and the flow is not fast in the channel which appears in the first two cases.

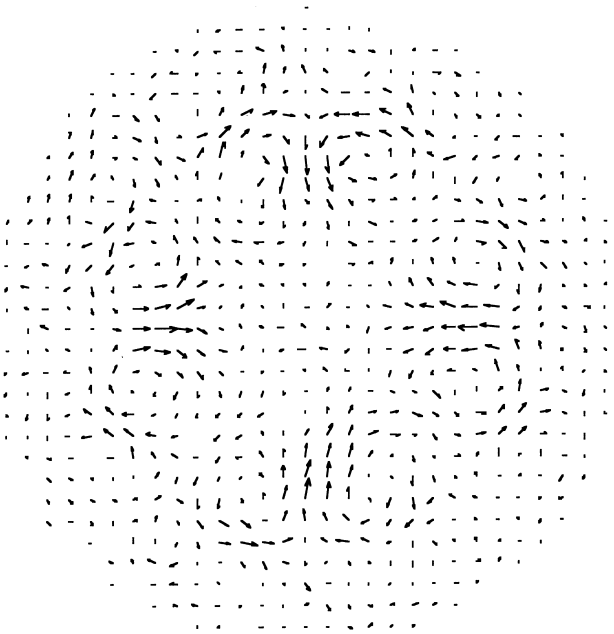

FIGURE 10. Slice through a sphere of radius 10, showing the spurious currents in the (x, y)-plane. The magnitude of the largest arrows is 0.1 lattice unit/time step, quite a large speed. The currents were averaged over 3000 time steps.

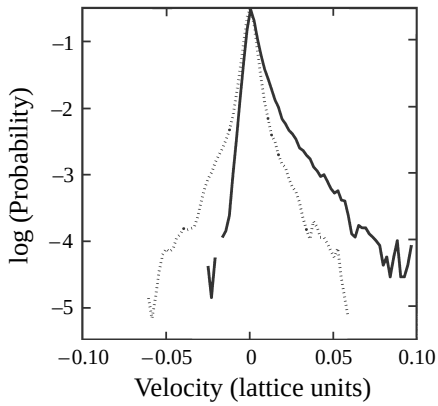

FIGURE 11. Histograms of the vertical component (solid line) and the horizontal components (dashed line) when pure fluid is forced. The forcing was 0.0004 momentum units per time step.

5.2.1. *Pure fluid*

When the pure fluid is forced (figure 11), the distributions are much wider than in the unforced case, with the horizontal components symmetric about zero and the z-component strongly skewed towards positive velocities. The horizontal component histograms are wide because the fluid must flow through channels that are not vertical. The width of these distributions indicates the range of horizontal velocities, and hence indicates the extent to which the flow paths are scattered from the vertical. It might be possible to derive a statistical tortuosity from these data. The z-component distribution is also wide, but with almost no occurrences of negative (downward)

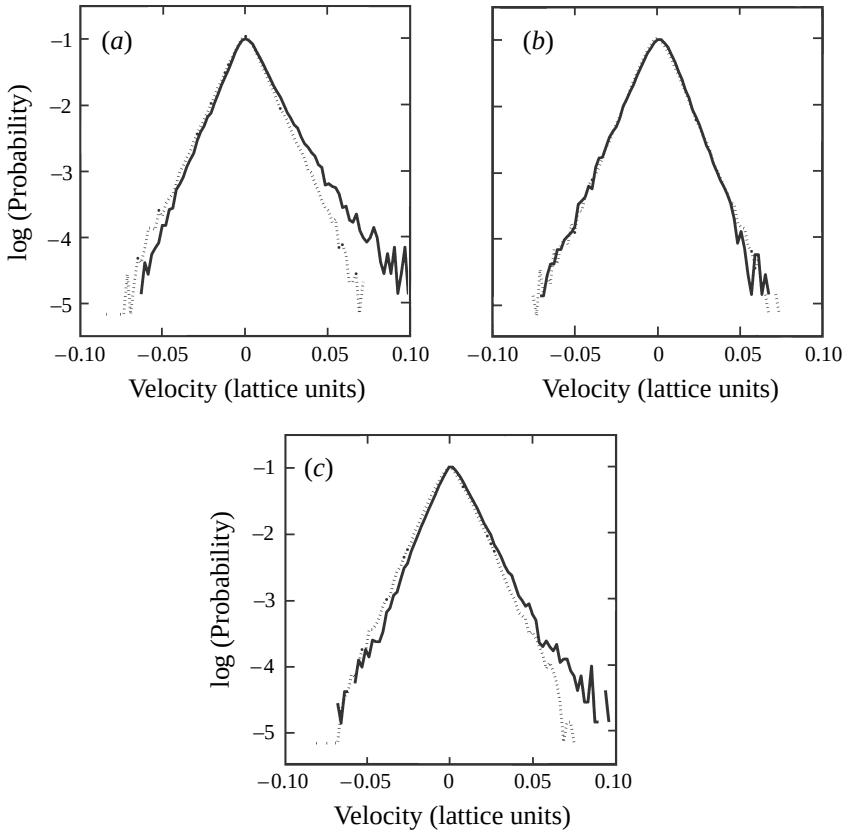

FIGURE 12. Distributions of velocity depend on which fluid is forced. Although the horizontal components (dotted lines) in each graph still look like the unforced histogram (figure 9*b*), the *z*-components (solid lines) are different to a greater or lesser degree as a result of the forcing. When the non-wetting fluid is forced (*a*), the distribution has a tail at high velocities; whereas when the wetting fluid is forced (*b*), there is no tail, but the most probable velocity is slightly larger than zero. When both fluids are forced (*c*), the histogram has properties of both.

velocity. This distribution reflects the range of channel widths through which the flow passes, from trapped regions that do not flow at all to a few wide channels with little resistance to flow. The fact that there is almost no negative velocity indicates that the fluid is moving coherently through the medium.

5.2.2. *Mixtures*

Histograms for wetting forced, non-wetting forced, and uniformly forced mixtures are shown in figure 12. In contrast to the difference between forced and unforced pure fluid, forcing does not significantly change the distributions of the horizontal components in the mixtures. The width of these distributions is apparently determined by the spurious currents. However, the *z*-component has a distinctly skewed distribution in the cases of non-wetting forced (figure 12*a*) and uniformly forced (figure 12*c*) mixtures. This is analogous to the skew observed when the pure fluid was forced. In contrast, when the wetting fluid is forced, the *z*-component is not visibly skewed (figure 12*b*).

The overall shape of the velocity distribution emphasizes the outlying data at extremes of velocity. We can complement these observations by examining the peak

of the distribution, the most common velocity. When the non-wetting fluid is forced (figure 12*a*), the most common velocity is zero, but when the wetting fluid is forced (figure 12*b*), the most common velocity is slightly larger than zero. Thus, while the non-wetting fluid is mostly stationary when it is forced (as evidenced by its most common velocity being zero), it has a few locations with a high velocity, resulting in an asymmetric distribution; while when the wetting fluid is forced, there is slight motion almost everywhere, with no large velocities anywhere. When both species are forced (figure 12*c*), the most common velocity is slightly larger than zero, and the distribution is somewhat asymmetrical: it is a combination, as one might expect, of the cases when wetting or non-wetting fluids are forced.

5.3. *Skewness*

The spatially and temporally complex velocity fields described above are reminiscent of those observed in turbulence studies. Statistical descriptions of turbulence have been found useful (see, for example, Tennekes & Lumley 1972), so we compute the moments of our observed distributions. In particular, the third moment or *skewness* turns out to be salient. The first moment of a distribution is the familiar mean, and the second moment is the variance; listing all three:

$$\left.\begin{aligned} \text{mean} = \bar{x} &= \sum_i^{\Omega} x_i/\Omega, \\ \text{variance} = \sigma^2 &= \sum_i^{\Omega} (x_i - \bar{x})^2/\Omega, \\ \text{skewness} = S &= \sum_i^{\Omega} (x_i - \bar{x})^3/\sigma^3\Omega. \end{aligned}\right\} \tag{5.1}$$

The skewness is a measure of the asymmetry of the distribution about the mean, scaled by the variance. For concreteness, consider a set of data $\{x_i\}, i \in \{1, \Omega\}$ which consists of n data with value 2α, $2n$ data with value $-\alpha$, and and $\Omega - 3n$ data with value zero. The mean of this simple distribution is zero by construction; its variance and skewness are

$$\sigma^2 = 6(n/\Omega)\alpha^2$$

and

$$S = (\Omega/6n)^{1/2}.$$

Note that the variance can be changed independently of the skew within this family of distributions by changing α. Hence, while the variance reports the width of the distribution, the skewness is independent of the width and reports instead that the outliers on one side of the mean are further away than the outliers on the other. In addition, the magnitude of the skewness is related to the number of outliers: the more outliers, the smaller the skewness will be. Hence a large skewness indicates that there are a few large values in the distribution. Since the distributions in this case refer to velocities distributed in space, the skewness indicates the extent to which directed flow is localized in a few places with large velocities. A flow field with a large skew is thus well-organized or focused into a few channels with large flow. Table 4 shows these moments for each of the cases illustrated in figures 9, 11 and 12.

Consistent with the histograms, the skewness of the transverse velocity component distributions is always small, since the distributions of these components are always

Forcing	$\bar{w}/10^{-3}$	$\sigma^2/10^{-4}$	Skewness horizontal	Skewness vertical	Figure(s)
Unforced pure	0.005	0.003	0.02	0.02	9(*a*)
Unforced mixed	0.005	0.13	0.05	0.01	9(*b*)
Forced pure	3.24	0.049	0.05	4.2	8(*b*), 11
Forced mixed	2.28	0.14	0.08	0.52	12(*c*)
Non-wetting forced	2.12	0.18	0.09	0.85	8(*a*), 12(*a*)
Wetting forced	0.80	0.14	0.02	0.31	8(*c*), 12(*b*)

TABLE 4. The mean vertical velocity ($\bar{w}$), its variance (σ^2) and the skewness of both horizontal and vertical velocities for the distributions illustrated in figures 9, 11 and 12. The last column indicates the figure in which the corresponding histogram (and optionally, the locations of large vertical velocity) is plotted. The skew is dimensionless, the mean is in lattice units per time step and the variance in (lattice units per time step)^2.

symmetric. The z-component skewness is largest for the pure fluid flow, which is by eye the most asymmetric distribution; the skewness also distinguishes between the non-wetting forced distribution (more skew) and the uniformly forced distribution (less skew), although this comparison is hard to make by inspection. Finally, this calculation detects some asymmetry in the wetting forced case, where there was not apparently any in the histogram. The calculated skewness is thus a sensitive measure of the distribution.

The fourth moment, or *kurtosis*, of the distributions can also be computed. The kurtosis, like the skewness, is large when there are a few, large outliers, so it could also be a measure of the organization of the flow. But whereas the skewness picks out asymmetry in the distribution, the kurtosis treats outliers symmetrically. Since the spurious currents are symmetric, the kurtosis includes them, which obscures the signal we want to see. The skewness, on the other hand, more selectively indicates the presence of directed flow.

We have thus found a simple statistical parameter which correlates with the apparent organization of the flow. But why should some cases be more organized than others? The explanation lies in the spatial distribution of the wetting and non-wetting fluids. The non-wetting fluids avoid the pore surfaces, forming smooth bubbles in the larger pores of the medium, while the wetting fluid clings to the walls. This spatial partitioning leads to a dynamic partitioning between the two fluids. When the wetting fluid is forced, its constricted and distributed geometry leads to a slow, distributed flow with little long-range coherence. Because the wetting fluid has no large channels through which it can pass (the large channels being filled with non-wetting fluid), it has few locations where it can have a large velocity and many where it must have a small velocity; thus the skew of this velocity distribution is small and the most common velocity is non-zero. On the other hand, the non-wetting fluid selectively occupies the larger pores, and when the forcing is strong enough to drive this fluid through the medium, almost all of its flow will take place in a few large channels where the velocity will be high. Meanwhile, much of the non-wetting fluid is trapped by the constricted medium and unable to move: hence the non-wetting fluid flow will have a larger skewness, reflecting the wider range and stronger localization of the non-wetting flow, while the most common velocity is zero. When both species are forced uniformly, then the wide channels will be filled with rapidly flowing non-wetting fluid while the narrower channels will contain slowly flowing wetting fluid,

so the overall flux will be larger while the flow will be less localized than in the non-wetting forced case. For this reason, the skewness of the uniformly forced case is smaller than that of the non-wetting forced case and the most common velocity is non-zero.

6. Conclusions

We have presented a series of results on the flow of two-phase mixtures in porous sandstone, based on lattice-gas simulations of flow in authentic pore-space geometry obtained by X-ray microtomography. We have made the first calculation of fluid–fluid coupling in simulated flow in a real geometry and we have presented the first test of the reciprocity of fluid–fluid coupling in two-fluid flow in porous rock: every previous determination, by experiment or by simulation, has been in an artificial medium.

In addition, we have demonstrated that the new lattice-gas model correctly simulates the subtle phenomenon of drag past a fluid sphere, despite the spurious currents at the interface; and that despite the finite spatial resolution, we could qualitatively reproduce results from experiments on two-phase flow through sandstone. These tests and others reported elsewhere (Olson & Rothman 1995; Appert *et al.* 1995) give us confidence that the combination of our model and tomographic reconstruction of rock geometry is indeed an appropriate tool with which to study the complex flow that takes place in porous rock.

Finally, we considered ways to characterize the complexity of the flow, and found that the third moment or *skewness* of the spatial velocity distribution was a convenient parameter. The systems which had apparently well-organized flow – the pure fluid and the mixture with non-wetting forcing – had a large skewness, while less organized flows had smaller skewness. The skewness is particularly appropriate because it seems to be relatively unaffected by the spurious currents which arise at interfaces in the lattice-gas model. It is also a parameter which might be conveniently measured in laboratory experiments.

We believe that the rock geometry used for this research is sufficiently generic that the properties we have observed should be qualitatively reproduced for other geometries and larger scales. However, this must be demonstrated by further calculations. In particular, we will need more examples before we can understand the observed symmetry of the coupling coefficients. One of the long-standing problems in the industrial modelling of porous flow is the question of how to use small-scale laboratory (or simulation) results to predict the outcome of much larger field or factory processes. While direct calculation and experiments may yield useful scaling relations for uniform media, real rocks are usually fractured and much of the fluid flow is dominated by the flow in fractures. A closer analysis of the statistics of flow in the porous rock should offer means to predict how the flow in rock will couple to flow in the space between rocks, which might then offer better ways to scale laboratory measurements up to the field.

The authors wish to thank Eirik Flekkøy and Steve Pride for valuable discussions on the topic of Onsager reciprocity; and Stéphane Zaleski and Denis Gueyffier for pointing out an error in our analysis of the Stokes drag calculations. Acknowledgment is made to the donors of The Petroleum Research Fund, administered by the ACS, for partial support of this research. The research was also supported in part by the sponsors of the MIT Porous Flow Project and by NSF grants EAR-9218819 and 9418039-EAR.

REFERENCES

Adler, P. M., Jacquin, C. G. & Quiblier, J. A. 1990 Flow in simulated porous media. *Intl J. Multiphase Flow* **16**, 691–712.

Adler, C., d'Humières, D. & Rothman, D. H. 1994 Surface tension and interface fluctuations in immiscible lattice gases. *J. Phys. Paris* **4**, 29–46.

Appert, C., Olson, J. F., Rothman, D. H. & Zaleski, S. 1995 Spinodal decomposition in a three-dimensional, two-phase, hydrodynamic lattice gas. *J. Statist. Phys.* **81**, 181–197.

Auzerais, F. M., Dunsmuir, J., Ferréol, B., Olson, J., Ramakrishnan, T. S., Rothman, D. H. & Schwartz, L. M. 1996 Transport in sandstone; a study based on three-dimensional microtomography. *Geophys. Res. Lett.* **23**, 705–708.

Avram, D. G. & Payatakes, A. C. 1995 Flow regimes and relative permeabilities during steady-state two-phase flow in porous media. *J. Fluid Mech.* **293**, 207–236.

Bacri, J. C., Chaouche, M. & Salin, D. 1990 Modèle simple de permeéabilités relatives croisées. *C. R. Acad. Sci. Paris* **311**, 591–597.

Bear, J. 1972 *Dynamics of Fluids in Porous Media.* Dover.

Bentsen, R. G. 1994 An investigation into whether the nondiagonal mobility coefficients which arise in coupled, two phase flow are equal. *Transport in Porous Media* **14**, 23–32.

Bentsen, R. G. & Manai, A. A. 1993 On the use of conventional cocurrent and countercurrent effective permeabilities to estimate the four generalized permeability coefficients which arise in coupled, two-phase flow. *Transport in Porous Media* **11**, 243–262.

Buckles, J., Hazlett, R., Chen, S., Eggert, K., Grunau, D. & Soll, W. 1994 Toward improved prediction of reservoir flow performance. *Los Alamos Science* **22**, 112–121.

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

Darcy, H. 1956 *Les Fontaines Publiques de la Ville de Dijon.* Victor Dalmot, Paris.

Doi, M. & Ohta, T. 1991 Dynamics and rheology of complex interfaces I. *J. Chem. Phys.* **95**, 1242–1248.

Frisch, U., Hasslacher, B. & Pomeau, Y. 1986 Lattice gas automata for the Navier-Stokes equations. *Phys. Rev. Lett.* **56**, 1505.

Genabeek, O. van & Rothman, D. H. 1996 Macroscopic manifestations of microscopic flows through porous media: Phenomenology from simulation. *Ann. Rev. Earth Planet. Sci.* **24**, 63–88.

Gennes, P. G. de 1983 Theory of slow biphasic flows in porous media. *Physico-Chem. Hydrodyn.* **4**, 175.

Ginsbourg, I. & Adler, P. M. 1994 Boundary flow condition analysis for the three-dimensional lattice Boltzmann model. *J. Phys. II, Paris* **4**, 191.

Ginsbourg, I. & d'Humières, D. 1996 Local second-order boundary method for lattice-Boltzmann models. *J. Statist. Phys.* (submitted).

Goode, P. A. & Ramakrishnan, T. S. 1993 Momentum transfer across fluid-fluid interfaces in porous media: a network model. *AIChE J.* **39**, 1124–1134.

Gueyffier, D. & Zaleski, S. 1996 Personal communication.

Gunstensen, A. K. 1992 Lattice-Boltzmann studies of multiphase flow through porous media. PhD thesis, MIT.

Gunstensen, A. K. & Rothman, D. H. 1993 Lattice-Boltzmann studies of immiscible two-phase flow through porous media. *J. Geophys. Res.* **98**, 6431–6441.

Hasimoto, H. 1959 On the periodic fundamental solutions of the Stokes equations and their application to viscous flow past a cubic array of spheres. *J. Fluid Mech.* **5**, 317–328.

Kalaydjian, F. 1990 Origin and quantification of coupling between relative permeabilities for two-phase flows in porous media. *Transport in Porous Media* **5**, 215–229.

Kinney, J. H., Breunig, T. M., Starr, T. L., Haupt, D., Nichols, M. C., Stock, S. R., Butts, M. D. & Saroyan, R. A. 1993 X-ray tomographic study of chemical vapor infiltration processing of ceramic composites. *Science* **260**, 789–792.

Ladd, A. J. C. 1994 Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 2. Numerical results. *J. Fluid Mech.* **271**, 311–339.

Lamb, H. 1932 *Hydrodynamics*, 6th Edn. Cambridge University Press.

Lindquist, W. B., Lee, S.-M., Coker, D. A., Jones, K. W. & Spanne, P. 1996 Medial axis analysis

of void structure in three-dimensional tomographic images of porous media. *J. Geophys. Res.* **101**,8297–8310.

Mannseth, T. 1991 Commentary on 'origin and quantification of coupling between relative permeabilities for two-phase flows in porous media' by F. Kalaydjin. *Transport in Porous Media* **6**, 469–471.

Olson, J. F. 1995 Two-fluid flow in sedimentary rock: Complexity, transport and simulation. PhD thesis, MIT.

Olson, J. F. & Rothman, D. H. 1995 Three-dimensional immiscible lattice gas: application to sheared phase separation. *J. Statist. Phys.* **81**, 199 -222.

Onsager, L. 1931*a* Reciprocal relations in irreversible processes I. *Phys. Rev.* **37**, 405–426.

Onsager, L. 1931*b* Reciprocal relations in irreversible processes II. *Phys. Rev.* **38**, 2265–2279.

Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. 1994 *Numerical Recipes*, 2nd Edn. Cambridge University Press.

Reif, F. 1965 *Fundamentals of Statistical and Thermal Physics*. McGraw-Hill.

Rothman, D. H. & Zaleski, S. 1994 Lattice-gas models of phase separation: interfaces, phase transitions, and multiphase flow. *Rev. Mod. Phys.* **66**, 1417–1479.

Rothman, D. H. & Zaleski, S. 1997 *Lattice-Gas Automata: Simple Models of Complex Hydrodynamics*. Cambridge University Press.

Sallès, J., Thovert, J. F. & Adler, P. M. 1993 Reconstructed porous media and their application to fluid flow and solute transport. *J. Contam. Hydrol.* **13**, 3–22.

Sallès, J., Thovert, J. F. & Adler, P. M. 1994 Transports in reconstructed porous media. In *Characterization of Porous Solids III* (ed. J. Roquerol, F. Rodríguez-Reinoso, K. S. W. Sing & K. K. Unger), pp. 211–223. Elsevier.

Sangani, A. S. 1987 Sedimentation in ordered emulsions of drops at low Reynolds numbers. *Z. Angew. Math. Phys.* **38**, 542–556.

Sheidegger, A. E. 1974 *The Physics of Flow Through Porous Media*. University of Toronto Press.

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

Soll, W., Chen, S., Eggert, K., Grunau, D. & Janecky, D. 1994 Application of the lattice-Boltzmann technique to multi-fluid flow in popous media. In *Computational Methods in Water Resources X* (ed. A. Peters), pp. 991–999. Kluwer.

Spanne, P., Thovert, J. F., Jacquin, W. B., Lindquist, W. B., Jones, K. W. & Adler, P. M. 1994 Synchroton computed microtomography of porous media: topography and transports. *Phys. Rev. Lett.* **73**, 2001–2004.

Tennekes, H. & Lumley, J. L. 1972 *A First Course in Turbulence*. MIT Press.

Yao, Y., Frykman, P., Kalaydjian, F., Thovert, J. F. & Adler, P. M. 1993 High-order moments of the phase function for real and reconstructed model porous media: a comparison. *J. Colloid Interface Sci.* **156**, 478–490.

Zarcone, C. 1994 Étude du couplage visqueux en milieu poreux: mesure des perméabilités croisées. PhD thesis, L'Institut National Polytechnique de Toulouse.

Zarcone, C. & Lenormand, R. 1994 Détermination expérimentale du couplage visqueux dans les écoulements diphasique en milieu poreux. *C. R. Acad. Sci. Paris* II **318**, 1429–1435.