Some studies of fluid mixing and transport

by Matthew David Finn, MMath

Thesis submitted to The University of Nottingham
for the degree of Doctor of Philosophy, September 2003

For my family.

Acknowledgements

First and foremost I would like to thank my supervisors, Steve and Helen. In addition to lending your impressive expertise, and endless supply of constructive comments, I am extremely grateful for your patience in reading so much draft writing, and for acting as therapists when necessary. May I also express gratitude to my family and friends, who have offered me continual support and encouragement.

Thanks to Rebecca Kinton for a tour of the Kidney Dialysis unit at Nottingham City Hospital. Thank you to Mike Clifford, Hassan Aref and Philip Boyland for discussions on the batch stirring device. I am particularly indebted to Nick Trefethen for generously donating his time and expertise in complex-variable techniques and Matlab; without his help, much of the work in this thesis could not have been undertaken.

This research was supported by a grant from the Engineering and Physical Sciences Research Council.

Contents

Abstract

In this thesis we study four problems with potential biological and industrial applications which rely on fluid mixing and transport.

The problem of simultaneous ultrafiltration, diffusion and osmosis across a membrane separating two fluids is studied, numerically and asymptotically, as a model for an artificial kidney dialyser. Couplings between the different transport mechanisms prove significant in determining overall transport rates. Our model appears to be the first to treat the three transport mechanisms in a spatially structured framework, and shows that previous, spatially averaged models can overestimate transport rates. Our results can be used to optimise dialyser geometry and to profile dialysis sessions.

The remainder of this thesis concerns some fundamentals of fluid mixing and mixer design. Techniques for assessing the quality of fluid mixing are reviewed, and applied to a two-dimensional laminar chaotic flow. We find no outright optimum mixing method across the range of measures, suggesting that 'sieving' a collection of mixing methods according to increasingly complicated mixing measures may fail to identify a global optimum.

'Topological chaos' appears to allow good mixing stretch rate to be built-in to batch mixer design, avoiding the need to tune the mixer parameters, provided a correct flow topology is created. We show that the theoretical stretch rate predictions are achieved quite tightly, in practice in a significant fraction of the flow domain; we investigate the practicalities of topologically chaotic mixers.

Finally, we discuss whether topological chaos may also apply to three-dimensional static mixer design, in a braided pipe mixer, in which pipe flow is mixed around carefully designed twisted inner pipes. We expect such a device to mix well if the inner pipes have appropriate topology. However, we demonstrate how three-dimensional flow features can undermine mixing performance.

Chapter 1

Introduction

Fluid mixing and transport are important in many biological and industrial scenarios. Biological examples include artificial lung ventilation of premature infants [117] and artificial kidney dialysis for kidney failure victims [135]. The Chemical Engineering industry alone, where mixing is used to blend products or facilitate chemical reactions, is worth at least £42 billion in sales to the United Kingdom economy each year [39].

The widespread occurrence and applicability of mixing and transport is the motivation for this thesis, in which we study aspects of fluid mixing and transport with potential biological and industrial applications. On the biological side, a particular problem of interest is the mixing and transport of blood solutes in kidney dialysis machines [36]. With a view to industrial applications, we begin by examining some fundamentals of fluid mixing, such as how to optimise a mixing protocol against a variety of mixing measures [66], before investigating how to apply the relatively new concept of topological chaos to the design of chaotic, laminar-flow mixers [20].

1.1 Artificial kidney dialysis

The primary role of the kidneys is to regulate the water and solute content of blood. Kidneys are also responsible for removing wastes from the body. Knep-

per and Rector [77] give a detailed physiological explanation of kidney function. Silverthorn [135] has also written a description of the kidney, with beautiful illustrations, from the whole kidney anatomy down to the microscopic filtering units called nephrons where regulation takes place. A primitive mathematical model of nephron function was formulated by Hoppensteadt and Peskin [63].

When a person suffers total kidney failure, either acutely (e.g. food poisoning) or chronically (e.g. diabetes), artificial dialysis is needed to perform the kidney function. Most commonly, dialysis involves connecting a patient's blood supply to an external machine, where it is regulated and cleaned, and then passed back into the body. The patient would typically endure three or four hospital visits per week, connected to the machine for up to four hours each session. There are alternative methods of dialysis, such as peritoneal dialysis, where a patient's peritoneal cavity membrane is used as a dialyser. While it can be used continuously, peritoneal dialysis is less efficient than a kidney dialysis machine. A description of the various forms of kidney dialysis is given by Daugirdas, Blake and Ing [36].

In Chapter 2 we study the dialyser, the key component of an artificial kidney dialysis machine, which allows removal of solutes and water from a patient's blood by exchange across a semi-permeable membrane separating the blood from an artificial dialysate fluid contained within the machine (see Figure 1.1). We model the dialysis process by studying transport of fluid and solute across a membrane separating two fluid channels. Whilst the mixing involved is very simple, there is a non-trivial feedback between motions of fluid and solute when the effects of ultrafiltration (a pressure-gradient driven fluid flux), osmosis (a solute-gradient driven fluid flux) and solute diffusion are considered together.

Whilst ultrafiltration, diffusion and osmosis have been studied before in various contexts [1, 5, 6, 31, 32, 51, 78, 79, 80, 81, 82, 109, 110, 114, 115, 116, 118, 136, 138, 144, 150, 152], our model appears to be the first to allow an investigation of the combined effects of ultrafiltration, osmosis and diffusion upon solute and water fluxes, and in particular the extent to which these differ from a straightforward superposition of individual effects. We also investigate whether there is an optimum channel geometry in which solute and water transport rates are maximised, and whether flow inlet conditions may be tuned to select indepen-

Figure 1.1: An artificial kidney fibre bundle, containing thousands of hollow membrane fibres. Blood is passed along these fibres, whilst another fluid, the dialysate, is passed outside the fibres. Solutes and water are transported across millions of microscopic pores in the fibre walls.

dently solute and water fluxes. Our expressions for water and solute fluxes can be used as a basis for creating optimal profiling schemes for artificial kidney dialysis machines with regard to treatment time and patient comfort.

1.2 Mixing measures for Stokes flow

Mixing performance is important in many situations. For example, in an industrial chemical reaction, the details of the method by which ingredients are blended can greatly influence both the amount and quality of the product [30], and it is the role of the engineer to optimise the mixer design. A flow regime commonly exploited in the polymer manufacturing industry [46, 154], and studied in this thesis, is slow viscous flow [2], or Stokes flow, in which stretching and folding of fluid creates chaotic particle paths. (Slow flow is employed so as not to damage the delicate polymer molecules.)

Optimising the settings for a mixer, for a given purpose, is an extremely difficult problem because chaotic flow simulations (see Figure 1.2) are computationally expensive and therefore exhaustive parameter-space searches are usually intractable. Furthermore, there are many techniques for quantifying different aspects of mixing, requiring varying amounts of computational effort, and these have been used in a variety of applications [25, 107, 111]. A heuristic method called numerical *sieving* [66] is a technique that attempts to circumvent this difficulty. The method is employed to find a global optimum design according to a chosen mixing measure when an exhaustive search exclusively using this measure is infeasible. Instead, a large initial set of candidate designs is sieved according to other, computationally cheaper measures, leaving a subset of candidates from which to select an optimum design according to the measure of interest. Sieving relies on understanding the relationship between different mixing measures, but this relationship has been the subject of very little research. We shall see that a lack of correlation between different measures can undermine the effectiveness of sieving in determining optimum mixer design.

In Chapter 3 we review a selection of mixing measures which commonly appear in

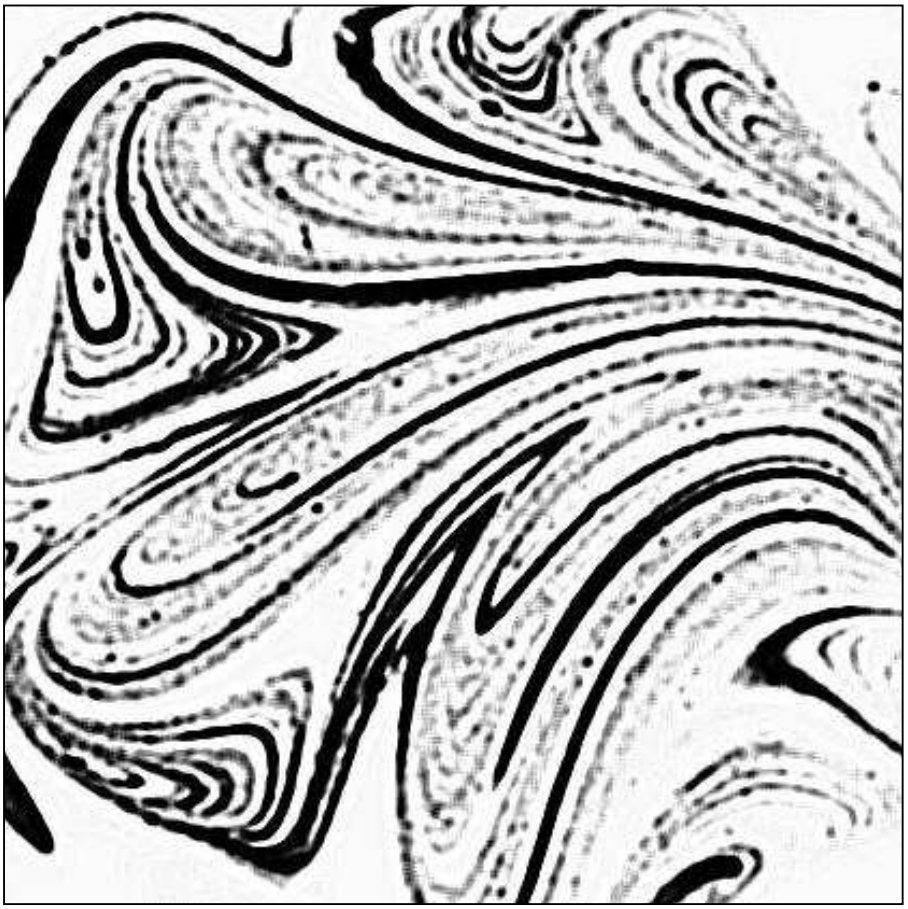

Figure 1.2: Photograph of striations created by stirring dye into golden syrup using a planetary mixer. The chaotic nature of the stretching and folding within the flow means that the striations very rapidly form an intricate structure, making computational simulation of the flow very difficult.

Figure 1.3: A typical industrial planetary mixer. Fluid in a vat is mixed by one or more bladed impellers that rotate as they are moved around.

the literature, including abstract dynamical systems techniques, statistical measures and measures with an immediate physical interpretation. We investigate the extent to which these measures correlate for the batch stirring device (BSD_1) [47], a two-dimensional chaotic, laminar, Stokes flow mixer [9], which is a simple model of a planetary mixer [27] (see Figure 1.3). The BSD_1 consists of a circular cylindrical vat of fluid, in which stirring is achieved by moving a single rod of circular cross-section through the fluid to generate a two-dimensional flow (the subscript on 'BSD_1' indicates that a single rod is used). The BSD_1 is significant in providing one of the first flow fields for which the velocity field in a time-dependent geometry may be determined exactly. As far as we are aware, this is the first systematic investigation into the correlations between a wide range of different mixing measures.

Each path taken by the stirring rod defines a *stirring protocol*. We compare mixing measures for a large number of different BSD_1 stirring protocols. Due to

the large number of protocols considered, we concentrate on measures that can be evaluated in an automated fashion without human intervention. The extent to which measures are correlated allows us to answer questions about how to optimise a protocol according to a given measure, or a given group of measures. Strong correlations would lend support to the concept of sieving, which relies on computationally cheap mixing measures being used as proxies for more expensive measures. By contrast, weak correlations would suggest that sieving should be applied with caution, if indeed at all.

1.3 Topological chaos in two dimensions

One weakness of the procedure described in Chapter 3 for mixer optimisation is that mixing performance can depend sensitively on the tuning of geometrical parameters (e.g. the size of a stirring impeller), and on the fluid rheology. The variety of mixing measures in use also means that best performance is not well defined. However, very recently, an elegant idea has emerged which potentially allows one particular mixing measure, material stretch rate, to be robustly built-in to a mixer. The concept is that of Topological Chaos, as described in a remarkable theoretical and experimental paper by Boyland, Aref and Stremler [20].

Topologically chaotic flows are a subset of chaotic flows in which the flow topology alone is responsible for achieving a guaranteed minimum material stretch rate, in some part of the domain, regardless of the geometrical parameters or exact fluid properties (viscosity, compressibility, rheology, etc).

One requirement for topologically chaotic flow is that a mixing device must contain at least three moving parts [20]. Surprisingly, therefore, very few industrial batch mixing devices are equipped to exploit topological chaos. Those that are appear to employ it almost accidentally, as they are not marketed as possessing any special topological flow properties. A very well known industrial device that does use topological chaos effectively is the antiquated four-pronged taffy pulling machine, shown in action in Figure 1.4.

Almost all mixing devices which have so far been modelled mathematically, in-

Figure 1.4: A taffy pulling machine, consisting of two motorised arms, each with two protruding taffy covered rods. The arms rotate so that the prong trajectories intersect causing rapid stretching of the taffy, and folding of air into the mixture.

cluding the BSD_1 studied in Chapter 3, are therefore incapable of producing topological chaos. However, such flows are readily generated in a generalised model of the batch stirring device, which is the subject of Chapter 4. We study the BSD_m, in which a circular cylindrical vat of fluid is stirred by m independently-moving stirring rods.

In order to perform meaningful simulations in the chaotic flow, accurate expressions for the velocity field are needed because numerical errors grow exponentially in time. Although we cannot determine an exact solution for the BSD_m velocity field, we are able to develop a computationally cheap, spectrally accurate (essentially exact), analytical expression for the streamfunction, which allows high-precision simulations of tracer advection experiments.

Following a summary of the work of Boyland *et al.* [20], we address the question of whether their topological ideas may be useful in practice for guiding mixer design, paying particular attention to whether the predicted material stretch rates are achieved in a physically useful area of the mixing domain, an important question which is not answered by topological chaos theory. Our modelling assumptions are validated by comparing numerical dye-advection simulations with simple experiments in a BSD_m with either three or four stirring rods. Our discussion addresses the issues of the energy input required for mixing and some practicalities of manufacturing a topologically chaotic mixer.

1.4 Topological chaos in three dimensions

Although batch mixers are used in industry, continuous-throughput mixing in three dimensions is more widely used, and is the type of flow found in biological systems (e.g. the lungs). The addition of an extra spatial dimension to the flow means that static mixing is possible, i.e. no time dependence is required in the velocity field in order to achieve chaotic particle advection [72]. Time-dependence can be, and commonly is, employed [85, 94], but many successful patented static mixer designs exist (e.g. by Kenics, Mixtec, TAH Industries and Sulzer–Chemtech). These static mixers generally consist of a circular cylindrical

Figure 1.5: Two static mixing devices. Shown at top is a Sulzer Chemtech mixer, with its square cross-section insert. At bottom is a twisted insert from a miniature Kenics static mixer, that causes stretching and folding of the fluid as it is pumped along the pipe.

pipe containing a judiciously designed insert that causes fluid in the cylinder to be 'cut and folded' in the cross-section as it is pumped axially, as illustrated in Figure 1.5.

In Chapter 5 we study a static mixer called the braided pipe mixer (BPM) inspired by the two-dimensional, topologically chaotic mixers of Chapter 4 and originally suggested by Boyland *et al.* [20]. The BPM consists of a circular cylindrical pipe containing smaller twisted pipes that cause the fluid to mix as it is driven along by an imposed pressure gradient. This may be thought of as a model for dialysate flow in a dialysis fibre bundle containing wavy or twisted fibres (see Figure 1.1 and Chapter 2), as a catheterised artery [3, 35, 68], or as a prototype industrial mixing device.

In such static mixers the axial coordinate occupies the role played by time in the two-dimensional batch mixer (BSD_m). The correspondence is not exact, however, because different fluid particles travel at different axial speeds, and some even

come to rest on the internal pipes. One important question that we are able to address, through direct comparison with the work in Chapter 4, is whether the topological ideas for two-dimensional flows may be applied to improve mixing in steady three-dimensional flows. Specifically, we examine the extent to which the stretch rate results generalise, enabling practical static mixer design to be influenced by topological considerations.

1.5 Layout of thesis

A more thorough review of the literature pertaining to the problems described above may be found in the relevant chapters, which are arranged as follows. In Chapter 2 we present our work on mixing and transport mechanisms in artificial kidney dialysis machines. In Chapter 3 [49] we discuss a selection of the wide variety of mixing measures in the literature and investigate the extent to which correlations between them can affect the optimisation technique of numerical sieving. Our analysis of the two-dimensional topologically chaotic batch stirring device is presented in Chapter 4 [50]. The three-dimensional braided pipe mixer is studied in Chapter 5 [48]. Detailed discussion of each topic is presented at the conclusion of each chapter, and our final concluding remarks are given in Chapter 6.

Chapter 2

Ultrafiltration, diffusion and osmosis in a channel dialyser

2.1 Introduction

Healthy humans have two kidneys, which form part of the urinary system, and whose primary function is to regulate the water content of, and the concentrations of solutes in, the blood. Blood is supplied to the kidneys from the aorta, along the renal artery, where it is then cleaned and passed back into the circulation along the renal vein and into the inferior vena cava. Typically blood enters the kidneys at a rate of 125 ml/min; there it traverses a branching structure of tubules ending at around one million microscopic filtering units, called nephrons. A schematic of a nephron is shown in Figure 2.1. At the beginning of each nephron is a Bowman's capsule which separates red blood cells and large proteins from a filtrate containing small molecular solutes (e.g. sodium, potassium and calcium) and waste. Solute and water content in this filtrate are regulated during passage along the proximal tubule and the loop of Henle, and in the distal tubule. An ingenious counter-current flow system operates, so that even though water cannot actively be pumped by the nephron its motion can accurately be controlled by solute transport and osmosis. Much of the water and solute is reabsorbed in peritubular capillaries, where it remains in circulation. The waste that is removed from the filtrate (urine) passes into a collecting duct. From here

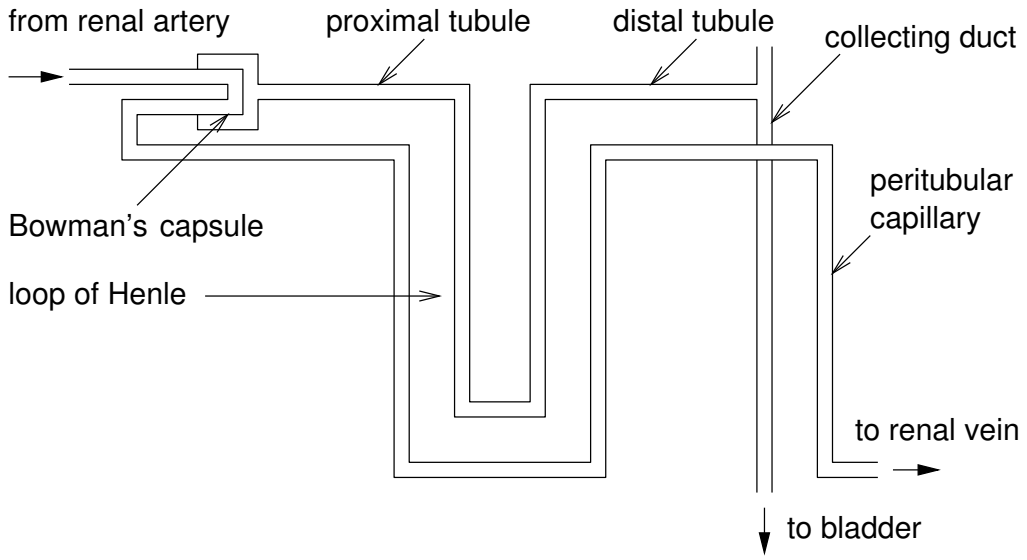

Figure 2.1: Schematic of a nephron in a human kidney. Each kidney contains around one million of these microscopic filtering units.

it is directed along two ureters to the bladder, where it is expelled from the body. Detailed physiological descriptions of kidney function have been given by Knepper and Rector [77] and Silverthorn [135]. A simple one-dimensional model of sodium ion regulation in the nephron has been presented by Hoppensteadt and Peskin [63], accounting for blood plasma dynamics and active (pumping) and passive (osmotic) motion of sodium.

Healthy human kidneys are tremendously efficient, and it is possible to live perfectly normally with just half of one kidney functioning. Unfortunately, a minority of people suffer serious kidney failure, either acutely (e.g. E-coli food poisoning or trauma) or chronically (e.g. diabetes or high blood-pressure) and they are reliant on a machine to perform their kidney function [36, 135].

Modern artificial kidneys, or dialysers, function along the same principles as, but in a much simpler way than, human nephrons. The vital component of a dialyser is an array of blood-carrying tubes (or channels) surrounded by a bath containing another fluid (dialysate); blood and dialysate are separated by a thin semi-permeable membrane. (Of course, kidney machines consist of much more than this, but the remainder of a typical machine is essentially controls, pumps and measuring equipment.) Microscopic pores in the membrane allow

water and solutes to be exchanged between the blood and dialysate. Three key transport mechanisms occur across the membrane: ultrafiltration (UF), solute diffusion (DF) and osmosis (OS). Ultrafiltration is the passage of water through the pores driven by a fluid pressure difference across the membrane. Solutes, such as sodium, calcium, potassium, waste products and foreign substances, can be removed from the blood by diffusion through the pores in the membrane. Osmosis describes the tendency of water to migrate through the membrane, up concentration gradients of solute, in order to equilibrate solute concentrations. (These three processes are all important in human kidney function.)

Patients requiring artificial dialysis may have to attend hospital up to four times per week, spending up to four hours connected to a kidney machine per session. (There are other, less common, alternative forms of dialysis, such as peritoneal dialysis, in which a patient's peritoneal cavity membrane is used a dialyser. This is less efficient than haemodialysis, but it can be performed continuously. This and other continuous dialysis methods are described in [36].) It would be of great benefit to the patient and the hospital staff if dialysis session times could be reduced. However, speeding up dialysis sessions is not a trivial task, since patients can experience serious trauma if their solute and water levels are altered too abruptly. State-of-the-art kidney machines are capable of profiling dialysis sessions to bring blood solute concentrations to a specified level during a session, but these profiling methods are still relatively crude and it is not well understood if the goals are being achieved in the shortest time with minimal patient discomfort.

There is a wealth of literature treating transfers of solute and fluid between two flows across a dividing membrane. Transport of solute by DF alone is well documented, and there exist many explicit exact expressions for the flux of solute by diffusion between two fluid flows [31, 32, 51, 78, 79, 80, 81, 82, 109, 110, 136, 138, 144, 150] (some of these articles are written in terms of mathematically equivalent heat-exchange problems, of interest in engineering applications). Simultaneous UF and DF have been studied [1, 152]. Osmotic transport effects are usually neglected in medically-oriented papers, since in artificial kidneys ultrafiltrative and diffusive fluxes are believed to be dominant. However, osmosis provides a further transport mechanism, which can, in principle, be of use in controlling

and profiling a dialysis session. Furthermore, it yields an interesting mathematical problem due to the non-linear coupling between the motions of solute and water, and has received much attention [5, 6, 114, 115, 116, 118] in this and other contexts. Osmosis is vital to real kidney function: since nephrons cannot actively move water, they instead rely on transporting solutes and the passive motion of water by osmosis. In previous studies of osmosis, authors have mainly devoted their attention to flow in wide channels, in which concentrations remain constant in the bulk flow, but where concentration boundary layers form near the membrane. However, in the context of artificial dialysis it is more appropriate to consider narrow channels, where solute concentrations vary across the whole width of the channel.

In addition to analyses of the dialyser-type problems described above there also exist simple studies of patient–dialyser interaction. Baigent *et al.* [15, 16] have investigated session profiling using a three-compartment model which looks at exchanges between a patient's body, that patient's blood and the dialysate fluid in the kidney machine. However, these spatially-averaged models do not account for the important interplay between DF, UF and OS at the dialyser membrane, nor for the concentration polarisation [38] that occurs there; these are analysed below. To make real improvements in profiled dialysis it will be necessary to understand spatial and temporal distributions of solutes in both the patient and the dialysis machine, but here we examine only the latter. Although specific recommendations for profiled dialysis are beyond the scope of this chapter, we hope that our work will eventually become part of a patient–dialyser profiling scheme. Our main aim is accurately to predict dialyser performance based on fixed design properties and variable in-vivo input conditions so that in future work the inverse problem of specifying inputs in order to achieve prescribed patient–dialyser interactions can be studied.

The outline of the rest of this chapter is as follows: in Section 2.2 we develop a simple dialyser model which allows us to study the non-trivial effect of interactions between UF, DF and OS on global solute and water transport rates. The model is nondimensionalised and simplified in Section 2.3 using parameter estimates from a typical clinical dialyser. We present a selection of numerical

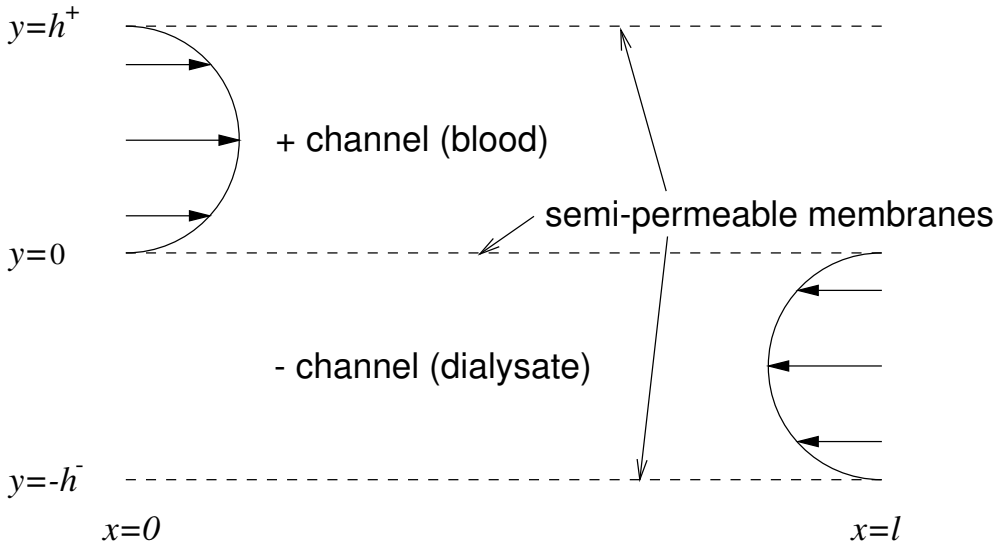

Figure 2.2: Geometry of the parallel-channel dialyser.

simulations of solute transport and fluid flow in Section 2.4 to guide our analytical progress. The construction of asymptotic solutions based on the long, thin nature of the artificial kidney fibres is then discussed in Section 2.5. Initially results for separate ultrafiltration, diffusion and osmosis are compared with numerical simulations; we then turn to the more physiologically interesting effects of coupling these processes together. Analytical solutions are presented where possible; otherwise the asymptotic problem is solved numerically. In Section 2.6 we calculate the solute and water fluxes from the completed asymptotic solution. Then in Section 2.7 we investigate the change in these fluxes with varying membrane permeabilities, input conditions, and dialyser geometry. Our conclusions are given in Section 2.8.

2.2 Dimensional model

In this section we develop a simple channel model for water and solute transport between a patient's blood and the dialysate in an artificial kidney dialysis machine.

Solute and water exchange in a modern dialyser occur across a very large membrane surface area (of the order 1m^2 [36]) between thousands of narrow cylindrical

blood-carrying fibres and a surrounding bath of dialysate fluid. For simplicity, however, we choose to work in a rectangular channel geometry, whereby blood and dialysate are transported along straight parallel two-dimensional channels, separated by plane permeable membranes. This set-up mimics an older parallel-plate dialyser design [36], although our model should still capture the essential features of transport in a cylindrical geometry. Our model geometry is depicted in Figure 2.2. Two rectangular channels of finite length l are shown. Blood and dialysate occupy the top (+) and bottom (−) channels, respectively, and the fluids are separated by a permeable membrane which runs along the x-axis. The + and − channels have heights h^+ and h^-, respectively and are assumed to form part of a periodic array,[1] whereby the boundary $y = h^+$ is identified with $y = -h^-$. Throughout this chapter, superscripts of + and − denote physical quantities associated with the corresponding channel.

We specify that blood flows from left to right in the + channel, but dialysate flow in the − channel might be prescribed in either direction (so-called 'counter-flow' is illustrated in Figure 2.2). For our present model we assume that a single solute (e.g. albumin, calcium, creatinine, glucose, phosphate, potassium, sodium or urea) is dissolved in the two fluids (henceforth called solvents) although the model can readily be extended to the case of multiple solutes. Typically the solute is more concentrated in the blood than in the dialysate.

While the transient dynamics of the system [79, 82] may be relevant to a more dynamic profiling methodology than is currently used, in the following analysis only steady-state solutions are considered. We justify this approximation on the basis of parameter estimates which indicate that, under typical conditions, concentrations in a dialyser reach their steady state within a minute. (Of course this is only a quasi-steady state since the concentrations of solutes in a patient's blood will change gradually over the four-hour period of dialysis.)

[1] As an alternative to considering this periodic array of channels, the channel walls $y = h^+$ and $y = -h^-$ can be considered to be impermeable to both solutes and to solvent, so that we are considering just a single membrane (i.e. a parallel-plate dialyser). We do not present the corresponding analysis for this case, but note that only trivial alterations to the wall boundary conditions are required and that the changes to subsequent results are readily found.

A further significant approximation we make, since we are concerned primarily with transport mechanisms, is to treat both the blood and dialysate as Newtonian fluids. In a clinical setting this is not the case, since blood contains red blood cells, which are approximately 10μm in diameter; this is comparable with the 100μm diameter of tubes typically used in a modern dialyser. In principle the model could be adapted to accommodate a more realistic blood rheology.

The analysis that follows is potentially applicable to flow in kidneys too, even though the loops of Henle, where exchange takes place (see Figure 2.1), are an order of magnitude thinner than in artificial kidneys (around 10μm [146]). This is because in the Bowman's capsule large cells (e.g. red blood cells) and molecules are filtered out and hence do not form part of the filtrate that passes into the loop of Henle.

We denote the fluid velocity components in the x and y directions by u and v, respectively, the fluid pressure by p and the concentration of solute by c. Then the solvent motion in each channel is governed by the steady-state Navier–Stokes and continuity equations

$$\begin{aligned} uu_x + vu_y &= -\frac{p_x}{\rho} + \nu(u_{xx} + u_{yy}), && (2.1)\\ uv_x + vv_y &= -\frac{p_y}{\rho} + \nu(v_{xx} + v_{yy}), && (2.2)\\ u_x + v_y &= 0, && (2.3) \end{aligned}$$

where ν is the kinematic viscosity of the solvent and ρ is solvent density (both assumed to be independent of solute concentration) and the solute concentration satisfies the advection–diffusion equation

$$uc_x + vc_y = D(c_{xx} + c_{yy}), \tag{2.4}$$

where D is the solute diffusivity. In practice the density, viscosity and diffusivities would differ between blood and dialysate, but for simplicity in this initial model the two solvents are taken to have identical physical properties.

By assuming a periodic array of blood and dialysate channels, we specify the symmetry conditions

$$u_y^+ = v^+ = 0 \qquad \text{at } y = \tfrac{1}{2}h^+, \tag{2.5}$$

$$c_y^+ = 0 \qquad \text{at } y = \tfrac{1}{2}h^+, \tag{2.6}$$

$$u_y^- = v^- = 0 \qquad \text{at } y = -\tfrac{1}{2}h^-, \tag{2.7}$$

$$c_y^- = 0 \qquad \text{at } y = -\tfrac{1}{2}h^-. \tag{2.8}$$

Since we model our solvents as viscous fluids we apply the no-slip conditions

$$u^+ = u^- = 0 \quad \text{at } y = 0. \tag{2.9}$$

At the channel inlets and outlets we prescribe a fully developed Poiseuille flow, with a uniform concentration of solute. Hence in the + channel we set the conditions

$$c^+ \to c_L^+, \quad u^+ \to -6w_L^+ y(y-h^+)/{h^+}^3, \quad v^+ \to 0 \qquad \text{as} \quad x \to -\infty, \tag{2.10}$$

$$c_x^+ \to 0, \quad u_x^+ \to 0, \quad v^+ \to 0 \qquad \text{as} \quad x \to +\infty, \tag{2.11}$$

where w_L^+ is the total volume flow rate in the + channel.

Note that for the dialysate (the − channel) the prescription of corresponding boundary conditions depends on whether there is co-flow or counter-flow. In the former case, the boundary conditions are

$$c^- \to c_L^-, \quad u^- \to -6w_L^- y(y+h^-)/{h^-}^3, \quad v^- \to 0 \qquad \text{as} \quad x \to -\infty, \tag{2.12}$$

$$c_x^- \to 0, \quad u_x^- \to 0, \quad v^- \to 0 \qquad \text{as} \quad x \to +\infty; \tag{2.13}$$

in the latter,

$$c^- \to c_R^-, \quad u^- \to -6w_R^- y(y+h^-)/{h^-}^3, \quad v^- \to 0 \qquad \text{as} \quad x \to +\infty, \tag{2.14}$$

$$c_x^- \to 0, \quad u_x^- \to 0, \quad v^- \to 0 \qquad \text{as} \quad x \to -\infty. \tag{2.15}$$

Here w_L^- and w_R^- are the total volume flow rates of dialysate for co-flow and counter-flow, respectively. In what follows we shall frequently consider co- and counter-flow cases together, writing such expressions as '$c^- \to c_{L/R}^-$ as $x \to \mp\infty$', with the obvious interpretation.

Most importantly we specify that the flux of solute j_{solute} and flux of solvent j_{solvent} across the membrane are given by

$$\begin{aligned} j_{\text{solute}} &= v^+c^+ - Dc_y^+ = v^-c^- - Dc_y^- \\ &= -k_{\text{df}}(c^+ - c^-) \quad \text{on } y = 0, \end{aligned} \tag{2.16}$$

$$\begin{aligned} j_{\text{solvent}} &= v^+ = v^- \\ &= -k_{\text{uf}}(p^+ - p^-) + k_{\text{os}}(c^+ - c^-) \quad \text{on } y = 0, \end{aligned} \tag{2.17}$$

where the coefficients $k_{\rm df}$, $k_{\rm uf}$ and $k_{\rm os}$, respectively, describe the permeability of the membrane to diffusion of solutes with low molecular weights, and to pressure-driven and osmosis-driven solvent fluxes. We assume the membrane is permeable only in the interval $0 < x < l$.

Since the solvent boundary condition (2.17) involves a pressure driven flux, in addition to conditions (2.5)–(2.17) we must also specify the pressure p at a single point in both of the channels (at the channel inlets, say). Then equations (2.1)–(2.4) together with the boundary conditions (2.5)–(2.17) form a closed system.

We now estimate the parameters of relevance to the model. Ultrafiltrative removal rates of water from blood vary between dialyser models, but a typical figure for overall transfer is around $10^{-11}\mathrm{m}^3\,\mathrm{s}^{-1}$ per Pascal pressure difference [36]. Parallel-plate dialysers typically have a total permeable-membrane area of around $1\,\mathrm{m}^2$ [36], giving an estimated ultrafiltrative permeability of $k_{\rm uf} = 10^{-11}\mathrm{m}\,\mathrm{s}^{-1}\,\mathrm{Pa}^{-1}$. To estimate $k_{\rm df}$ we begin by noting that the effectiveness of clinical dialysers at removing blood solutes is usually specified in terms of 'clearance'. A typical clearance value for small-molecule solutes in a parallel plate dialyser is $250\,\mathrm{ml}\,\mathrm{min}^{-1}$ for a blood flow rate of $300\,\mathrm{ml}\,\mathrm{min}^{-1}$ [36], meaning that 250/300 of the total solute is removed. The permeability $k_{\rm df}$ is the solute mass flux per unit membrane area per unit concentration difference. Using the blood inlet concentration as a representative concentration difference, the permeability $k_{\rm df}$ can be expressed as the inlet flow rate multiplied by the clearance divided by the membrane area. A typical total blood flow rate through a dialysis machine is of the order $3\times10^{-6}\,\mathrm{m}^3\,\mathrm{s}^{-1}$ (while dialysate flow rate is typically twice as great at $6\times10^{-6}\,\mathrm{m}^3\,\mathrm{s}^{-1}$). Thus we estimate that $k_{\rm df} = 3\times10^{-6}\times250/300/1 = 2.5\times10^{-6}\,\mathrm{m}\,\mathrm{s}^{-1}$. In an artificial kidney the movement of water by osmosis is usually neglected compared to the active movement by ultrafiltration. Hence it is difficult to estimate $k_{\rm os}$ from any clinical dialyser data, so in our numerical simulations we choose values for $k_{\rm os}$ that are large enough to demonstrate the qualitative effects of osmosis. Real kidneys, however, are unable actively to transport water, and rely on controlled passive motion by osmosis; thus in modelling a real kidney we would use larger values for $k_{\rm os}$ than would be applicable to an artificial kidney.

In addition to dialyser permeability constants we will need to know the dialyser

geometry and fluid property parameters. In a typical parallel-plate dialyser $l \sim 20 \times 10^{-2}\,\mathrm{m}$ and $h^+, h^- \sim 100 \times 10^{-6}\,\mathrm{m}$ [131]. When the blood flow is distributed amongst the ($\sim 10^4$) tubules in a hollow-fibre dialyser the corresponding flow rate is of the order $4 \times 10^{-10}\,\mathrm{m^3\,s^{-1}}$ per tubule, corresponding to a typical blood velocity $\xi = 0.05\,\mathrm{m\,s^{-1}}$.

Though the fluid properties of blood and dialysate differ, for simplicity we use $\nu = 1 \times 10^{-6}\,\mathrm{m^2\,s^{-1}}$ (viscosity of water [2]), $\rho = 1000\,\mathrm{kg\,m^{-3}}$ (density of water) and $D = 1.5 \times 10^{-9}\,\mathrm{m^2\,s^{-1}}$ (diffusivity of salt in water [57]) in both channels, noting that in reality blood is more dense and viscous than water.

2.3 Nondimensional model

In order to determine which terms are dominant, and hence to guide construction of asymptotic solutions, we first nondimensionalise (2.1)–(2.4).

Before rescaling, it is convenient to define the dimensionless parameters

$$\delta = \frac{h^+}{l}, \quad H = \frac{h^-}{h^+}, \quad \mathrm{Re} = \frac{\xi h^+}{\nu}, \quad \mathrm{Pe} = \frac{\xi h^+}{D}, \quad \epsilon = \delta \mathrm{Pe}. \tag{2.18}$$

In (2.18) the aspect ratio of the upper channel is given by δ, H is the ratio of the two channel heights, and Re and Pe are the flow Reynolds and Péclet numbers, respectively. Data presented in Section 2.2 can be used to determine all the scaling parameters. We estimate that $H \sim 1$, $\mathrm{Re} \sim 1$, $\mathrm{Pe} \sim 10^3$ and $\delta \sim 10^{-4}$, which gives $\epsilon \sim 10^{-1}$. The lower-case dimensional variables in equations (2.1)–(2.4) are now scaled according to

$$X = \frac{x}{l}, \quad Y = \frac{y}{h^+}, \quad U^\pm = \frac{u^\pm}{\xi}, \quad V^\pm = \frac{v^\pm}{\delta \xi}, \quad P^\pm = \frac{h^+ \delta p^\pm}{\rho \nu \xi}, \quad C = \frac{c}{\zeta}, \tag{2.19}$$

where upper-case letters represent dimensionless variables. The typical streamwise velocity ξ was given in Section 2.2, and we are free to choose ζ, which represents a typical solute concentration. With these changes of variable, equations (2.1)–(2.4) become

$$\mathrm{Re}\delta(UU_X + VU_Y) = -P_X + \delta^2 U_{XX} + U_{YY}, \tag{2.20}$$

$$\mathrm{Re}\delta^3(UV_X + VV_Y) = -P_Y + \delta^4 V_{XX} + \delta^2 V_{YY}, \tag{2.21}$$

$$U_X + V_Y = 0, \tag{2.22}$$

$$\epsilon(UC_X + VC_Y) = \delta^2 C_{XX} + C_{YY}. \tag{2.23}$$

We neglect terms in equations (2.20)–(2.23) containing positive powers of the small parameter δ. The system then reduces to the classical lubrication equations [2]

$$P_X = U_{YY}, \tag{2.24}$$

$$P_Y = 0, \tag{2.25}$$

$$U_X + V_Y = 0, \tag{2.26}$$

$$C_{YY} = \epsilon(UC_X + VC_Y), \tag{2.27}$$

which in the next section will be solved asymptotically using ϵ as a small parameter.

For completeness we also present here the rescaled boundary conditions (2.5)–(2.17). Note that, because of neglected terms, the system (2.24)–(2.27) is of lower order than (2.20)–(2.23) and therefore we impose fewer conditions.

The symmetry conditions

$$U_Y^+ = V^+ = 0 \quad \text{at } Y = \tfrac{1}{2}, \tag{2.28}$$

$$C_Y^+ = 0 \quad \text{at } Y = \tfrac{1}{2}, \tag{2.29}$$

$$U_Y^- = V^- = 0 \quad \text{at } Y = -\tfrac{1}{2}H, \tag{2.30}$$

$$C_Y^- = 0 \quad \text{at } Y = -\tfrac{1}{2}H, \tag{2.31}$$

may be maintained, along with the no-slip condition

$$U^+ = U^- = 0 \quad \text{at } Y = 0. \tag{2.32}$$

However, it will be seen in Section 2.5 that we are no longer able to specify the velocity profile at entry, but only the overall flow rates

$$W^+ = W_L^+ \text{ at } X = 0, \tag{2.33}$$

$$W^- = W_{L/R}^- \text{ at } X = 0/1 \tag{2.34}$$

at the channel inlets, which are given by

$$W^+(X) = \int_0^1 U^+ \mathrm{d}Y, \quad W^-(X) = \int_{-H}^0 U^- \mathrm{d}Y. \tag{2.35}$$

(Recall that in (2.34) the case $W^- = W_L^-$ at $X = 0$ applies for co-flow, while $W^- = W_R^-$ at $X = 1$ for counter-flow.) There is a corresponding restriction on our specification of the solute concentration at the channel inlets. This is discussed later, in Section 2.5.1. The important membrane boundary conditions (2.16) and (2.17) transform to give

$$\begin{aligned} J_{\text{solute}} &= V^+C^+ - C_Y^+/\epsilon = V^-C^- - C_Y^-/\epsilon \\ &= -K_{\text{df}}(C^+ - C^-) \quad \text{on } Y = 0, \end{aligned} \tag{2.36}$$

$$\begin{aligned} J_{\text{solvent}} &= V^+ = V^- \\ &= -K_{\text{uf}}(P^+ - P^-) + K_{\text{os}}(C^+ - C^-) \quad \text{on } Y = 0, \end{aligned} \tag{2.37}$$

where the dimensionless membrane permeability coefficients are

$$K_{\text{uf}} = \frac{\rho\nu}{h^+\delta^2}k_{\text{uf}}, \quad K_{\text{df}} = \frac{1}{\xi\delta}k_{\text{df}}, \quad K_{\text{os}} = \frac{\zeta}{\xi\delta}k_{\text{os}}. \tag{2.38}$$

Using the clinical parameter estimates given earlier, and the dimensional permeability estimates from Section 2.2, we find that the dimensionless ultrafiltrative and diffusive permeabilities are given by $K_{\text{uf}} \sim 0.01$ and $K_{\text{df}} \sim 0.5$, respectively. In the absence of a reliable estimate for k_{os}, we will later choose values of K_{os} to illustrate the effects of osmosis.

In the remainder of this chapter we study solutions to the system (2.24)–(2.27). To avoid repetition we now state the 'standard' parameter values that we use in the majority of the chapter:

$$\begin{gathered} W_L^+ = 1, \quad C_L^+ = 2, \quad P_L^+ = 0, \quad H = 1, \\ \begin{cases} W_L^- = 2, \quad C_L^- = 1, \quad P_L^- = 0 & \text{co-flow}, \\ W_R^- = -2, \quad C_R^- = 1, \quad P_R^- = 0 & \text{counter-flow}. \end{cases} \end{gathered} \tag{2.39}$$

Here C_L^+ and $C_{L/R}^-$ are the inlet solute concentrations in the $+$ and $-$ channels, respectively. The fluid pressures at the inlets to the $+$ and $-$ channels are P_L^+ and $P_{L/R}^-$ respectively.

2.4 Numerical solutions

In general, the full system (2.1)–(2.17) cannot be solved analytically. However, we can compute numerical solutions to examine the key features of the model in

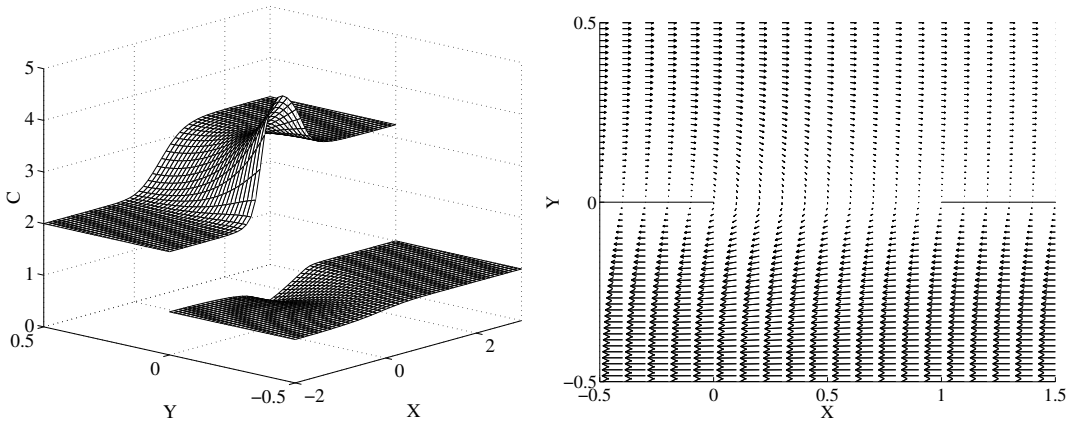

Figure 2.3: Numerical solute concentration and velocity field plots with ultrafiltration only ($K_{\text{uf}} = 0.01$, $K_{\text{df}} = 0$, $K_{\text{os}} = 0$) across the membrane. Equations (2.40)–(2.42) were solved using the standard counter-flow parameter set (2.39) and $\delta = 1$, $\text{Re} = 0$ and $\text{Pe} = 10$. The solute concentration is shown in the surface plot. The upper and lower surfaces represent the blood (+) and dialysate (-) channels, respectively. The corresponding fluid velocity field is shown in the arrow plot. The membrane separating the two channels is permeable only in the interval $0 < X < 1$.

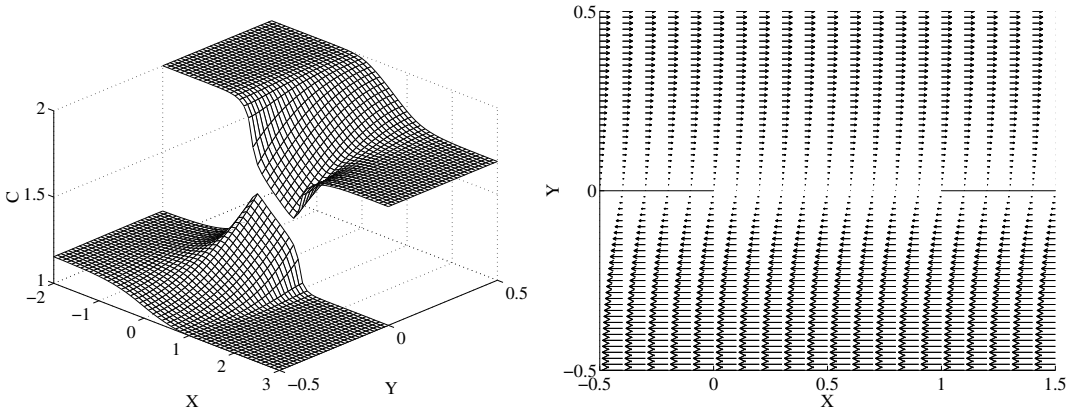

Figure 2.4: Numerical solute concentration and velocity field plots with diffusion only ($K_{\text{uf}} = 0$, $K_{\text{df}} = 1$, $K_{\text{os}} = 0$) across the membrane. Equations (2.40)–(2.42) were solved with the counter-flow parameters (2.39) and $\delta = 1$, $\text{Re} = 0$ and $\text{Pe} = 10$. The membrane separating the two channels is permeable only in the interval $0 < X < 1$.

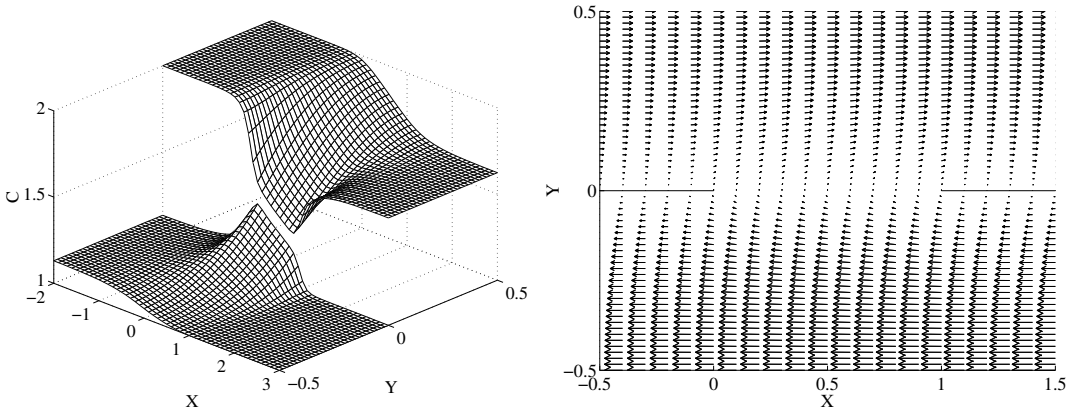

Figure 2.5: Numerical solute concentration and velocity field plots with osmosis only ($K_{\rm uf} = 0$, $K_{\rm df} = 0$, $K_{\rm os} = 1$) across the membrane. Equations (2.40)–(2.42) were solved with the parameters (2.39) and $\delta = 1$, $\mathrm{Re} = 0$ and $\mathrm{Pe} = 10$. The membrane separating the two channels is permeable only in the interval $0 < X < 1$.

order to guide our analysis of the reduced system in Section 2.3.

Using a successive-over-relaxation finite-difference scheme, we solve numerically the full streamfunction–vorticity formulation of equations (2.20)–(2.23), namely

$$\delta^2 \Psi_{XX} + \Psi_{YY} = -\Omega, \tag{2.40}$$

$$\delta^2 \Omega_{XX} + \Omega_{YY} = \delta \mathrm{Re}(\Psi_Y \Omega_X - \Psi_X \Omega_Y), \tag{2.41}$$

$$\delta^2 C_{XX} + C_{YY} = \epsilon(\Psi_Y C_X - \Psi_X C_Y), \tag{2.42}$$

where $U = \partial\Psi/\partial Y$ and $V = -\partial\Psi/\partial X$. (The pressure has been eliminated from the system by taking the curl of the Navier–Stokes equations (2.20) and (2.21).) Our numerical method is adapted from [119] to cope with the additional advection–diffusion equation (2.42).

Illustrative counter-flow velocity profiles and corresponding solute concentration profiles are shown in Figures 2.3, 2.4 and 2.5. These plots show the individual effects of ultrafiltration ($K_{\rm uf} = 0.01$), solute diffusion ($K_{\rm df} = 1$) and osmosis ($K_{\rm os} = 1$), respectively, using the standard parameters in (2.39) with $\delta = 1$, $\mathrm{Re} = 0$ and $\mathrm{Pe} = 10$. Note that the physically unrealistic values of the parameters used for these plots are chosen in order to emphasise the qualitative features in the figures; in Section 2.5 and beyond, we compute solutions based on the physical

parameters given in Section 2.3.

These figures demonstrate the variation of solute concentration across the width of the channels. Concentration polarisation [38] can be seen, meaning that solute concentration at the membrane is different to the mean concentration in the transverse direction. This variation is ignored in existing one-dimensional cross-sectionally averaged (well-mixed) dialysis models [74]. We show below that when spatial variation in the y-direction is neglected the cross-membrane solute and solvent fluxes can be significantly under- or overestimated (see Figures 2.12, 2.13 and 2.14 later).

2.5 Asymptotic solutions

To explore the combined effects of transport by UF, DF and OS, as illustrated in the previous section, we now outline the construction of two-dimensional asymptotic solutions, which exploit the long thin nature of the tubes that are used in dialysis machines.

2.5.1 General analysis

In this section we begin by constructing a general solution for the asymptotic flow rate and concentration profile from (2.24)–(2.27) assuming ϵ (proportional to the channel aspect ratio) to be a small parameter. (Note that we will use $\epsilon = 0.3$ in all of the plots which follow.[2]) The general form of the solution can be found without making reference to the transport mechanisms that act at the membrane. In subsequent sections, specific solutions are then constructed for particular membrane properties.

[2]When computing numerical solutions to the full equations (2.40)–(2.42) to compare with the asymptotic solutions, numerical problems associated with the small value $\delta \sim 10^{-4}$ can arise. To obtain the value of $\epsilon = 0.3$ used for the asymptotic solution plots, we use the parameters $\delta = 0.1$, Re $= 0$ and Pe $= 3$ for our corresponding numerical simulations. These parameters do not create any numerical problems, and still maintain a negligible value for δ^j for $j > 2$, an assumption made in Section 2.3 in deriving the reduced asymptotic model.

Equation (2.25) tells us that the pressures $P^{\pm}$ in both channels are independent of Y, so that (2.24) and then (2.26) can be fully integrated to give

$$U^+ = \tfrac{1}{2}P_X^+\mathcal{U}(Y) = -6W^+\mathcal{U}(Y), \tag{2.43}$$

$$U^- = \tfrac{1}{2}H^2P_{\bar{X}}^-\mathcal{U}(-Y/H) = -6W^-\mathcal{U}(-Y/H)/H, \tag{2.44}$$

$$V^+ = \tfrac{1}{2}W_X^+\mathcal{V}(Y), \tag{2.45}$$

$$V^- = -\tfrac{1}{2}W_{\bar{X}}^-\mathcal{V}(-Y/H), \tag{2.46}$$

where

$$\mathcal{U}(Y) = Y(Y-1), \tag{2.47}$$

$$\mathcal{V}(Y) = 4Y^3 - 6Y^2 + 1, \tag{2.48}$$

after applying boundary conditions (2.28), (2.30) and (2.32).

It remains now to solve (2.27) for the concentration profile. To make progress, we seek solutions for $C^{\pm}$ and $W^{\pm}$ in terms of the regular asymptotic expansions

$$C^{\pm}(X,Y) = C_0^{\pm}(X,Y) + \epsilon C_1^{\pm}(X,Y) + \epsilon^2 C_2^{\pm}(X,Y) + \cdots, \tag{2.49}$$

$$W^{\pm}(X) = W_0^{\pm}(X) + \epsilon W_1^{\pm}(X) + \epsilon^2 W_2^{\pm}(X) + \cdots, \tag{2.50}$$

(recall that $\epsilon \ll 1$). The expansion (2.50) for W can now be inserted into (2.43)–(2.46), and these in turn, together with (2.49), substituted into (2.27). In this way, we obtain a set of equations for $C_i^{\pm}$ ($i = 0, 1, 2, \ldots$) by equating terms of the same powers in ϵ. At order ϵ^0 we obtain

$$C_{0_{YY}}^+ = C_{0_{YY}}^- = 0, \tag{2.51}$$

which implies (after applying (2.29) and (2.31)) that the solute concentrations $C_0^{\pm} = C_0^{\pm}(X)$ are uniform across both channels. The physical interpretation is that diffusion across the channel is much more rapid than diffusion along the channel because of the small channel aspect ratio δ, so the concentration profile is, at leading order only, homogeneous across the channel.

Now considering terms at order ϵ^1 we obtain

$$C_{1_{YY}}^+ = -6W_0^+C_{0_X}^+\mathcal{U}(Y), \tag{2.52}$$

$$C_{1_{YY}}^- = -6W_0^-C_{0_X}^-\mathcal{U}(-Y/H)/H, \tag{2.53}$$

with the solution, after applying boundary conditions (2.29) and (2.31),

$$C_1^+ = -\tfrac{1}{2}W_0^+ C_{0_X}^+ \mathcal{C}(Y) + F^+(X), \tag{2.54}$$

$$C_1^- = -\tfrac{1}{2}W_0^- C_{0_X}^- H\mathcal{C}(-Y/H) + F^-(X). \tag{2.55}$$

where

$$\mathcal{C}(Y) = Y^4 - 2Y^3 + Y. \tag{2.56}$$

The functions of integration, $F^+(X)$ and $F^-(X)$ are determined later when we consider the properties of the membrane that separates the two channels.

By considering terms at order ϵ^2 in (2.27) we obtain a further pair of second-order equations for $C_2^\pm$. These equations are readily solved, subject to (2.29) and (2.31), to give

$$\begin{aligned} C_2^+ = {} & k_8^+(X)Y^8 + \ldots + k_2^+(X)Y^2 + k_0^+(X) + \{\tfrac{17}{280}{W^+}^2 C_{0_{XX}}^+ \\ & + \tfrac{17}{140}W_0^+ W_{0_X}^+ C_{0_X}^+ - \tfrac{1}{2}(W_0^+ F_X^+ + W_1^+ C_{0_X}^+)\}Y, \end{aligned} \tag{2.57}$$

$$\begin{aligned} C_2^- = {} & k_8^-(X)Y^8 + \ldots + k_2^-(X)Y^2 + k_0^-(X) - \{\tfrac{17}{280}H{W^-}^{-2} C_{0_{XX}}^- \\ & + \tfrac{17}{140}HW_0^- W_{0_X}^- C_{0_X}^- - \tfrac{1}{2}(W_0^- F_X^- + W_1^- C_{0_X}^-)\}Y, \end{aligned} \tag{2.58}$$

where the $k_i^\pm$ are complicated functions of X, not presented here. It will be seen below that only the coefficients of Y in $C_2^\pm$ are important in determining the functions $F^\pm$.

In principle we can continue constructing higher-order terms in (2.49) by balancing terms at higher powers of ϵ in (2.27), however the algebra quickly becomes tedious so it is omitted here. Also we note that it may be inconsistent to construct higher-order corrections in ϵ, since we have already neglected Reynolds number (Re) and aspect ratio (δ) effects in obtaining (2.24)–(2.27). Hence from here onward we specify quantities up to first order in ϵ only.

Thus far we have made no reference to the solute and solvent flux boundary conditions (2.36) and (2.37). It is these conditions that we now use in order to determine the remaining unknown functions $W_0^\pm(X)$, $W_1^\pm(X)$, $C_0^\pm(X)$ and $F^\pm(X)$; this will complete our solution up to terms in ϵ^1.

In the analysis that follows we shall consider various asymptotic limits for the three transport mechanisms and so it proves convenient to write the three mem-

brane permeability constants as

$$K_{\rm uf} = K_{\rm uf0} + \epsilon K_{\rm uf1}, \tag{2.59}$$

$$K_{\rm df} = K_{\rm df0} + \epsilon K_{\rm df1}, \tag{2.60}$$

$$K_{\rm os} = K_{\rm os0} + \epsilon K_{\rm os1}. \tag{2.61}$$

We substitute the expansions (2.49), (2.50) and (2.59)–(2.61) into equations (2.36) and (2.37). In addition we differentiate equation (2.37) with respect to X in order to eliminate $P^{\pm}$ in favour of $W^{\pm}$ by using (2.43) and (2.44). By collecting together terms at order ϵ^0 and ϵ^1 we obtain the eight equations

$$W^+_{0XX} = 24K_{\rm uf0}(W^+_0 - W^-_0/H^3) + 2K_{\rm os0}(C^+_{0X} - C^-_{0X}), \tag{2.62}$$

$$W^-_{0XX} = -W^+_{0XX}, \tag{2.63}$$

$$(W^+_0 C^+_0)_X = -2K_{\rm df0}(C^+_0 - C^-_0), \tag{2.64}$$

$$(W^-_0 C^-_0)_X = 2K_{\rm df0}(C^+_0 - C^-_0), \tag{2.65}$$

$$\begin{aligned} W^+_{1XX} &= 24K_{\rm uf0}(W^+_1 - W^-_1/H^3) + 24K_{\rm uf1}(W^+_0 - W^-_0/H^3) \\ &\quad + 2K_{\rm os0}(F^+_X - F^-_X) + 2K_{\rm os1}(C^+_{0X} - C^-_{0X}), \end{aligned} \tag{2.66}$$

$$W^-_{1XX} = -W^+_{1XX}, \tag{2.67}$$

$$\begin{aligned} (F^+W^+_0)_X &= \tfrac{17}{140}(W^{+2}_0 C^+_{0X})_X - (W^+_1 C^+_0)_X \\ &\quad - 2K_{\rm df0}(F^+ - F^-) - 2K_{\rm df1}(C^+_0 - C^-_0), \end{aligned} \tag{2.68}$$

$$\begin{aligned} (F^-W^-_0)_X &= \tfrac{17}{140}H(W^{-2}_0 C^-_{0X})_X - (W^-_1 C^-_0)_X \\ &\quad + 2K_{\rm df0}(F^+ - F^-) + 2K_{\rm df1}(C^+_0 - C^-_0) \end{aligned} \tag{2.69}$$

for the eight unknown functions $W^{\pm}_0$, $W^{\pm}_1$, $C^{\pm}_0$ and $F^{\pm}$. Note the perfect derivatives in (2.62)–(2.69); these indicate which quantities are conserved in the absence of any fluxes through the membrane (i.e. when $K_{\rm uf} = K_{\rm df} = K_{\rm os} = 0$).

Because we introduced an extra derivative in equations (2.62), (2.63), (2.66) and (2.67) by eliminating the pressure $P^{\pm}$, we require four further conditions to determine the model fully. By substituting (2.45) and (2.46) for $V^{\pm}$ into (2.37) we find these to be

$$W^+_{0X} = +2K_{\rm os0}(C^+_0 - C^-_0) - 2K_{\rm uf0}(P^+_0 - P^-_0) = -W^-_{0X}, \tag{2.70}$$

$$\begin{aligned} W^+_{1X} &= +2K_{\rm os0}(F^+_0 - F^-_0) + 2K_{\rm os1}(C^+_0 - C^-_0) - 2K_{\rm uf1}(P^+_0 - P^-_0) \\ &\quad - 2K_{\rm uf0}(P^+_1 - P^-_1) = -W^-_{1X}. \end{aligned} \tag{2.71}$$

These may be considered as additional boundary conditions, fixing the cross-membrane pressure difference, to be applied at either end of the channel, as appropriate to the co- or counter-flow problems.

The flow rates $W_0^{\pm}$ and $W_1^{\pm}$ at the inlets are given by

$$W_0^+ = W_L^+, \quad W_1^+ = 0 \qquad \text{at } X = 0, \tag{2.72}$$

$$W_0^- = W_{L/R}^-, \quad W_1^- = 0 \qquad \text{at } X = 0/1, \tag{2.73}$$

and the leading-order inlet concentrations are chosen to be

$$C_0^+ = C_L^+ \qquad \text{at } X = 0, \tag{2.74}$$

$$C_0^- = C_{L/R}^- \qquad \text{at } X = 0/1. \tag{2.75}$$

Finally, boundary conditions on F^+ and F^- are derived by specifying that the solute inlet fluxes to the $+$ and $-$ channels are exactly $W_L^+ C_L^+$ and $W_{L/R}^- C_{L/R}^-$, respectively. To see how this is achieved, consider the solute flux that enters the $+$ channel at $X = 0$; by integrating the expansion for U^+C^+ across the channel inlet we find this to be

$$\int_0^1 U^+C^+ \mathrm{d}Y = W_0^+C_0^+ + \epsilon(W_0^+F^+ + W_1^+C_0^+ - \tfrac{17}{140}{W_0^+}^2 C_{0_X}^+) + \mathcal{O}(\epsilon^2). \tag{2.76}$$

We have already imposed that $W_1^+ = 0$ at $X = 0$ and so, in order that the order-ϵ contribution to the flux is zero, we choose

$$F^+ = \tfrac{17}{140} W_0^+ C_{0_X}^+ \quad \text{at } X = 0. \tag{2.77}$$

By performing a similar calculation in the $-$ channel we find that

$$F^- = \tfrac{17}{140} H W_0^- C_{0_X}^- \quad \text{at } X = 0/1, \tag{2.78}$$

where, as usual, we apply the condition at $X = 0$ for co-flow and at $X = 1$ for counter-flow.

The asymptotic problems at order ϵ^0 and ϵ^1 are now fully specified. It remains to construct solutions for $W_0^{\pm}$, $W_1^{\pm}$, $C_0^{\pm}$ and $F^{\pm}$ from the eight equations (2.62)–(2.69) subject to (2.70)–(2.75), (2.77) and (2.78).

In general it is not possible to find an exact solution to equations (2.62)–(2.69). It is, however, possible to solve them in certain special cases, and these are

discussed in the next three sections. Numerical solutions to (2.62)–(2.69), on the other hand, can be computed for any combination of membrane permeability constants (see Section 2.5.5); they complement the exact asymptotic solutions to (2.62)–(2.69) where these are available, and allow us to extend our results to cases where exact solutions are unobtainable.

2.5.2 Ultrafiltration-dominated flow

Consider the case of strong ultrafiltration, so that $K_{\mathrm{uf0}} > 0$, $K_{\mathrm{uf1}} = K_{\mathrm{df}} = K_{\mathrm{os}} = 0$. In this case the solution to (2.62) and (2.63) is

$$W_0^+ = +\alpha_1 \cosh(\lambda_{\mathrm{u}} X) + \alpha_2 \sinh(\lambda_{\mathrm{u}} X) + (\beta_1 X + \beta_2)/(1+H^3), \quad (2.79)$$

$$W_0^- = -\alpha_1 \cosh(\lambda_{\mathrm{u}} X) - \alpha_2 \sinh(\lambda_{\mathrm{u}} X) + (\beta_1 X + \beta_2)H^3/(1+H^3), \quad (2.80)$$

where $\lambda_{\mathrm{u}}^2 = 24K_{\mathrm{uf0}}(1+1/H^3)$ and α_1, α_2, β_1 and β_2 are constants, determined by the boundary conditions (2.70), (2.72) and (2.73). For co-flow the constants are readily found to be

$$\alpha_1 = (H^3 W_{0_L}^+ - W_{0_L}^-)/(1+H^3), \quad (2.81)$$

$$\alpha_2 = -2K_{\mathrm{uf0}}(P_{0_L}^+ - P_{0_L}^-)/\lambda_{\mathrm{u}}, \quad (2.82)$$

$$\beta_1 = 0, \quad (2.83)$$

$$\beta_2 = W_{0_L}^+ + W_{0_L}^-. \quad (2.84)$$

When applying the counter-flow boundary conditions (2.70) it is necessary to integrate W^- along the channel in order to relate the inlet pressures, which are specified at opposite ends. By considering (2.44) we see that $P_X^- = -12W^-/H^3$, and by integrating this equation over the length of the $-$ channel it follows that

$$\begin{aligned} P_{0_L}^- - P_{0_R}^- &= 12\int_0^1 W_0^- \mathrm{d}X/H^3 \\ &= \frac{12}{H^3}\left\{-\frac{\alpha_1}{\lambda_{\mathrm{u}}}\sinh\lambda_{\mathrm{u}} + (1-\cosh\lambda_{\mathrm{u}})\frac{\alpha_2}{\lambda_{\mathrm{u}}} + \frac{(\beta_1/2+\beta_2)H^3}{1+H^3}\right\} \end{aligned} \quad (2.85)$$

This expression can be used to write $P_{0_L}^-$ in terms of $P_{0_R}^-$ and allow the boundary conditions (2.70) to be applied at the left-hand end of the channels, $X = 0$, as with co-flow. The constants α_1, α_2, β_1 and β_2 are then readily determined, but are not presented here as their algebraic expressions are too cumbersome.

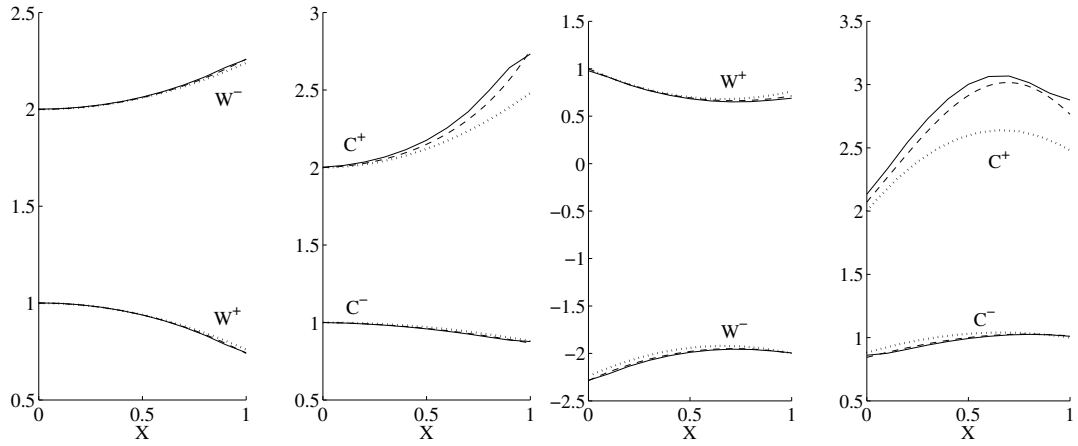

Figure 2.6: Flow rate and concentration plots with only ultrafiltrative flux of solvent ($K_{\rm uf} = 0.02$, $K_{\rm df} = 0$, $K_{\rm os} = 0$) through the membrane. The left-hand pair of plots correspond to co-flow; the right-hand pair to counter-flow. The dotted lines show the solution obtained by assuming $K_{\rm uf} = \epsilon K_{\rm uf1}$ is small, the dashed lines show the asymptotic solution with $K_{\rm uf} = K_{\rm uf0} = \mathcal{O}(1)$, and the solid lines show the corresponding full numerical solution.

The order-ϵ^1 flow rates $W_1^\pm$ satisfy (2.66) and (2.67), which are of the same form as for the order-ϵ^0 flow rates $W_0^\pm$; thus $W_1^\pm$ are given by expressions of the form (2.79) and (2.80), although the constants of integration corresponding to α_1, α_2, β_1 and β_2 are different, and are determined by the boundary conditions (2.71), (2.72) and (2.73).

This completes the solution for the flow rates. Now we can return to find the concentration profile. The functions $C_0^\pm$ and $F^\pm$ are easily determined: since no solute can pass through the membrane, the solute conservation equations (2.64), (2.65), (2.68) and (2.69) can be integrated trivially (subject to (2.74), (2.75), (2.77) and (2.78)) to give

$$C_0^+ = W_L^+ C_L^+ / W_0^+, \tag{2.86}$$

$$C_0^- = W_{L/R}^- C_{L/R}^- / W_0^-, \tag{2.87}$$

$$F^+ = \tfrac{17}{140} W_0^+ C_{0_X}^+ - W_1^+ C_0^+ / W_0^+, \tag{2.88}$$

$$F^- = \tfrac{17}{140} H W_0^- C_{0_X}^- - W_1^- C_0^- / W_0^-. \tag{2.89}$$

Illustrative plots of the flow rate and concentration functions, $W^\pm$ and $C^\pm$ are

given in Figure 2.6 for both co- and counter-flow. It can be seen that for co-flow the flow rates are monotonic functions of X and hence the solvent flux is always negative (directed towards the $-$ channel); by contrast, in a counter-flow regime the flow rates are not monotonic, implying that the direction of the solvent flux can change along the length of the membrane.

These solutions are physically meaningful only if neither of the flow rates $W^{\pm}$ changes sign along the length of the membrane, otherwise there would be a stagnation region in one channel where $W^{\pm} = 0$, and an undesirable backflow, with fluid entering one of the channels from both ends. Our mathematical solution also breaks down in this case, accumulation of solute leading to an infeasible infinite steady-state concentration in the stagnation region. In terms of the total solvent flux Q_{solvent} across the membrane $Y = 0$ (which is derived later in Section 2.6), the maximum flux sensibly permitted is $|Q_{\text{solvent}}| = \frac{1}{2}W_L^+$ (the factor of one half is due to the equal flux across the membranes at $Y = 0$ and at $Y = 1$).

2.5.3 Diffusion-dominated flow

In this case we assume that the ultrafiltrative and osmotic permeabilities of the membrane are small so that $K_{\text{uf0}} = K_{\text{os0}} = 0$.

At leading-order, (2.62) and (2.63) tell us that W^+ and W^- are constants equal to their inlet values, given by (2.72) and (2.73), respectively. This is consistent with there being zero solvent flux at leading order across the membrane if ultrafiltration and osmosis effects are weak. Substituting these values into equations (2.64) and (2.65) generates two coupled linear ordinary differential equations for $C_0^{\pm}$. For the co-flow problem with boundary conditions (2.74) and (2.75) the corresponding solution is

$$C_0^{\pm} = \frac{W_L^+ C_L^+ + W_L^- C_L^- \pm W_L^{\mp}(C_L^+ - C_L^-)\exp(-\lambda_{\text{d}} X)}{W_L^+ + W_L^-}, \tag{2.90}$$

where $\lambda_{\text{d}} = 2K_{\text{df0}}(W_L^+ + W_L^-)/(W_L^+ W_L^-)$. For counter-flow the solution is

$$C_0^+ = \frac{W_L^+ C_L^+ + \gamma W_R^- C_R^- + W_R^-(C_L^+ - C_R^-)\exp\left(\lambda_{\text{d}'}(1-X)\right)}{W_L^+ + \gamma W_R^-}, \tag{2.91}$$

$$C_0^- = \frac{W_L^+ C_L^+ + \gamma W_R^- C_R^- - W_L^+(C_L^+ - C_R^-)\exp\left(\lambda_{\text{d}'}(1-X)\right)}{W_L^+ + \gamma W_R^-}, \tag{2.92}$$

where $\lambda_{d'} = 2K_{df0}(W_L^+ + W_R^-)/(W_L^+ W_R^-)$ and $\gamma = \exp(\lambda_{d'})$. Using these solutions for $W_0^\pm$ and $C_0^\pm$ we continue to solve equations (2.66) and (2.67) subject to (2.71)–(2.73). In the case of co-flow the solution is

$$\begin{aligned} W_1^+ = -W_1^- &= K_{uf1}\left\{ 12\frac{{W_L^+}^2 + W_L^+ W_L^-(1-1/H^3) - {W_L^-}^2/H^3}{W_L^+ + W_L^-}X^2 \right. \\ &\quad \left. - 2(P_L^+ - P_L^-)X \right\} \\ &\quad - \frac{K_{os1}}{K_{df0}}\frac{W_L^+ W_L^-}{W_L^+ + W_L^-}(C_L^+ - C_L^-)\left(\exp(-\lambda_d X) - 1\right), \end{aligned} \quad (2.93)$$

while for counter-flow the expressions are

$$\begin{aligned} W_1^+ &= \frac{W_L^+ W_R^-}{W_L^+ + \gamma W_R^-}\frac{K_{os1}}{K_{df0}}(C_L^+ - C_R^-)\gamma(1 - \exp(-\lambda_{d'}X)) \\ &\quad + K_{uf1}\left\{-2(P_L^+ - P_R^-)X + 12W_L^+ X^2 - 12W_R^- X(X-2)/H^3\right\} \end{aligned} \quad (2.94)$$

$$\begin{aligned} W_1^- &= \frac{W_L^+ W_R^-}{W_L^+ + \gamma W_R^-}\frac{K_{os1}}{K_{df0}}(C_L^+ - C_R^-)(-1 + \exp(-\lambda_{d'}(X-1))) \\ &\quad - K_{uf1}(X-1) \\ &\quad \left\{-2(P_L^+ - P_R^-) + 12W_L^+(X+1) - 12W_R^-(X-1)/H^3\right\}. \end{aligned} \quad (2.95)$$

Finally, we determine $F^\pm$ by substituting for $W_0^\pm$, $C_0^\pm$ and $W_1^\pm$ in equations (2.68) and (2.69). These equations are easily solved exactly using the techniques described above, but the solutions are algebraically cumbersome and so are not presented here.

Illustrative flow rate and concentration functions, $W^\pm$ and $C^\pm$, are plotted in Figure 2.7 for both co- and counter-flow. The $\mathcal{O}(\epsilon)$ correction terms to $C^\pm$ indicate that the concentration is depleted below (raised above) its 'inlet' value in the $+$ $(-)$ channel, a result which is consistent with the full numerical simulations in Figure 2.4. We note also that more solute is removed from the $+$ channel using counter-flow than using co-flow conditions [74]. By employing a counter-flow regime, a much shorter membrane is needed to achieve a given solute flux, than with co-flow.

2.5.4 Osmotically-driven flow

Osmotic flux through the permeable membrane is not usually considered to be important in a clinical setting, where ultrafiltration and diffusion are the dominant

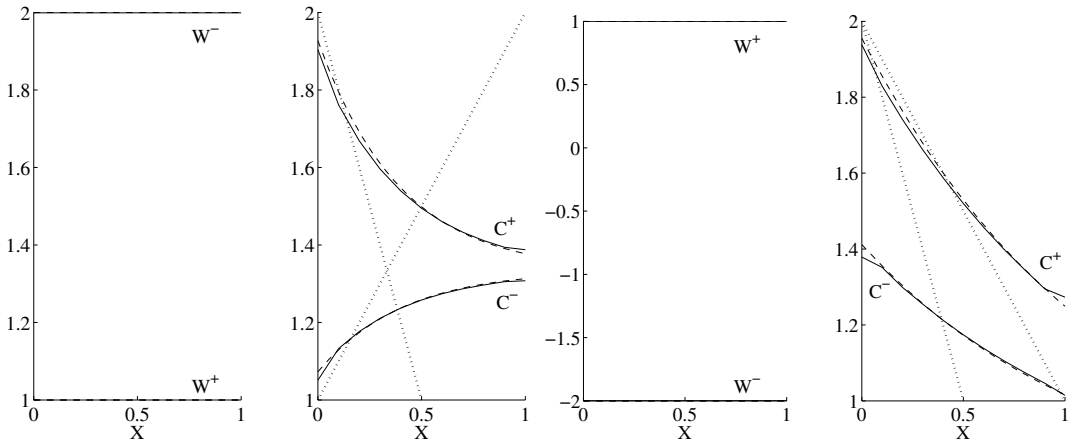

Figure 2.7: Concentration plots with only diffusive solute flux ($K_{\rm uf} = 0$, $K_{\rm df} = 1$, $K_{\rm os} = 0$) through the membrane. The left-hand and right-hand pair of plots correspond to co-flow and counter-flow, respectively. The dotted lines show the (linear) solution obtained by assuming $K_{\rm df} = \epsilon K_{\rm df1}$ is small, the dashed lines show the asymptotic solution with $K_{\rm df} = K_{\rm df0} = \mathcal{O}(1)$, and the solid lines show the corresponding full numerical solution.

effects [36]. However, osmotic flux may be exploited, for the purpose of dialysis session profiling, by adding large-molecular-weight solutes to the dialysate, to increase its osmotic potential in order to remove more water from the blood. The osmosis problem is the most interesting from a mathematical point of view as it involves a non-linear coupling between solvent and solute motion at the membrane. Furthermore, the type of analysis presented here can be applied to real kidneys, which rely on osmosis for water movement, since water cannot actively be pumped around inside the nephrons [135].

In the case of strong osmosis, where $K_{\rm os0} > 0$ in equations (2.62)–(2.69) but $K_{\rm os1} = K_{\rm df} = K_{\rm uf} = 0$, we can solve only part of the problem analytically. Equations (2.62), (2.64) and (2.65) for W_0^- and $C_0^\pm$ can be solved to give these quantities parametrically in terms of W_0^+ as

$$C_0^+ = \frac{W_L^+ C_L^+}{W_0^+}, \quad C_0^- = \frac{W_{L/R}^- C_{L/R}^-}{W_T - W_0^+}, \quad W_0^- = W_T - W_0^+, \tag{2.96}$$

where W_T is the total flow rate in both channels. (For co-flow $W_T = W_L^+ + W_L^-$, but for counter-flow W_T remains to be determined.)

This leaves the first-order non-linear differential equation

$$W_{0_X}^{+} = 2K_{\text{os}0}\frac{W_L^+ C_L^+ (W_T - W_0^+) - W_{L/R}^- C_{L/R}^- W_0^+}{(W_T - W_0^+)W_0^+} \tag{2.97}$$

to determine W_0^+, which can readily be solved by separation of variables to give the implicit solution

$$\begin{aligned} 4\,\alpha_2^3 K_{\text{os}0} X &= \alpha_2(2\alpha_1 + \alpha_2(W_0^+ - 2W_T))W_0^+ - \alpha_2(2\alpha_1 + \alpha_2(W_L^+ - 2W_T))W_L^+ \\ &\quad + 2\alpha_1(\alpha_1 - W_T\alpha_2)\log\left(\frac{\alpha_1 - \alpha_2 W_0^+}{\alpha_1 - \alpha_2 W_L^+}\right), \end{aligned} \tag{2.98}$$

where $\alpha_1 = W_L^+ C_L^+ W_T$ and $\alpha_2 = W_L^+ C_L^+ + W_{L/R}^- C_{L/R}^-$.

The boundary condition $W_0^+ = W_L^+$ at $X = 0$ has been imposed in deriving (2.98). By solving (2.98) numerically, the solution $W_0^+(X)$ may be continued smoothly from $X = 0$ to $X = 1$. This continuation procedure requires special care in its numerical implementation since, for a given value of X, (2.98) may have multiple roots for W_0^+; it is, however, a straightforward matter to determine the physically meaningful root that yields a smooth function $W_0^+(X)$.

Since it is not possible to write $W_0^\pm$ and $C_0^\pm$ explicitly, it is not possible to find explicit analytical expressions for the higher-order corrections $W_1^\pm$ and $F^\pm$. However, the remaining ordinary differential equations (2.66)–(2.69) may be integrated numerically. (For counter-flow we obtain a two-point boundary value problem, and it is necessary to use a technique such as shooting.)

Flow rate and concentration functions, $W^\pm$ and $C^\pm$, are plotted in Figure 2.8 for both co- and counter-flow. Note that in this co-flow regime with purely osmotic flux at the membrane, both the flow rates $W^\pm$ and the concentrations $C^\pm$ have nearly equilibrated at $X = 1$ (and would more nearly do so if the membrane were extended); when other transport mechanisms across the membrane are incorporated, as we consider later, this is no longer true.

2.5.5 General membrane properties

The three transport mechanisms UF, DF and OS can be combined in the mathematical model in many different ways, with either strong or weak contributions

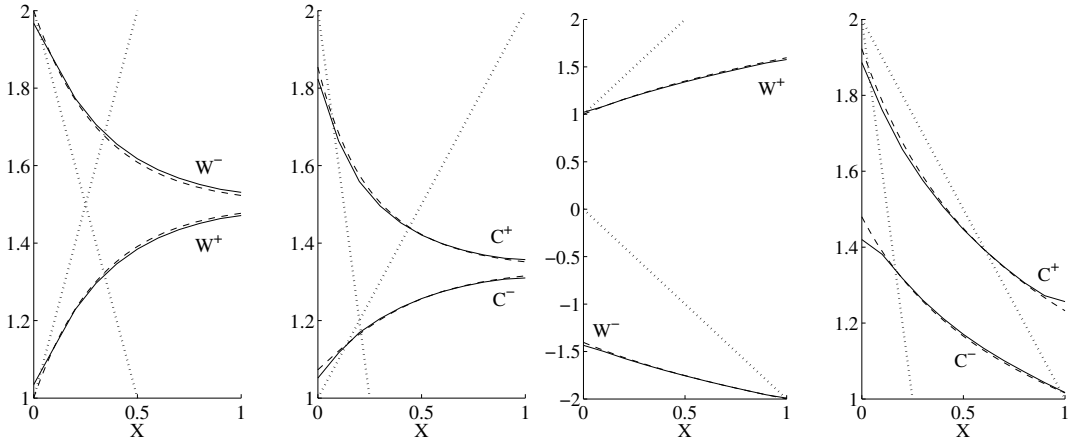

Figure 2.8: Flow rate and concentration plots with only osmotic solvent flux ($K_{\rm uf} = 0$, $K_{\rm df} = 0$, $K_{\rm os} = 1$) through the membrane. The left-hand pair of plots correspond to co-flow; the right-hand pair to counter-flow. The dotted lines show the (linear) solution obtained by assuming $K_{\rm os} = \epsilon K_{\rm os1}$ is small, the dashed lines show the asymptotic solution with $K_{\rm os} = K_{\rm os0} = \mathcal{O}(1)$, and the solid lines show the corresponding full numerical solution.

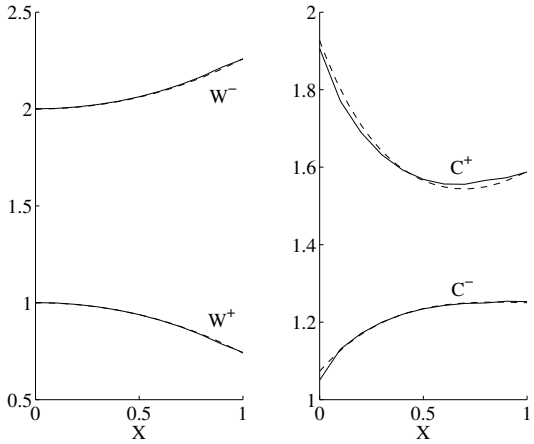

Figure 2.9: Flow rate and concentration plots with ultrafiltrative and diffusive fluxes occurring together ($K_{\rm uf} = 0.02$, $K_{\rm df} = 1$, $K_{\rm os} = 0$) for co-flow. Dashed and solid lines correspond to the $\mathcal{O}(1)$-permeability asymptotic solution, and the full numerical solution, respectively.

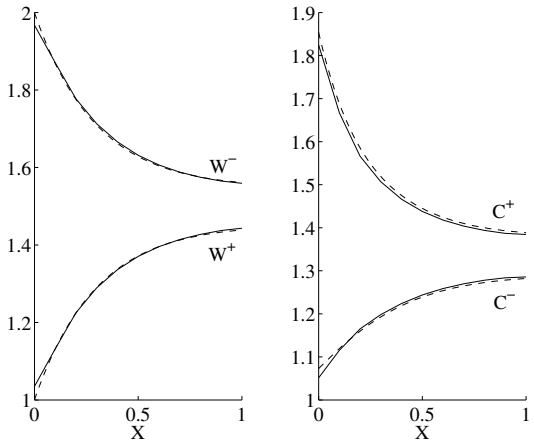

Figure 2.10: Flow rate and concentration plots with ultrafiltrative and osmotic fluxes occurring together ($K_{\text{uf}} = 0.02$, $K_{\text{df}} = 0$, $K_{\text{os}} = 1$) for co-flow. Dashed and solid lines correspond to the $\mathcal{O}(1)$-permeability asymptotic solution, and the full numerical solution, respectively.

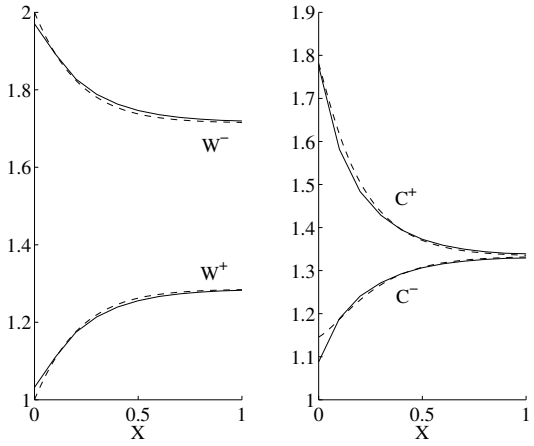

Figure 2.11: Flow rate and concentration plots with diffusive and osmotic transport occurring together ($K_{\text{uf}} = 0$, $K_{\text{df}} = 1$, $K_{\text{os}} = 1$) for co-flow. Dashed and solid lines correspond to the $\mathcal{O}(1)$-permeability asymptotic solution, and the full numerical solution, respectively.

from each transport type (though not all these regimes may arise in biological contexts). As can be seen from Figures 2.6–2.8, the linear solutions found by assuming small, $\mathcal{O}(\epsilon)$ permeabilities do not correctly capture the form of the true solutions, whereas the more involved asymptotic analyses with $\mathcal{O}(1)$ permeabilities agree better. This is also apparent in Figures 2.9, 2.10 and 2.11 which show co-flow solutions with combinations of UF, DF and OS.

In Figure 2.9 with UF and DF fluxes it can be seen that for small X, near the channel inlet, the solution is diffusion dominated (cf. Figure 2.7); then for larger X ultrafiltration has a greater influence on the ultimate solute concentrations at the exit from the channels (cf. Figure 2.6).

Ultrafiltrative and osmotic fluxes compete with each other in Figure 2.10; here the flow rates in the $+$ and $-$ channels tend to distinct values at the channel exits. This is in contrast to the situation with OS only (in Figure 2.8) where the flow rates equilibrate.

In Figure 2.11 the flow rates and concentrations are shown for diffusive and osmotic transport through the membrane. Here the two transport mechanisms compete, with solute and solvent driven in opposite directions across the membrane. The concentrations equilibrate relatively quickly, and this attenuates the osmotic flux, so that the flow rates are bounded away from each other as X becomes larger (again in contrast to Figure 2.8, where the flow rates come closer to convergence).

Later, in Section 2.7.1, we discuss the effect of coupling together UF, DF and OS on overall solute and solvent removal rates. In the next section we describe how these fluxes are calculated from the asymptotic solutions described above.

2.6 Calculation of flux

2.6.1 General considerations

The total fluxes Q_{solute} and Q_{solvent} of solute and solvent through the membrane are of physiological interest, and are defined by

$$Q_{\text{solute}} = \int_0^1 J_{\text{solute}}\,\mathrm{d}X, \tag{2.99}$$

$$Q_{\text{solvent}} = \int_0^1 J_{\text{solvent}}\,\mathrm{d}X. \tag{2.100}$$

Although we already have expressions for J_{solute} and J_{solvent}, defined in (2.36) and (2.37), we can avoid explicit evaluation of the integrals in (2.99) and (2.100) by noting the following conservation arguments.

We begin by considering the flux of solvent. According to (2.26), the velocity field (U, V) is divergence free and so we can use the divergence theorem to write

$$\oint_{\partial\Omega} (U, V) \cdot \boldsymbol{n}\,\mathrm{d}S = 0, \tag{2.101}$$

where $\partial\Omega$ is the perimeter of the + channel, and $\boldsymbol{n}$ is the unit outward normal vector to the + channel. By expanding this integral and using the symmetry $V^+|_{Y=1} = -V^+|_{Y=0}$, it is easily shown that

$$Q_{\text{solvent}} = \tfrac{1}{2}\left[W^+\right]_{X=0}^{X=1}. \tag{2.102}$$

We can perform a similar procedure with the advection–diffusion equation (2.27) to obtain an expression for Q_{solute}. By integrating (2.27) over the area of the + channel, using the continuity equation (2.26), and the symmetries $V^+|_{Y=1} = -V^+|_{Y=0}$ and $C^+|_{Y=1} = C^+|_{Y=0}$, it can be shown that

$$Q_{\text{solute}} = \tfrac{1}{2}\int_0^1 \left[U^+C^+\right]_{X=0}^{X=1}\,\mathrm{d}Y. \tag{2.103}$$

This provides a particularly convenient expression for Q_{solute}, since we have already derived an expression, (2.76), for this integral.

2.6.2 Special case calculations of the flux

If the membrane permeability properties admit explicit solutions for the functions $W_0^{\pm}$, $C_0^{\pm}$, $W_1^{\pm}$and $F^{\pm}$ then we can substitute these into (2.102) and (2.103) to obtain explicit expressions for the total solute and solvent fluxes, as we now describe.

In Section 2.5.3 we calculated the solution with strong diffusion and weak ultrafiltration and osmosis; substitution of the co- and counter-flow solutions in (2.102) and (2.103) yields general flux expressions, but these are too cumbersome to present here. In the case of strong diffusion only, there is a flux of solute (but $Q_{\text{solvent}} = 0$), given for co-flow by

$$\begin{aligned} Q_{\text{solute}} &= \tfrac{1}{2} W_L^+ W_L^- \frac{C_L^+ - C_L^-}{W_L^+ + W_L^-} (\exp(-\lambda_{\text{d}}) - 1) \\ &\quad + \tfrac{17}{70} \epsilon K_{\text{df0}}^2 (1+H)(C_L^+ - C_L^-) \exp(-\lambda_{\text{d}}), \end{aligned} \tag{2.104}$$

and for counter-flow by

$$Q_{\text{solute}} = \frac{1}{2} \frac{W_L^+ W_R^-}{W_L^+ + \gamma W_R^-} (1-\gamma)(C_L^+ - C_R^-) + \mathcal{O}(\epsilon), \tag{2.105}$$

where we recall from Section 2.5.3 that $\lambda_{\text{d}} = 2K_{\text{df0}}(W_L^+ + W_L^-)/(W_L^+ W_L^-)$, $\lambda_{\text{d}'} = 2K_{\text{df0}}(W_L^+ + W_R^-)/(W_L^+ W_R^-)$ and $\gamma = \exp(\lambda_{\text{d}'})$. The $\mathcal{O}(\epsilon)$ correction term in (2.105) is algebraically cumbersome and not presented here. Variation of Q_{solute} (and the other fluxes calculated below) with the various model parameters is discussed in Section 2.7.

For strong ultrafiltration only, the solute flux is zero and the solvent flux is given by

$$Q_{\text{solvent}} = \frac{1}{2} \left(\alpha_1 (\cosh \lambda_{\text{u}} - 1) + \alpha_2 \sinh \lambda_{\text{u}} + \frac{\beta_1}{1+H^3} \right), \tag{2.106}$$

with $\alpha_1, \alpha_2, \beta_1$ and λ_{u} as defined in Section 2.5.2.

In the case of weak permeabilities only ($K_{\text{uf0}} = K_{\text{df0}} = K_{\text{os0}} = 0$), the solute and solvent fluxes are given by

$$Q_{\text{solute}} = -\epsilon K_{\text{df1}} (C_L^+ - C_L^-), \tag{2.107}$$

$$\begin{aligned} Q_{\text{solvent}} &= \epsilon K_{\text{uf1}} \left\{ 6(W_L^+ - W_L^-/H^3) - (P_L^+ - P_L^-) \right\} \\ &\quad + \epsilon K_{\text{os1}} (C_L^+ - C_L^-) \end{aligned} \tag{2.108}$$

for co-flow and by

$$Q_{\text{solute}} = -\epsilon K_{\text{df1}}(C_L^+ - C_R^-), \tag{2.109}$$

$$Q_{\text{solvent}} = \epsilon K_{\text{uf1}}\left\{6(W_L^+ + W_R^-/H^3) - (P_L^+ - P_R^-)\right\} + \epsilon K_{\text{os1}}(C_L^+ - C_R^-) \tag{2.110}$$

for counter-flow. Thus for a combination of *weak* transport mechanisms the solute and solvent fluxes are a linear superposition of the appropriate individual fluxes, so that the processes do not interact at this order.

For strong diffusion coupled with weak osmosis, the solute flux can be calculated explicitly but is too complicated to state here. The solvent flux for co-flow is

$$Q_{\text{solvent}} = \tfrac{1}{2}\epsilon\frac{K_{\text{os1}}}{K_{\text{df0}}}\frac{W_L^+ W_L^-}{W_L^+ + W_L^-}(C_L^+ - C_L^-)\left(1 - \exp(-\lambda_{\text{d}})\right), \tag{2.111}$$

while for counter-flow the flux is

$$Q_{\text{solvent}} = -\tfrac{1}{2}\epsilon\frac{K_{\text{os1}}}{K_{\text{df0}}}\frac{W_L^+ W_R^-}{W_L^+ + \gamma W_R^-}(C_L^+ - C_R^-)(1 - \gamma). \tag{2.112}$$

Expressions for the solute flux when weak ultrafiltration is coupled to strong diffusion can be calculated but they are also too unwieldy to present here. The solute motion in this case has no influence on the solvent flux, so that it is still given by (2.108) and (2.110).

While other combinations of weak and strong transport mechanisms can, in principle, be studied, it is not always possible to construct explicit expressions for the corresponding fluxes; in such cases fluxes can be estimated by substituting numerical calculations of $W_0^\pm$, $C_0^\pm$, $W_1^\pm$ and $F^\pm$ into equations (2.102) and (2.103).

2.6.3 Large-permeability flux limits

In this section we derive explicit expressions for the limiting values of the total solute and solvent fluxes when the membrane is highly permeable (through one of the three mechanisms UF, DF and OS).

As the ultrafiltrative membrane permeability increases, the physical limit to the flux of solvent is the flow rate of fluid into the channel that is being drained. As

we have discussed in Section 2.5.2 our model assumptions break down before this extreme is reached.

With only diffusive and osmotic fluxes, expressions for the large-permeability solute and solvent fluxes can be found without solving the full problem. Consider the boundary condition (2.36) in the limit $K_{\mathrm{df}} \to \infty$. In a physically well behaved system, where the solute flux remains finite, it must be the case that C^+ and C^- are (approximately) equal immediately either side of the membrane. Using this fact in combination with simple solute conservation arguments and the inlet conditions, it is straightforward to show that the large-permeability solute flux for co-flow is

$$\lim_{K_{\mathrm{df}}\to\infty} Q_{\mathrm{solute}} = -\tfrac{1}{2} W_L^+ W_L^- \frac{C_L^+ - C_L^-}{W_L^+ + W_L^-}, \tag{2.113}$$

and for counter-flow is

$$\lim_{K_{\mathrm{df}}\to\infty} Q_{\mathrm{solute}} = -\tfrac{1}{2} W_L^+ (C_L^+ - C_R^-). \tag{2.114}$$

These expressions agree with the diffusive solute fluxes (2.104) and (2.105) when the limit $K_{\mathrm{df}} \to \infty$ is taken. If we now take the inlet solute concentrations in the lower channel (C_R^- and C_L^-) to be the same in the two cases, then the ratio of the counter-flow limit (2.114) to the co-flow limit (2.113) is $1 + W_L^+/W_L^- > 1$, demonstrating that the maximum flux achievable is always greater in a counter-flow system [74].

Applying similar considerations to (2.37), we find that in the limit $K_{\mathrm{os}} \to \infty$ the osmotic solvent fluxes are

$$\lim_{K_{\mathrm{os}}\to\infty} Q_{\mathrm{solvent}} = \tfrac{1}{2} W_L^+ W_L^- \frac{C_L^+ - C_L^-}{W_L^+ C_L^+ + W_L^- C_L^-} \tag{2.115}$$

and

$$\lim_{K_{\mathrm{os}}\to\infty} Q_{\mathrm{solvent}} = \tfrac{1}{2} W_L^+ (C_L^+/C_R^- - 1), \tag{2.116}$$

for co- and counter-flow, respectively. Again, taking the same inlet concentrations in the two cases, we find that the ratio of these quantities is $1 + W_L^+ C_L^+/(W_L^- C_L^-) > 1$ so the maximum flux achievable is always greater in a counter-flow system [74].

2.7 Effects of changing various parameters

In this section we investigate the physiologically important effects on the total solute and solvent fluxes (Q_{solute} and Q_{solvent}) of varying the model parameters. We begin by examining the effects of changing the three membrane permeabilities (K_{uf}, K_{df} and K_{os}); later we discuss effects of changes in dialyser geometry and the fluid inlet conditions.

2.7.1 Effects of changes in permeability

We have calculated the relationships between solute and solvent flux and the ultrafiltrative, diffusive and osmotic membrane permeabilities in Section 2.6. These variations are plotted in Figures 2.12, 2.13 and 2.14.

Figure 2.12 shows that the ultrafiltrative flux does not depend on whether there is a co- or counter-flow regime in the small-K_{uf} limit; however, as the ultrafiltration permeability is increased, the counter-flow flux becomes significantly larger than that for co-flow. In Figures 2.13 and 2.14 we see that in the small-permeability limits the diffusive and osmotic fluxes also do not depend on whether there is a co- or counter-flow regime; again for larger permeabilities the counter-flow fluxes are significantly larger than for co-flow. It is worth noting that the leading-order asymptotic flux predictions, corresponding to a well-mixed analysis [74], always overestimate the diffusive and osmotic fluxes, while the corrected $\mathcal{O}(\epsilon^0) + \mathcal{O}(\epsilon^1)$ fluxes presented here agree significantly better with the numerical predictions.

Of much greater physiological interest is the variation of solute and solvent fluxes when a combination of ultrafiltration, diffusion and osmosis occur together. Using numerical solutions to the asymptotic problem, we have plotted in Figures 2.15, 2.16 and 2.17 the flux variation with combined UF/DF, UF/OS and DF/OS using our standard co-flow[3] parameter values.

[3]We use co-flow for these simulations, rather than the more common physiological counter-flow, because solving the ordinary differential equations (2.62)–(2.69) numerically for general counter-flow problems would require shooting with several parameters. Although this can in principle be done, it is a rather more sophisticated numerical approach which would not allow such a comprehensive exploration of parameter space, given available computational resources.

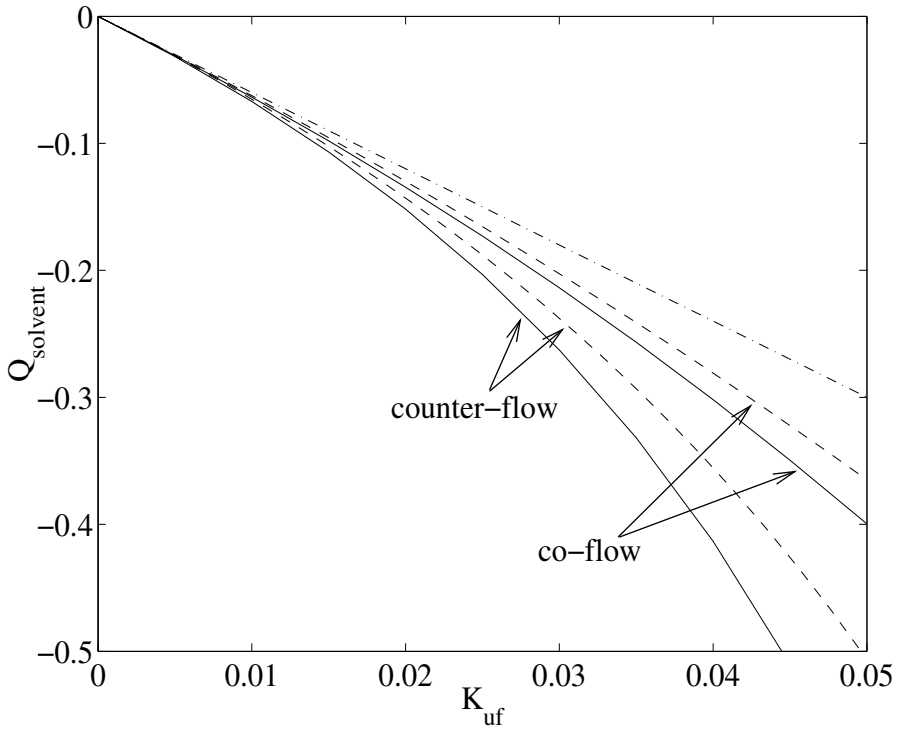

Figure 2.12: Variation of the solvent flux Q_{solvent} with the ultrafiltrative membrane permeability K_{uf}. Results for both co- and counter-flow are labelled. The dash–dotted line shows the flux according to the small-K_{uf} analysis (for both co-flow and counter-flow). The dashed line shows the flux according to the full asymptotic analysis. Numerical calculations are represented by the solid line. Note that there is a physical limit to this flux of $-\frac{1}{2}W_L^+ = -0.5$, whereby the + channel is emptied of solvent.

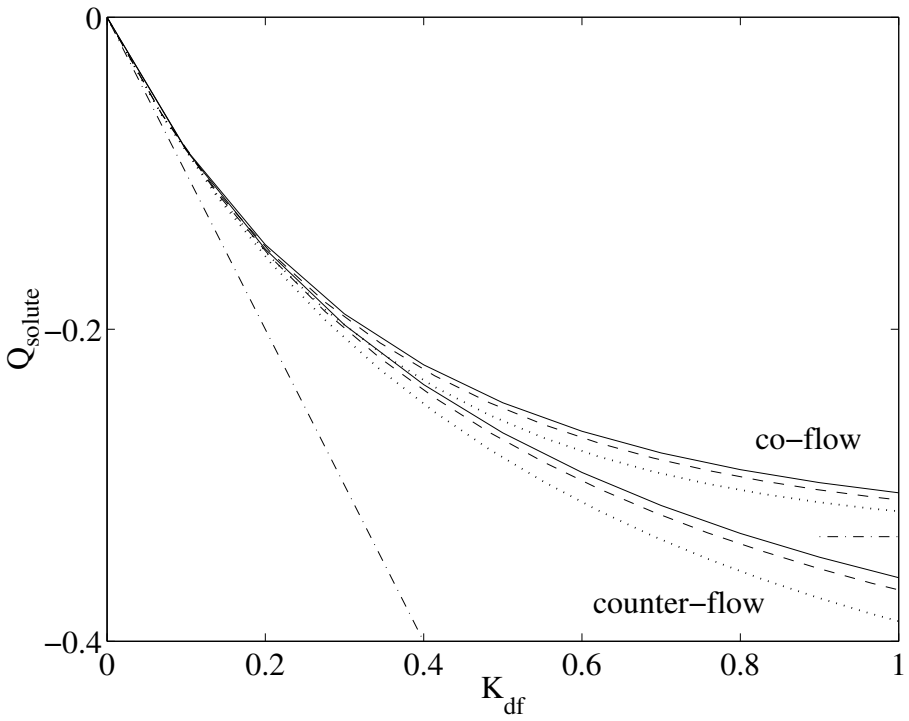

Figure 2.13: Variation of the solute flux Q_{solute} with the diffusive membrane permeability K_{df}. Results for both co- and counter-flow are labelled. The dotted lines show the leading-order asymptotic contribution to the flux. The dashed lines show the asymptotic flux, with the $\mathcal{O}(\epsilon)$ correction included. Numerical calculations are depicted by the solid lines. The dash–dotted lines show the flux in the small-K_{df} and also the co-flow large-K_{df} limit. (The large-K_{df} counter-flow limit of $Q_{\text{solute}} \to -0.5$ is not shown.)

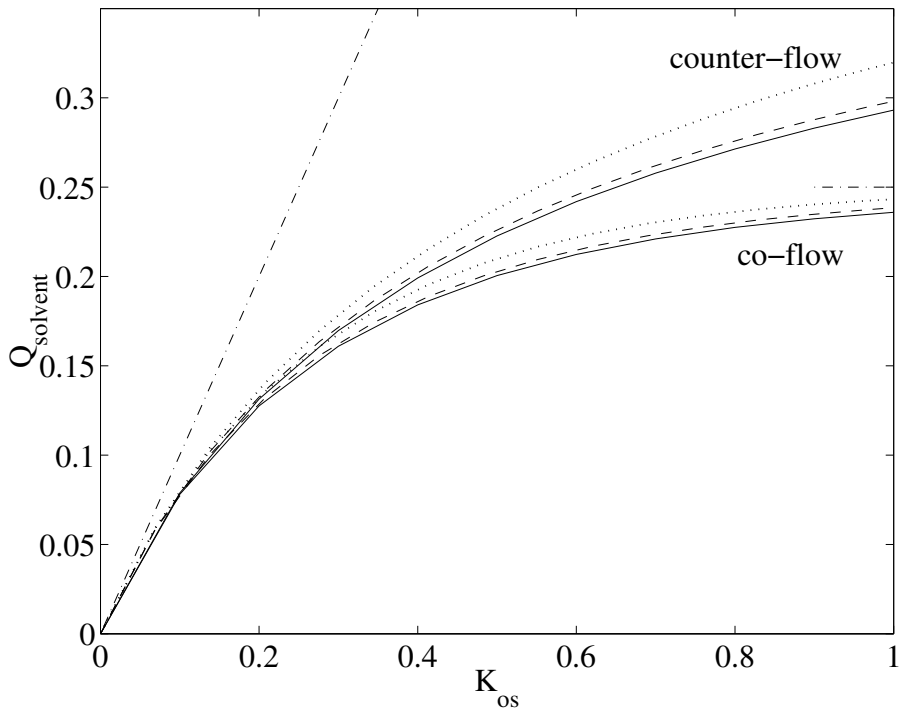

Figure 2.14: Variation of the solvent flux Q_{solvent} with the osmotic membrane permeability K_{os}. Results for both co- and counter-flow are labelled. The dotted lines denote the leading-order asymptotic contribution to the flux. Dashed lines show the asymptotic flux, with the $\mathcal{O}(\epsilon)$ correction included. Numerical calculations are shown by the solid lines. The dash–dotted lines show the flux in the small-K_{os} and the co-flow large-K_{os} limit. (The counter-flow large-K_{os} limit of $Q_{\text{solvent}} \to 0.5$ is not shown.)

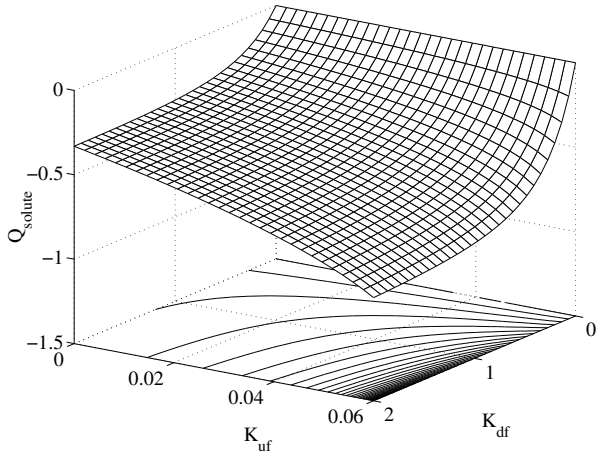

Figure 2.15: Plot showing the variation of solute flux Q_{solute} with K_{uf} and K_{df} (with $K_{\text{os}} = 0$). All other parameters are fixed to their default co-flow values. (Note that the solvent flux is independent of K_{df}, and so is not shown.) Contours of constant flux are shown beneath the surface.

Figure 2.15 shows that an ultrafiltrative solvent flux enhances the diffusive solute flux by concentrating (diluting) the solute in the + (−) channel. The large-K_{df} limit of Q_{solute} is considerably enhanced over that given by (2.104) with diffusive flux alone. Since the solvent flux is unaffected by the motion of the solute, it is not plotted here.

In contrast, Figure 2.16 shows that there is a competition when UF and OS occur together, giving an overall reduction in solvent flux. Depending on the relative sizes of K_{uf} and K_{os}, the net solvent flux can be positive or negative. In a clinical setting it is normal for ultrafiltrative effects to dominate, resulting in a negative flux that removes fluid from the + (blood) channel. A similar competition is evident in Figure 2.17 between diffusive solute flux and osmotic solvent flux: here an increase in K_{os} reduces the jump in solute concentration across the membrane and so leads to a reduced solute flux; an increase in K_{df} has a similar effect on solute concentration and so correspondingly reduces the osmotic flux.

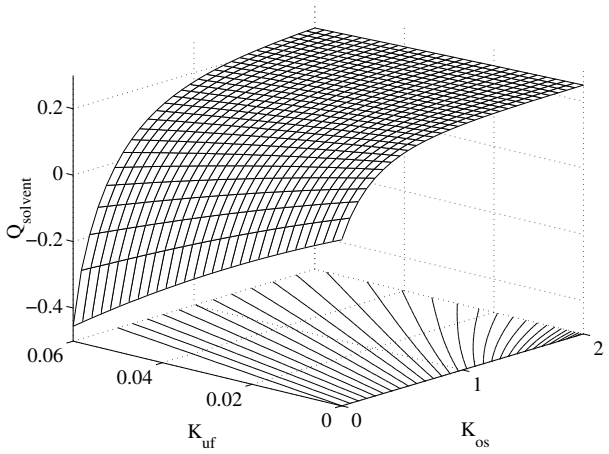

Figure 2.16: Plot showing the variation of solvent flux Q_{solvent} with K_{uf} and K_{os} (with $K_{\text{df}} = 0$). All other parameters are fixed to their default co-flow values. (Note that the solute flux is always zero in this regime, and so is not plotted here.) Contours of constant flux are shown beneath the surface.

2.7.2 Effects of changes in geometry

The geometry of our simple dialyser is characterised in our dimensionless model with just the single parameter H (the ratio of the heights of the $-$ and $+$ channels). For a more realistic model, based on cylindrical tubes, rather than channels, H would characterise the packing arrangement of the blood and dialysate tubes.

We have calculated the variation of solute and solvent flux with H. In doing so, the maximum inlet flow speed is held fixed in the $-$ channel by choosing $W^-_{L/R}$ to be proportional to H; all other parameters are fixed at their default values. Plots for ultrafiltrative, diffusive and osmotically-permeable membranes are shown in Figures 2.18, 2.19 and 2.20, respectively.

For the ultrafiltration membrane results, plotted in Figure 2.18, as the $-$ channel becomes narrower (i.e. H is decreased), any influx of solvent causes a large increase in the channel flow speed, resulting in a more rapid pressure drop down the channel, consistent with an increased flux. This feedback increases as H decreases until a threshold value is reached at which the physical flux limit $Q_{\text{solvent}} = -\frac{1}{2}W^+_L = -0.5$ is met, and the $+$ channel is completely drained. By contrast, as H is increased, the co-flow and counter-flow problems give rise to the

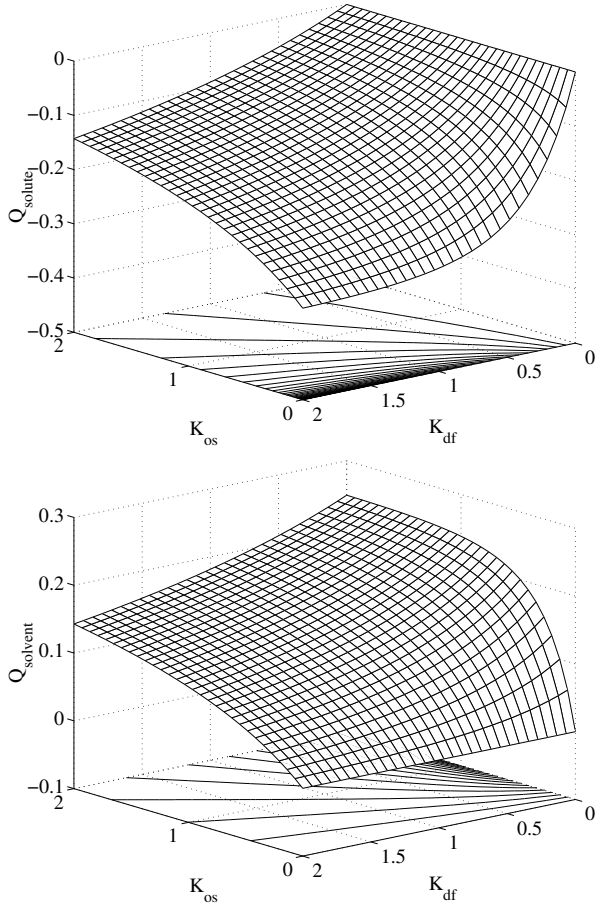

Figure 2.17: Plots showing the variation of solute flux Q_{solute} and solvent flux Q_{solvent} with K_{df} and K_{os} (with $K_{\text{uf}} = 0$). All other parameters are fixed to their default co-flow values. Contours of constant flux are shown beneath the surfaces.

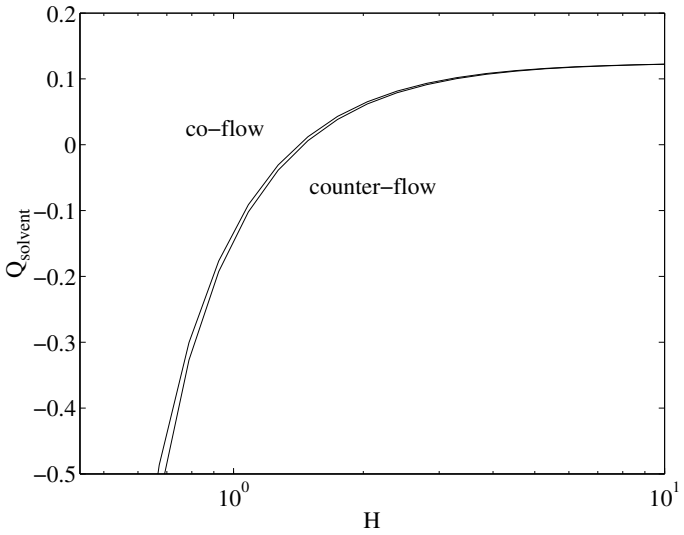

Figure 2.18: Variation of ultrafiltrative solvent flux Q_{solvent} with the channel height ratio H, according to the $\mathcal{O}(1) + \mathcal{O}(\epsilon)$ asymptotic solution. Counter-flow is shown to yield a larger flux for all values of H.

same limiting value for the solvent flux. The reason for this is clear if we recall that it is the maximum inlet flow speed that is held fixed in the lower channel as H is varied; thus when H is large, the flow velocity near the membrane in the $-$ channel is small and both co-flow and counter-flow problems approach that in which the lower channel is replaced by a large, stationary bath of fluid.

Figure 2.19 shows the variation in Q_{solute} with H when the diffusive flux (only) is acting. This plot has several features of note. First, as indicated above, the distinction between co-flow and counter-flow problems is lost in the limit as $H \to \infty$. Second, for large H, as H is decreased, the solute flux $|Q_{\text{solute}}|$ increases. The reason for this trend is that, while the limit $H \to \infty$ corresponds to stationary fluid in the $-$ channel, reducing H increases the flow speed near the membrane, removing excess solute, thereby improving solute flux. Third, for small H, as H is decreased, $|Q_{\text{solute}}|$ decreases, simply due to the correspondingly decreasing volume of fluid in the $-$ channel available to remove the solute. By considering the second and third of these points together, we see that there is an optimum channel depth ratio, at which $|Q_{\text{solute}}|$ is maximised. This optimum value of H can be found by solving $\mathrm{d}Q_{\text{solute}}/\mathrm{d}H = 0$ from (2.104) and (2.105) and remembering that $W^-_{L/R} \propto H$.

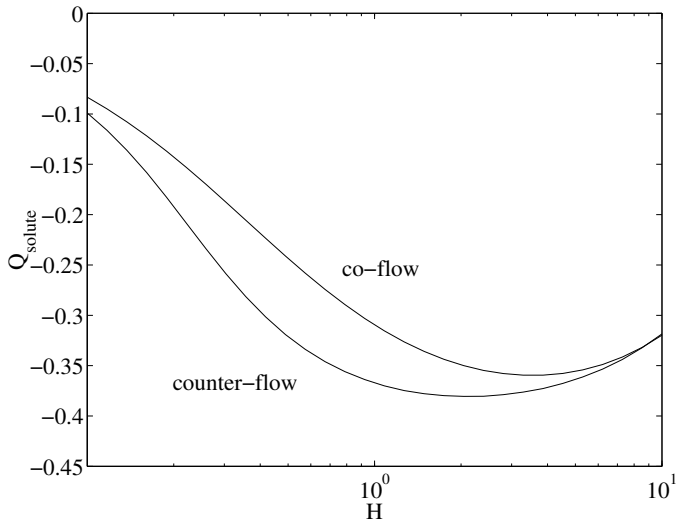

Figure 2.19: Variation of diffusive solute flux Q_{solute} with the channel height ratio H, according to the $\mathcal{O}(1) + \mathcal{O}(\epsilon)$ asymptotic solution.

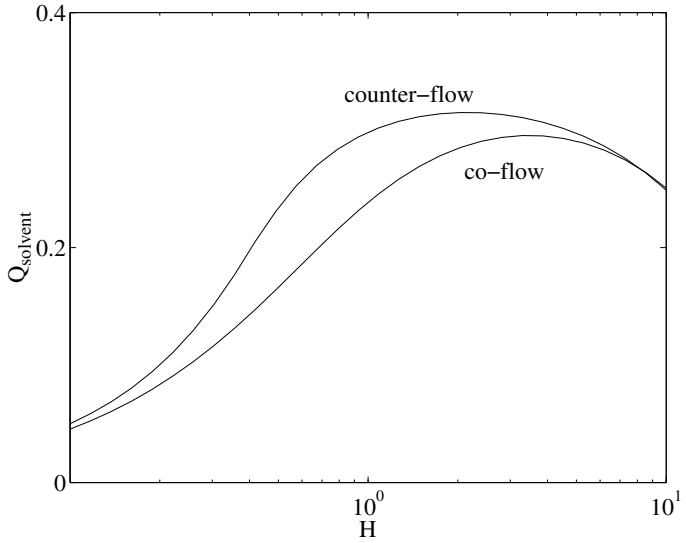

Figure 2.20: Variation of osmotic solvent flux Q_{solvent} with the channel height ratio H, calculated using the $\mathcal{O}(1) + \mathcal{O}(\epsilon)$ asymptotic solution.

Similar behaviour is displayed in Figure 2.20 with only osmotic flux through the membrane. The solvent flux is always larger for counter-flow than for co-flow, for a given value of H. The small-H decay of the solvent flux is readily understood, as follows. For small H, when the $-$ channel contains little fluid, the effect of a given solvent flux upon the solute concentration there is more pronounced than when H is larger. The solute concentration in the $-$ channel therefore quickly rises to that in the $+$ channel, causing further solvent flux to cease. In the limit as $H \to 0$ this effect is amplified so that the osmotic solvent flux decays to zero. An argument similar to that given above for diffusive flux explains why, at large H, the flux increases as H decreases, and hence why there is an optimum value of H at which the osmotic flux is maximised.

2.7.3 Effects of changes in inlet conditions

When there is competition between UF, DF and OS, the behaviours of both the solute and solvent fluxes are nontrivial. It is of clinical interest to determine whether, for fixed values of the membrane permeability constants, it is possible to achieve specific values for the two fluxes simply by adjusting the flow inlet conditions. In a clinical setting the solute composition of the blood is fixed, and it is not wise to alter the flow rate of blood from the patient, since this can cause additional discomfort (although the blood flow rate may vary between patients). However, the flow rate W_L^- and solute concentration C_L^- of the dialysate may be varied, in vivo, by controlling the dialysate pump speed and the dilution ratio of the pure water added to the dialysate concentrate, respectively.

Using a co-flow regime, and three different sets of membrane permeability constants, we have plotted in Figures 2.21, 2.22 and 2.23 the variation of solute and solvent fluxes with the inlet parameters W_L^- and C_L^-.

Figure 2.21 shows the variation of solute flux for a membrane allowing ultrafiltrative and diffusive fluxes only. (The solvent flux is not plotted since it does not vary with C_L^-.) It can be seen that if the solvent flow rate is low and the solute concentration high in the $-$ channel then the solute flux is positive (into the $+$ channel), while faster flow and lower solute concentration result in solute being

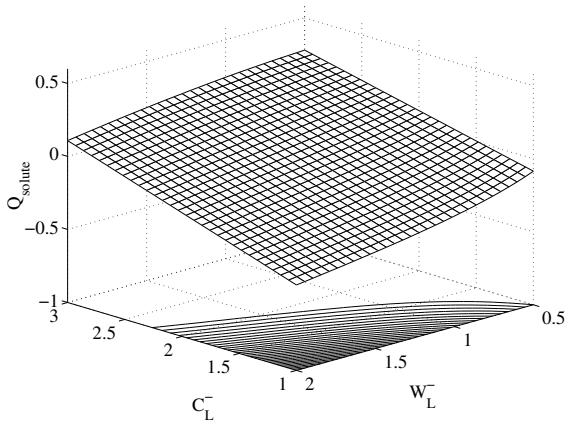

Figure 2.21: Plot showing the variation of Q_{solute} with the $-$ channel inlet conditions W_L^- and C_L^- when ultrafiltration and diffusion occur ($K_{\text{uf}} = 0.05$ and $K_{\text{df}} = 1$) but osmosis is neglected ($K_{\text{os}} = 0$). All other parameters are fixed to their default co-flow values in (2.39). Contours of constant flux are shown beneath the surface.

removed from the $+$ channel.

Figure 2.22 shows the corresponding solvent flux for a membrane allowing ultrafiltration and osmosis only. In this case the solute flux is zero. The total solvent flux Q_{solvent} can be either positive or negative depending on the $-$ channel inlet conditions.

When the permeable membrane allows simultaneous diffusive and osmotic fluxes only, as shown in Figure 2.23, both Q_{solute} and Q_{solvent} exhibit variation with W_L^- and C_L^-. These fluxes are in fact proportional to one another (as is evident from the boundary conditions (2.36) and (2.37)) so that the contours shown in Figure 2.23 are the same shape for both the solute and solvent fluxes. The fluxes change sign as the inlet conditions are varied. With $C_L^- < C_L^+$, diffusive solute flux is negative and osmotic solvent flux is positive, whilst the situation is reversed when $C_L^- > C_L^+$. The magnitudes of these fluxes can be increased by increasing the flow rate W_L^-.

The next stage in investigating selection of solute and solvent fluxes is to consider dialysers where UF, DF and OS occur simultaneously. Understanding the extra interactions in these cases may allow a wide range of desired solute and solvent

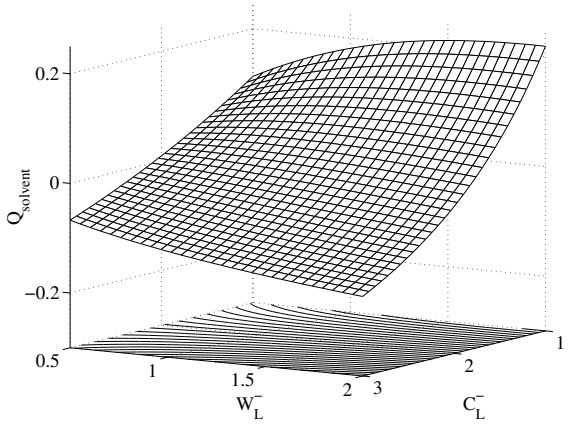

Figure 2.22: Plot showing the variation of solvent flux Q_{solvent} with W_L^- and C_L^- when ultrafiltration and osmosis occur ($K_{\text{uf}} = 0.05$ and $K_{\text{os}} = 1$) but solute diffusion is neglected ($K_{\text{df}} = 0$). All other parameters are fixed to their default co-flow values in (2.39). Contours of constant flux are shown beneath the surface.

fluxes to be achieved. A systematic inversion of membrane-parameter–flux information would then reveal the necessary values of the dialyser properties and inlet parameters needed to attain any given pair of solute and solvent fluxes for a clinical dialysis session. However, the many parameters involved in such a study prohibit a thorough discussion here.

2.8 Conclusions

Kidneys remove excess (or waste) solutes and water from the blood. Artificial dialysis machines perform this function for humans with kidney failure or loss. Patient dialysis sessions can be profiled, in the grossest sense, by choosing the correct dialyser fibre-bundle for a patient's needs. Extreme cases are the removal of water only (an ultrafiltrative dialyser) or removal of blood solutes only (a diffusive dialyser), though it would be normal to require a combination of both. Further, more sophisticated, control is achieved by using a computer to monitor patient solute and fluid levels, and adjusting the composition and flow rate of the dialysate fluid to 'optimise' removal rates.

These methods of profiling clinical dialysis sessions are widely used, based on

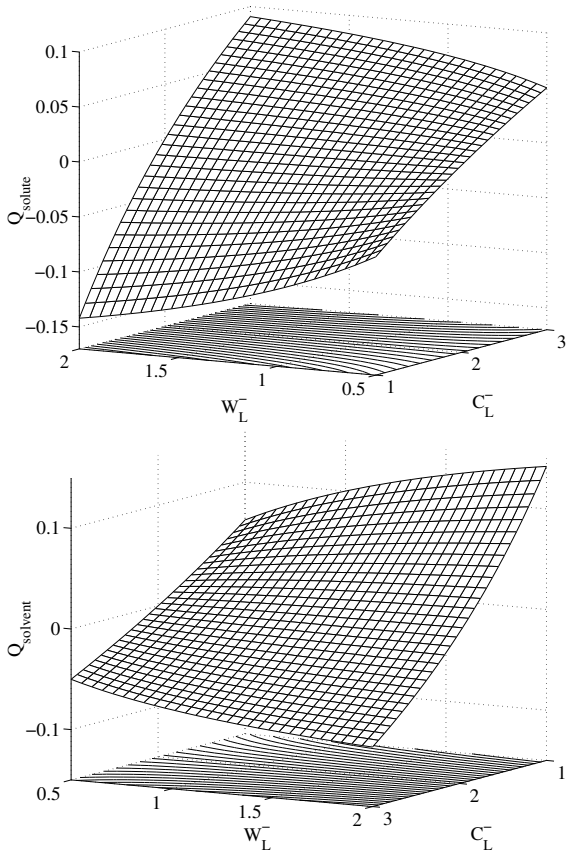

Figure 2.23: Plots showing the variation of Q_{solute} and Q_{solvent} with W_L^- and C_L^- when solute diffusion ($K_{\text{df}} = 1$) and osmotic solvent flux ($K_{\text{os}} = 1$) occur but ultrafiltration is neglected ($K_{\text{uf}} = 0$). All other parameters are fixed to their default co-flow values in (2.39). Contours of constant flux are shown beneath the surfaces. Note the different orientations of the axes in the two plots, to improve visualisation of the surface plots.

empirical data, but they are not fully understood. Improvements in dialyser design and performance can conceivably be made if we improve our understanding of the links between machine set-up and usage on the one hand and the removal rate of substances from the patient on the other. While we have not attempted to model interactions with the patient's body, the model developed in this chapter permits an understanding of the interactions that occur in a dialysis machine itself.

Using realistic values for the properties of a dialyser fibre-bundle, we were able to identify two small model parameters, one geometrical (related to the aspect ratio of the fibres) and the other related to the ratio of advective and cross-fibre diffusive fluxes of solute. A corresponding asymptotic solution for the model was constructed in Section 2.5 and used in Section 2.6 to calculate the physiologically important overall solute and solvent transfer rates between the channels carrying blood and dialysate.

Plots of flux versus membrane permeability in Section 2.7.1 show that spatially-averaged (well-mixed) one-dimensional analyses consistently overestimate the removal rate of solute and solvent for the cases of diffusion and osmosis, whereas the asymptotic theory developed here gives improved estimates of these fluxes (though the size of the discrepancy depends on the model parameters and may be quite small for clinical values). In addition we have confirmed that counter-flow regimes allow greater fluxes for a given geometry, even when transport types are combined.

Of particular interest is the effect of combining pairs of membrane properties, so that transport can occur by a combination of ultrafiltration, diffusion and osmosis. Non-trivial interactions have been shown to occur and competition between the processes can either attenuate or enhance solute and solvent fluxes compared with values when the processes occur independently. Such behaviour is not evident in existing lumped-parameter, spatially-averaged models [15, 16].

Our investigation into the variation of flux with dialyser geometry (Section 2.7.2) showed that the relative widths of the channels in a parallel-plate dialyser influence the rates of solute and fluid transport between a patient's blood and the

dialysate fluid. Just as we predict an optimum ratio of channel widths (when considering diffusive and osmotic fluxes), we anticipate that there will exist some corresponding optimum packing arrangement for the fibres in a clinical dialyser (see Figure 1.1). Further studies of this nature could aid dialyser design.

Finally in Section 2.7.3 we investigated the dependence of fluxes on parameters that would be variable in a clinical situation. We found that independent variation of solute and solvent flux is possible by varying the permeability properties of the membrane. Future studies along similar lines could act as the basis for solving the inverse problem of selecting dialyser parameters to achieve desired solute and solvent removal rates. (Obviously, the inverse problem would be a significant challenge for a general membrane where ultrafiltration, diffusion and osmosis are all present.)

There are many ways in which our model can be generalised to accommodate additional features of kidney dialysis. For example, a multiple-species model could account for interactions between different blood solutes (e.g. albumin, calcium, creatinine, glucose, phosphate, potassium, sodium and urea). A complicated coupling between the solute motions would result from the osmotic flux of water taking place in response to the *total* osmotic potential, combined with different diffusive permeabilities of the membrane to individual solutes [36]. Our model could also be adapted to account for non-Newtonian flow of blood in small scale structures, and the different rheology of dialysate and blood [55].

Other techniques can be used to provide additional flexibility in solving the profiling problem. One such method is to inject supplementary solutes of large molecular weight into the dialysate to increase its osmotic potential, thus enhancing the osmotic flux of water from blood into the dialysate. If the solute molecular size is sufficiently large, it is prevented from escaping through the microscopic membrane pores into the blood. A generalised multiple-solute-species analysis of the kind presented here would allow this to be studied. Ultimately, such an analysis could then be coupled with a patient–dialyser model, providing a safe, accurate and flexible way of controlling dialysers, improving dialysis speed whilst maximising patient comfort.

Chapter 3

Mixing measures for a two-dimensional chaotic Stokes flow

3.1 Introduction

Achieving effective fluid mixing is important in many areas, both in industry (e.g. food production [120] or controllable manufacture of polymer blends and composites [154]) and in biological scenarios (e.g. ventilation [21, 117], digestion [90], drug manufacture [133] and dentistry [33]). In this chapter we investigate a range of common measures used to assess the quality of fluid mixing. In doing so it is useful to be aware of a formal distinction between the processes of stirring and mixing: the latter refers to homogenising a scalar concentration through diffusion, whilst stirring refers to the process of advecting a fluid in order to increase the scalar concentration gradient, thereby enhancing the rate of homogenisation by mixing [42]. In the remaining chapters of this thesis we are mainly interested in the advective stirring process, however, to avoid monotonous prose, we refer to this process as 'stirring' or 'mixing' interchangeably.

On an industrial scale, the stirring necessary to achieve good mixing can be enormously expensive, and this justifies the many recent mathematical, numerical

and experimental investigations into the quantification and optimisation of the stirring/mixing process (see, e.g. [113, 153]). Of course, there are many practical and economic constraints that may be placed on the mixer (e.g. speed of operation, expense and energy usage), and many different criteria by which the resulting mixtures may be judged (e.g. uniformity of mixing, and yield of a chemical reaction), depending on the application.

Laminar mixing is desirable in many applications, such as polymer production and dentistry, where turbulent mixing is either infeasible due to the high viscosity of the working fluid, or damaging to the fluids being mixed. Further, it is now well established that even simple mixing devices, capable of generating Stokes flows (slow flows) with chaotic Lagrangian particle paths, can mix viscous fluids effectively [9, 10]. Stokes flow mixing is also common in many biological scenarios, such as the small scale lung airways [59], or blood flow in capillaries. We concentrate on mixing in the Stokes flow regime in the remainder of this thesis (we also consider inviscid flow in Chapter 4).

Throughout this chapter we consider an idealised two-dimensional mixer [47], in which flow is chaotic, but laminar; the device approximates the mixing of a tank of fluid with a long stick of circular cross-section. This serves as an idealisation of many industrial planetary mixing devices, in which a vat of fluid is stirred by a gyrating impeller (see Figure 1.3). Later, in Chapter 4, we study a generalised device with $m > 1$ stirring sticks.

We examine, in detail, the effectiveness of a large number of 'stirring protocols', which correspond to varying both the path of the stick through the tank and the stick diameter, and do so according to a number of different measures. The central question that we aim to answer is whether one protocol or a small subset of protocols can be identified as somehow mixing 'best' across a range of mixing measures. In order to make our task feasible, we have required that all mixing measures be readily implementable in automatic fashion in a computer code.

We are concerned here with computational mixing measures rather than formal, mathematical ones (see, e.g. [34]). A variety of such tools exist for systematically measuring mixing quality, almost all of which are based on the Lagrangian ap-

proach of tracking passive, non-diffusive, inertia-free tracer particles in the flow.

The prediction of mixing quality directly from the Eulerian velocity field, and the relationship between Eulerian and Lagrangian measures has been discussed by Yannacopoulos *et al.* [151] and King *et al.* [75]. In two model three-dimensional flows, these papers show a proportionality between an effective axial diffusion coefficient (calculable by Lagrangian particle tracking) and two functions of the Eulerian velocity field related to flow symmetries. Whether the results may be applied more generally to other flows, with different symmetries, has yet to be demonstrated. It is an important question, since it often easier to compute quantities directly from the Eulerian velocity field than to perform expensive particle advection simulations.

Broadly speaking, mixing measures may be classified as dynamical-systems, statistical or physical measures; while such a classification is useful for some purposes, it is of course rather arbitrary. The first class of measures that we study relies on the observation that, since the mixing device operates in time-periodic fashion, it is convenient to work with the mapping that takes tracer particles to their new locations after one period of operation of the mixing device. The dynamical-systems techniques that we use describe certain features of this mapping, and include: iterated mappings [11, 62, 65, 72]; periodic-point, manifold and symmetry analysis [52, 53, 56, 72, 83, 88, 95], and Liapunov-exponent calculation [64, 65, 106]. The second set of measures that we use are concerned with the end-product of a sequence of mappings, rather than with individual mappings themselves. These statistical techniques include: Kolmogorov [89] and mixing-variance [60, 62, 126] tests of the uniformity of the mixing; diffusive return percentage tests [12], indicating the degree to which the mixing can be undone, and intermaterial area density [58, 106, 103, 101], related to the production of thin striations. The third, and final, class of measures that we consider are more physically-based, and include: the energy input required to achieve a given state of mixing [83, 111, 122], the growth rate of an interface across which mixing can take place [93], and the residence time of tracer in the mixer [72].

We have not attempted to incorporate every mixing measure in the literature into the present study. Other common mixing measures that are not included in our

simulations include the persistence of strain [40], related to material stretching, and the scale of segregation [84], related to the sizes of structures of the mixture. Some other significant measures also lie beyond the scope of the present work, for example, the yield of a chemical reaction taking place in the flow, which for a multi-stage reaction is known to be significantly affected by the details of the mixing [28, 30, 96, 104, 105, 113, 130].

In previous evaluations of the mixing effectiveness of Stokes flow devices, a common approach has been to apply a single measure to a given set of stirring protocols, in order to find the optimum protocol. A more refined strategy is to apply a 'sieve' [66], whereby a succession of ever more sophisticated (and computationally expensive) measures are used, at each stage refining the choice of protocol, until an optimum protocol or set of protocols is established. There is no suggestion in [66] that any such optimum protocol need exist; nevertheless, the approach requires that different mixing measures broadly agree on their ranking of mixing quality. This latter issue, while crucial to the 'sieve' approach, seems to have received little direct attention in the literature; it is, however, a significant concern here. We investigate the extent to which different mixing measures agree, and in particular whether such agreements can lead to a saving in computational time, with a cheap measure serving as a substitute for a more expensive one. We also highlight the discrepancies that may exist between mixing measures, since these indicate the difficulties inherent in simultaneous optimisation of different mixing properties, and hence in the implementation of 'sieving'.

Since this is, to our knowledge, the first systematic study of the correlations between a range of different mixing measures, we focus on a flow where the exact velocity field is known. This provides some significant computational advantages over a flow for which the velocity field itself must be computed at each time step, in terms of both accuracy and CPU time. It also allows us to evaluate thoroughly a large number of stirring protocols.

The mixer used here is the batch stirring device (BSD_1) [47], whose velocity field is expressible in closed form involving finitely many terms.[1] The BSD_1 is nearly

[1] In [47] the BSD_1 is called the 'translating, rotating mixer (TRM)'. We adopt the name BSD_1 in this thesis for consistency with Chapter 4, where we consider a generalised BSD_m,

Dynamical-systems measures	Iterated mapping	IM
	Liapunov exponent	LE
Statistical measures	Diffusive deviation	DD
	Intermaterial density	ID
	Kolmogorov statistic	KS
	Lamellar widths	LW
	Mixing variance	MV
Physical measures	Energy usage	EU
	Interface stretch	IS
	Residence time	RT

Table 3.1: The mixing measures used in this chapter, with the abbreviations used in discussions.

unique among Stokes flow mixing devices in having the desirable feature of a geometry that changes with time. There are two notable exceptions to the rule of fixed geometry. One is the baffled cavity mixer [67] in which flow in a rectangular cavity is driven by sliding motion of one of the walls, while the geometry of the flow domain changes due to time-varying intrusion of baffles. A second variable-geometry device is the cavity transfer mixer [149], comprising two cavities, where flow is driven by sliding motion of one wall, but a time variation in geometry is introduced by relative sliding of the two cavities.

The BSD_1 is introduced in Section 3.2, and a set of candidate stirring protocols is constructed by independently varying three of the mixer parameters. There follows in Section 3.3 a description of each of the mixing measures that we use. However, it should be emphasised that this chapter is not intended as a review of mixing measures, and so these descriptions are necessarily selective; we present sufficient detail to justify the measures and to explain how they are computed in our particular example. Of course, we recognise that the various measures each involve some arbitrariness in their implementation. In order to provide a reference point for later discussion and to preclude confusion due to the abundance of different measures, we summarise the measures used in Table 3.1. Calculations of

studied by Boyland, Aref and Stremler [20].

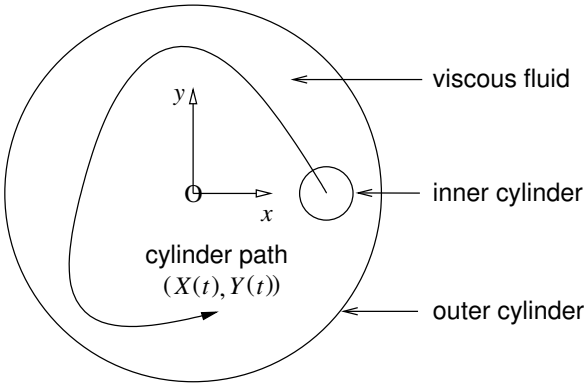

Figure 3.1: The batch stirring device (BSD_1) [47]. The coordinate system is chosen such that the origin $(x, y) = (0, 0)$ lies on the outer cylinder axis.

the various mixing measures are described in Section 3.4. Finally, in Section 3.5, we draw some, necessarily tentative, conclusions for 'optimal' selection of the stirring protocol.

3.2 Stirring protocols for the BSD_1

The batch stirring device (BSD_1) is shown in Figure 3.1. This two-dimensional mixer consists of two parallel circular cylinders, one inside the other. The annular region between the two cylinders is filled with a viscous fluid, and mixing is achieved by translating the inner cylinder through the fluid domain, so mimicking the intuitive method of stirring a tank of fluid using a stick (the outer cylinder is the 'tank', the inner cylinder a 'stick'). Although in principle more complex flow regimes can be generated by additionally allowing the inner and outer cylinders to rotate about their respective axes [47], this additional feature is not considered here, in the interests of limiting the number of parameters to be investigated.

The BSD_1 is a generalisation of the well-studied eccentric annular mixer (EAM), in which the inner cylinder does not translate, but fluid motion is driven only by rotation of the inner and/or outer cylinders about their respective axes [11, 12, 111, 145]. The exact velocity field for the BSD_1 in the Stokes flow regime has been presented by Finn and Cox [47]; it is readily implemented for numerical simulations, and receives no further discussion here. The availability of an exact

expression for the velocity field significantly improves our ability accurately to simulate the mixing, at modest computational expense. The more difficult problem that arises when the velocity field is not known a priori does not lend itself to such a detailed study as presented here, with current computing capabilities. A further benefit arising from our knowledge of the exact velocity field is that the power input required to drive the BSD_1 can also be determined exactly [47]; such information is clearly of great importance in an industrial setting, where it provides a constraint on the mixing, and it receives detailed discussion later.

3.2.1 The 252 stirring protocols

A stirring protocol is represented by the closed curve $(X(t), Y(t))$ along which the axis of the inner cylinder moves, the motion repeating with period T. The protocols that we consider correspond to cycloids of the form

$$(X, Y) = (r_1 \cos(2\pi t') + (0.8 - r_1)\cos(2\pi n t'), r_1 \sin(2\pi t') + (0.8 - r_1)\sin(2\pi n t'))a, \tag{3.1}$$

where $t' = t/T$, a is the radius of the outer cylinder, and the parameters r_1 and n can be varied to generate different paths for the inner cylinder. In addition, we consider three different values for the radius ratio a_{in}/a, where a_{in} is the radius of the inner cylinder:

$$a_{\mathrm{in}}/a: \quad 0.05, 0.10, 0.15. \tag{3.2}$$

Figure 3.2 shows the 84 inner-cylinder paths considered in this chapter. They were obtained by choosing r_1 and n from the following lists:

$$r_1: \quad 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, \tag{3.3}$$

$$n: \quad -7, -6, -5, -4, -3, -2, 2, 3, 4, 5, 6, 7. \tag{3.4}$$

We note here the significance of the parameter n, in that each cycloid has $|n-1|$-th order rotational symmetry. By taking all possible combinations from (3.2), (3.3) and (3.4), we thus generate 252 distinct protocols.

For brevity we use the notation $\mathcal{P}(i, j, k)$ to index the protocols, where $i = 1, 2, 3$ correspond, respectively, to the values of a_{in}/a given in (3.2), $j = 1, \ldots, 7$ correspond, respectively, to the values of r_1 given in (3.3) and $k = 1, \ldots, 12$ corre-

spond, respectively, to the values of n given in (3.4). For example, the protocol with $a_{\text{in}} = 0.15a$, $r_1 = 0.2$ and $n = 4$ is $\mathcal{P}(3, 2, 9)$.

3.3 Mixing measures

In this section we examine some standard mixing measures and describe the ten quantities that will be used to evaluate the mixing achieved by the protocols introduced in Section 3.2. With two exceptions, the measures are based on tracking the paths of passive, non-diffusive, inertia-less particles, whose positions $(x(t), y(t))$ satisfy

$$\dot{x} = u(x, y, t), \quad \dot{y} = v(x, y, t), \tag{3.5}$$

where (u, v) is the fluid velocity. (In the two remaining cases, the particles also diffuse, so that their advection according to (3.5) is supplemented by a random component of motion, described below in Section 3.3.4.) Depending on the mixing measure, the particles may initially be placed on a regular grid or form the interface to a blob: the four initial conditions that we use are a coarse grid, a fine grid, the interface to a rectangular blob and a linear (and hence infinitely thin) 'blob' (see Figure 3.3). The coarse grid is a regular 10×10 array of particles, occupying the right rectangle whose lower-left hand corner is at $(-0.7a, -0.7a)$ and whose upper-right hand corner is at $(0.6a, 0.7a)$, where the xy-origin is on the axis of the outer cylinder; the fine grid is a regular 100×100 array of particles, occupying the same region. As we explain below, the grid that is adopted depends largely on the computational demands of the particular mixing measure. The rectangular blob interface is defined by placing 100 equally spaced particles along each side of the rectangle used for the coarse and fine grids. Finally, an infinitely thin linear 'blob' is defined by placing 10000 equally spaced particles along the right-hand edge of the same rectangle. The sizes and initial locations of the grids and blobs have been chosen so that they are always caught in the path of the inner cylinder.

We now describe each of the mixing measures in turn.

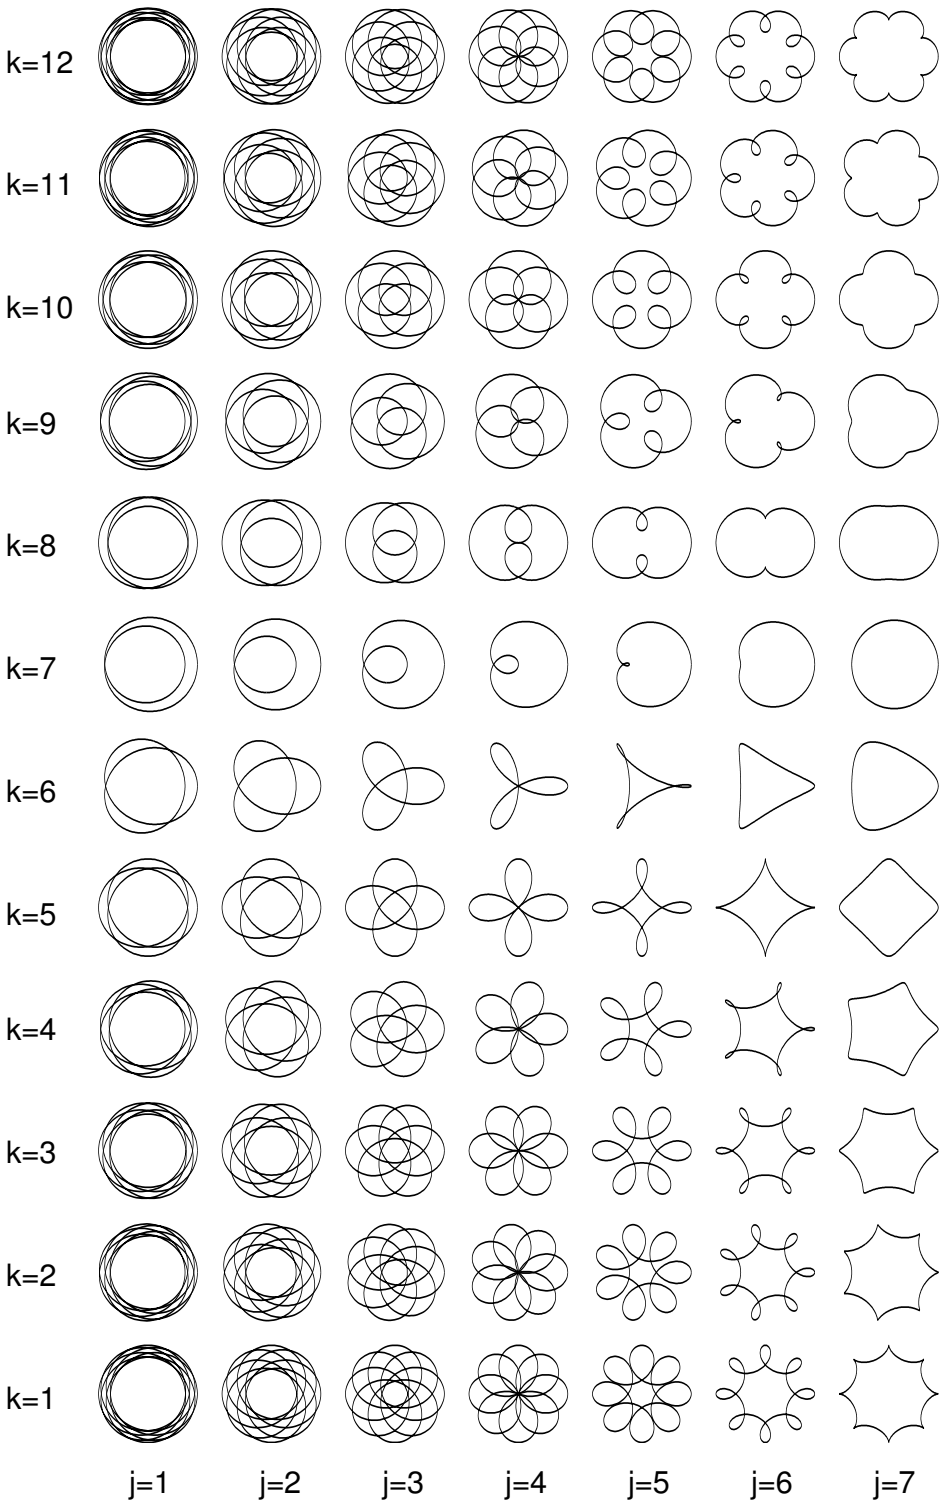

Figure 3.2: Motion of the inner cylinder axis for the stirring protocols $\mathcal{P}(*, j, k)$. The values of j and k are shown for each protocol. (The value of i is irrelevant for this diagram since it reflects the radius of the inner cylinder, not shown here, and does not influence its path.)

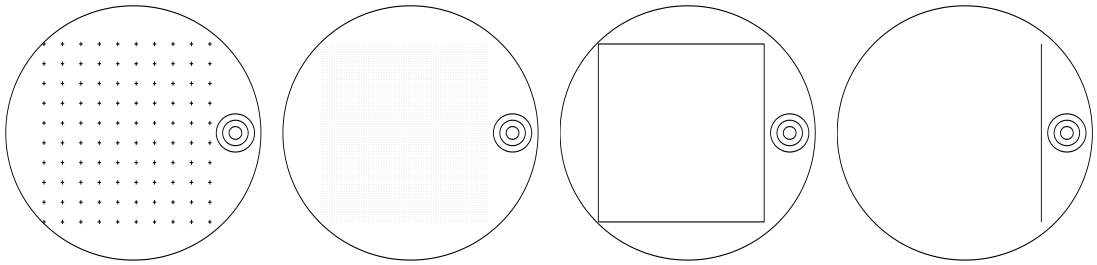

Figure 3.3: The reference initial conditions used to compute the various mixing measures. The three concentric circles indicate the three inner cylinder radii used in this investigation. From left to right are shown the 10×10 coarse grid, the 100×100 fine grid, the rectangular blob and the linear 'blob'.

3.3.1 Dynamical-systems measures

It is convenient, when evaluating the various stirring protocols, to consider instead of the flow (3.5) the period-one mapping $\mathcal{M}(\mathcal{P}) : (x_0, y_0) \mapsto (x_1, y_1)$, which maps the position (x_0, y_0) of a particle at time $t = 0$ to its position (x_1, y_1) one period later, at time $t = T$. Although the velocity field is known exactly, this mapping requires numerical computation, and is not available in closed form. In view of the periodic nature of the BSD_1 motion, $\mathcal{M}(\mathcal{P})$ also maps the position of a particle at $t = mT$ to its position at $t = (m+1)T$, for any integer m. The iterated mapping and Liapunov exponent measures that we describe below are defined in terms of $\mathcal{M}(\mathcal{P})$.

Iterated mapping measure of chaoticity (IM)

The simplest and most common dynamical systems tool for evaluating mixing quality is the iterated mapping plot, or Poincaré section [11, 62, 65, 72]. This is constructed by taking a small number of tracer particles (here the coarse grid, containing 100 particles) and plotting their positions under repeated application of the mapping $\mathcal{M}$, up to $\mathcal{M}^m$, for some large integer m (here 1023). Regions of poor mixing are revealed as islands of integrable quasi-periodic motion, which surround the elliptic periodic points of $\mathcal{M}$ [111]. Stretching of fluid elements is slow (algebraic in time) in these islands, and no fluid exchange can take place

across island boundaries, in the absence of diffusion. By contrast, effective mixing takes place in the chaotic sea, with exponential-in-time separation of initially neighbouring particles. Boundaries between islands and chaotic sea may have an intricate fractal structure, containing infinitely many smaller islands on all spatial scales. This structural complexity gives rise to computational difficulties in distinguishing between genuinely chaotic and merely quasi-periodic particle paths involving long periods. However, at least informally, iterated mappings provide a good visual identification of the degree to which the stirring domain is chaotic.

To systematise the evaluation of the degree of chaoticity in the iterated mapping, we proceed as follows. For a given protocol $\mathcal{P}$ we track the motion of the coarse grid of particles, recording the particle positions under the application of $\mathcal{M}^m$ for $m = 0, \ldots, 1023$ (where $m = 0$ corresponds to the initial locations of the particles). This procedure generates 100 times-series of the form $\{(x_0, y_0), (x_1, y_1), \ldots, (x_{1023}, y_{1023})\}$, whose periodicity, quasi-periodicity or chaoticity can be determined by taking the discrete Fourier transform (DFT) of the time series of the x- and y-components (see Figure 3.4). If the DFT appears noisy (i.e. if the peak amplitude in the DFT is comparable in size to the mean amplitude) then we declare the corresponding particle path to be chaotic. If, on the other hand, an amplitude peak exists in the DFT that is significantly greater than the mean (we use a factor of 20 as our threshold ratio) then the corresponding particle path is declared quasi-periodic (exceptionally it may be periodic). Of course, points may be mis-diagnosed by this procedure; in particular, quasi-periodic paths involving very long periods may be mistakenly declared chaotic, but for practical purposes we do not concern ourselves with such fine distinctions.

For a given protocol, we define the iterated mapping mixing measure IM to be the proportion of particles on the coarse grid that give rise to chaotic trajectories, as determined above. This measure provides an estimate of the percentage area of the BSD_1 domain that experiences chaotic advection; high values correspond to good mixing.

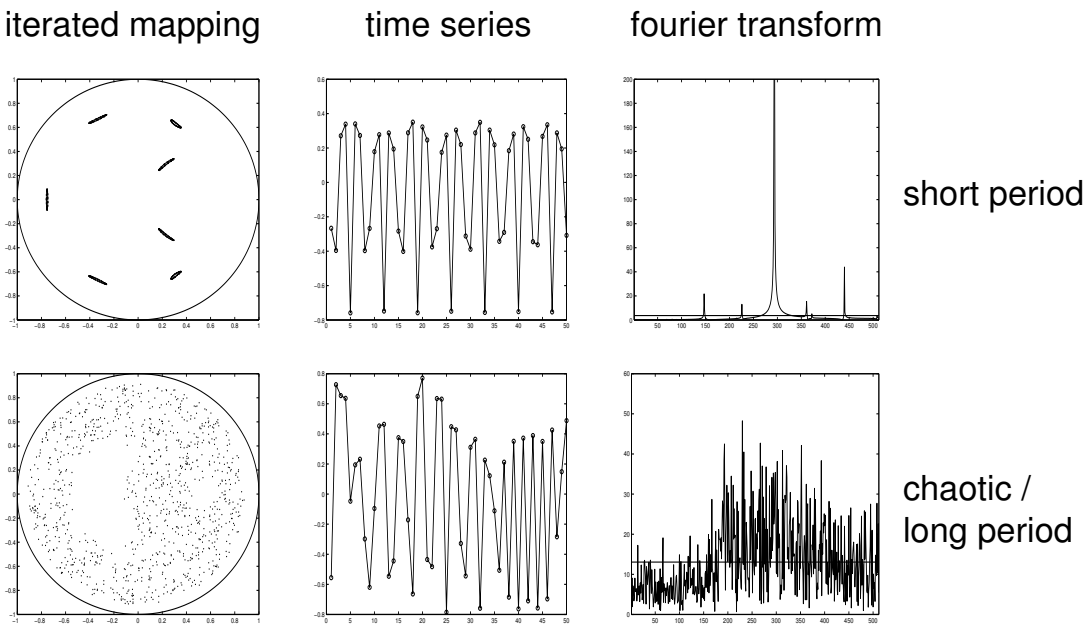

Figure 3.4: Detecting periodicity using a DFT. In the upper sequence of plots, a quasiperiodic trajectory gives rise to a DFT with well defined, sharp peaks. In the lower sequence, by contrast, an apparently chaotic trajectory yields a full spectrum in the DFT.

Liapunov exponent (LE)

In the chaotic sea, neighbouring particles separate exponentially in time, and a corresponding stretching-rate exponent, or Liapunov exponent, can be defined as follows. We first imagine placing an infinitesimal material line element $\mathrm{d}\boldsymbol{x}$ at a point $\boldsymbol{x}$, extending from $\boldsymbol{x}$ to $\boldsymbol{x} + \mathrm{d}\boldsymbol{x}$; then the Liapunov exponent λ at $\boldsymbol{x}$ is

$$\lambda = \lim_{n\to\infty} \limsup_{|\mathrm{d}\boldsymbol{x}|\to 0} \frac{1}{nT} \log \frac{|\mathcal{M}^n(\boldsymbol{x} + \mathrm{d}\boldsymbol{x}) - \mathcal{M}^n(\boldsymbol{x})|}{|\mathrm{d}\boldsymbol{x}|}. \tag{3.6}$$

Roughly speaking, an infinitesimal material line element of initial length l_0 in a chaotic region will have length $l_n \sim l_0 \exp(\lambda n T)$ at a time nT later. If the flow domain is globally chaotic then ergodicity implies that λ is independent of the starting position $\boldsymbol{x}$ so that a single Liapunov exponent can be used to characterise the protocol. However, in general, we expect a spatial distribution of values of the Liapunov exponent [124].

For the purpose of evaluating our stirring protocols, we compute the *finite-time* Liapunov exponent, obtained from (3.6) by replacing ∞ with a moderate number N of mappings (here $N \sim 50$), and using a small, but finite, line element (here $|\mathrm{d}\boldsymbol{x}| \sim 10^{-5}a$), rather than an infinitesimal one (although in principle the evolution of the latter could be determined by considering the gradient flow [111]). We compute the exponent at all points on the coarse grid and take the mean exponent to be our mixing measure LE; high values of LE correspond to good mixing. Of course, material lines lying within a periodic island experience sub-exponential stretching, and have correspondingly small finite-time Liapunov exponents; this tends to punish those stirring protocols that give rise to large areas of (poorly mixed) islands. Finally, we note that there are other, more sophisticated, related mixing measures, e.g. based on the *distribution* of stretching values, or on Ottino's mixing efficiency [111], which compares local elemental stretching rates against a theoretical upper bound for stretching, but these are not considered here.

3.3.2 Statistical measures

In contrast to the dynamical systems tools described in Section 3.3.1, statistical mixing measures are based on the final outcome of the mixing rather than on the

details of the evolution of tracer particles in the flow.

Mixing variance (MV)

A simple and intuitive statistical mixing measure is the mixing variance. To motivate this measure, we note that perfect mixing would distribute material uniformly over the entire fluid domain, so that in any partition of the domain into regions of equal area, each region would contain an identical amount of material, and thus the variance of the tracer distribution would be precisely zero. Any other distribution of the same material would lead to a non-zero variance, and so the variance of the tracer distribution gives an indication of the mixing quality.

Specifically, we define the mixing variance as follows. For a given protocol $\mathcal{P}$ we take the linear blob as an initial condition and evolve it for ten periods of the stirring protocol, i.e. we apply the mapping $\mathcal{M}^{10}$. Using the rectangular blob to provide the corner points, we then construct a 9×9 array of equal-sized rectangles, and compute the number of particles ultimately found in each of these rectangles. The sample variance of these numbers is calculated, giving the mixing variance measure MV; low values of MV are indicative of good mixing.

Kolmogorov statistic (KS)

The Kolmogorov test is an alternative measure of the difference between an actual distribution of stirred tracer particles and a desired (usually uniform) distribution, as, for example, applied by Ling [89] to viscous flow in a rectangular domain stirred by periodic sliding of two of the domain walls. Ling considered rectangles of different sizes, with one corner fixed at some given point in the fluid domain, and recorded the number of particles found in each rectangle. The maximum deviation in this number from that expected under ideal, uniform mixing provided the mixing measure. (Under ideal, uniform mixing, any rectangle should ultimately contain an amount of mixed material in proportion to its area.)

A similar Kolmogorov statistic may be defined for the BSD_1. As with the MV

measure introduced above, we take as initial tracer locations the linear blob shown in Figure 3.3, and apply $\mathcal{M}^{10}$. We then consider a large number (say, 100) of concentric circles, centred at the BSD_1 origin. For perfectly uniform mixing, a circle of radius r should contain $N_{\text{uniform}} = r^2 N_{\text{total}}/a^2$ particles, where $N_{\text{total}} = 10000$ is the total number of particles used for the simulation. (We ignore the presence of the inner cylinder for the sake of a simpler calculation because its area is small compared to that of the 'test' circle and so has little effect on KS.) If the actual number of particles found in a given circle is N_{actual} then the scaled deviation from uniformity is

$$\frac{|N_{\text{uniform}} - N_{\text{actual}}|}{N_{\text{total}}}. \tag{3.7}$$

The maximum value of this quantity over the set of circles is recorded, and provides the Kolmogorov statistic KS. As for the MV measure, small values of KS correspond to more uniform (i.e. better) mixing.

Intermaterial interface density (ID)

When an interface between two chemical species evolves in the chaotic region of the BSD_1 flow, it grows exponentially in length and rapidly becomes highly contorted, by the well known stretching and folding mechanisms of chaotic flows. Thin striations (or stripes) of the chemicals are formed. The way in which a stretched interface is distributed in space is highly non-uniform [101], and the density with which the interface occupies the flow domain varies over many orders of magnitude. Correspondingly, it has been observed [58, 83, 101] that the distribution of striation widths varies over a similarly wide range. This correspondence is natural since the stretching of the interface in one direction is accompanied by a compression of the striations in another, due to incompressibility of the fluid. (In practice, the production of striations leads, through diffusion, to effective intermingling of the species involved, and hence to good mixing, although, of course, tracer particles do not diffuse in these simulations.) Therefore the uniformity with which a material interface is distributed throughout the mixer provides a measure of mixing quality. In keeping with previous nomenclature [101], we call this measure the 'intermaterial interface density'.

We define the ID measure as follows. For each protocol, we evolve the linear blob shown in Figure 3.3 for two periods. We take the same 9×9 array of rectangles as for the MV measure, and compute the length of interface contained within each rectangle after the mixing has taken place. Since chaotic advection causes exponential stretching of the interface, we must be careful to ensure that the interface is at all times adequately resolved. This requires particular care, even during just two periods of the flow. To do so, we use the dynamic particle insertion algorithm of Krasnopolskaya et al. [83] accurately to resolve the interface into a smooth curve as it is stretched. Our method of computing the length of interface contained within a given rectangle is then as follows: since the interface is well resolved and the dimensions of each rectangle greatly exceed the resolution threshold of the interface, we merely count the number of particles used to represent the interface fragment within each rectangle, after application of the particle insertion algorithm. This gives a reasonably accurate and easy-to-compute approximation to the length of interface contained within the rectangle. The intermaterial interface density ID is then the variance of the distribution of lengths, scaled by the squared mean length of interface in each rectangle. Thus the measure ID is concerned with the uniformity of interface distribution, rather than with its absolute length (for that, see the interface stretch measure below).

3.3.3 Physical measures

We now describe some mixing measures which may be categorised as physically based. Indeed, the first of these measures is not concerned with the results of the mixing process, only with the energy input to the device. The remaining two measures described in this section are, however, more directly concerned with mixing quality.

Energy usage (EU)

A crucial consideration in large-scale industrial mixers is the energy expended to achieve a certain degree of mixing. Many authors [83, 111] have examined the energetic requirements of laminar mixing devices, sometimes using crude substi-

tutes for energy, such as the distance moved by the boundaries that drive the motion. For the BSD_1, however, as for some other Stokes flow mixing devices, an exact expression is available for the power input required to drive the device [47], and so we can compute its energy requirements directly.

If we wish to minimise the energy requirements of the mixer then we need to examine how best to vary the speed of the inner cylinder as it traces out its prescribed path $\mathcal{P}$ in a given time T. The total energy required to drive the mixer can be written as

$$\mathcal{E} = \mu a^3 T^{-2} \int_0^T P(s, \dot{s}) \, \mathrm{d}t, \tag{3.8}$$

where μ is the kinematic viscosity of the fluid and P is a dimensionless function related to the power input [47]. Here $s(t)$ denotes arc length along the path $(x, y) = (X(t), Y(t))$ indicated in (3.1). We wish to choose $s(t)$ to minimise the energy integral (3.8), subject to the constraints

$$s(0) = 0, \quad s(T) = \Lambda, \tag{3.9}$$

where Λ is the inner-cylinder path length

$$\Lambda \equiv \oint_{\mathcal{P}} \mathrm{d}s.$$

Since the integrand in (3.8) is autonomous, it follows that $\mathcal{E}$ is minimised when the Beltrami identity

$$P(s, \dot{s}) - \dot{s} \frac{\partial P(s, \dot{s})}{\partial \dot{s}} = -\kappa \tag{3.10}$$

is satisfied for some constant κ. For Stokes flow, the power function can be written in the form $P(s, \dot{s}) = P(s)\dot{s}^2$, where $P(s)$ is a known function [47], and thus (3.10) reduces to

$$P(s)\dot{s}^2 = \kappa, \tag{3.11}$$

so that the motion must be carried out in such a way as to maintain a constant power input to the device. In order to satisfy the constraints in (3.9), κ must take the value

$$\kappa = \frac{1}{T^2} \left(\int_0^T \sqrt{P(s, \dot{s})} \, \mathrm{d}t \right)^2 = \frac{1}{T^2} \left(\int_0^\Lambda \sqrt{P(s)} \, \mathrm{d}s \right)^2,$$

and hence the speed of the inner cylinder for minimum-energy motion is

$$\dot{s} = \frac{\int_0^\Lambda \sqrt{P(s)}\,\mathrm{d}s}{\sqrt{P(s)}\,T}.$$

Finally, from (3.8) the minimum energy usage measure EU is given by

$$\mathrm{EU} = \min \mathcal{E} = \mu a^3 \kappa T^{-1}. \tag{3.12}$$

While EU is not itself directly a measure of the mixing quality, it is useful to correlate the energy input with genuine mixing measures, to provide an informal measure of mixing efficiency.

Interface stretch (IS)

An obvious physical manifestation of effective mixing is a rapid growth in the length of a finite material interface placed in the flow. The interface stretch exponent λ_{is}, also known as *topological entropy*, quantifies the exponential rate of growth of such an interface as it is stretched by the flow [83, 93]: a finite interface of initial length l_0 has length $l_n \sim l_0 \exp(\lambda_{\mathrm{is}} nT)$ at time nT, for large nT. This quantity is related conceptually to the Liapunov exponent λ, although the latter provides a measure of the *local* stretch, whereas the interface stretch exponent gives a global stretching measure [4, 7]. By its nature, the interface stretch exponent gives greater weight to regions of high stretch than does the Liapunov exponent (because regions of locally high stretch preferentially generate new interface precisely in regions of high stretch), and so $\lambda_{\mathrm{is}} > \lambda$ [7]. We define the interface stretch measure IS to be the final perimeter of the linear blob after two periods of the flow. As in our calculation of ID, we are careful to maintain adequate resolution of the growing interface.

Lamellar width variance (LW)

In many applications, mixing is accompanied by chemical reaction between the various species. At present, even on a powerful workstation it remains a significant computational challenge to resolve simultaneous chemical reaction, advection and

diffusion in a two-dimensional chaotic flow. However, such simulations may be necessary in order accurately to predict the yield of a multi-stage reaction, since it is known that fine details of the reactant distribution can have a surprisingly great influence on the final yield [30, 102, 153]. While an accurate simulation of advection–diffusion–reaction in the BSD_1 is beyond the scope of this study, we recognise that the statistics of the length scales that develop will influence the progress and the yield of any multi-stage reaction that takes place in the flow, and that the details may be highly specific to individual reaction schemes. Thus we introduce a mixing measure based on the distribution of lamellar widths, defined as follows. We imagine the linear blob to be the interface between two chemical reactants. This interface is tracked for two periods of the mixer, interpolating where necessary as with IS and ID [83]. In the course of the mixing, a large number of highly contorted striations are generated, and we measure their widths by calculating the intersection points of the interface with the y-axis. The distance between successive intersections gives an estimate of the lamellar widths generated by the flow. Our numerical code uses cubic interpolation [29] to generate accurate values for the intersection points between the interface and the y-axis. Once the lamellar widths have been found, we compute the variance of their distribution, scaled by the square of the mean lamellar width, to give the mixing measure LW. (Whether one considers good mixing to correspond to high or low values of LW may depend on the details of the chemical reactions imagined to take place in the flow.)

Note that we use a straight line segment as the initial interface, whereas a closed curve should be used to compute genuine lamellar widths. However, we find that this consideration makes little difference to the LW measure, scaled as indicated here; furthermore, any influence of the nature of the initial interface is mitigated by our later presentation of the results as *rankings* of the various measures.

3.3.4 Effects of diffusion

Thus far we have used the term 'mixing' to describe what is essentially a process of pure stirring (i.e. mechanical stretching of an interface in a laminar flow), but

it is worthwhile, from the point of view of practical application, to consider the combined effects of stirring and particle diffusion. Accordingly, our final two mixing measures, described below, concern diffusive tracer particles.

Residence time (RT)

In a chemically reactive system, the reaction may take place not only in the bulk of the flow, but also through boundary effects. One source of these, appropriate in some applications, is diffusion through the boundary (e.g. in models for kidney dialysis, as in Chapter 2 or heat exchangers). A second source is a chemical reaction, often catalysed, at the boundary that removes chemicals from the system. In either case, the action of mixing is crucial in continually introducing new fluid to the neighbourhood of the boundary, and thereby significantly enhancing these removal effects. We therefore introduce a mixing measure related to the residence time; the more effective the mixing, the shorter will be the average residence time of a tracer particle. Of course, this measure should not be confused with the more common usage of the term 'residence time' to indicate the duration of a particle's sojourn in a continuous-throughput mixer [87].

For this mixing measure, we suppose that the tracer particles, in addition to being advected by the flow, also diffuse, with diffusivity D, and consider diffusion through either wall to be the sole mechanism of removal (i.e. there is no chemical reaction there). Then the advection equations (3.5) are replaced by the stochastic differential equations

$$\mathrm{d}x = u(x, y, t)\,\mathrm{d}t + \mathrm{d}\mathcal{N}_x, \quad \mathrm{d}y = v(x, y, t)\,\mathrm{d}t + \mathrm{d}\mathcal{N}_y, \tag{3.13}$$

where $\mathrm{d}\mathcal{N}_x$ and $\mathrm{d}\mathcal{N}_y$ are independent, identically distributed, Gaussian random variables with mean zero and variance $2D\,\mathrm{d}t$.

To model diffusion of tracer particles through the BSD_1 walls, we solve (3.13), at each time step removing from the numerical simulations any particle that happens to have diffused through a wall. We then adopt a straightforward definition of the residence time mixing measure RT by recording the fraction of particles remaining in the BSD_1 at the end of a given time (here two periods of the flow). We use the fine grid of points as the initial condition and, in order to obtain results that

significantly distinguish between the various protocols, we somewhat arbitrarily take the diffusivity to be $D = 10^{-3}a^2T^{-1}$.

Of course, this is a rather crude measure, and other, more sophisticated, versions of the residence time measure can easily be conceived. For instance, in practice, the residence time of a diffusive particle initially placed at some given spatial location will depend on the precise details of the random component of its path, so it may be of interest to consider the distribution of residence times for a set of particles initially placed at the same point in the flow domain. Similarly, the mean residence time of diffusive particles will have some spatial distribution, which could readily be computed. It may also be of interest to examine whether there are regions of the boundary from which particles are preferentially removed from the flow domain.

Diffusive deviation (DD)

Our final mixing measure serves to quantify the extent to which the mechanical action of stirring enhances the intermingling of species brought about by diffusion. To motivate the measure we begin by noting that in the absence of diffusion the motion of tracer particles in Stokes flow is time-reversible: if the mixer is first run forwards and then the boundary motion is reversed, any tracer particle will return precisely to its initial location. However, when $D \neq 0$ this return is not exact, and the particle will ultimately lie some distance from its initial location. Furthermore, in a chaotic flow, this deviation can be significantly greater than that predicted by simple diffusion alone, because the sensitivity of the chaotic flow to initial conditions enhances the diffusion [13, 24].

To compute the diffusive deviation mixing measure, we allow each particle in the fine grid of points to evolve according to (3.13), with $D = 10^{-5}a^2T^{-1}$, for two periods of the flow forwards in time, followed by two periods backwards in time. The particles diffuse during both forwards and backwards motions of the boundaries. We then calculate the mean distance of each particle from its initial position, giving the mean diffusive deviation mixing measure DD. (The value of D used here is smaller than for the RT measure purely for computational

convenience – we find that for this reduced value of the diffusivity we need to take no special precautions to prevent particles erroneously escaping from the mixer.)

A related 'return percentage' measure was studied by Aref and Jones [12] for flow in the eccentric annular mixer (a special case of the BSD_1 where the cylindrical walls are rotated but there is no translation of the inner stirrer). They took an initially circular blob of particles, evolved it for a given number of periods of their flow, then reversed the motion. In the absence of diffusion, all particles returned to their initial locations inside the blob. However, with non-zero diffusion, the return percentage was diminished, and significantly so in chaotic flow, even for small values of the diffusion coefficient. The tremendous enhancement to the diffusion was due to the exponential growth of the blob boundary during the flow, and consequently greatly enhanced diffusive particle escape from the blob. Although we have taken only a single value of D, the likely trends with varying D might reasonably be inferred from the work of Aref and Jones [12].

3.4 Results

In the previous section we defined ten mixing measures that can be evaluated numerically without any user intervention (see Table 3.1 for an aide-memoire of the various measures). We have computed each of these mixing measures for the 252 protocols described in Section 3.2. Clearly, there is a degree of arbitrariness in some of the measures, largely due to a particular choice of parameters (such as the number of periods of the flow over which to calculate the measure, initial conditions, or the particle diffusivity, where this is non-zero). Furthermore, for certain of the measures we have chosen a crude, but simple, version to calculate, despite the availability of more sophisticated versions. While the details of our results may change if the measures are defined slightly differently, we expect our broad conclusions to be qualitatively unaffected.

Before describing our results, we make some preliminary comments concerning correlations between the mixing measures, and about energy usage. We then

move on to describe general trends in our results, before highlighting observations regarding specific protocols and measures.

3.4.1 Correlations between mixing measures

In an attempt to find optimal protocols, it is useful to understand the extent to which the various mixing measures are pairwise correlated. Since the measures have very different distributions (some, like RT, vary only slightly across the range of protocols, while others, like IS, vary by many orders of magnitude, for instance), we evaluate their degree of correlation using a distribution-free test based on the respective rankings of the 252 protocols. We use the straightforward Spearman rank correlation coefficient [108], as follows. To calculate the rank correlation coefficient between two mixing measures, we first independently rank the set of 252 protocols according to each of the two measures. Thus the i-th protocol is assigned integers m_i and n_i that are its respective rankings according to the two measures (the worst protocol having rank 1, the best protocol rank 252, in this case). Exceptionally, protocols may have precisely the same value of the mixing measure, in which case they are ranked equally. Using the two sets of rankings, we calculate the Spearman rank correlation coefficient

$$C = \frac{\sum(m_i - \bar{m})(n_i - \bar{n})}{\sqrt{\sum(m_i - \bar{m})^2 \sum(n_i - \bar{n})^2}}, \tag{3.14}$$

where $\bar{m} = \sum m_i/252 = 253/2$ and $\bar{n} = \sum n_i/252 = 253/2$. By construction, $-1 \leq C \leq 1$. Large values of $|C|$ indicate that the two measures in question are well correlated, while small values of $|C|$ indicate that the measures concerned do not agree on their assessment of 'good' and 'bad' protocols.

The occurrence of a strong correlation between a pair of measures is of particular interest when the measures concerned have different computational requirements (e.g. correlations between IS and EU, which are expensive and cheap to compute, respectively). If a good correlation exists, we might attempt to use the computationally cheaper measure as a substitute for the more expensive measure, thereby saving computational effort when evaluating a given set of protocols.

Equally, it is instructive to be aware of the existence of poor correlation between

two measures. This indicates that the distributions are not simply related and therefore it may be possible to find a protocol which has desired values for both measures (e.g. a low energy usage but a high Liapunov exponent). A consequence concerns the idea of using a 'sieve' [66] systematically to eliminate poor protocols, through the application of a succession of mixing measures, each more sophisticated and computationally intensive than the last. If two mixing measures used at different stages in the sieving process are poorly correlated, this may lead to the premature elimination of a protocol that would in fact score well with the later measures, simply because it did not also score well with some prior, less sophisticated measure, uncorrelated with the measure of true interest. A further shortcoming of the sieving approach is that it may eliminate protocols that perform well overall, if they do so only in some average sense.

We note that although the use of a *rank* correlation allows us to compare measures with very different (or unknown) distributions, it necessarily leads to results that depend on the set of protocols chosen for study. Our expectation is that the set of protocols is sufficiently large (252) and the protocols sufficiently varied that our broad conclusions hold, not withstanding this caveat.

3.4.2 Energy usage

In order to appreciate the role of energy in our considerations, we note that for a given BSD_1 protocol, operating in the Stokes flow regime, and mixing non-diffusive particles, the outcome depends only upon the path taken by the inner cylinder, and not upon the speed of operation of the device. However, the power input to the BSD_1 scales with the square of the speed of the inner cylinder, and hence the *energy* requirement for execution of a given protocol scales in proportion to this speed (since the period of operation varies in inverse proportion to the speed). In other words, the overall energy requirement varies in inverse proportion to the period of time over which the mixing takes place. A further consequence of operation in the Stokes flow regime is that if we fix an identical period T for all stirring protocols then those with greatest inner-cylinder path length Λ will tend to require the greatest energy input $\mathcal{E}$ per period, simply because they have the

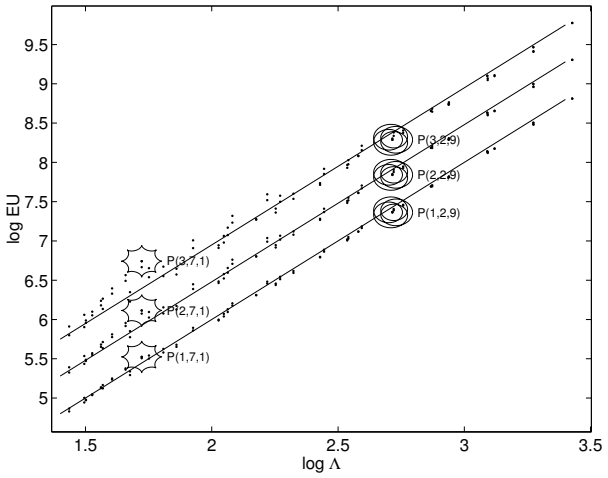

Figure 3.5: Logarithmic plot of mixing measure EU, from (3.12), against inner-cylinder path length Λ for each of the 252 stirring protocols. Trend lines with slope 2 are also shown.

greatest average speed of the inner cylinder. In fact, with T fixed in this way, the energy required per period is roughly proportional to Λ^2. However, as Figure 3.5 demonstrates, there is some spread about this broad trend. Figure 3.5 also shows that in general the energy requirements of the protocols $\mathcal{P}(3,*,*)$ exceed those of the protocols $\mathcal{P}(2,*,*)$, which in turn exceed those of $\mathcal{P}(1,*,*)$. The reason for this ordering is straightforward: the protocols fix the path of the inner cylinder axis, and the two cylinders approach one another more closely when the inner cylinder radius is greater, thereby incurring greater penalties in terms of viscous drag.

3.4.3 General trends

Our results, described in detail below, lead us to the conclusion that it is not possible to choose a single 'best' protocol from the 252 protocols examined, since the various measures do not agree in their relative rankings of these protocols. Although there is a strong correlation between some pairs of measures, other pairs are only weakly correlated. However, it is possible to make some general comments, as follow below. In this subsection we shall discount EU (which does not genuinely measure the quality of mixing, although it is important in any

discussions of mixing efficiency) and LW (which is only poorly correlated with all other measures; it is also unclear whether large or small values of this quantity are desirable until an application is specified). We therefore discuss here only the remaining eight mixing measures.

Rather arbitrarily we have chosen to focus on the twenty best protocols for each mixing measure. For DD, ID and RT, at least eighteen of the top twenty protocols have $2 \leq j \leq 4$ and $1 \leq k \leq 3$. A similar, but slightly less marked, preference for protocols in this region of parameter space is exhibited by KS, LE and MV. With IM there is a greater spread among top-twenty protocols, although all have $2 \leq j \leq 6$. These results indicate a strong preference for what might be termed 'flowery' protocols. A simple explanation for this is that the flowery protocols tend to have the longest path length Λ and therefore they do more stirring during one complete cycle of the protocol than do other protocols. However, upon further reflection, we see the true story is rather more subtle and that in fact the strong preference for protocols with $k \leq 4$ is intriguing because (see Figure 3.2) to the eye these protocols look similar to those with $k \geq 10$, whereas the latter protocols do not perform as well as the former (e.g. $\mathcal{P}(*, 3, 12)$ seem as convoluted as, say, $\mathcal{P}(*, 3, 2)$, but are not ranked as highly). In contrast, the top twenty for the final mixing measure, IS, stand alone in including a few protocols from the top half of Figure 3.2.

Some resolution to this puzzle can be obtained by considering Figure 3.6, which shows a qualitative picture in the form of iterated mapping plots for protocols of the form $\mathcal{P}(1, *, *)$ (i.e. those with cylinder radius ratio $a_{\text{in}}/a = 0.05$). From this figure it is clear that the 'best' protocols, with $2 \leq j \leq 4$ and $1 \leq k \leq 3$, have much smaller periodic islands (and therefore exhibit better mixing) than corresponding protocols with $k \geq 10$. Since all measures (except IS) involve some averaging over the fluid domain, the islands exact a penalty on the $k \geq 10$ protocols, which is why they do not appear in the top-twenty lists. The iterated mapping plots also indicate why IS does not discriminate against the $k \geq 10$ protocols: the initial line whose stretch is calculated lies entirely in the chaotic region, and so the existence of significant regular regions does not materially affect this mixing measure.

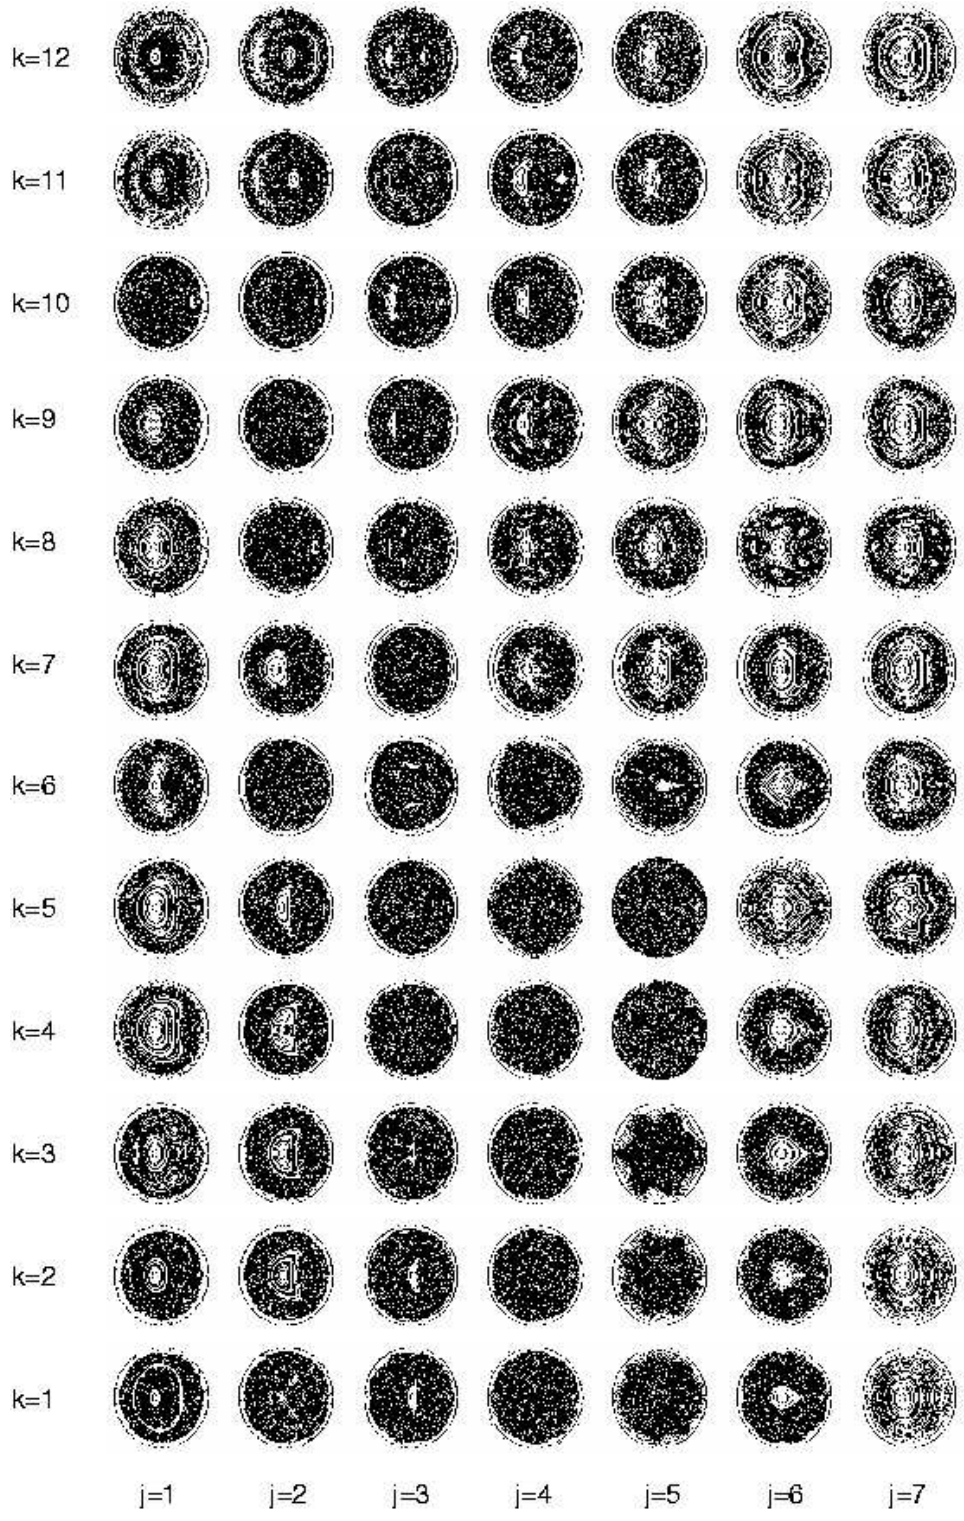

Figure 3.6: Iterated mapping plots for protocols $\mathcal{P}(1, *, *)$. This figure may be compared directly with Figure 3.2.

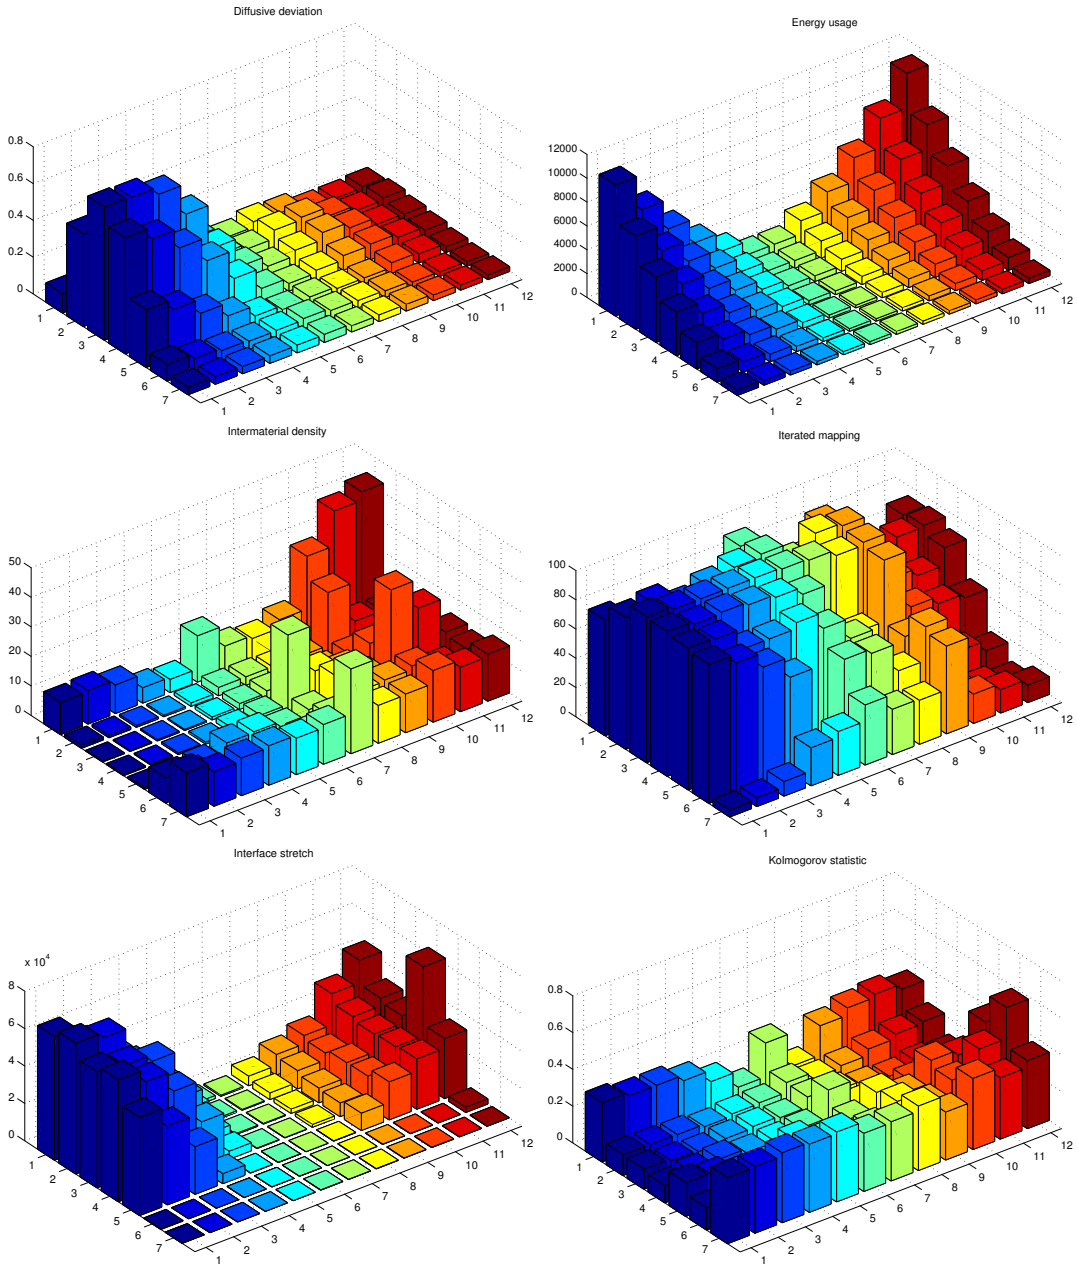

Figure 3.7: Plots of the ten mixing measures for each of the protocols $\mathcal{P}(2, *, *)$.

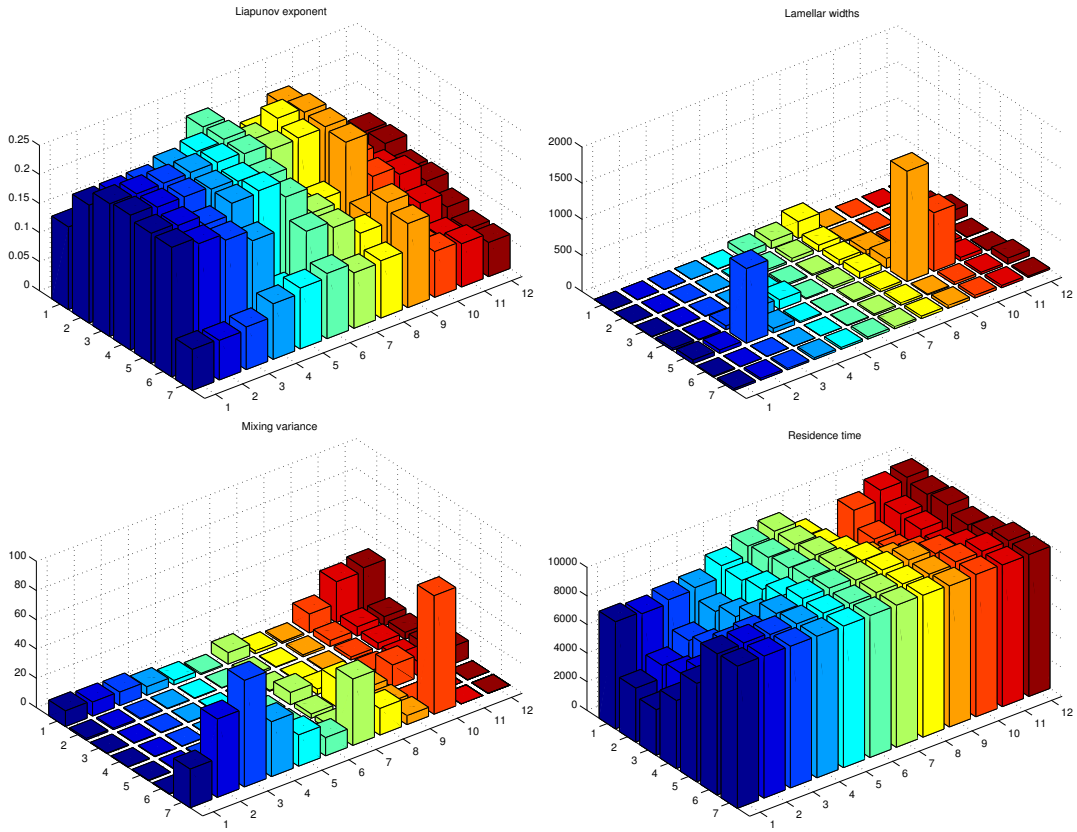

3.4.4 Detailed results

Figure 3.7 shows in detail data for the individual mixing measures, for the protocols $\mathcal{P}(2, *, *)$. As might be expected from Figure 3.6, many mixing measures show wide variation between the various protocols. For some measures, such as IS, there is a clear clustering of good protocols, while for others, such as MV, the good protocols are more isolated and scattered.

Plots of the pairwise joint distributions of the ranked mixing measures are shown in Figure 3.8; the corresponding rank correlation coefficients can be found in Table 3.2. To understand Figure 3.8, consider as an example the plot in column DD and row EU. This plot contains 252 data points, with coordinates (m_i, n_i), for $i = 1, \ldots, 252$. The m_i and n_i are the rankings of the i-th protocol according to the mixing measures DD and EU, respectively. The remaining plots in the figure are defined similarly. More detailed versions of these plots are shown in Figures 3.9, 3.10 and 3.11, which, in addition, illustrate the stirring protocols that correspond to the various data points in Figure 3.8.

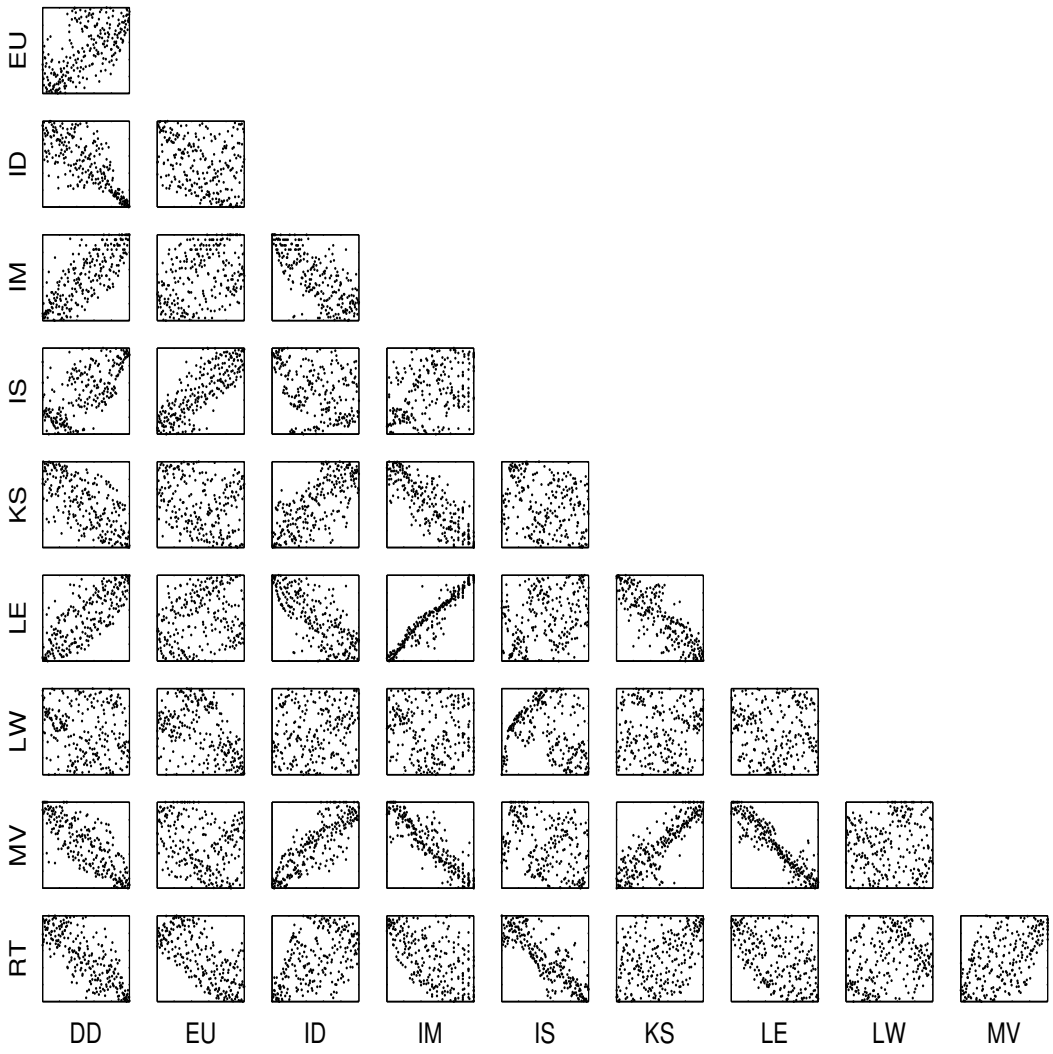

Figure 3.8: Plots of the joint distributions of the ranked mixing measures.

	DD	EU	ID	IM	IS	KS	LE	LW	MV	RT
DD	1.00	0.64	-0.78	0.79	0.68	-0.67	0.77	-0.18	-0.80	-0.80
EU	0.64	1.00	-0.31	0.37	0.81	-0.25	0.24	-0.44	-0.27	-0.70
ID	-0.78	-0.31	1.00	-0.74	-0.37	0.72	-0.75	0.15	0.84	0.52
IM	0.79	0.37	-0.74	1.00	0.40	-0.83	0.95	-0.12	-0.90	-0.57
IS	0.68	0.81	-0.37	0.40	1.00	-0.26	0.32	-0.23	-0.36	-0.85
KS	-0.67	-0.25	0.72	-0.83	-0.26	1.00	-0.84	0.06	0.87	0.42
LE	0.77	0.24	-0.75	0.95	0.32	-0.84	1.00	-0.05	-0.93	-0.52
LW	-0.18	-0.44	0.15	-0.12	-0.23	0.06	-0.05	1.00	0.04	0.27
MV	-0.80	-0.27	0.84	-0.90	-0.36	0.87	-0.93	0.04	1.00	0.56
RT	-0.80	-0.70	0.52	-0.57	-0.85	0.42	-0.52	0.27	0.56	1.00

Table 3.2: Correlation coefficient C between pairs of mixing measures.

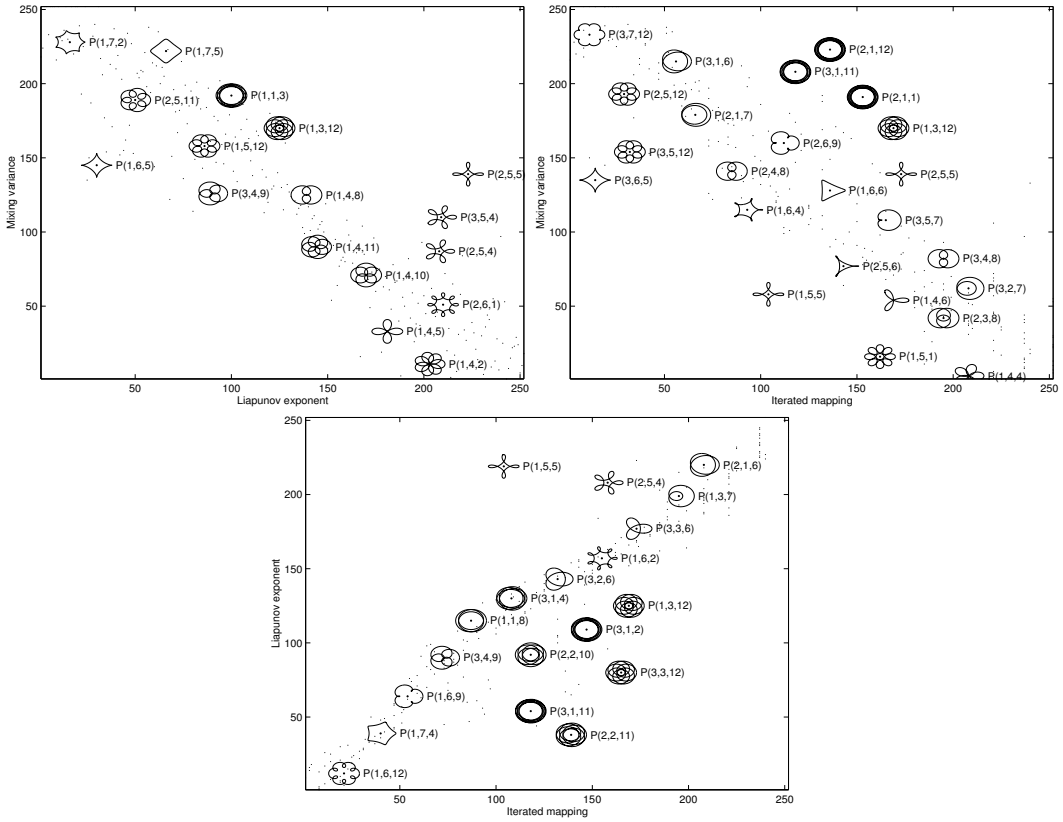

Figure 3.9: Joint distribution plots for strongly correlated measures, illustrating some of the protocols that correspond to various data points.

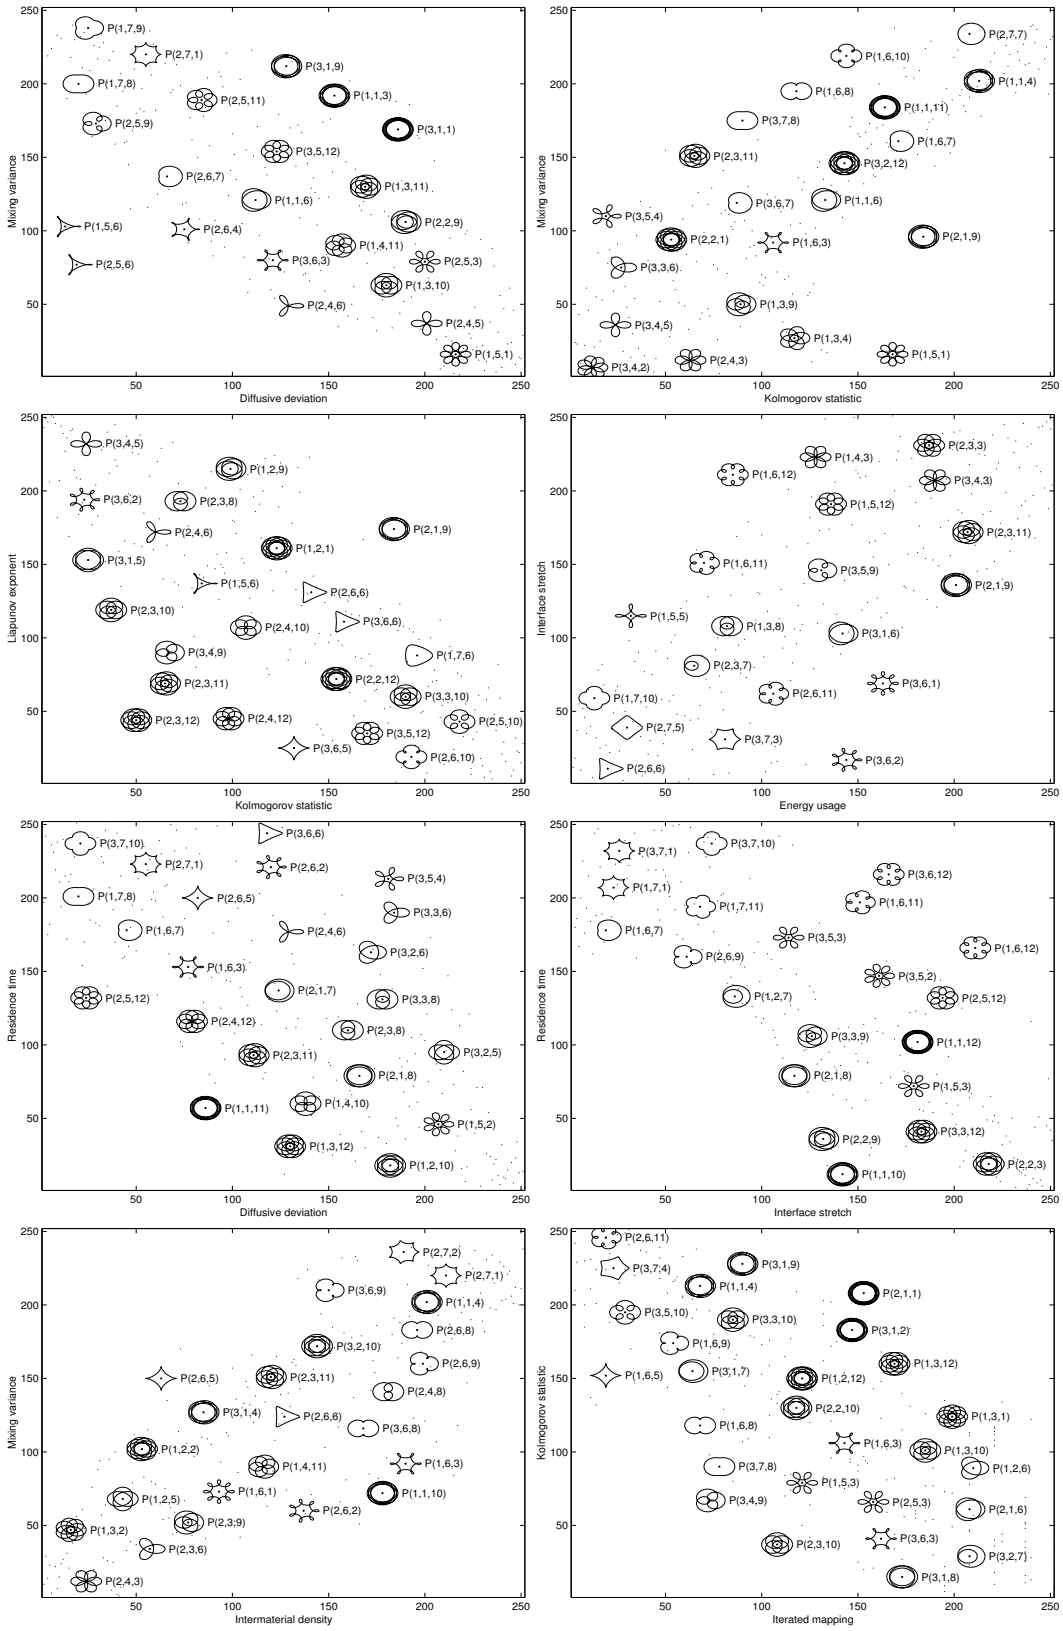

Figure 3.10: Joint distribution plots for well correlated measures, illustrating some of the protocols that correspond to various data points.

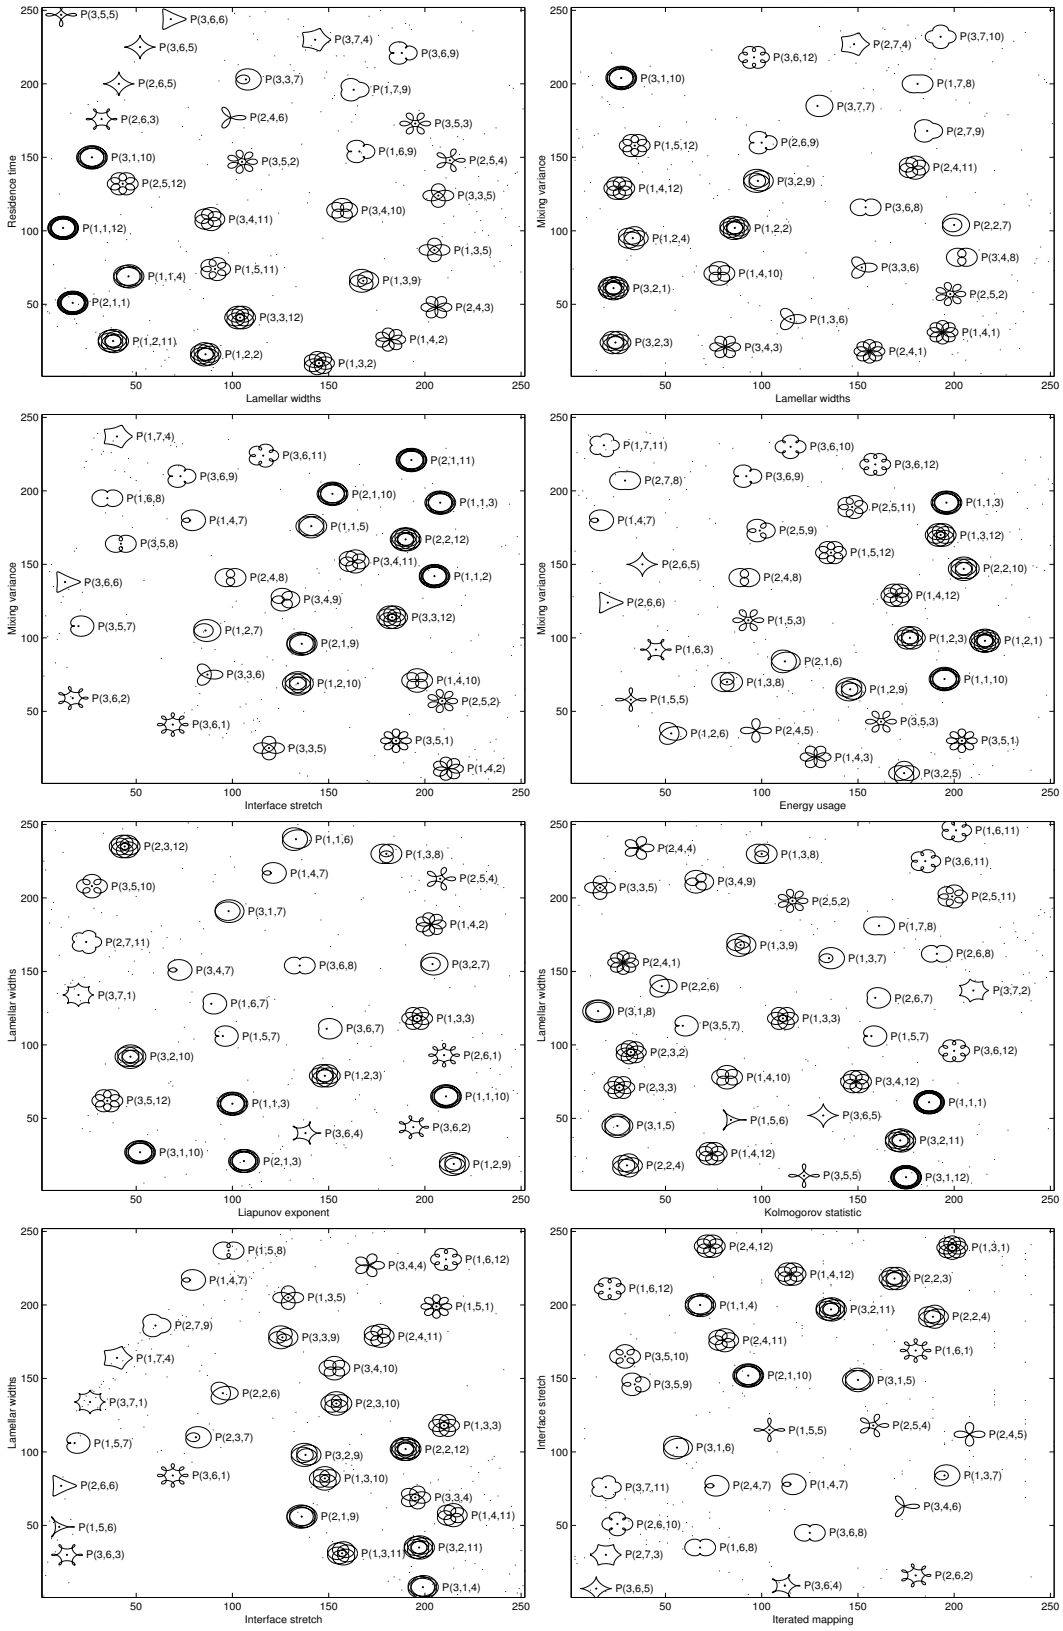

Figure 3.11: Joint distribution plots for weakly correlated measures, illustrating some of the protocols that correspond to various data points.

In the present study, the strongest correlations ($|C| \geq 0.9$), evident in Figure 3.8 and Table 3.2, are found between iterated mapping chaoticity, Liapunov exponent and mixing variance mixing measures. Correlation plots illustrating the protocols responsible for some of the data points are shown in Figure 3.9. MV is significantly simpler to compute than IM or LE, in terms of storage and time, respectively. Thus, for the protocols investigated here, one might reasonably substitute MV for either IM or LE, depending on one's computational constraints. These strong correlations accord with intuition, since each measure heavily penalises those protocols with large regular regions. We emphasise our earlier point that superficially similar protocols, e.g. $\mathcal{P}(2,1,7)$ and $\mathcal{P}(3,2,7)$, may lie near opposite ends of the spectrum of mixing quality (see Figure 3.9).

The next strongest correlations ($0.8 \leq |C| < 0.9$) exist between (MV,KS), (LE,KS), (MV,DD), (IS,EU), (RT,DD), (RT,IS), (MV,ID) and (KS,IM). When we consider these pairings (see Figure 3.10), we note that there is generally a clustering of 'like' protocols (e.g. those we have described above as 'flowery'). However, in some cases there is a sensitive dependence of results on the exact protocol. For example, in the (IS,EU) plot, long flowery protocols, such as $\mathcal{P}(1,6,12)$, $\mathcal{P}(1,4,3)$ and $\mathcal{P}(1,6,11)$ clearly perform well. Clustering of like protocols is also evident in the (RT,IS) plot, where simple protocols like $\mathcal{P}(3,7,1)$ perform poorly while broader, flowery protocols such as $\mathcal{P}(2,2,3)$ perform better. In the (RT,LW) plot similar clustering is evident. In the (KS,IM) plot the simple protocols exhibit a greater range of mixing quality, and greater sensitivity to protocol details is apparent. For example, the similar protocols $\mathcal{P}(3,2,7)$ and $\mathcal{P}(3,1,7)$ achieve good and poor mixing, respectively.

In a previous paper [47] we observed from a sample of three protocols the counter-intuitive result that patterns producing highly chaotic iterated mappings tended to generate relatively poor interface stretch (and vice versa). However, the corresponding joint distribution plot of the iterated mapping and interface stretch measures for the larger sample of 252 protocols shown in Figure 3.11 suggests this is not always the case. It can be seen that, while simpler patterns generally produce a poor interface stretch, the degree of chaoticity of the iterated mapping can depend sensitively upon the pattern shape. A simple thought experiment

shows why we should not expect interface stretching necessarily to provide much information about uniformity of mixing. Consider the BSD_1 in the limit as the scaled inner cylinder radius a_{in}/a becomes small (imagine using a wire rather than a stick to stir a tin of paint): although the inner cylinder may still be dragged repeatedly through a material line and thereby stretch it considerably, the thin cylinder transports little of the bulk of the fluid. Thus we should expect a relatively poor correlation between IS and other measures of mixing uniformity such as IM, MV and KS.

As well as showing sensitive dependence upon protocol parameters, some pairs of measures demonstrate a poor correlation ($|C| \leq 0.4$) – see Figure 3.11 for examples. Notably, the energy input, EU, does not correlate well with ID, IM, KS, LE or MV. This indicates that some low energy protocols may perform well according to these mixing measures. For example, both the $\mathcal{P}(1,2,6)$ and $\mathcal{P}(1,5,5)$ protocols have a desirable low mixing variance and also require relatively little energy input. Desirable low energy protocols for IM, KS and LE could be $\mathcal{P}(1,3,6)$, $\mathcal{P}(1,4,6)$ and $\mathcal{P}(1,5,5)$, respectively.

The relatively poor correlation between interface stretch and Liapunov exponent ($C = 0.32$) may seem counter-intuitive. However, this is an artifact of the crude manner in which we compute the average Liapunov exponent: we do not distinguish between points in the chaotic region (with positive Liapunov exponent) and points in a regular region (for which the finite-time Liapunov exponent is small). This technique tends to depress the computed average Liapunov exponent for flows with a significant area occupied by regular regions, whereas the interface chosen for the stretch calculation lies (by design) largely in the chaotic region, regardless of the particular protocol. In a more sophisticated implementation of our code, it would, of course, be possible to average over only points designated as lying in the chaotic region (using techniques from the iterated mapping chaoticity measure to identify such points).

3.5 Conclusions

We have described above the numerical evaluation of a range of mixing measures for a large number of stirring protocols in the batch stirring device. Before attempting to extrapolate from these results to the problem of protocol selection, we make some general comments.

First, we note that the preference for protocols with $2 \leq j \leq 4$ and $1 \leq k \leq 3$ suggests a certain robustness to the exact details of the protocol. This robustness is of particular practical importance because in deriving the velocity field for our simulations we have made the Stokes flow approximation, effectively setting the Reynolds number Re of the flow to be zero. Of course, in any real experiment performed at small, but finite, Reynolds number, corrections to this velocity field of $O(\mathrm{Re})$ will be required. The robustness of the mixing quality measures to protocol choice does not necessarily imply a similar insensitivity to finite-Reynolds number effects. To investigate this we have performed experimental dye advection in the BSD_1 at various Re between 10^{-3} and 10^{1} [27]. We find that as the Reynolds number increases, regular islands in the flow (where mixing is noticeably poor) shrink, leading to uniform mixing over a greater fraction of the domain, as shown in Figure 3.12. That is, the iterated mapping chaoticity measure IM increases with increasing Reynolds number for the protocol considered. Hobbs and Muzzio [62] have studied Reynolds number effects numerically in the Kenics static mixer (see Figure 1.5), finding that for $\mathrm{Re} < 10$ and $\mathrm{Re} > 1000$ their flow is globally chaotic, but for $\mathrm{Re} \sim 100$ significant regular islands of poor mixing appear.

The measures adopted in this chapter were selected because they allow for automated computation. There is inevitably some arbitrariness in our sample of initial tracer particle locations. However, we should bear in mind that many other measures are inherently analytical, requiring significant human intervention. For example, a major class of measures ignored here concerns the location of periodic points and an understanding of the structure of their associated invariant manifolds [111]. Although periodic points can be found systematically, the requisite numerical algorithms [37] are computationally intensive, and extremely good

Figure 3.12: From left to right: iterated mapping plot, showing islands in numerical simulation of Stokes flow; large time dye-advection experiments in glycerol, with $\mathrm{Re} = 0.1, 1, 3, 10$.

initial guesses for their locations are sometimes required. Elliptic fixed points, which are found at the centre of periodic islands, are comparatively easy to discern (provided the island is visible to the eye). Hyperbolic points, on the other hand, reside in the chaotic sea and are more difficult to identify. Furthermore, although protocol symmetries can be exploited to find periodic points [52, 53, 56, 88, 95], this process often requires significant human intuition, and even then is not guaranteed to locate all such points; it is thus difficult to program robustly. It would, nevertheless, be instructive to investigate the extent to which the automated computational mixing measures contemplated here agree with human analyses and intuitions about real-life stirring devices.

Another important computational practicality is the extent to which the mixing measures considered here could feasibly be evaluated when the velocity field is not known. We have already undertaken a significant computational endeavour in evaluating such a large number of protocols; the additional time required to compute the velocity field for each protocol in a mixer with no expression for the velocity could render similar investigations intractable.

Finally, we note that, from the point of view of protocol selection, it would also be useful to study the inverse problem to that studied here, i.e. rather than to determine which of a given set of stirring protocols performs best, according to a given mixing measures, to design a protocol to maximise a target mixing measure, subject to constraints on others (such as energy usage). Although such an undertaking is beyond our present scope, optimisation studies have been attempted on other mixers, such as lid-driven cavity flows [84, 143], the cavity

transfer mixer [148] and a twin screw extruder [84] by approximating the stirring dynamics sufficiently to make computations practicable.

3.5.1 Implications for protocol selection

We now summarise some of our results as they influence the selection of 'optimal' stirring protocols.

First we emphasise that the best protocols cannot be selected 'by eye': only detailed calculation seems able to distinguish the mixing capabilities of $\mathcal{P}(2,3,1)$ (a good protocol) from those of $\mathcal{P}(2,3,12)$ (a poor one), or $\mathcal{P}(2,6,2)$ (a reasonable protocol) from $\mathcal{P}(2,6,12)$ (a bad one).

We have also found that a useful rule of thumb for eliminating some of the poorest protocols is that for good mixing the inner cylinder should move through a large region of the flow domain (e.g. $\mathcal{P}(2,3,1)$), rather than on a more limited, nearly circular path (e.g. $\mathcal{P}(2,7,7)$). This rule alone though is not sufficiently sophisticated to facilitate selection of the best protocols: for example, presumably it would not discriminate between the relative strengths of $\mathcal{P}(2,4,1)$ and $\mathcal{P}(2,4,12)$. Exploration of the entire domain by the mixer impeller is often exclaimed as proof of perfect mixing in brochures for industrial batch mixers. However, there are exceptions to the rule, as for instance with $\mathcal{P}(2,6,1)$ (which mixes well despite the path of the inner cylinder, which explores relatively little of the flow domain) and $\mathcal{P}(2,5,11)$ (which mixes poorly despite the more extensive path of the inner cylinder).

Another design rule of thumb might be simply to choose a protocol with as great a path length Λ as possible, on the basis that the more you stir the better the mixing. To a certain extent this is the case, the rank correlation coefficients of Λ with DD, IS and RT being 0.64, 0.89 and -0.83, respectively. (The correlation between path length and EU is of course the highest: $C = 0.94$.) However, the rank correlation coefficients of Λ with the remaining six mixing measures are all less than 0.4 in magnitude. We therefore conclude that the variations in mixing effectiveness between various protocols are not simply a matter of

some protocols having longer paths than others; there is much to be gained from judiciously selecting the manner in which the inner cylinder moves, as well as merely maximising the length of its path.

Chapter 4

Topological chaos in two-dimensional inviscid and viscous mixers

4.1 Introduction

In the previous chapter we have seen that optimising a mixer against a given mixing quality measure represents a significant challenge. This is mainly due to the sensitivity of chaotic flows to parameter changes and also due to computational limitations preventing exhaustive parameter-space searches. It is desirable, therefore, to find a method of designing a mixer that has a certain degree of mixing quality built-in, which is robust with respect to changes in fluid properties, and also allows additional tuning of various mixer parameters.

A significant recent theoretical advance, which allows good mixing to be built in, concerns the concept of 'topological chaos' [18, 19, 20, 92, 123]: without reference to any computation or exact flow details, Boyland *et al.* [20] have demonstrated, in an unusual blend of ad-hoc experimentation and abstract mathematics, that flows with the topology of certain braids (e.g. like plaited hair) are guaranteed to provide effective mixing. By topological chaos, we mean a complexity that cannot be removed by continuous deformation of the fluid region *unless* that deformation

involves appropriate changes to the topology of the region. According to the theory described by Boyland *et al.* [20], a *quantitative* material stretch rate can be predicted, depending only on the topology of the flow. Remarkably, the theory requires only fluid continuity, and is independent of the geometrical parameters or exact fluid properties (viscosity, compressibility, rheology, etc).

However, a key feature that is not predicted by the theoretical considerations of Boyland *et al.* [20] is the area of the domain in which the stretch rate is attained. Here we provide numerical and experimental results that support the observations of Boyland *et al.*, that the chaotic region is commensurate with the region of fluid through which the stirring elements move during the operation of their device.

In this chapter we study the BSD_m, a generalisation of the batch stirring device described in Chapter 3, in which an arbitrary number m of cylindrical rods of circular cross-section move independently to stir a fluid. In order to generate topological chaos in such a device, it is necessary to have three or more stirring rods [20]. As we have seen for the BSD_1 in the previous chapter, the generation of chaotic flow with fewer stirring rods is perfectly possible, but quality mixing relies on a careful tuning of the system parameters, and is not guaranteed by the considerations of topological chaos theory. In this sense, the known velocity fields for Stokes flow with two or fewer stirring elements are inadequate; such flows include, in addition to the BSD_1, the double journal bearing flow [66, 93], the rotated arc mixer [127], and the partitioned pipe mixer [85, 94] used to model the Kenics static mixer (Figure 1.5).

We construct the velocity field for two-dimensional flow in the BSD_m, for either an inviscid fluid or Stokes flow in a viscous fluid; in each case, we consider first the simpler model of an unbounded flow and then the more analytically sophisticated, but physically more realistic, model of bounded flow inside a cylindrical tank. When there is only a single stirring element, it is possible to construct the exact velocity field for Stokes flow in closed form [47, 145], using finitely many terms, by means of image systems of singularities [14, 132]; however, we have found that attempts to derive the corresponding velocity field with multiple stirring elements lead to insurmountable analytical difficulties associated with a fractal set of singularities. The alternative series solution adopted here

is based on an approach due to Price, Mullin and Kobine [121] and Prof. L. N. Trefethen (personal communication). We construct series whose terms individually satisfy the governing equations exactly, and which, provided the series are carefully chosen, satisfy the boundary conditions with spectral accuracy when truncated for numerical purposes. This technique is fast and easy to implement numerically and allows accurate simulations to be performed in a time-dependent flow. Vikhansky [142] has computed some corresponding flows with three stirring rods using the 'immersed boundary method': these illustrate the chaotic motion of fluid particles in the BSD_3, but a corresponding level of accuracy is difficult to achieve in his simulations, or indeed any which use a finite-element or finite-difference approach.

We test the topological chaos stretch predictions given by Boyland *et al.* [20] by numerically calculating stretch rates for both inviscid and viscous models of their mixing device (the BSD_3 with three stirring rods), and comparing them with the theoretical stretch rates. We also validate our numerical dye-advection simulations by comparing them with corresponding experiments performed in a simple rig.

The structure of the chapter is as follows. In Section 4.2 we describe the BSD_m, and present a fast algorithm for accurately finding the corresponding velocity field. In Section 4.3 we summarise key results from topological stretch rate theory, and in Section 4.4 we present numerical stirring simulations which allow comparison with the abstract theory of Boyland *et al.* [20]. In Section 4.5 we compare experiments and numerics to validate our results. We then discuss practical issues that need to be considered when constructing topological mixers. Discussion and conclusions are given in Sections 4.6 and 4.7.

4.2 Mathematical models for the BSD_m

We begin this section by describing the BSD_m [20]. Then we present a fast, accurate method for constructing the velocity field in this device. Four increasingly sophisticated models, involving inviscid or viscous fluid, in either an infinite or

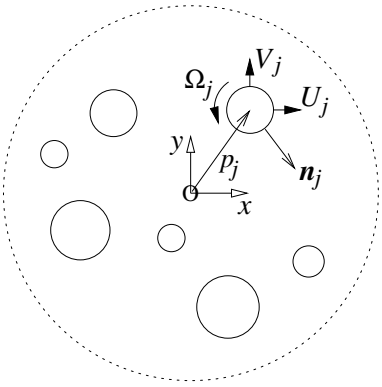

Figure 4.1: Geometry of the BSD_m. Here x and y are Cartesian coordinates. Two-dimensional flow is driven by m circular cylinders, each respectively centred at $z = p_j$ (where $z = x + \mathrm{i}y$), with radius a_j, translating with velocity (U_j, V_j) and (for the viscous model only) rotating about its axis with angular velocity Ω_j. The unit normal of the j-th cylinder is $\boldsymbol{n}_j$. The fluid may be either inviscid or viscous, and may be either infinite in extent or bounded by a circular cylinder centred at $z = 0$ with radius a.

a finite flow domain, are discussed in turn. We start with the simplest case of unbounded, inviscid flow to illustrate the technique, and progress to the most analytically sophisticated bounded, viscous model, upon which we shall concentrate attention later.

The BSD_m geometry is shown in Figure 4.1. Flow is two-dimensional, in the xy-plane. For both the inviscid and viscous models, it proves convenient to work with complex coordinates $z = x + \mathrm{i}y$ and $\bar{z} = x - \mathrm{i}y$, where x and y are Cartesian coordinates as indicated in Figure 4.1. We assume that there are m infinitely long circular cylinders placed in the fluid, with their axes at locations $z = p_j(t)$ and with radii a_j $(j = 1, \ldots, m)$. Each cylinder may move independently with a prescribed velocity $U_j + \mathrm{i}V_j = \mathrm{d}p_j/\mathrm{d}t$ and, in the case of the viscous model only, may rotate about its axis with angular velocity $\Omega_j(t)$. Although in practice the fluid domain is finite, we shall initially model the m cylinders as lying in an infinite fluid; by the method of images, we then refine the model so that the flow domain is finite, the cylinders being enclosed in the circular domain $|z| \leq a$.

In order to simulate mixing in the BSD_m numerically, we shall need to track with

high accuracy the motion of inertialess tracer particles as they are advected with the fluid. It is therefore advantageous to have available an exact expression for the fluid velocity field, which allows evaluation of the velocity at an arbitrary point in the flow domain (rather than, say, a finite-difference approximation, where some interpolation is in general required). Unfortunately, for $m > 1$, such an exact expression is not available in closed form, so instead we calculate a series solution involving terms that individually satisfy the appropriate governing equation exactly. The series coefficients are then determined by computing the least-squared-error in the boundary conditions, subject to certain constraints, detailed later. The component solutions are carefully chosen to give overall spectral accuracy with only a small number of terms (Prof. L. N. Trefethen, personal communication). The least-squares optimisation is carried out in Matlab, which provides a convenient environment in which to implement our algorithm, since it handles complex arithmetic and provides the fast and robust `lsqnonlin` function in the Optimisation Toolbox.

Mixing is achieved in the BSD$_m$ by moving the m cylinders and, hence, varying the geometry of the flow domain, in a time-periodic fashion. In order to avoid repeated recalculation of the same velocity fields, it is efficient first to compute the velocity field at a large number of time steps, and then to store the results of these calculations for ultimate use in simulations of tracer advection. In this respect it is significant that our method of solution is not only faster and more accurate than alternative finite-element or finite-difference techniques, but it requires less storage, since we need only the series coefficients, rather than the velocity field at a full set of grid points.

4.2.1 Inviscid, irrotational model

Our first, and simplest, model for the BSD$_m$ is of incompressible, inviscid, irrotational flow. Here it is most convenient to work with the complex potential $w = \phi + \mathrm{i}\psi$. The functions ϕ and ψ are the velocity potential and streamfunction, respectively, and are each harmonic functions [2]. This is guaranteed if $w = f(z)$ for any analytic function f, the corresponding fluid velocity field (u, v) being

obtained from

$$u - \mathrm{i}v = \frac{\mathrm{d}w}{\mathrm{d}z}. \tag{4.1}$$

Unbounded flow

For the unbounded, inviscid model, the boundary conditions on the cylinders are

$$(u, v) \cdot \boldsymbol{n}_j = (U_j, V_j) \cdot \boldsymbol{n}_j \quad \text{on} \quad |z - p_j| = a_j \quad (j = 1, \ldots, m), \tag{4.2}$$

where $\boldsymbol{n}_j$ is the unit normal vector of the j-th cylinder (see Figure 4.1).

Following [121], we look for a complex potential in the form

$$w = \sum_{j=1}^{m} \left\{ b_{j,1} \log(z - p_j) + \sum_{k=2}^{n} b_{j,k} (z - p_j)^{1-k} \right\}, \tag{4.3}$$

where $b_{j,k}$ are mn unknown constants, corresponding to the sum of a logarithm and a Laurent series centred on the axis of each of the cylinders. (The exact solution is obtained in the limit $n \to \infty$.) The velocity field resulting from this ansatz, according to (4.1), is readily verified to satisfy

$$u - \mathrm{i}v = \sum_{j=1}^{m} \left\{ b_{j,1} (z - p_j)^{-1} + \sum_{k=2}^{n} b_{j,k} (1 - k)(z - p_j)^{-k} \right\}. \tag{4.4}$$

As this expression stands, the fluid velocity decays to zero at infinity. Since there can be no net source of fluid inside any internal cylinder, each $b_{j,1}$ must be pure imaginary. Furthermore, by the Kelvin circulation theorem [2], each $b_{j,1}$ must remain constant during the operation of the device; we choose $b_{j,1} = 0$, which corresponds to starting the system from rest, although clearly other fixed circulations could be imposed.

The remaining $m(n-1)$ coefficients in (4.3) may then be found by minimising the squared error in the boundary conditions (4.2) on the cylinders [17]; the accuracy of the method is illustrated in Table 4.1 below. A sample of streamline plots for this model of the BSD_m is shown in Figure 4.2.

Bounded flow

In this section we construct a bounded, inviscid model for the BSD_m. In addition to the boundary conditions specified by (4.2), we also impose the no-penetration

condition that

$$(u, v) \cdot \boldsymbol{n} = 0 \quad \text{on} \quad |z| = a, \tag{4.5}$$

where $\boldsymbol{n}$ is the unit normal of the bounding cylinder.

We can automatically satisfy this extra boundary condition by using the Milne-Thomson circle theorem [2, 99], which states that if the complex potential $w = f(z)$ has singularities that lie in $|z| < a$, then the complex potential

$$w = f(z) + \overline{f(a^2/\bar{z})} \tag{4.6}$$

has the same set of singularities inside $|z| < a$ (except possibly for an additional singularity at the origin $z = 0$) and also satisfies the boundary condition $\mathrm{Im}\, w = 0$ on $|z| = a$, equivalent to (4.5). In most cases any extraneous singularities at the origin can be removed by repeated applications of the circle theorem [47].

Taking $f(z)$ to be the complex potential (4.3) for the unbounded problem, we apply (4.6) to generate the new potential

$$\begin{aligned} w &= \sum_{j=1}^{m} \left\{ b_{j,1} \log(z - p_j) + \sum_{k=2}^{n} b_{j,k} (z - p_j)^{1-k} \right\} \\ &\quad + \sum_{j=1}^{m} \left\{ \bar{b}_{j,1} \left(\log(z - q_j) + \log(-\bar{p}_j/a) \right) + \sum_{k=2}^{n} \bar{b}_{j,k} \left(\frac{a^2}{z} - \bar{p}_j \right)^{1-k} \right\}, \end{aligned} \tag{4.7}$$

which automatically satisfies the boundary condition (4.5), where

$$q_j = \frac{a^2}{\bar{p}_j}.$$

As explained for the unbounded case, each $b_{j,1}$ must be purely imaginary and must remain constant during the operation of the device.

The velocity field corresponding to (4.7) can then be found using (4.1), and satisfies

$$\begin{aligned} u - \mathrm{i}v &= \sum_{j=1}^{m} \left\{ b_{j,1} (z - p_j)^{-1} + \sum_{k=2}^{n} b_{j,k} (1 - k)(z - p_j)^{-k} \right\} \\ &\quad + \sum_{j=1}^{m} \left\{ \bar{b}_{j,1} (z - q_j)^{-1} + \sum_{k=2}^{n} \bar{b}_{j,k} (1 - k) \frac{q_j}{z} \left(-\frac{\bar{p}_j}{z} \right)^{1-k} (z - q_j)^{-k} \right\}. \end{aligned} \tag{4.8}$$

As in Section 4.2.1, the coefficients $b_{j,k}$ $(k > 1)$ are found by minimising the squared error in the boundary conditions (4.2); see Table 4.1. By construction,

the condition (4.5) on the outer cylinder is satisfied automatically. A sample of streamline plots for this model is shown in Figure 4.2.

Although we shall not implement such a scheme here, we note that use of the Milne-Thomson circle theorem may be avoided by adopting instead of (4.7) an alternative series expansion, which considers the outer boundary cylinder on the same footing as the internal cylinders. The solution is then posed as the sum of $m+1$ series, respectively centred on the axes of all $m+1$ cylinders (m internal and one external). The boundary condition (4.5) is then subject to squared-error minimisation, rather than being forced exactly by the structure of the solution. This technique is, however, rather more algebraically involved than that described above, principally because of the more pronounced role of the singularity at the origin, and of the special care required to remove it.

4.2.2 Viscous model

Having begun with the more straightforward, though less physical, inviscid model, we now turn to our model for the BSD_m involving an incompressible, Newtonian viscous fluid. We consider Stokes flow, at vanishing Reynolds number, which is specified in terms of the streamfunction $\psi(z,\bar{z})$, which satisfies the biharmonic equation $\nabla^4\psi = 0$ [86]. The general solution to this equation is [99]

$$\psi = \bar{z}f(z) + g(z) + c.c., \tag{4.9}$$

for any analytic functions f and g, where *c.c.* means the complex conjugate of the preceding terms. The resulting velocity field may be obtained from

$$u + \mathrm{i}v = -2\mathrm{i}\frac{\partial\psi}{\partial\bar{z}}. \tag{4.10}$$

Unbounded flow

For the viscous model in an unbounded fluid domain, our task is to find functions f and g consistent with the no-slip boundary conditions

$$u + \mathrm{i}v = U_j + \mathrm{i}V_j + \mathrm{i}\Omega_j(z - p_j) \quad \text{on} \quad |z - p_j| = a_j \quad (j = 1, \ldots, m). \tag{4.11}$$

Note that, in addition to the translational velocities of the stirring rods, their rotational velocities Ω_j now also influence the flow field. Although rotation of a rod about its axis does not influence the topology of the flow, it does provide an additional degree of freedom for tuning parameters to optimise mixing performance.

Motivated by the exact solution for the BSD_1 with a single moving cylinder in a bounded domain [47] (as we have employed in Chapter 3), we construct a trial streamfunction with a rotlet and Stokeslet singularity at the centre of each of the m cylinders. By analogy with our results from the previous section, we also include in our expressions for f and g series in $(z-p_j)$ for $j=1,\ldots,m$, and, in addition, a term representing a constant translational velocity, ζ, at infinity (to be determined as part of the solution). We are thus led to the expression

$$\begin{aligned}\psi &= \bar{z}\sum_{j=1}^{m}\left\{\tfrac{1}{2}b_{j,1}\log|z-p_j| + \sum_{k=2}^{n} b_{j,k}(z-p_j)^{1-k}\right\} \\ &\quad + \sum_{j=1}^{m}\left\{\tfrac{1}{2}c_{j,1}\log|z-p_j| + \sum_{k=2}^{n} c_{j,k}(z-p_j)^{1-k}\right\} + \tfrac{1}{2}\mathrm{i}\zeta\bar{z} + c.c. \quad (4.12)\end{aligned}$$

Using (4.10), we deduce that the velocity field associated with the streamfunction (4.12) is given by

$$\begin{aligned}u+\mathrm{i}v &= -2\mathrm{i}\sum_{j=1}^{m}\left\{\tfrac{1}{4}b_{j,1}\left[\log|z-p_j|^2 + \frac{\bar{z}}{\bar{z}-\bar{p}_j}\right] + \bar{b}_{j,1}\frac{z}{\bar{z}-\bar{p}_j}\right. \\ &\quad \left. + \sum_{k=2}^{n}\left[b_{j,k}(z-p_j)^{1-k} + \bar{b}_{j,1}(1-k)z(\bar{z}-\bar{p}_j)^{-k}\right]\right\} \\ &\quad - 2\mathrm{i}\sum_{j=1}^{m}\left\{\tfrac{1}{2}c_{j,1}\frac{1}{\bar{z}-\bar{p}_j} + \sum_{k=2}^{n}\bar{c}_{j,k}(1-k)(\bar{z}-\bar{p}_j)^{-k}\right\} + \zeta \quad (4.13)\end{aligned}$$

For general choices of the constants $b_{j,k}$, the velocity field diverges at infinity. We force the limit $u+\mathrm{i}v \to \zeta$ at infinity by applying the condition

$$\sum_{j=1}^{m} b_{j,1} = 0. \quad (4.14)$$

(In general one cannot find a solution of the form (4.12) for which $\zeta = 0$ [69, 70, 147].) The squared error in the boundary conditions (4.11) is then minimised to determine the various coefficients in (4.12); see Table 4.1 and Figure 4.2.

Bounded flow

We now consider our final model, of bounded, viscous flow in the BSD_m. Here, in addition to the boundary conditions (4.11) specified on the m inner cylinders, we also require

$$u + \mathrm{i}v = 0 \quad \text{on} \quad |z| = a, \tag{4.15}$$

so that there is no slip on the outer cylinder.

Our approach towards accommodating the external boundary is similar to that used in Section 4.2.1 for inviscid flow. We can satisfy the boundary condition (4.15) automatically by applying the circle theorem for the biharmonic equation [14, 47, 54, 97, 132] to the series (4.12). This theorem states that if the streamfunction $\psi = \bar{z}f(z) + g(z) + c.c.$ has singularities in $|z| < a$, then the streamfunction

$$\begin{aligned}\psi &= \bar{z}f(z) + g(z) - z\overline{f(a^2/\bar{z})} - \overline{g(a^2/\bar{z})} + (a^2 - |z|^2)\left\{\overline{f'(a^2/\bar{z})} + z^{-1}\overline{g'(a^2/\bar{z})}\right\} \\ &\quad + c.c. \end{aligned} \tag{4.16}$$

has the same set of singularities in $|z| < a$ (except perhaps for a further singularity at $z = 0$) and, in addition, satisfies (4.15). The undesirable, unphysical singularity induced at the origin by direct implementation of (4.16) can, in most cases, be removed by further applications of the circle theorem [47, 54]. Following the procedure given by [47], we calculate from (4.16) that the modified trial streamfunction is

$$\begin{aligned}\psi &= \sum_{j=1}^{m}\left\{b_{j,1}\left[\frac{1}{4}\bar{z}\log\frac{q_j|z-p_j|^2}{p_j|z-q_j|^2} + \frac{|z|^2-a^2}{4|p_j|^2|z-q_j|^2}\left(\bar{p}_j\bar{z}(z+p_j) - a^2(\bar{z}+\bar{p}_j)\right)\right]\right. \\ &\quad \left. + \sum_{k=2}^{n} b_{j,k}\left[\bar{z}(z-p_j)^{1-k} + ((p_j+(k-1)z)\bar{z} - a^2k)(-p_j)^{-k}\bar{z}^k(\bar{z}-\bar{q}_j)^{-k}\right]\right\} \\ &\quad + \sum_{j=1}^{m}\left\{c_{j,1}\left[\frac{1}{4}\log\frac{|q_j||z-p_j|^2}{|p_j||\bar{z}-\bar{q}_j|^2} + \frac{(|z|^2-a^2)(|z|^2-|q_j|^2)}{4a^2|z-q_j|^2}\right]\right. \\ &\quad \left. + \sum_{k=2}^{n} c_{j,k}\left[(z-p_j)^{1-k} + ((p_j+(k-1)z)\bar{z} - a^2k)(-p_j)^{-k}\bar{z}^{k-1}(\bar{z}-\bar{q}_j)^{-k}\right]\right\} \\ &\quad + c.c. \end{aligned} \tag{4.17}$$

In order successfully to remove the singularity at the origin, we must impose the constraint

$$\mathrm{Im} \sum_{j=1}^{m} c_{j,1} = 0; \tag{4.18}$$

that is, the sum of the rotlet strengths must be real.

It may be shown, after a great deal of tedious algebra, that (4.17) generates through (4.10) a velocity field $u + \mathrm{i}v$ given by

$$\begin{aligned}
u + \mathrm{i}v \;=\; & -2\mathrm{i}\sum_{j=1}^{m}\Bigg\{\tfrac{1}{4}b_{j,1}\log\frac{q_j|z-p_j|^2}{p_j|z-q_j|^2} \\
& \quad + \tfrac{1}{4}b_{j,1}\left[\frac{z}{p_j}+\frac{z}{z-q_j}+\frac{\bar{p}_j}{\bar{z}-\bar{p}_j}-\frac{a^4(z-p_j)}{p_j^3(\bar{z}-\bar{q}_j)^2}-\frac{\bar{q}_j}{\bar{z}-\bar{q}_j}\right] \\
& \quad + \tfrac{1}{4}\bar{b}_{j,1}\left[\frac{z^2}{\bar{p}_j(z-q_j)}+\frac{z}{\bar{z}-\bar{p}_j}-\bar{q}_j\frac{z-p_j}{(\bar{z}-\bar{q}_j)^2}-\frac{z}{\bar{z}-\bar{q}_j}\right] \\
& \quad + \sum_{k=2}^{n}\Bigg[b_{j,k}(z-p_j)^{1-k} \\
& \quad + \bar{b}_{j,k}(k-1)z\left((-\bar{p}_j)^{-k}z^k(z-q_j)^{-k}-(\bar{z}-\bar{p}_j)^{-k}\right) \\
& \quad + b_{j,k}(-p_j)^{-k}\bar{z}^k(\bar{z}-\bar{q}_j)^{-1-k} \\
& \qquad \left(\frac{a^4k^2}{p_j\bar{z}}+(p_j+(k-1)z)(\bar{z}-(k+1)\bar{q}_j)\right)\Bigg]\Bigg\} \\
& - 2\mathrm{i}\sum_{j=1}^{m}\Bigg\{c_{j,1}\frac{(a^2-|z|^2)\,(a^6-|p_j|^2(2a^2-p_j\bar{z})(a^2-\bar{p}_jz+|z|^2))}{2a^2(a^2-\bar{p}_jz)(\bar{z}-\bar{p}_j)(a^2-p_j\bar{z})^2} \\
& \quad + \sum_{k=2}^{n}\Bigg[c_{j,k}k(k-1)\bar{q}_j(-p_j)^{-k}(a^2-|z|^2)\bar{z}^{k-2}(\bar{z}-\bar{q}_j)^{-1-k} \\
& \quad - \bar{c}_{j,k}(k-1)\left((\bar{z}-\bar{p}_j)^{-k}-(-\bar{p}_j)^{-k}z^k(z-q_j)^{-k}\right)\Bigg]\Bigg\}. \qquad (4.19)
\end{aligned}$$

As in the previous three models, the coefficients (here $b_{j,k}$ and $c_{j,k}$) may be determined by minimising the squared error in the boundary conditions (4.11) on the inner cylinders, subject to (4.18); see Table 4.1 and Figure 4.2.

Although we shall not implement this feature here, we note that a rotation of the outer cylinder about its axis with angular velocity Ω_{out} can readily be modelled by adding a rotation term $\mathrm{i}\Omega_{\mathrm{out}}z$ to the right-hand side of (4.19) before minimising the squared error in the boundary conditions on the internal cylinders.

4.2.3 Features of the four models

Typical streamline plots for the BSD_m with two, three, four and five stirring rods are shown together in Figure 4.2, according to the four models presented above (inviscid or viscous, unbounded or bounded). It is not our intention here

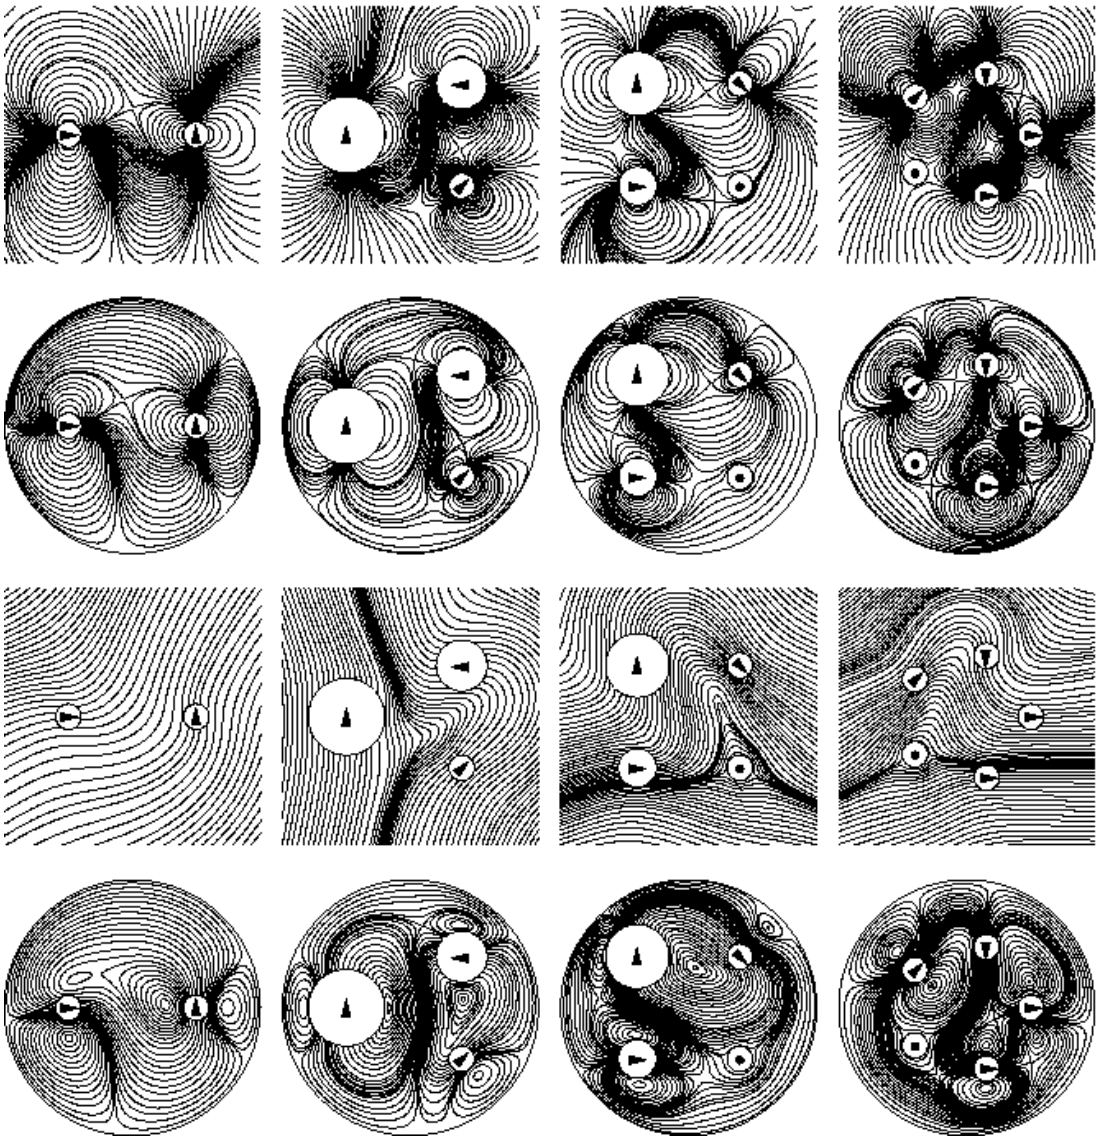

Figure 4.2: Illustrative plots showing instantaneous streamlines for the four BSD_m models. From top to bottom, the rows correspond to unbounded inviscid, bounded inviscid, unbounded viscous and bounded viscous flow. Each column corresponds to identical cylinder positions and velocities (with directions indicated by arrows), allowing the four flow regimes to be compared. In each case, all cylinders have zero angular velocity about their respective axes. For the inviscid plots, zero circulation is imposed around each cylinder.

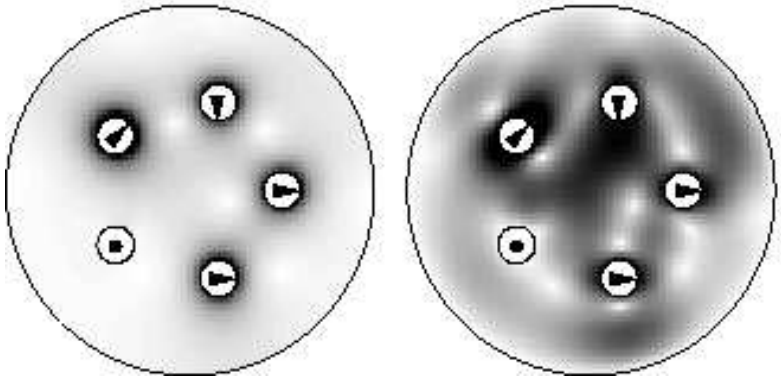

Figure 4.3: Grey-scale plots of the fluid speed, for the bounded inviscid (left) and viscous (right) models of the BSD_5. Stationary fluid is depicted in white, the most rapidly moving fluid in black, with an intermediate linear grey-scale. The fluid motion is more localised to the vicinity of the moving cylinders in the inviscid case than in the viscous regime.

to undertake a discussion of the possible flow topologies (see [121] for such a discussion with $m = 2$); however we make some general comments.

First we comment on the significant qualitative differences between the flows in the unbounded inviscid and viscous regimes (the square plots in Figure 4.2), associated with the different far-field behaviour of the velocity fields. In the inviscid case the velocity field decays quickly to zero, whereas in the viscous case the velocity field tends more slowly to a non-zero value.

In the bounded models, the question of the behaviour at infinity does not arise, yet the inviscid and viscous flows (the circular plots in Figure 4.2) remain quite distinct. Figure 4.3 shows a plot of the fluid speed for the two bounded five-cylinder flows in Figure 4.2. We note that in the inviscid case the fluid motion is more localised, because the fluid slips around the moving cylinders; by contrast, in the viscous scenario the fluid motion is more uniform, with the fluid speed being of the same order over much of the domain. One consequence of this difference is that material stretching in the inviscid BSD_m is also more localised than in the viscous BSD_m; consequently, it is more problematic to perform accurate numerical simulation of the advection of a fluid interface, for example, in the inviscid case. A second significant difference between the inviscid and viscous flows is the

nature of the stagnation points. In an inviscid flow, stationary points are always hyperbolic, because the streamfunction satisfies Laplace's equation ($\nabla^2\psi = 0$), and so cannot have internal extrema, whereas the viscous streamfunction satisfies the biharmonic equation ($\nabla^4\psi = 0$), which allows both elliptic and hyperbolic stationary points.

It seems clear from Figure 4.2 (and similar results that we have computed for other cylinder configurations, not presented here) that if the internal cylinders execute a given motion to stir the fluid then at any instant the streamlines will be quite different according to the inviscid and viscous models. This makes the results of [20], that a material stretch rate can be found, independent of the flow field, all the more remarkable.

4.2.4 Accuracy of the flow field

In Table 4.1 we illustrate the error in the computed velocity field due to truncation of the series (4.4), (4.8), (4.13) and (4.19) after n terms. Since the governing equations are, by construction, satisfied exactly, we record in the table the maximum error in the boundary conditions: (4.2) and (4.5) for inviscid flow, and (4.11) and (4.15) for viscous flow. We find that in all four cases the error decreases rapidly with n, with approximately one extra digit of precision gained for each extra term. We find that the largest error in the velocity field over the entire domain (compared with an 'exact' solution obtained using the 'effectively infinite' value $n \approx 15$) is always located on an internal cylinder boundary. In Figure 4.4 we plot the magnitude of the coefficients $b_{j,k}$ and $c_{j,k}$, from (4.19), against the index k for a typical BSD_3 solution with $n = 10$ (the coefficients are scaled by the cylinder velocity V). The exponential decay of the coefficients with k is indicative of spectral accuracy. For the purpose of numerical stirring simulations, we find that taking $n = 10$ terms reduces the error in the velocity field below that introduced by time discretisation (described later). The limit of machine precision in Matlab (16 decimal places) is typically reached with around $n \approx 15$ terms. By comparison, Vikhanksy [142] reports errors of $O(10^{-4})$ for his finite-difference immersed-boundary simulations, although his method allows

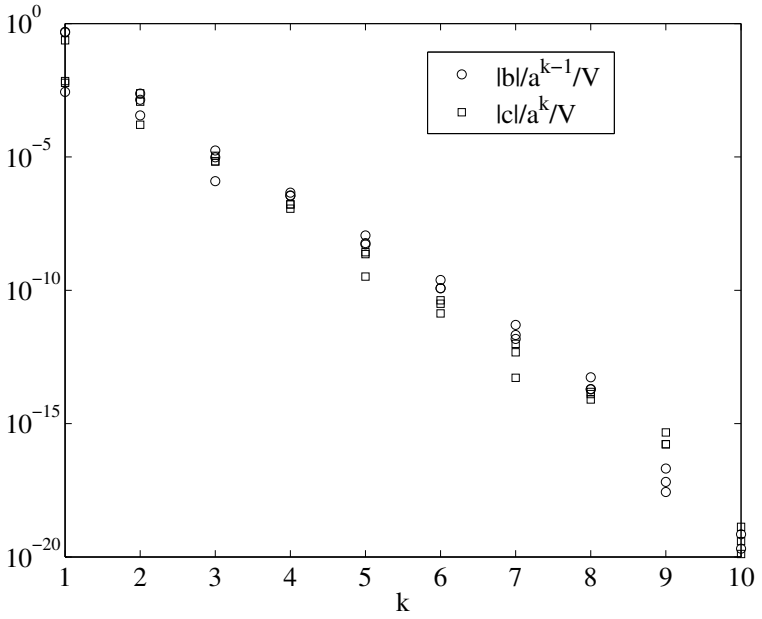

Figure 4.4: Exponential decay of the coefficients $b_{j,k}$ and $c_{j,k}$ with k. For definiteness, the values plotted correspond to the BSD_3 simulations described in Section 4.4 at the instant when the cylinders lie along a diameter, though the decay illustrated is typical for all times.

extension to finite Reynolds number, which ours does not.

It is clear from the table that, in our implementation, the time taken to compute the coefficients in the various series solutions increases rapidly with the number of terms taken. This is because for most of our computations, purely for analytical convenience, we used Matlab's *nonlinear* least-squares optimisation function `lsqnonlin`, whereas it is perfectly feasible to cast the problem in a more efficient manner using *linear* optimisation. Indeed, we have used the latter formulation to check our results from the former in some cases.

4.2.5 Energetics

An important consideration when designing a mixing device is the energy input that will be required to operate it. In this section we calculate the energy input required to operate the viscous BSD_m. (For our inviscid model, there is no drag on the cylinders, and consequently the energy input is zero.)

model	n	3	5	7	9
unbounded, inviscid	time	1.2	4.0	8.5	17
	error	4×10^{-3}	2×10^{-4}	1×10^{-5}	4×10^{-7}
bounded, inviscid	time	1.9	8.9	22	48
	error	4×10^{-3}	2×10^{-4}	9×10^{-6}	4×10^{-7}
unbounded, viscous	time	4.4	18	49	110
	error	3×10^{-2}	6×10^{-4}	2×10^{-5}	5×10^{-7}
bounded, viscous	time	11	33	84	170
	error	4×10^{-2}	1×10^{-3}	2×10^{-5}	9×10^{-7}

Table 4.1: Computation time (in seconds) and maximum error in boundary conditions, for the four BSD_m models on a 1.5GHz Pentium 4 computer. Using the $m = 5$ cylinder configuration from Figure 4.2, we compute the maximum error in the boundary conditions (4.2) and (4.5), or (4.11) and (4.15), scaled by the maximum cylinder speed, when truncating the appropriate series to n terms.

The complex-variable formulation used in Section 4.2 affords us a convenient method for calculating the instantaneous power input to the BSD_m, using the calculus of residues. As detailed by Finn and Cox [47], who have studied the single-cylinder BSD_1 (see Chapter 3), for a streamfunction in the form (4.9), the force $\mathcal{F} = \mathcal{F}^x + \mathrm{i}\mathcal{F}^y$ and moment $\mathcal{M}$ (about the origin $z = 0$) on any closed contour $\mathcal{C}$ in the flow domain are [99]

$$\mathcal{F} = 4\mu \left[-f(z) + z\overline{f'(z)} + \overline{g'(z)} \right]_{\mathcal{C}}, \quad \mathcal{M} = 4\mu \, \mathbb{I}\mathrm{m} \left[z\bar{z}\overline{f'(z)} + \bar{z}\overline{g'(z)} - \overline{g(z)} \right]_{\mathcal{C}}, \tag{4.20}$$

where $[h]_{\mathcal{C}}$ represents the change in the function h upon one anticlockwise tour of $\mathcal{C}$, and μ is the dynamic viscosity of the fluid.

By taking the contour $\mathcal{C}$ to be the boundary of the j-th cylinder, and noting that only the logarithmic terms in the streamfunction (4.17) make nonzero contributions to the expressions in (4.20), we calculate that the force $\mathcal{F}_j$ exerted by the j-th cylinder on the fluid is

$$\mathcal{F}_j = 4\pi\mu\mathrm{i}\, b_{j,1}, \tag{4.21}$$

and the moment $\mathcal{M}_j$ about the axis of the j-th cylinder exerted by that cylinder

on the fluid is

$$\mathcal{M}_j = -4\pi\mu\, c_{j,1} - \mathbb{I}\mathrm{m}(\mathcal{F}_j \bar{p}_j). \tag{4.22}$$

By considering the overall equilibrium of the BSD_m, or by using the outer boundary as the contour $\mathcal{C}$, we find that the corresponding force $\mathcal{F}_{\mathrm{out}}$ and moment $\mathcal{M}_{\mathrm{out}}$ about the origin exerted by the outer cylinder on the fluid are

$$\mathcal{F}_{\mathrm{out}} = -4\pi\mu \mathrm{i} \sum_{j=1}^{m} b_{j,1}, \quad \mathcal{M}_{\mathrm{out}} = 4\pi\mu \sum_{j=1}^{m} c_{j,1}. \tag{4.23}$$

The total power input $\mathcal{P}$ to the BSD_m is then

$$\mathcal{P} = \mathcal{M}_{\mathrm{out}}\Omega_{\mathrm{out}} + \sum_{j=1}^{m} \left(\mathcal{F}_j^x U_j + \mathcal{F}_j^y V_j + \mathcal{M}_j \Omega_j \right). \tag{4.24}$$

4.3 Theoretical background of topological chaos

In this section we begin by reviewing some basics of braid notation, which provides a framework for subsequent analysis of stretch rates for flows in the BSD_m with various topologies. We examine a number of different ways of moving the internal cylinders (these are termed the 'stirring protocols') and evaluate their effectiveness according to the interface stretch exponent (IS) measure defined in the previous chapter. We also consider the physically important issue of energy usage (EU) by the various protocols.

As we have discussed in the Chapter 3, the aim in previous studies of chaotic advection has often been to maximise measures such as material stretch rate by appropriate tuning of the system parameters. However, the theoretical results of Boyland *et al.* [20] allow calculation of a material stretch rate in a BSD_m flow generated by the motion of $m \geq 3$ cylinders (with particular emphasis on the case $m = 3$), *without explicit reference to the underlying equations governing the fluid motion.* While some tuning of the parameters may lead to improvements beyond this predicted stretching a certain minimal stretch rate is guaranteed. It should be emphasised that, at present, the great generality of the theory comes at a cost: the size of the region in which good mixing is achieved is unknown, and typically depends on the governing fluid equations; according to the theory, it may even have zero measure. It is our goal, therefore, to investigate the quality

of mixing due to various stirring protocols, according to the four models outlined above, with particular emphasis on the size of the chaotic regions.

4.3.1 Braid notation

We assume that the stirring motion is periodic in time, with period T. At the start of any period, at $t = pT$, and at the end, at $t = (p+1)T$, the axes of the m internal cylinders lie on the diameter $y = 0$ of the outer cylinder. During the course of the stirring, pairs of adjacent cylinders are interchanged. Boyland *et al.* [20] have described how such motion of m cylinders, when considered appropriately, has the topology of a physical braid on m strings, which in turn corresponds to a mathematical braid [140]. To see how the mathematical theory of braids may be applied to the mixing problem, we begin by assigning to each interchange a braid letter σ_j or σ_j^{-1}, as shown in Figure 4.5. The subscript j indicates which two (adjacent) cylinders are interchanged, from $j = 1$, which represents the left-most pair, to $j = m - 1$, which represents the right-most pair. A superscript of -1 indicates that the cylinders orbit one another in an anticlockwise sense; no superscript indicates a clockwise orbit.

The sequence of interchanges during one complete period is the stirring protocol, and can be characterised by a braid word, which is a string of braid letters. Interchanges are made corresponding to each braid letter in sequence, reading from right to left. For example, the braid word $\sigma_1\sigma_2^{-1}$ means that we first interchange the second and third cylinders in an anticlockwise sense, then interchange the first and second cylinders in a clockwise sense. A braid word may also be represented schematically by a braid diagram, as shown in Figure 4.5. Here each line represents one of the m internal cylinders and time progresses from the bottom to the top of the diagram; each crossing of the lines represents an interchange of neighbouring cylinders.

Once we have selected the stirring protocol, and, in particular, the exact paths that the cylinders take during their interchanges, we are in a position to simulate the corresponding fluid mixing, by tracking the advection of passive tracer particles in the flow. Since the stirring protocol is repeated periodically, we find

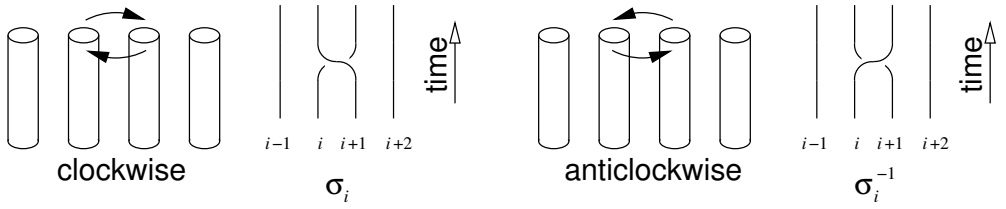

Figure 4.5: Braid diagrams illustrating the exchange in position of two adjacent cylinders. Neighbouring cylinders may interchange positions in either a clockwise or anticlockwise sense, characterised by braid letters σ_i or σ_i^{-1}, respectively. In the braid diagram corresponding to a clockwise exchange σ_i, the $i+1$-th strand crosses over the i-th strand, whereas for the braid letter σ_i^{-1}, the $i+1$-th strand crosses beneath the i-th strand.

it efficient to store the series coefficients for the velocity field at a large number of time steps during the period T, so that they can be computed just once, prior to the simulation of the particle advection. Other simplifications in the model allow us to reduce the computational and storage requirements further: if the cylinders are equal in radius and are placed symmetrically either side of the origin $z = 0$, and if 'reasonable' interchanges are used, then reflectional and time-reversal symmetries may be invoked, so that with m cylinders, only $\lceil \frac{1}{2}(m-1) \rceil$ 'half-interchanges' need be calculated (where $\lceil \xi \rceil$ means the least integer not less than ξ).

4.3.2 Prediction of stretch rate

Using Thurston–Nielson theory, Boyland *et al.* [20] have shown that the action of each braid letter in a stirring protocol is intimately linked with a matrix derived from its braid (Burau) matrix. In particular, for $m = 3$, a certain stretch rate given by the stirring action of a particular braid letter is predicted by the spectral radius (magnitude of the largest eigenvalue) of its associated matrix. Boyland *et al.* [20] explain how the braid matrix is constructed from its corresponding braid letter. When $m = 3$, the matrices corresponding to the braid letters σ_1, σ_1^{-1}, σ_2,

σ_2^{-1} are, respectively,

$$s_1 = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}, \quad s_1^{-1} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad s_2 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \quad s_2^{-1} = \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix}. \tag{4.25}$$

The matrix corresponding to the action of a braid word may be found be taking the matrix product of the component braid matrices in the same order as the braid letters appear in the word. For example, the protocol $\sigma_1\sigma_1\sigma_2^{-1}\sigma_2^{-1}$ has the braid matrix

$$s_1 s_1 s_2^{-1} s_2^{-1} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} 5 & -2 \\ -2 & 1 \end{bmatrix}, \tag{4.26}$$

which has spectral radius $\lambda_{\text{theory}} = 3 + 2\sqrt{2}$. If the BSD_3 is operated according to this protocol, there is at least one material line whose length $l(t)$ grows at least as rapidly as

$$l(t)/l(0) \sim \lambda_{\text{theory}}^{t/T}, \tag{4.27}$$

where T is the period. According to the theory, only the flow topology matters, and this is completely characterised by the braid word. The nature of the flow, for example, whether it is inviscid or viscous, incompressible or compressible, Newtonian or non-Newtonian, does not matter: while the stretch rate λ_{theory} is the same in each case, the size and location of the appropriate region in which the stretch rate is achieved depends on the details of the flow.

Using the models from Section 4.2, we are able to test the predictions of Boyland *et al.* [20] in numerical simulations of both inviscid and viscous models for the BSD_3. For the case $m = 3$, we have chosen five different test protocols, labelled A–E, as shown in Figure 4.6. These all have braid words containing four letters, but different theoretical stretching rates, as given in Table 4.2. Protocol A is special because its braid is topologically equivalent to the 'identity' braid, for which the crossings can be untied [20]. In Figure 4.7 we illustrate that three iterations of protocol A mimics the action of rotating the tank of fluid by two complete revolutions; correspondingly the third power of the braid matrix for protocol A is the identity matrix. Although we expect protocol A might stir effectively in practice, particularly if its physical parameters are chosen appropriately, its braid

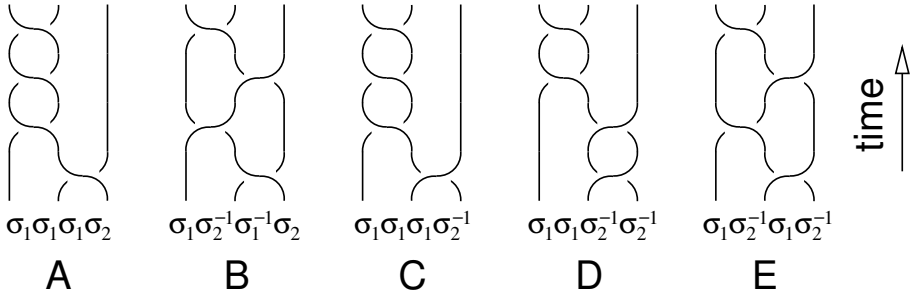

Figure 4.6: Braid diagrams and braid words for the five test protocols. Protocol A is 'finite-order' and does not have the correct topology to ensure effective stirring. Protocols B–E are 'pseudo-Anosov' and each have a (distinct) stretch rate; these protocols illustrate all possible theoretical stretch-rates corresponding to braid words of four letters.

Protocol	Braid word	λ_{theory}	$\lambda_{\text{numerical}}$
A	$\sigma_1\sigma_1\sigma_1\sigma_2$	—	2.2
B	$\sigma_1\sigma_2^{-1}\sigma_1^{-1}\sigma_2$	$\frac{1}{2}\left(3+\sqrt{5}\right) \approx 2.6$	2.8
C	$\sigma_1\sigma_1\sigma_1\sigma_2^{-1}$	$\frac{1}{2}\left(5+\sqrt{21}\right) \approx 4.8$	5.0
D	$\sigma_1\sigma_1\sigma_2^{-1}\sigma_2^{-1}$	$3+2\sqrt{2} \approx 5.8$	5.8
E	$\sigma_1\sigma_2^{-1}\sigma_1\sigma_2^{-1}$	$\frac{1}{2}\left(7+3\sqrt{5}\right) \approx 6.9$	7.1

Table 4.2: Table of stretch rates for the five test protocols A–E. Indicated are the theoretical values λ_{theory}, together with numerical results based on our bounded, viscous model for the BSD_3.

matrix has complex eigenvalues, and so the topological theory does not guarantee exponential stretching of a material interface. Protocols B–E do, however, have nontrivial topology; in fact, these protocols correspond to the four possible distinct material stretch rates that can be generated by four-letter braid words.

4.4 Numerical simulations of dye advection

We have carried out numerical simulations of passive dye advection for the five test protocols A–E. Initially the axes of the cylinders are equi-spaced along the line $y = 0$; the cylinders are taken to be of equal radius, symmetrically placed

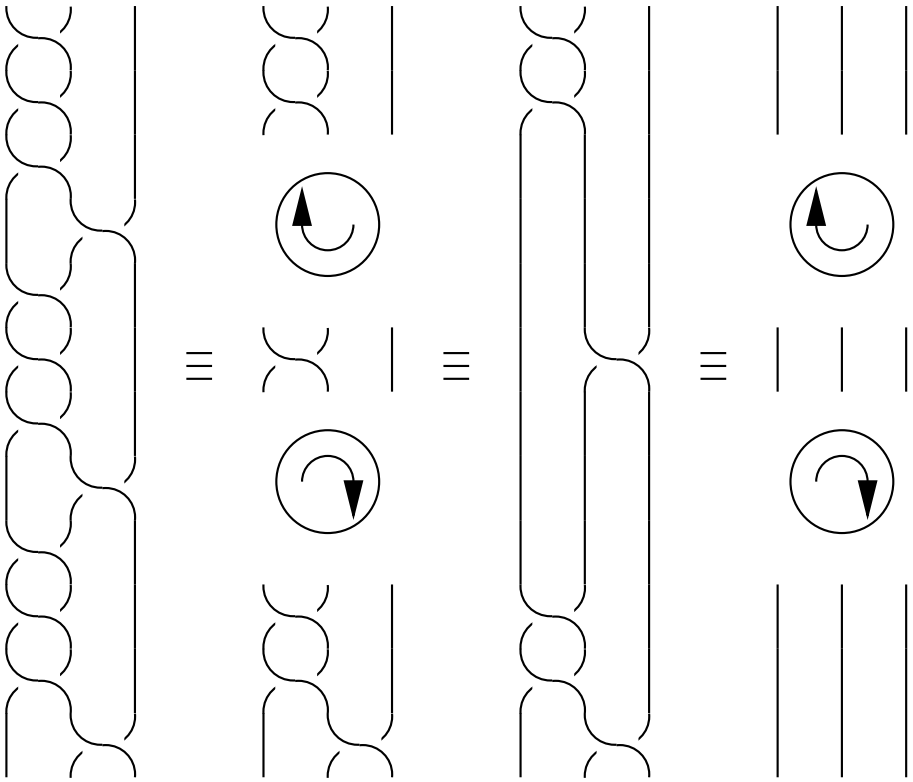

Figure 4.7: Diagrams showing how three applications of protocol A are topologically equivalent to two complete rotations of the tank. On the left is shown the braid diagram for three iterations of protocol A. In the neighbouring diagram, we have identified two instances of $\sigma_1\sigma_2\sigma_1$, corresponding to reversing the order of the stirring rods – a move performed by rotating the tank clockwise by 180°. The topological action of the rotations and the single braid letter σ_1 between them can be accomplished, without any tank rotations, by the braid letter σ_2, as shown in the third diagram. What is left can now be identified as topologically equivalent to two further clockwise 180° rotations, shown on the right.

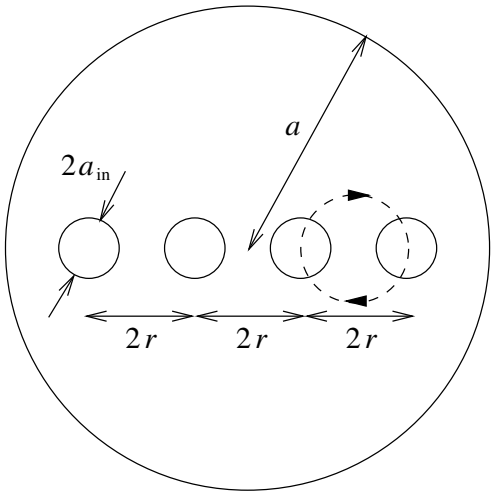

Figure 4.8: The BSD_m geometry at the start of a period, illustrated for $m = 4$. The inner cylinders all have radius a_{in}, and are equally spaced along the line $y = 0$, configured symmetrically. The distance between neighbouring cylinder axes is $2r$. When two adjacent cylinders are interchanged, they are moved along circular arcs of radius r centred on the midpoint between the two cylinders, either clockwise or anticlockwise.

about the origin, with the axes of neighbouring cylinders separated by a distance $2r$, as shown in Figure 4.8. We have chosen to investigate the parameter regime $a_{\mathrm{in}}/a = 0.1$ (where a_{in} is the radius of each cylinder) and $r/a = 0.25$. A streak of dye is initially placed along the perimeter of a rectangle joining the four points with coordinates $(\pm\frac{1}{4}a, \pm\frac{1}{2}a)$, and is then advected by the fluid.

Although the details of the cylinder motion during the interchanges between neighbours are unimportant for the theoretical prediction of the stretch rate, in order to perform our stirring simulations we need to specify the exact cylinder motions. When two adjacent cylinders are interchanged, we move them along circular arcs of radius r centred on the midpoint between their axes, so that the two cylinders remain at all times on opposite ends of a diameter. During this procedure, the cylinders themselves do not rotate about their axes.

Results for the bounded inviscid and viscous models of the BSD_3 are shown in Figures 4.9 and 4.10, respectively. One difference that is immediately apparent between the two sets of results is that the inviscid flow generates structures that

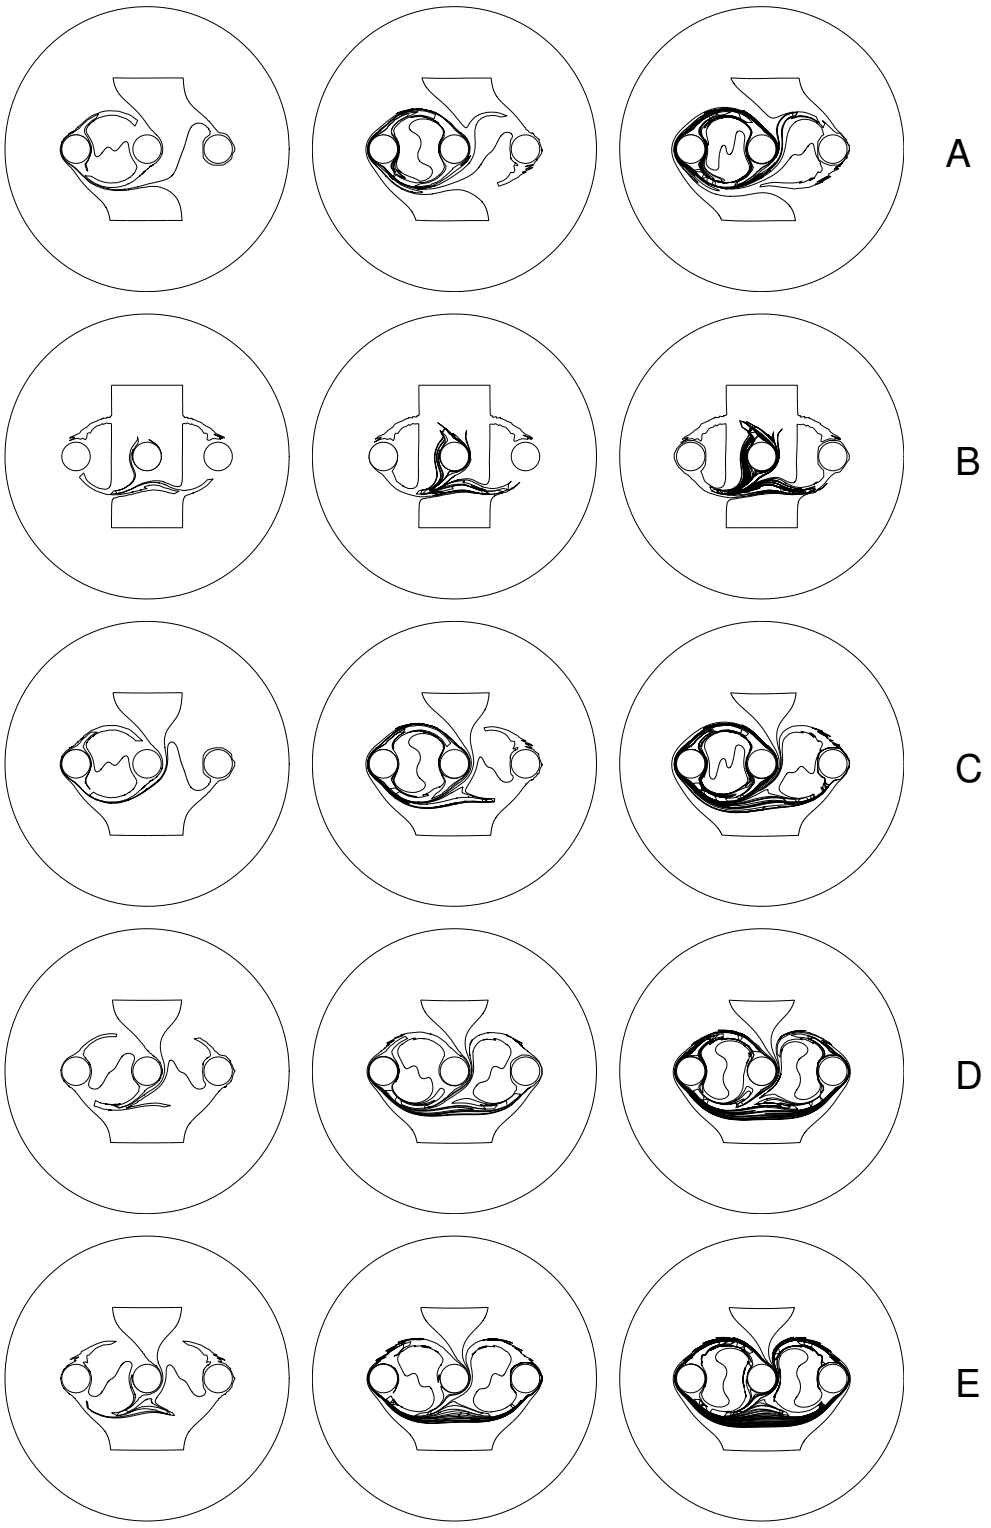

Figure 4.9: Numerical dye-advection simulations of the five test protocols A–E in the inviscid model for the BSD_3. The protocols are labelled by row. The columns show the dye location after (from left to right) 1, 2 and 3 periods of the flow.

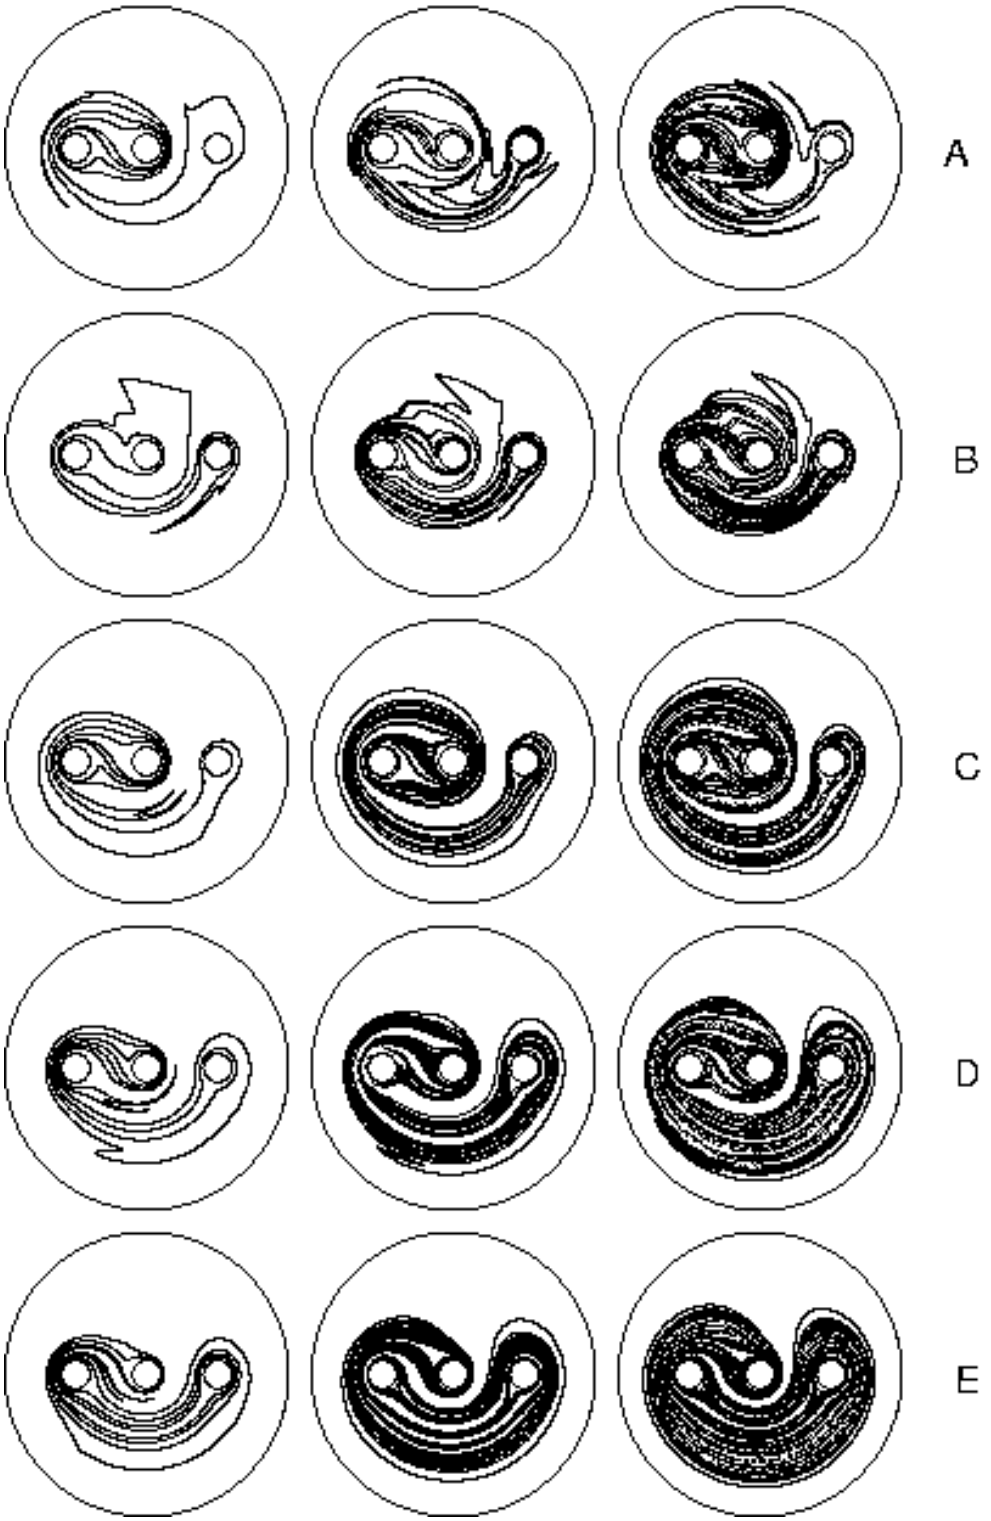

Figure 4.10: Numerical dye-advection simulations of the five test protocols A–E in the viscous model for the BSD_3. The protocols are labelled by row. The columns show the dye location after (from left to right) 1, 2 and 3 periods of the flow.

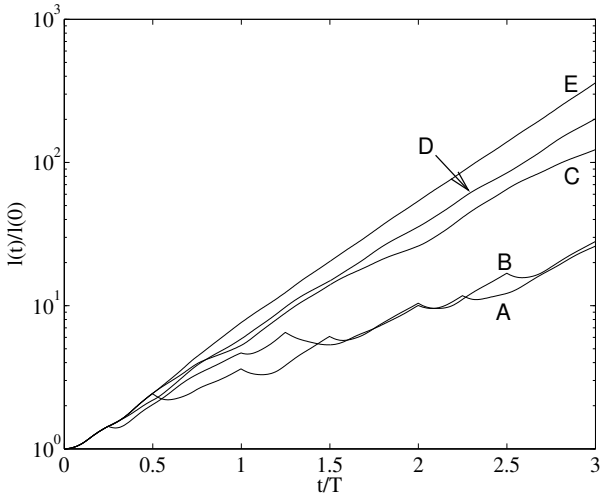

Figure 4.11: Dye-streak length $l(t)$ for the five test protocols A–E in the bounded viscous, model for the BSD_3.

are less smooth than in the viscous case. This feature is a consequence of the greater localisation of the inviscid flow, as described in Section 4.2.3. In each case, the region of good mixing is commensurate with the region through which the inner cylinders are moved (cf. [20]).

Another observation we make for the viscous case (Figure 4.10) is that after three periods under the actions of protocols C–E the dye has generally moved into the lower half of the domain. This is because the component braid letters σ_1 and σ_2^{-1} which comprise these three braid words cause fluid in the centre of the mixer to move downwards, and fluid on the left and right of the domain to move upwards. After many more periods, the dye in the lower half of the domain would precess around the edge of the mixer thus occupying both the upper and lower halves.

The time-variation in the length of the dye-streak is plotted in Figure 4.11, according to the bounded, viscous model for the BSD_3. From these results we have computed the corresponding stretch rate for each protocol, our results being shown in Table 4.2. There is close agreement between these and the theoretical stretch rates. In tracking the dye-streak, we have used the dynamic particle insertion algorithm of [83] to ensure that the exponentially growing curve remains well resolved. Of course, it should be noted that the initial dye-streak does not seem to lie entirely in the chaotic region. Therefore some section of it is likely

to stretch at a rate less than λ_{theory}. However, the influence of any such section on the numerical estimate of the stretch rate diminishes with time (since it occupies an increasingly small fraction of the total length of the dye-streak). Thus our results confirm the predictions of Boyland *et al.* [20], and indicate that the predicted stretch rates seem quite tight, at least for these protocols.

Corresponding calculations of the stretch rate for the inviscid model are not presented, because we have found it difficult to compute with accuracy the interface length in this case, due to the highly localised stretching. However, although we have less confidence in the inviscid results, they do yield stretch rates that are also consistent with those given in Table 4.2.

4.4.1 Effect of stirrer radius

We recall that one feature of the topological theory is that exact details of the mixer geometry have no influence on the predicted stretch rates. To illustrate this we plot in Figures 4.12 and 4.13 the results corresponding to the viscous BSD_3 simulations in Figures 4.10 and 4.11, respectively, but performed using the much smaller inner cylinder radii, with $a_{\text{in}}/a = 0.01$. We find that the stretch rates associated with Figure 4.13 are still in accordance with those in Table 4.2; however, by inspecting Figure 4.12, we see that the region of high stretch is more localised around paths along which the cylinders are moved than with the larger cylinders in Figure 4.10. Hence although the predicted stretch rates are maintained we find that mixing quality according to other measures (e.g. iterated mapping chaoticity) would be reduced. The dependence of stretch rate upon stirrer thickness is discussed further in Chapter 5, where we explore topological chaos in three dimensions.

4.4.2 Comparison with experiments of Boyland *et al.*

Finally, we present in Figure 4.14 numerical simulations of the bounded, viscous model corresponding to the experiments reported by Boyland *et al.* [20]. These results may be compared with their Figure 2, which shows the result of stirring

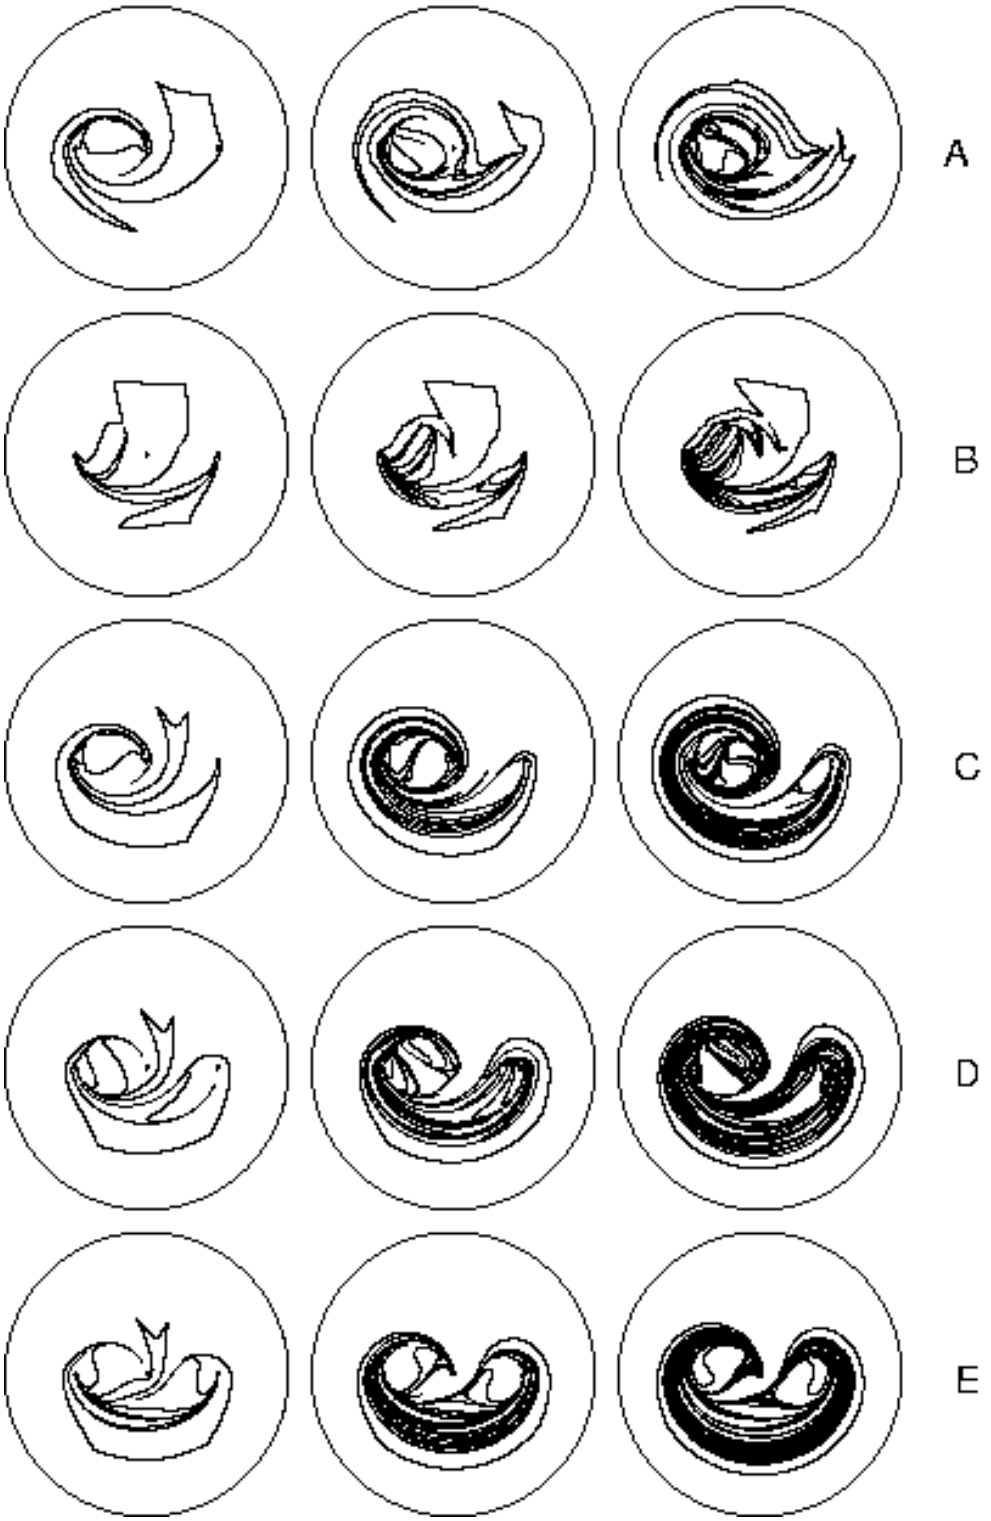

Figure 4.12: Numerical dye-advection simulations of the five test protocols A–E in the viscous model for the BSD_3 as in Figure 4.10, but with smaller inner cylinders of radii $a_{in}/a = 0.01$. The protocols are labelled by row. The columns show the dye location after (from left to right) 1, 2 and 3 periods of the flow.

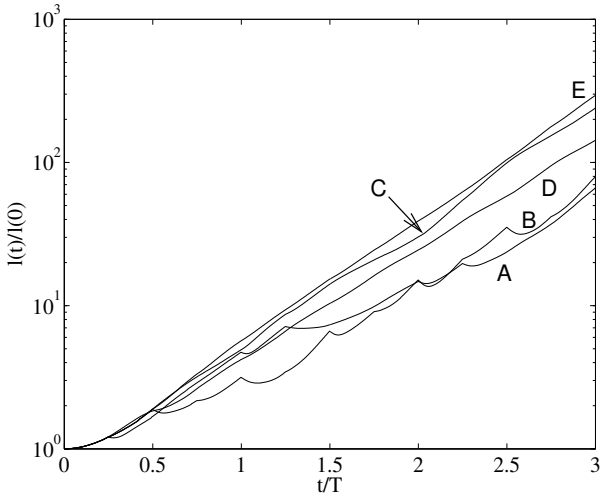

Figure 4.13: Dye-streak length $l(t)$ for the five test protocols A–E in the bounded viscous, model for the BSD_3, as in Figure 4.11, but with smaller inner cylinders of radii $a_{in}/a = 0.01$.

by motions $(\sigma_1\sigma_2)^p$ and $(\sigma_1^{-1}\sigma_2)^p$ in our notation (respectively, f^p and g^p in their notation), for $p = 0, 1, 2, 3, 4, 5, 6, 9$. The agreement is excellent; our plots show qualitatively the same features as the experimental pictures presented in [20] (the differences presumably being due to slight mis-matches between the cylinder radii and the initial placement of the dyed fluid). The bounded, viscous model thus seems to provide an excellent numerical test-bed in which to investigate topologically chaotic flows further.

4.5 Experimental results

Further validation of our bounded, viscous model (discussed in Section 4.2) may be obtained by comparing numerical simulations with our own experimental dye-advection results, for the case $m = 4$.

We have built a simple experimental BSD_4 rig. The outer cylinder has radius $a = 7.4$cm and the inner cylinders have radius $a_{\text{in}} = 0.7$cm; the separation between the centres of the inner cylinders is $2r = 3.3$cm. These cylinders are interchanged manually, by means of a braiding fork, which may be lowered into

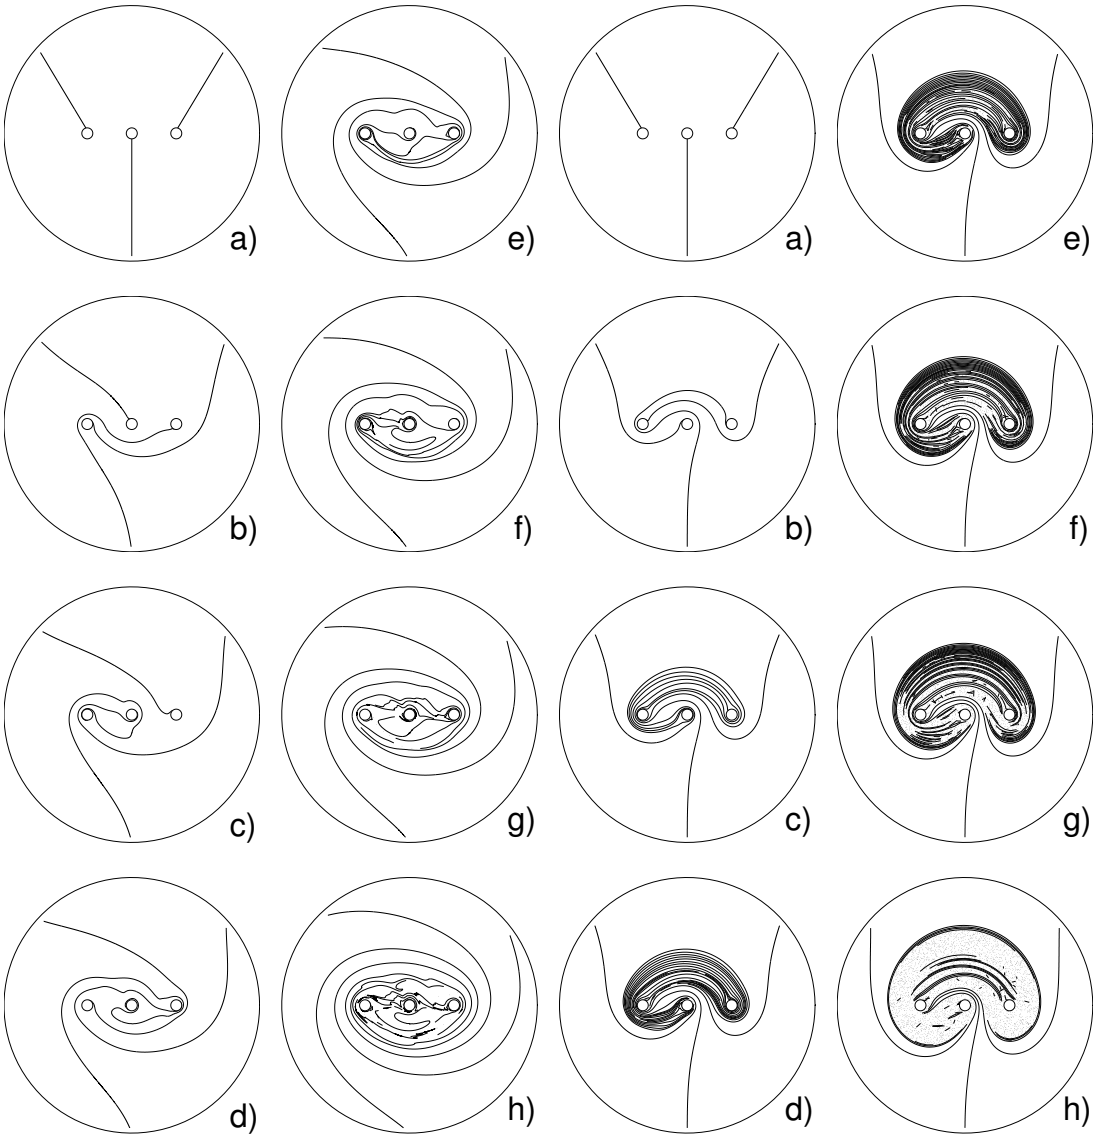

Figure 4.14: Numerical dye-advection simulations in the viscous model for the BSD_3 under motions corresponding to $(\sigma_1\sigma_2)^p$ (left) and $(\sigma_1^{-1}\sigma_2)^p$ (right). Here $a_{\text{in}}/a = 0.043$ and $r/a = 0.17$, roughly in accordance with the experimental conditions of Boyland *et al.* [20]. Configuration at (a) $p = 0$; (b) $p = 1$; (c) $p = 2$; (d) $p = 3$; (e) $p = 4$; (f) $p = 5$; (g) $p = 6$, and (h) $p = 9$.

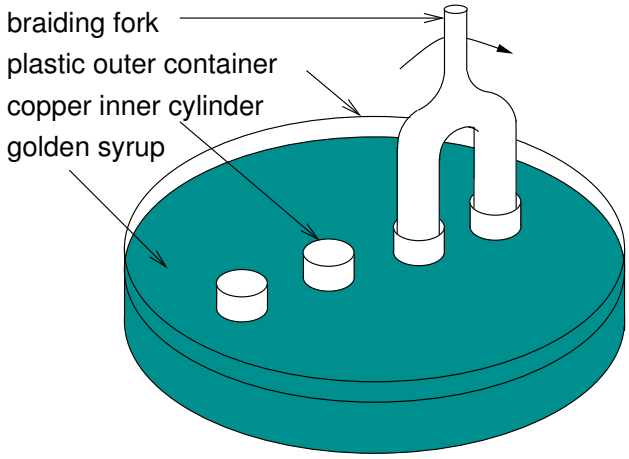

Figure 4.15: Schematic of our experimental apparatus, illustrating how adjacent cylinders are interchanged by inserting and rotating a braiding fork, which causes the cylinders to move along circular arcs centred on the mid-point between the cylinders.

any two adjacent cylinders and then rotated (see Figure 4.15). We note that, since the braiding fork does not apply a moment to the cylinders as it moves them, they rotate through some small amount about their own axes during an interchange; we neglect this (small) effect in our numerical simulations, as it is beyond our control in such a crude experiment.

As our working fluid we use golden syrup, which has a kinematic viscosity at room temperature of approximately $\nu = 1200\text{cm}^2\text{s}^{-1}$ [2]. For each interchange, we rotate the braiding fork slowly, completing a half-circle in a time of around $T_{\text{swap}} = 10\text{s}$. Using r as a characteristic length, and r/T_{swap} as a characteristic velocity, this gives a flow Reynolds number of the order $r^2/\nu T_{\text{swap}} \sim 10^{-4} \ll 1$, so we are indeed in the Stokes flow regime.

Figure 4.16 shows a comparison between numerical and experimental dye advection simulations using three different protocols. Although the experiments are admittedly crude, particularly in the initial placement of the dye, the intricacies of the striation patterns are reproduced remarkably well. In each case, the region of good mixing is again roughly the same size as that through which the inner cylinders are moved, consistent with the observations of [20].

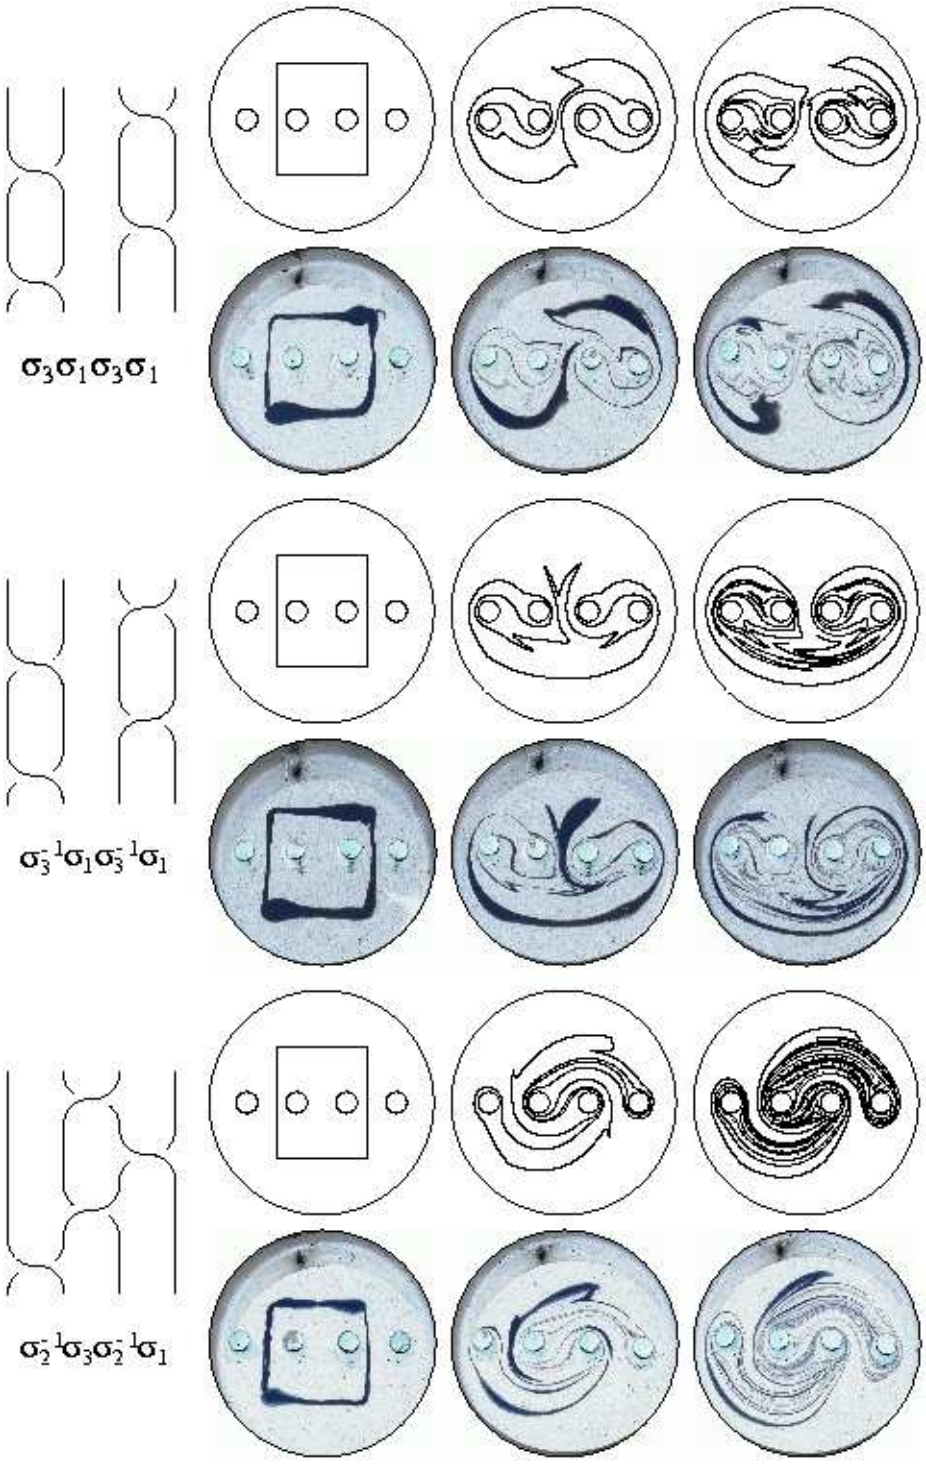

Figure 4.16: Left: braid diagrams for three protocols with $m = 4$ cylinders. Right: corresponding numerical simulations (top row) and experiments (bottom). The first column of results gives the initial location of the dye-streak. The next two columns show the position of the dye-streak after one and two periods of the flow, respectively.

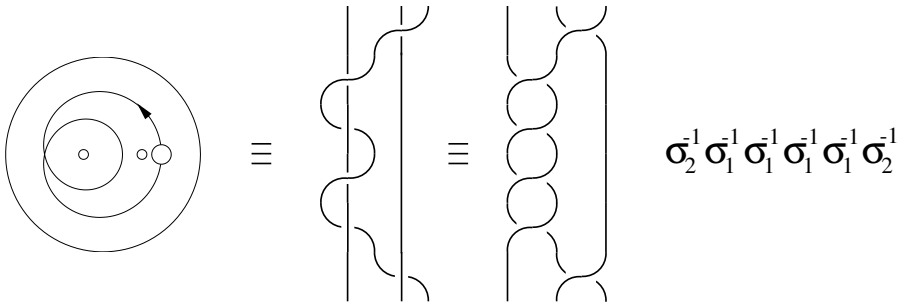

Figure 4.17: A BSD_3 protocol created by inserting two stationary cylinders into the BSD_1 [47].

4.5.1 Practicalities of mixer construction

Having described some simple small-scale experiments, we now discuss some of the practicalities of building a mechanical BSD_m. In designing such a device, one would need to decide between a gearing system or some kind of robotics for moving the cylinders. While the latter would certainly be more flexible, and allow the protocol to be changed easily, it is feasible, and simpler, to achieve effective stirring topologies using systems of gearing. For instance, in the three-cylinder BSD_3 discussed in Section 4.3.2, only one cylinder need be moved to create a non-trivial flow topology – the two remaining cylinders can be held stationary, provided that the moving cylinder loops around them appropriately. This significantly simplifies construction of the device, but it is not clear for general values of m whether the flow topology with the greatest $\lambda_{\mathrm{theory}}$ can always be achieved with $m-1$ stationary cylinders and one cylinder moving by means of simple gearing.

In Figure 4.17 we show a cycloidal path for one of the cylinders (radius $a_{\mathrm{in}} = 0.1a$), of the form

$$p_1/a = -0.1 + 0.2\exp(2\pi \mathrm{i}t/T) + 0.5\exp(4\pi \mathrm{i}t/T) \tag{4.28}$$

(recall that $z = p_1(t)$ is the location of the axis of the cylinder at time t), which can be generated using a simple gearing arrangement [47]. These parameters mimic the $\mathcal{P}(2,2,7)$ protocol from Chapter 3 shown in Figure 3.2. Corresponding simulations of mixing in this BSD_1 are given below. By placing two further stationary cylinders (each of radius $a_{\mathrm{in}} = 0.05a$) with axes fixed at $p_2 = -0.2a$ and $p_3 = 0.4a$,

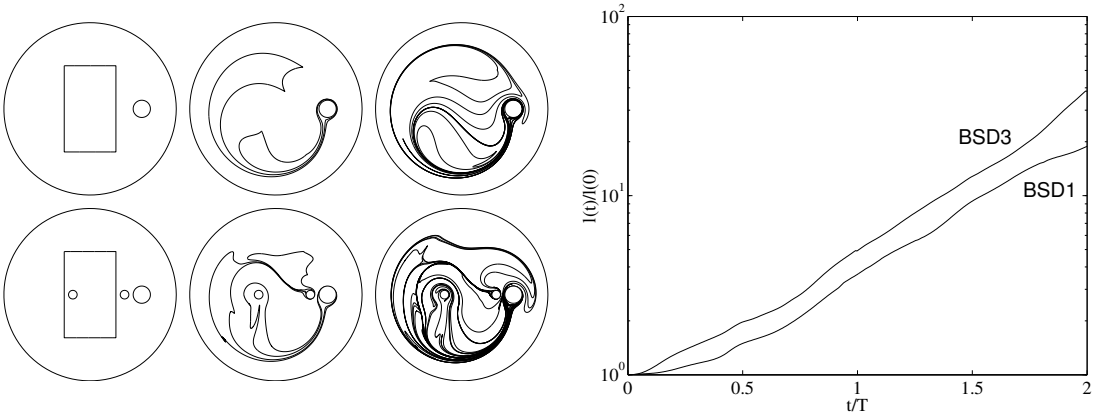

Figure 4.18: Stretch-rate comparison between the BSD_1 and the BSD_3. On the left are numerical dye-advection plots for the BSD_1 (top row) and BSD_3 (bottom row) after 0, 1 and 2 periods. Apart from the presence of two additional stationary cylinders in the BSD_3, the two stirring motions are otherwise identical. On the right are plotted the corresponding dye-streak lengths $l(t)$ (cf. Figure 4.11).

we obtain a BSD_3 flow with the topology of the braid $\sigma_2^{-1}\sigma_1^{-1}\sigma_1^{-1}\sigma_1^{-1}\sigma_1^{-1}\sigma_2^{-1}$ (see Figure 4.17). According to the theory outlined in Section 4.3.2, this protocol generates a stretch rate of $\lambda_{\text{theory}} = 3 + 2\sqrt{2} \approx 5.8$, whereas in the absence of the two stationary cylinders, the theory offers no corresponding value for the BSD_1. One might therefore expect material stretch rates to be significantly greater in the BSD_3 than in the BSD_1: Figure 4.18 shows a stretch-rate comparison between the BSD_1 and BSD_3 protocols. We find that λ_{BSD_3} slightly exceeds λ_{BSD_1} but, by calculating the energy input, we find, not surprisingly, that it is more expensive to operate the mixer with the additional cylinders in place (approximately 400% more energy is required for the BSD_3 than the BSD_1). The quantitative details of the improvement in stretch rate and the excess energy consumption depend sensitively on the size and location of the additional cylinders. However computing the most energy efficient mixer configurations (especially when involving more stationary cylinders) is a difficult problem that is beyond the scope of the present study.

4.6 Discussion

We have confirmed through numerical simulations that the theory of braids appears to provide a fairly tight prediction for the stretch rate in various BSD_3 protocols, according to both inviscid and viscous models for the flow field. As Boyland *et al.* [20] point out, the theory does not predict the size of the region in which good stretching is achieved: from our numerical simulations and experiments, we confirm their observation that this region seems commensurate with that through which the cylinders are moved during their interchanges.

Of course, as we have discussed in Chapter 3, the stretch rate of a finite material line element is not the only measure of mixing quality, and so one might reasonably ask whether the results described above provide any insight into the effectiveness of the various protocols according to other mixing measures.

We recall that in Chapter 3 we found a strong correlation between material stretch rate and energy usage ($|C| = 0.85$) for the BSD_1. However, it is unlikely that such a strong correlation would be observed for the BSD_3, with three stirring rods, because, by the symmetry of the motions described in Section 4.3.2, we find that the same energy is required to execute any of the braid letters σ_1, σ_1^{-1}, σ_2, σ_2^{-1}, and hence to execute any of the protocols A–E. Despite this equality of energy usage, these protocols exhibit very different stretch rates — the correlation between interface stretch and energy usage is broken by the flow topology.

Hence, if an empirical correlation exists between material stretch rate λ_{is} and other mixing measures, it necessarily depends on flow topology. The problem of finding such a relationship is important since the (apparently quite tight) estimates on the stretch rate are independent of the exact fluid properties and flow field and so a strong link to another mixing measure would allow similarly tight predictions for that measure too, based purely on topological considerations. Unfortunately the additional computational resources required for numerical simulations of the BSD_m prohibit as thorough an investigation as was undertaken with the BSD_1 in Chapter 3 and this must be left for a future project.

We now make some further considerations regarding the energetics of the BSD_m.

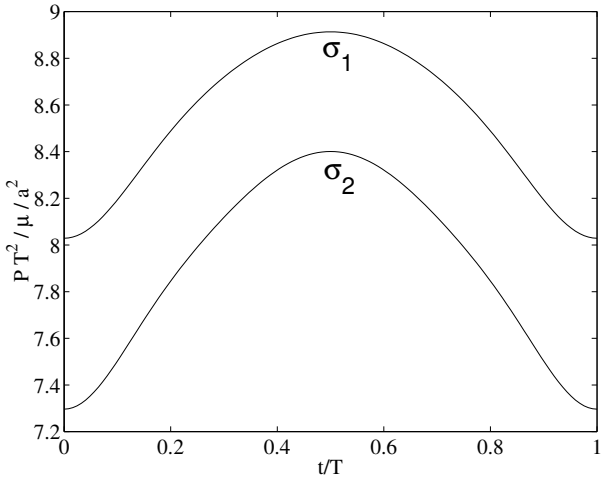

Figure 4.19: Power input calculated from (4.24) for the cylinder exchanges σ_1 and σ_2 in the four-cylinder BSD_4 (here the time T is that required to perform a single exchange). A greater power input is required for σ_1 than for σ_2, since in the former case the relevant cylinders are closer to the outer boundary.

Although we have focused on computing the energy input required only for the protocols investigated in this chapter, three broader issues now arise, which would influence a more ambitious project in which one were to attempt to find some optimal stirring protocol (using least energy, say).

First, we note that once we have selected a protocol with a given topology and have to engineer a given material stretch rate, it remains to choose such parameters as a_{in} and r, for example. A simple thought experiment shows that these can significantly affect the energy requirements of the BSD_m, for if the internal cylinders lie close together then the power input is large (and is infinite when two moving cylinders touch) and the region of good mixing is presumably correspondingly small. Alternatively, when the internal cylinders lie far apart, their proximity to the outer boundary becomes significant, and once again the power input becomes large (infinite if any moving internal cylinder touches the outer boundary). We anticipate that for a given set of internal cylinders there exists some optimum way to execute a given braid motion, for minimum energy input.

Secondly we note that for the BSD_m with $m > 3$, cylinder exchanges taking place closer to the outer boundary (e.g. σ_1 or σ_{m-1}) require more energy than those

further away, so that the choice of braid word itself can influence the energy usage. As an example, we demonstrate in Figure 4.19 the different power requirements during the execution of the braid letters σ_1 and σ_2 with $m = 4$ cylinders.

The third consideration, alluded to in Section 4.2.2, is that in addition to translating each of the internal cylinders, we may also choose to rotate them about their axes (or, indeed, the external cylinder about its axis), in an attempt to reduce the power input during the cylinder exchanges.

Finally, we note that the series solution approach implemented here is readily adapted for stirring elements whose cross-section is not circular. Spectral accuracy in the series for the velocity field can be maintained by applying an appropriate local conformal mapping from each non-circular cross-section to a circle (Prof. L. N. Trefethen, personal communication).

To be precise, if $w_j(z)$ maps the boundary of the stirring element j to the unit circle $|w| = 1$, then the streamfunction (4.17) is recast into the form

$$\begin{aligned} \psi &= \bar{z} \sum_{j=1}^{m} \left\{ \tfrac{1}{2} b_{j,1} \log |w_j(z)| + \sum_{k=2}^{n} b_{j,k} (w_j(z))^{1-k} \right\} \\ &\quad + \sum_{j=1}^{m} \left\{ \tfrac{1}{2} c_{j,1} \log |w_j(z)| + \sum_{k=2}^{n} c_{j,k} (w_j(z))^{1-k} \right\} \\ &\quad + \text{image system} + c.c., \end{aligned} \tag{4.29}$$

where 'image system' represents the images of the singularities inside the circle $|z| = a$ according to the biharmonic circle theorem [14, 47, 54, 97, 132]. The expression (4.29) is of the form (4.9), so this expression satisfies the governing biharmonic equation exactly.

We have in this manner carried out preliminary simulations for stirring elements of both smooth (elliptical) and non-smooth (flat, plate-like) cross-section (cf. [128, 129] for the former). A streamline plot with two moving plates and one stationary plate is shown in Figure 4.20. The mappings used in this case are Joukowski transformations [2] of the form

$$w(z) = \zeta + \sqrt{\zeta^2 - 1}, \quad \zeta = \mathrm{e}^{-\mathrm{i}\theta}(z - p)/l, \tag{4.30}$$

which take plates of length $2l$, centred at $z = p$, and inclined at an angle θ to the

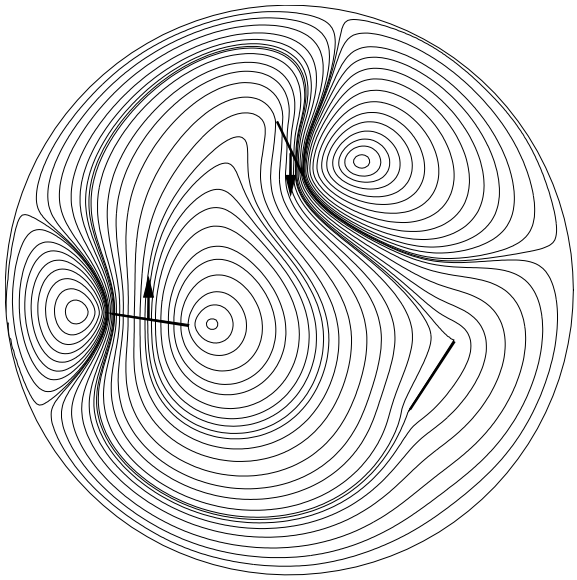

Figure 4.20: Streamline pattern for two moving plates and one stationary plate.

circle $|w| = 1$. Special care must be taken to ensure that appropriate branch cuts are used when transforming between the z and w domains.

In principle, more complicated maps could be used, such as Schwarz–Christoffel transformations, which would allow stirring elements with an arbitrary polygonal cross-section [41].

4.7 Conclusions

We have examined the generation of 'topological chaos' in a batch stirring device (BSD_m) [20], through numerics and experiment. The device consists of a number of stirring rods whose positions are successively interchanged in a 'plaiting' motion, the topology of which is alone sufficient to determine a material stretch rate that can be achieved in a corresponding two-dimensional chaotic fluid flow. This bound is independent of the fluid properties, and of the exact cylinder motions used to achieve the interchanges.

In any practical implementation of the BSD_m, however, the theory does not indicate the actual stretch rate that can be achieved, nor does it indicate the extent of the flow domain in which the greatest stretching takes place. To address such

issues, we have proposed four models for the fluid velocity field in the BSD$_m$, with particular emphasis on the most sophisticated, in which the BSD$_m$ is described by Stokes flow in a finite domain. Through a complex series for the instantaneous streamfunction, we have been able numerically to simulate stirring protocols with various topologies. Our algorithm is simpler, faster, more accurate, and requires less storage than an alternative finite-element or finite-difference approach [142]. Furthermore, it allows accurate evaluation of the velocity field at any point in the domain, without the necessity of interpolation between values known at relatively few grid points. We have applied our algorithm to perform dye-advection and stretch-rate calculations, and our results agree well with the theoretical predictions given by Boyland *et al.* [20], in particular, concerning the size of the chaotic region. A validation of our numerical code against experiments performed in a simple, inexpensive rig showed excellent agreement with the bounded, viscous model. Although our experiments involved a crude manual interchange of the stirring rods, we have discussed how one might construct an automated BSD$_m$. Many effective stirring protocols can be realised with a single stirring rod that loops around other rods, which are fixed, and this kind of motion can be achieved using only simple gearing.

As yet, there is no corresponding theory to predict a priori how a given protocol will score against a mixing measure other than its rate of material line stretch. We have already discussed in Chapter 3 the important question of whether, in general, there is a quantifiable link between stretch rates and other measures. It is not clear whether topological chaos could play a key role in establishing this link; however, the ability of braid theory to bound stretch rates allows us to some extent to circumvent expensive computer simulations; this is an important step, and these theoretical results may well provide the foundation on which a future study involving a wider range of other mixing measures can be based.

Chapter 5

Topological chaos in three-dimensional mixers

5.1 Introduction

We have seen in Chapter 4 that topological chaos (TC) theory may be used to predict fluid stretch rates [18, 20, 123, 142] in two-dimensional batch mixers. Although a practical disadvantage of the TC theory is that it gives no indication of the area in which high stretch rate is achieved, we have seen that the regions of effective mixing are commensurate with the region through which the stirring elements move [20]. Furthermore, even thin stirring elements appear to mix effectively, provided they move with the topology of a nontrivial mathematical braid.

In this chapter we broaden our exploration of mixer design to three-dimensional mixers. Chaotic advection is readily achieved in three-dimensional mixing devices and *can* lead to effective stirring [9, 10, 11, 111, 112] provided flow parameters are tuned carefully. Indeed, in three dimensions no time-dependence of the Eulerian velocity field is required in order to create chaotic particle trajectories [72], and so we examine in this chapter three-dimensional *static* mixers, which contain no moving parts. Generally, a static mixer consists of a pipe containing a specially designed insert that causes transverse mixing of the passing fluid flow, as

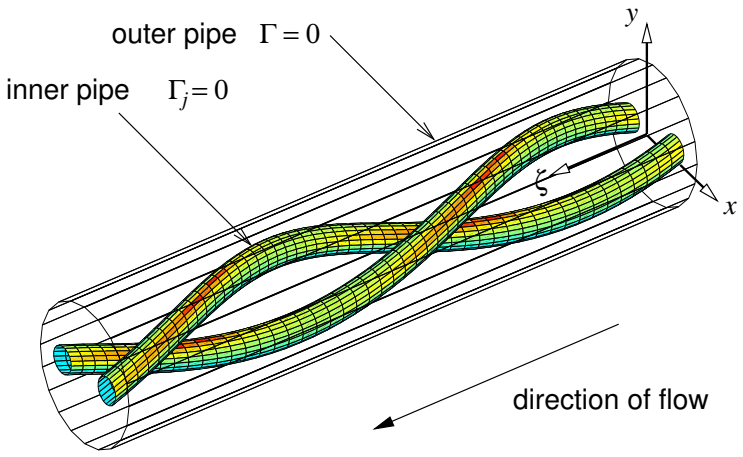

Figure 5.1: Geometry of the Braided Pipe Mixer.

illustrated in Figure 1.5. This mixing occurs at a cost, however, since a greater pressure gradient is required to drive the flow than when the insert is absent.

We investigate the possibility, suggested by Boyland, Aref and Stremler [20] and Mackay [92], of applying topological chaos to the design of a three-dimensional static mixer. Specifically, we consider a 'braided pipe mixer' (BPM) in which fluid is driven by a pressure gradient down a circular cylindrical pipe that contains several intertwining, solid pipes. As fluid flows past the internal pipes, it is forced to mix in the cross-section (see Figure 5.1). An analogy between the BPM and the two-dimensional BSD_m may be made, whereby the axial coordinate in the BPM plays the role that time occupies in the BSD_m. One might then expect that mixing in the BPM could be improved if the internal pipes form an appropriate mathematical braid. The correspondence between the BPM and BSD_m is not exact however, since particles in the BPM travel at different axial speeds depending on their proximity to walls; some particles can even come to rest on the braided internal pipes. Since the notion of applying TC insight to static mixer design is based on analogy rather than a sound theoretical basis, a thorough investigation is needed into whether TC can be exploited for static mixer design, or whether additional flow features not found in two-dimensional devices can undermine mixing quality in three dimensions.

Preliminary flow simulations have been performed by Vikhansky [141] in a 'topologically chaotic' static mixer consisting of a rectangular outer pipe, containing three braided pipes of circular cross-section. The flow field solution was computed numerically using an 'immersed' or 'virtual' boundary method in which no-slip conditions are not applied exactly on the circular boundaries, but by interpolation at the nearest neighbouring grid points. Vikhansky [141] reports that the mixer outperforms many leading commercial static mixers according to stretch rate and energy consumption measures, but suffers from the appearance of significant regular islands of poor mixing in a Poincaré section taken across the rectangular cross-section. The numerical method used also allowed the inclusion of inertial effects, up to a Reynolds number of 14, with second-order spatial accuracy; however, it is important to note that such a scheme is not ideal for studying the intricacies of chaotic advection, where exponential temporal growth of errors demands a more accurate description of the flow field.

Other analytical studies have been made of three-dimensional mixing in the Stokes flow regime, such as in the Eccentric Helical Annular Mixer [125], Kenics Static Mixer [60, 61, 62] (see Figure 1.5), Partitioned Pipe Mixer [85, 88, 94, 100], Rotated Arc Mixer [127], flow through twisted pipes [71, 72, 73, 130] and others [8, 56, 84]. However, none of the mixers in these papers is sufficiently flexible to allow an investigation of different flow 'topologies'.

We are therefore motivated in this chapter to develop an accurate representation of the velocity field in the BPM, in order to pursue the analogy with TC in three dimensions. We use our model to address two central questions: (i) whether the mixing is improved when the internal pipes form a nontrivial mathematical braid, and (ii) whether good mixing can be achieved when the internal pipes are very thin, just as effective mixing could be achieved in the BSD_m with thin stirring rods.

The chapter layout is as follows. In Section 5.2 we describe the Braided Pipe Mixer (BPM) and an algorithm for constructing the velocity field with spectral accuracy. Example computations are shown, and features of the solution are described. In Section 5.3 we extend the braid notation used in Chapter 4 to three dimensions, and describe how various BPM braids are constructed. These

are then evaluated by performing numerical dye advection simulations. In Section 5.4 we discuss features of the flow in the BPM, including the applicability of topological chaos, and the effects of stretching the internal pipes. Our conclusions are given in Section 5.5.

5.2 Mathematical model for the BPM

In this section we construct the velocity field for the Braided Pipe Mixer (BPM), shown in Figure 5.1. We work with the Cartesian coordinate system (x, y, ζ), as shown, with x and y as cross-sectional coordinates and ζ in the axial direction. It is also convenient to introduce the complex variables $z = x + \mathrm{i}y$ and $\bar{z} = x - \mathrm{i}y$. The mixer has no moving parts, and consists of a circular cylinder of radius a whose axis is aligned in the ζ direction and whose surface is defined by

$$\Gamma \equiv |z|^2 - a^2 = 0. \tag{5.1}$$

Contained inside the mixer are m twisted inner pipes of circular cross-section (not necessarily of fixed radius) defined by

$$\Gamma_j \equiv |z - p_j(\zeta)|^2 - a_j^2(\zeta) = 0 \quad j = 1, \ldots, m, \tag{5.2}$$

where $p_j(\zeta)$ determines how the j-th pipe twists down the device, and $a_j(\zeta)$ determines its radius in any given ζ cross-section; Figure 5.1 thus illustrates the case $m = 2$. (In practice one would need a means of fixing these inner pipes to the outer wall, such as wire stays, but we neglect these here.) Our attention in this chapter is focused on inner pipes that are periodic with wavelength L in the axial direction ζ. The mixer walls are referred to frequently in what follows, so for brevity we simply refer to them as Γ and Γ_j for the outer and inner pipes, respectively.

Incompressible viscous fluid occupies the domain between Γ and the Γ_j and flows steadily along the pipe at vanishing Reynolds number (Stokes flow). Mixing is achieved by forcing the fluid along the pipe, making it wind around the inner pipes as it progresses.

In a viscous flow regime the velocity field $\boldsymbol{u}$ satisfies the Stokes equation and the continuity equation

$$\mu\nabla^2\boldsymbol{u} = \boldsymbol{\nabla}p, \quad \boldsymbol{\nabla}\cdot\boldsymbol{u} = 0, \tag{5.3}$$

where $\boldsymbol{u}$, p and μ are the fluid velocity, pressure and dynamic viscosity, respectively. The no-slip condition $\boldsymbol{u} = \boldsymbol{0}$ applies on Γ and the Γ_j.

We prescribe the volume flow rate

$$Q = \iint_A w \,\mathrm{d}A, \tag{5.4}$$

where $w(x, y, \zeta)$ is the axial component of the fluid velocity and $A(\zeta)$ is the cross-sectional area of fluid, that is, the area enclosed between Γ and the Γ_j. Of course the flow rate Q must be the same across every cross-section A since the fluid is incompressible, and this condition is later used to determine the downstream pressure gradient, which varies along the length of the device. Alternatively, instead of specifying the flow rate, we could, equivalently, fix the pressure drop down the pipe, and compute from it the corresponding flow rate Q.

In a BPM having inner pipes with axial period L, the parameter $\epsilon = a/L$ is an indication of how twisted the pipes are. The velocity field for any ϵ can be calculated numerically [141]. To make analytical progress we focus on the regime $\epsilon \ll 1$ which corresponds to inner pipes with a long wavelength. We assume flow is driven by a pressure field in the form

$$p \sim \mu\left(\mathcal{P}(\zeta) + \epsilon\pi(x, y, \zeta)\right), \tag{5.5}$$

where the constraint that π should have zero average over any cross-section of constant ζ ensures that $\mathcal{P}(\zeta)$ is specified uniquely.

Here the flow is predominantly in the axial direction, enabling us to expand the velocity field as

$$\boldsymbol{u} \sim (\epsilon u, \epsilon v, w + \epsilon\tilde{w}), \tag{5.6}$$

and make the axial rescaling $Z = \epsilon\zeta$, so that the system (5.3) can be written

$$\begin{aligned}
\nabla_\perp^2 u + \epsilon^2 u_{ZZ} &= \pi_x, && (5.7)\\
\nabla_\perp^2 v + \epsilon^2 v_{ZZ} &= \pi_y, && (5.8)\\
\nabla_\perp^2 w + \epsilon\nabla_\perp^2\tilde{w} + \epsilon^2 w_{ZZ} + \epsilon^3\tilde{w}_{ZZ} &= \epsilon^2\pi_Z - P, && (5.9)\\
u_x + v_y + w_Z + \epsilon\tilde{w}_Z &= 0, && (5.10)
\end{aligned}$$

where $\nabla_\perp^2 = \partial^2/\partial x^2 + \partial^2/\partial y^2$. In (5.9), the pressure gradient function P is defined as

$$P(Z) = -\frac{\mathrm{d}\mathcal{P}}{\mathrm{d}\zeta}. \tag{5.11}$$

Accompanying these equations are the no-slip conditions $u = v = w = \tilde{w} = 0$ on Γ and Γ_j. At leading order, with $\epsilon \ll 1$, the governing equations reduce to give

$$\nabla_\perp^2 u = \pi_x, \tag{5.12}$$

$$\nabla_\perp^2 v = \pi_y, \tag{5.13}$$

$$\nabla_\perp^2 w = -P, \tag{5.14}$$

$$u_x + v_y + w_Z = 0, \tag{5.15}$$

together with constraint (5.4) of fixed volume flux.

In the next two sections we describe the steps taken to construct a spectrally accurate solution to equations (5.12)–(5.15).

5.2.1 Axial velocity field

We begin our solution by determining the leading-order axial velocity component w. Equation (5.14) may solved independently of u and v, and at each ζ gives the flow corresponding to a straight annular pipe with the local cross-section. The case of a single inner cylinder $m = 1$ has been studied before [3, 35], and an exact solution for steady viscous flow down a straight annular pipe has been derived [43, 91, 137]. For $m > 1$ no corresponding analytical solution is known, and we turn to the complex series approach employed for the two-dimensional batch stirring device discussed in Chapter 4.

The local pressure gradient term P is a parameter in the Poisson equation (5.14). Hence the solution at a given value of Z may be written in the form

$$w = P(Z)\left\{\tfrac{1}{4}(a^2 - z\bar{z}) + \hat{w}\right\}, \tag{5.16}$$

where $\hat{w}$ satisfies Laplace's equation

$$\nabla_\perp^2 \hat{w} = 4\hat{w}_{z\bar{z}} = 0 \tag{5.17}$$

and the boundary conditions

$$\hat{w} = \tfrac{1}{4}(z\bar{z} - a^2) \quad \text{on} \quad \Gamma_j, \tag{5.18}$$

$$\hat{w} = 0 \quad \text{on} \quad \Gamma. \tag{5.19}$$

Note that since the boundaries Γ_j vary with Z then $\hat{w}$ depends non-trivially on Z.

In its canonical form, (5.17) may be integrated easily. The most general real solution which satisfies (5.19) is [2, 98, 99]

$$\hat{w} = f(z) - f(a^2/\bar{z}) + c.c. \tag{5.20}$$

for any analytic function $f(z)$, where *c.c.* means the complex conjugate of all the preceding terms. By choosing f to be

$$f(z) = \sum_{j=1}^{m} \left\{ \alpha_{j,1} \left(\log(z - p_j) - \tfrac{1}{2} \log z \right) + \sum_{k=2}^{n} \alpha_{j,k} (z - p_j)^{1-k} \right\}, \tag{5.21}$$

that is, a logarithm plus n-term Laurent series centred at each p_j, as in Chapter 4, we generate a solution (5.20) that is single-valued and free of singularities in the domain provided that the $\alpha_{j,1}$ are real. Spectral accuracy is achieved using this form of solution, whereby the coefficients α decay according to $|\alpha_{j,k}| \sim C^k$ for some $C < 1$. The accuracy of the method allows for meaningful simulations of chaotic particle advection (see Section 5.3.2). The logarithmic singularity in $f(z)$ at the origin does not give rise to any corresponding singularity in $\hat{w}$ when $f(z)$ is substituted into (5.20).

The mn constants $\alpha_{j,k}$ are determined numerically by minimising the squared-error in the remaining boundary conditions (5.18) as discussed in the previous chapter [121]. The rapid convergence of the series in (5.21) means that the method is fast and requires storage only of the coefficients $\alpha_{j,k}$, and not of a grid of velocities as in a finite-element or finite-difference scheme. After the $\alpha_{j,k}$ are determined, we then compute P by imposing (5.4), as we now describe.

In substituting (5.16) into the flow rate expression (5.4), we find that evaluating the resulting integral is very involved, but can be done exactly. The key step is to transform the surface integral to a contour integral using the Milne-Thomson area theorem [99]

$$\iint_A w \, \mathrm{d}A = \frac{1}{2}\mathrm{i} \oint_{\partial A} \left(\int w \, \mathrm{d}z \right) \mathrm{d}\bar{z}. \tag{5.22}$$

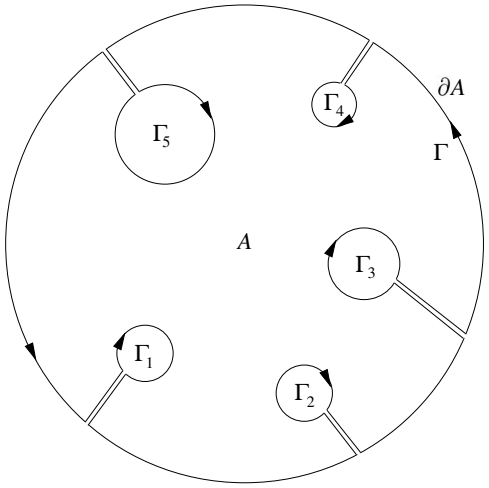

Figure 5.2: The contour ∂A used in the Milne-Thomson area theorem [99] to evaluate the axial flow rate.

The path ∂A traverses Γ in an anti-clockwise direction and each of the Γ_j in a clockwise sense, as illustrated in Figure 5.2. It is convenient to evaluate the contributions from each $\alpha_{j,k}$ term by term, and to consider each circle separately. When integrating the term corresponding to $\alpha_{j,k}$, the cases $k = 1$, $k = 2$ and $k > 2$ must be treated separately, and the integrals around the outer, j-th and non-j-th circular boundaries require slightly different approaches.

To demonstrate how the calculation is performed, consider the volume flux Q^* generated by the $\alpha_{j,1}$ term from (5.16) through the outer pipe Γ. According to (5.22) this is

$$\begin{aligned} Q^*/P &= \frac{1}{2}\mathrm{i}\alpha_{j,1}\oint_\Gamma \int \left[\log|z-p_j|^2 - \log|z-q_j|^2 - \log|p_j/a|^2\right] \mathrm{d}z\,\mathrm{d}\bar{z} \\ &= \frac{1}{2}\mathrm{i}\alpha_{j,1}\oint_\Gamma \Big[(z-p_j)\log(z-p_j) + z\log(\bar{z}-\bar{p}_j) \\ &\qquad - (z-q_j)\log(z-q_j) - z\log(\bar{z}-\bar{q}_j) \\ &\qquad + (p_j-q_j) - z\log|p_j/a|^2 + h(\bar{z})\Big]\,\mathrm{d}\bar{z}, \end{aligned} \tag{5.23}$$

where $q_j = a^2/\bar{p}_j$. In performing the inner integral, with respect to z, we have introduced $h(\bar{z})$, an arbitrary function of integration. We choose $h(\bar{z})$ to make the resulting integrand for the outer contour integral single valued, noting that, by Cauchy's theorem, the exact choice of h does not influence the calculated flow

rate, provided it does not introduce a singularity inside A. Then we have

$$\begin{aligned} Q^*/P &= \frac{1}{2}\mathrm{i}\alpha_{j,1}\oint_\Gamma \Big[(z-p_j)\left(\log(z-p_j)+\log(\bar{z}-\bar{p}_j)\right) \\ &\qquad -(z-q_j)\left(\log(z-q_j)+\log(\bar{z}-\bar{q}_j)\right) \\ &\qquad +(p_j-q_j)-z\log|p_j/a|^2\Big]\,\mathrm{d}\bar{z}. \end{aligned} \tag{5.24}$$

To make further progress we express the logarithm terms in (5.24) as Taylor expansions, giving

$$\begin{aligned} Q^*/P &= \frac{1}{2}\mathrm{i}\alpha_{j,1}\oint_\Gamma \Big[(z-p_j)\left(\log(z)-p_j/z-\cdots+\log(\bar{z})-\bar{p}_j/\bar{z}-\cdots\right) \\ &\qquad -(z-q_j)\left(\log(-q_j)-z/q_j-\cdots+\log(-\bar{q}_j)-\bar{z}/\bar{q}_j-\cdots\right) \\ &\qquad +(p_j-q_j)-z\log|p_j/a|^2\Big]\,\mathrm{d}\bar{z}. \end{aligned} \tag{5.25}$$

Note that care must be taken to write down the correct (convergent) form of Taylor expansion; on the contour Γ we have $|p_j/z|<1$ and $|z/q_j|<1$, since the singularities p_j and q_j lie inside and outside Γ, respectively. On the outer boundary Γ we have $z=a\mathrm{e}^{\mathrm{i}\theta}$, $\bar{z}=a\mathrm{e}^{-\mathrm{i}\theta}$ and $\mathrm{d}\bar{z}=-\mathrm{i}a\mathrm{e}^{-\mathrm{i}\theta}\,\mathrm{d}\theta$, and so we can rewrite (5.25) as

$$\begin{aligned} Q^*/P &= \frac{1}{2}\alpha_{j,1}\int_0^{2\pi}\Big[(a\mathrm{e}^{\mathrm{i}\theta}-p_j)\left(\log a^2-p_j\mathrm{e}^{-\mathrm{i}\theta}/a-\bar{p}_j\mathrm{e}^{\mathrm{i}\theta}/a-\cdots\right) \\ &\qquad -(a\mathrm{e}^{\mathrm{i}\theta}-q_j)\left(\log|q_j|^2-a\mathrm{e}^{\mathrm{i}\theta}/q_j-a\mathrm{e}^{-\mathrm{i}\theta}/\bar{q}_j-\cdots\right) \\ &\qquad +(p_j-q_j)-a\mathrm{e}^{\mathrm{i}\theta}\log|p_j/a|^2\Big]a\mathrm{e}^{-\mathrm{i}\theta}\,\mathrm{d}\theta \\ &= \alpha_{j,1}\pi\left[a^2\log\left(a^4/|p_jq_j|^2\right)+|p_j|^2-a^2\right]=\alpha_{j,1}\pi(|p_j|^2-a^2), \end{aligned} \tag{5.26}$$

which completes the calculation. (Note that the terms represented by dots in (5.25) and (5.26) do not contribute to the final result.)

The flow rate contributions from all the other terms in (5.16) may be treated in a similar way. When all the individual terms are collected together, we find the exact expression for the flow rate is

$$\begin{aligned} Q/P &= \tfrac{1}{16}\pi a^4-\sum_{j=1}^m \tfrac{1}{8}\pi a_j^2\left(a^2-|p_j|^2-\tfrac{1}{2}a_j^2\right) \\ &\quad +\sum_{j=1}^m\Big\{\tfrac{1}{2}\alpha_{j,1}\pi\Big[|p_j|^2-a^2 \\ &\quad -a_j^2\left(\log a_j^2-\log|p_j-q_j|^2-\log|p_j|^2+\log a^2-1\right) \\ &\quad -\sum_{s=1,s\neq j}^m a_s^2\left(\log|p_s-p_j|^2-\log|p_s-q_j|^2-\log|p_j|^2+\log a^2\right)\Big] \end{aligned}$$

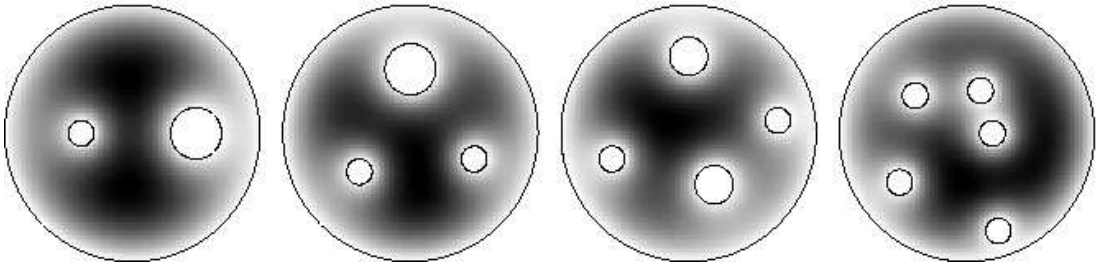

Figure 5.3: Illustrative grey-scale plots of the axial velocity component w in a cylinder containing between two and five inner pipes. Flow speed is shown by grey-level, with darker tones indicating greater flow speed.

$$
\begin{aligned}
&+ \alpha_{j,2}\pi\Big[- \bar{p}_j + \bar{p}_j a_j^2 (a^2 - |p_j|^2)^{-1} \\
&- \sum_{s=1, s\neq j}^{m} \Big(a_s^2 (p_s - p_j)^{-1} - \bar{p}_s a_s^2 (a^2 - p_j \bar{p}_s)^{-1} \Big) \Big] \\
&- \sum_{k=3}^{n} \alpha_{j,k}\pi \Big[- \bar{p}_j^{k-1} a_j^2 (a^2 - |p_j|^2)^{1-k} \\
&+ \sum_{s=1, s\neq j}^{m} \Big(a_s^2 (p_s - p_j)^{1-k} - \bar{p}_s^{k-1} a_s^2 (a^2 - p_j \bar{p}_s)^{1-k} \Big) \Big] \Big\} + c.c., \quad (5.27)
\end{aligned}
$$

where $c.c.$ means the complex conjugate of all the preceding terms. Since Q is specified, and all the other quantities are known, this enables us to determine P. Four sample solutions for w are shown in Figure 5.3. This figure illustrates that axial flow speed is small in the proximity of all the boundary walls. The retarding effect and the reduction in cross-sectional area caused by the presence of inner pipes means that a larger pressure gradient is required to drive a specified flow rate than that which would be needed in a Poiseuille flow through a pipe of radius a containing no inner pipes. A discussion of how the pressure gradient P depends on the locations p_j and radii a_j of the inner pipes is given in Section 5.2.3.

5.2.2 Transverse velocity field

Having found the axial velocity component w we now compute the transverse velocity components u and v. The remaining equations (5.12), (5.13) and (5.15) may be written as

$$\nabla_\perp^2 (u, v) \;=\; \nabla_\perp \pi, \qquad (5.28)$$

$$u_x + v_y = -w_Z. \tag{5.29}$$

We look for a solution $u = \phi_x + \psi_y$, $v = \phi_y - \psi_x$, in terms of a potential ϕ and streamfunction ψ, so that the velocity field is given by

$$u + \mathrm{i}v = 2\left(\phi_{\bar{z}} - \mathrm{i}\psi_{\bar{z}}\right), \tag{5.30}$$

and equations (5.28) and (5.29) become [2]

$$\nabla_\perp^2 \phi = -w_Z, \tag{5.31}$$

$$\nabla_\perp^4 \psi = 0, \tag{5.32}$$

subject to the no-slip condition

$$\phi_{\bar{z}} = \mathrm{i}\psi_{\bar{z}} \quad \text{on} \quad \Gamma \text{ and } \Gamma_j. \tag{5.33}$$

Since we have already computed w, the right-hand side of the Poisson equation (5.31) could be calculated numerically by finite-differencing w at neighbouring values of Z; however we can deduce w_Z more accurately as follows. We first differentiate (5.14) with respect to Z, leaving us to solve

$$\nabla_\perp^2 w_Z = -P_Z, \tag{5.34}$$

where P_Z can be computed by making just a single finite-difference approximation. To find boundary conditions on w_Z, consider two neighbouring points (x, y, Z) and $(x + \delta x, y + \delta y, Z + \delta Z)$ on a given boundary (say Γ, although the argument also applies for the Γ_j). Since

$$\Gamma(x + \delta x, y + \delta y, Z + \delta Z) = \Gamma(x, y, Z) + \delta x \Gamma_x + \delta y \Gamma_y + \delta Z \Gamma_Z + \cdots \tag{5.35}$$

$$w(x + \delta x, y + \delta y, Z + \delta Z) = w(x, y, Z) + \delta x w_x + \delta y w_y + \delta Z w_Z + \cdots \tag{5.36}$$

and since both Γ (by definition) and w (by no-slip) vanish on the boundary, it follows that to leading order in small quantities,

$$\delta x \Gamma_x + \delta y \Gamma_y + \delta Z \Gamma_Z = \delta x w_x + \delta y w_y + \delta Z w_Z = 0. \tag{5.37}$$

By considering, in turn, displacements with $\delta x = 0$ or $\delta y = 0$ it becomes apparent that $\Gamma_x w_Z - \Gamma_Z w_x = \Gamma_y w_Z - \Gamma_Z w_y = 0$ and by taking linear combinations of these results we can finally arrive at the conditions

$$w_Z = w_z \Gamma_{j_Z} / \Gamma_{j_z} \quad \text{on} \quad \Gamma_j, \tag{5.38}$$

$$w_Z = 0 \quad \text{on} \quad \Gamma. \tag{5.39}$$

The right-hand side of the boundary condition (5.38) is known, since from (5.16) and (5.20) we have

$$w_z/P \;=\; -\tfrac{1}{4}\bar{z} + f'(z) + a^2 z^{-2}\overline{f'(a^2/\bar{z})}, \tag{5.40}$$

where the derivative of f is easily calculated from (5.21), and also

$$\Gamma_{jZ}/\Gamma_{jz} = -p_{jZ} - \bar{p}_{jZ}\frac{z-p_j}{\bar{z}-\bar{p}_j} - 2\frac{a_{jZ}}{\bar{z}-\bar{p}_j}. \tag{5.41}$$

Again, we seek solutions to (5.34) in the form

$$w_Z \;=\; P_Z\left\{\tfrac{1}{8}(a^2 - z\bar{z}) + g(z) - g(a^2/\bar{z})\right\} + c.c. \tag{5.42}$$

$$g(z) \;=\; \sum_{j=1}^{m}\left\{\beta_{j,1}\left(\log(z-p_j) - \tfrac{1}{2}\log z\right) + \sum_{k=2}^{n}\beta_{j,k}(z-p_j)^{1-k}\right\} \tag{5.43}$$

with the constraint that the $\beta_{j,1}$ must be real, and determine the mn coefficients $\beta_{j,k}$ numerically by our least-squares algorithm. With the solution for w_Z determined we could go on to find ϕ by integrating (5.31). However, since we are interested only in the velocity field contribution from ϕ, given by (5.30), we need only integrate (5.31) once with respect to z to determine $\phi_{\bar{z}}$ rather than ϕ itself. We find that

$$2\phi_{\bar{z}} = -\tfrac{1}{2}\int w_Z \,\mathrm{d}z + \chi(\bar{z}), \tag{5.44}$$

where $\chi(\bar{z})$ is an arbitrary function. We must choose χ so that the velocity field resulting from ϕ is single-valued and after careful integration deduce

$$\begin{aligned}
2\phi_{\bar{z}}/P_Z \;=\; & -\tfrac{1}{8}(a^2 - \tfrac{1}{2}z\bar{z})z - \tfrac{1}{2}z\sum_{j=1}^{m}\sum_{k=2}^{n}\bar{\beta}_{j,k}(\bar{z}-\bar{p}_j)^{1-k} \\
& + \tfrac{1}{2}z\sum_{j=1}^{m}\left\{\beta_{j,1}\log(-p_j/a) + \sum_{k=2}^{n}\beta_{j,k}\left(\frac{a^2}{\bar{z}} - p_j\right)^{1-k}\right\} \\
& - \tfrac{1}{2}\sum_{j=1}^{m}\left\{\beta_{j,1}(z-p_j)(\log|z-p_j|^2 - 1)\right. \\
& \left. + \beta_{j,2}\log|z-p_j|^2 + \sum_{k=3}^{n}\beta_{j,k}\frac{(z-p_j)^{2-k}}{(2-k)}\right\} \\
& + \tfrac{1}{2}\sum_{j=1}^{m}\left\{\beta_{j,1}\left((z-q_j)(\log|z-q_j|^2 - 1) + z\log(-\bar{p}_j/a)\right)\right. \\
& + \sum_{k=2}^{n}\bar{\beta}_{j,k}(-\bar{p}_j)^{1-k}\left(z - q_j + (k-1)q_j\log|z-q_j|^2\right. \\
& \left.\left. + \sum_{l=2}^{k-1}\binom{k-1}{l}q_j^l\frac{(z-q_j)^{1-l}}{1-l}\right)\right\},
\end{aligned} \tag{5.45}$$

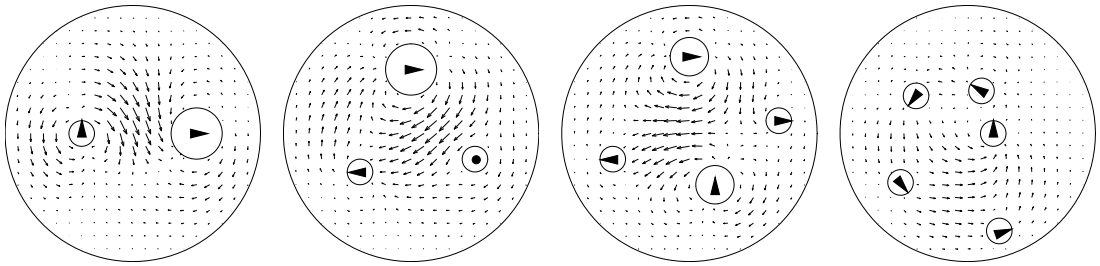

Figure 5.4: Example cross-sectional velocity fields (u, v) with between two and five inner pipes. In the flow domain, velocities are shown as arrows, with the length of the arrow indicating the speed. The arrows drawn inside the circles indicate $\mathrm{d}p_j/\mathrm{d}Z$, i.e. they show how the axis of each inner pipe varies as Z is increased.

where again $q_j = a^2/\bar{p}_j$. Now we must construct ψ to complete the description of the velocity field. The general solution to (5.32) is [99]

$$\psi = \bar{z}\kappa(z) + \lambda(z) + c.c., \tag{5.46}$$

giving rise to the velocity field contribution, through (5.30), of

$$-2i\psi_{\bar{z}} = -2i\left(\kappa(z) + z\overline{\kappa'(z)} + \overline{\lambda'(z)}\right). \tag{5.47}$$

A series solution for ψ is adopted, as discussed in the previous chapter, with κ and λ taking the form [26, 76]

$$\kappa = \sum_{k=1}^{n} \gamma_{0,k} z^{k-1} + \sum_{j=1}^{m} \left\{ \gamma_{j,1} \log(z - p_j) + \sum_{k=2}^{n} \gamma_{j,k} (z - p_j)^{1-k} \right\}, \tag{5.48}$$

$$\lambda = \sum_{k=2}^{n} \delta_{0,k} z^{k-1} + \sum_{j=1}^{m} \left\{ \bar{\gamma}_{j,1} z \log(z - p_j) + \delta_{j,1} \log(z - p_j) + \sum_{k=2}^{n} \delta_{j,k} (z - p_j)^{1-k} \right\}, \tag{5.49}$$

which correspond to a rotlet, Stokeslet and Laurent series at the centre of each Γ_j and also a Taylor series at the centre of Γ (the origin). The series in κ and λ enable us to satisfy (5.33) with spectral accuracy. The coefficients $\gamma_{j,k}$ and $\delta_{j,k}$ (with $j = 0, 1, \ldots, m$) are chosen to minimise the squared error in (5.33) on both Γ and the Γ_j. Sample cross-sectional velocity fields are shown in Figure 5.4.

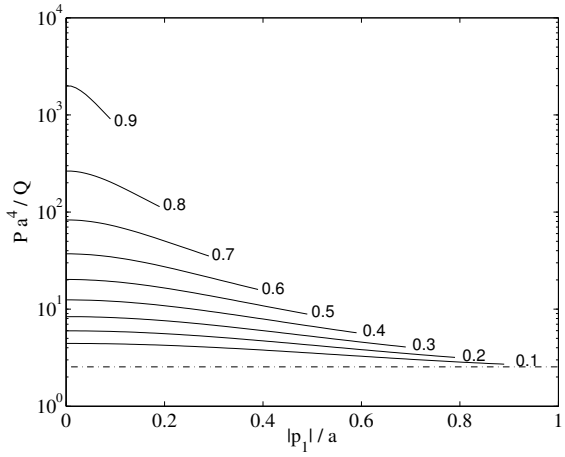

Figure 5.5: Plot of pressure gradient P against centre p_1 of the single inner pipe in a BPM with $m = 1$. The corresponding value of a_1/a is indicated beside each curve. The right-hand end point of each curve corresponds to $p_1 = a - a_1$, where the inner pipe is touching the outer cylinder. The limiting value $Pa^4/Q = 8/\pi$ for Poiseuille flow without an inner pipe is shown as a dashed line [91, 137].

5.2.3 Features of the solution

Having completed the description of the BPM velocity field we now describe some of the flow features. Figures 5.3 and 5.4 show example velocity fields at cross-sections of four mixers with between two and five inner pipes.

Figure 5.3 shows that the axial velocity is retarded in the proximity of the outer boundary Γ or any inner pipe Γ_j. The pressure gradient P required to produce a given flow rate Q depends not only on the cross-sectional area of the flow, but also on the locations p_j and radii a_j of the pipes.

For example, with $m = 1$ inner pipe, Shivakumar and Ji [134] have calculated how Q varies for a specified pressure gradient P (see also [91, 137]). As a check on our numerical implementation of the algorithm described in Section 5.2 we have successfully recalculated their results. In Figure 5.5 we plot the pressure gradient P required to produce a flow rate Q as a function of position p_1 and radius a_1 of the inner pipe. Our results confirm the findings of Shivakumar and Ji, that the pressure gradient P is greatest when the inner and outer pipes are

concentric.[1] The minimum of P occurs when $|p_1| = a - a_1$ so that the inner and outer pipes are touching.

With $m > 1$ inner pipes it becomes more difficult to determine which arrangement of pipes will minimise the pressure gradient P required to obtain a given flow rate Q, since the parameter space is many-dimensional. For example, with $m = 2$, there are already five parameters governing the cross-sectional geometry, namely $|p_1|/a$, $|p_2|/a$, a_1/a and a_2/a and $\arg(p_2/p_1)$. We make the general observation that if the inner pipes are placed in proximity, close to the outer boundary, then their overall retarding effect on the axial flow is reduced, meaning that the pressure gradient P in the axial direction is reduced. Further remarks on pressure variation are made in Section 5.3.

Flow in the cross-section occurs only when the inner pipes are twisted so that $\mathrm{d}p_j/\mathrm{d}Z \neq 0$ for some j. Figure 5.4 shows that fluid moves in the cross-section to avoid encountering the pipes further downstream, giving flow patterns that loosely mimic those seen in the two-dimensional batch stirring device from Chapter 4. In three dimensions, however, the cross-sectional flow is described by (5.30) in terms of both a streamfunction ψ and velocity potential ϕ, whereas the BSD_m flow is described in terms of a streamfunction only. The velocity field contribution from the potential ϕ has a non-zero divergence given by (5.31) which allows cross-sectional flow behaviour not seen in the two-dimensional BSD_m. Specifically, local source- or sink-like behaviour is possible, so that instead of fluid 'following' the twisting motion of the inner pipes, it is preferential for fluid to 'escape' in the axial direction, with an accompanying change in axial flow speed.

5.2.4 Accuracy of the flow field

In this section we expand on our claim above that the form of solution used allows the velocity to be determined accurately while requiring only a modest number

[1]Note that the form of solution discussed in Section 5.2 cannot be applied numerically if any pipe is perfectly concentric with the outer pipe, so that $p_j = 0$ for some j. This is because as the singularity location $p_j \to 0$ the corresponding image singularities present in (5.20) and (5.45) diverge to infinity. However, this singularity is removable, and may be avoided in our numerical implementation by placing the relevant pipe a small distance from the origin.

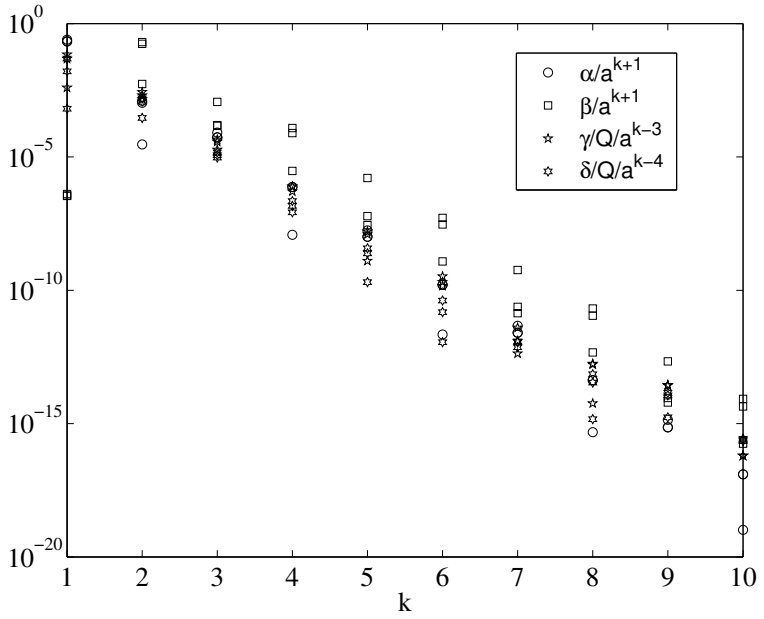

Figure 5.6: Exponential decay of the coefficients $\alpha_{j,k}$, $\beta_{j,k}$, $\gamma_{j,k}$ and $\delta_{j,k}$ with k. For definiteness, the values plotted correspond to the $Z = 0$ cross-section of the σ_1 braid letter described in Section 5.3.1, though the decay illustrated is typical for all cross-sections.

of terms to be kept in each series. In Figure 5.6 we plot the magnitude of the coefficients α, β, γ and δ, from (5.21), (5.43), (5.48) and (5.49), respectively, against the index k for a typical BPM with $m = 3$ pipes. The coefficients decay exponentially, indicating spectral accuracy. The times taken to determine α, β, γ and δ using the least-squares algorithm for various n are given in Table 5.1. Typically, the maximum error in the boundary conditions $\boldsymbol{u} = \mathbf{0}$ can be reduced to $10^{-6}Q/a^2$ with as few as $n = 10$ terms per series.

5.3 Investigating topological chaos

In this section we describe numerical dye advection simulations to examine the effect of changing the topology of the inner pipes in the braided pipe mixer. Boyland *et al.* [20] have conjectured that non-trivially braided inner pipes will mix fluid more effectively than other pipe arrangements (at least according to a stretch rate measure) and so we test this conjecture here. We then test, given a particular

	n	3	5	7	9	11
w	time	0.12	0.38	0.92	1.9	3.4
	error	1×10^{-4}	2×10^{-5}	6×10^{-7}	9×10^{-9}	1×10^{-9}
u, v	time	0.69	3.3	9.9	23	47
	error	7×10^{-2}	2×10^{-3}	2×10^{-4}	2×10^{-5}	1×10^{-6}

Table 5.1: Computation time (in seconds) and error in boundary conditions for the $m = 3$ inner pipe configuration in Figures 5.3 and 5.4. We compute the maximum error in the no-slip conditions for u, v and w in expressions (5.16) and (5.30) when truncating their series solutions to n terms. The errors are scaled by the typical fluid velocity Q/a^2. All calculations were performed on a 1.5GHz Pentium 4 computer.

braid topology, whether the quality of mixing is significantly influenced by varying other geometrical parameters. Recall that in the two-dimensional batch stirring device in Chapter 4, where the flow was topologically chaotic, we found that even with stirring rods of very small radius we could achieve effective mixing over a useful area of the domain. We address in this section the corresponding question of whether a BPM containing narrow pipes (therefore offering very little flow resistance beyond pure Poiseuille flow) can be a practically useful mixing device.

5.3.1 Braid construction

A simple way to create different braid topologies is to construct a BPM from a sequence of component units, each of which transposes two adjacent inner pipes. At the start and end of each unit the cross-section has the same geometry, with the pipe centres lying along a diameter; between these cross-sections two pipes twist around each other. In this way, each unit therefore corresponds to a braid letter σ, as described in Chapter 4. The mixer is then characterised by its particular sequence of component units corresponding to a braid word.

For simplicity, and to allow comparison with the results in Chapter 4, we study a BPM with $m = 3$ inner pipes. We also consider axially periodic mixers where a sequence of units corresponding to four braid letters is repeated indefinitely.

Braid letter	p_1/a	p_2/a	p_3/a
σ_1	$-\frac{1}{4}-\frac{1}{4}e^{i\pi Z/a}$	$-\frac{1}{4}+\frac{1}{4}e^{i\pi Z/a}$	$\frac{1}{2}$
σ_1^{-1}	$-\frac{1}{4}-\frac{1}{4}e^{-i\pi Z/a}$	$-\frac{1}{4}+\frac{1}{4}e^{-i\pi Z/a}$	$\frac{1}{2}$
σ_2	$-\frac{1}{2}$	$\frac{1}{4}-\frac{1}{4}e^{i\pi Z/a}$	$\frac{1}{4}+\frac{1}{4}e^{i\pi Z/a}$
σ_2^{-1}	$-\frac{1}{2}$	$\frac{1}{4}-\frac{1}{4}e^{-i\pi Z/a}$	$\frac{1}{4}+\frac{1}{4}e^{-i\pi Z/a}$

Table 5.2: Inner pipe shapes for a BPM component unit $0 \leq Z \leq a$ with $m = 3$, for each of the four different component braid letters σ_1, σ_1^{-1}, σ_2 and σ_2^{-1}.

Figure 5.7: Illustrations of the four BPM component units σ_1, σ_1^{-1}, σ_2 and σ_2^{-1}, showing the twisting of the inner pipes. The outer pipe surrounding all the inner pipes is not shown.

With $m = 3$ there are four possible crossings of adjacent pipes (two pairs of pipes, which may each be twisted 'anticlockwise' or 'clockwise'). Clearly there is a great deal of freedom in specifying the shape of the pipes. For simplicity we design each unit so that the axes of the two crossing pipes follow helices whilst third pipe is straight, with a constant cross-section. Expressions that we adopt for the p_j for the four possible braid letters σ_1, σ_1^{-1}, σ_2 and σ_2^{-1} are given in Table 5.2 generating the pipe shapes shown in Figure 5.7. Each of these units has length a (in terms of Z), so a four-letter braid word gives rise to a BPM with axial period $4a$.

Although our formulation in principle allows us to consider pipes with varying radii $a_j(Z)$ this has no influence on flow topology, so for simplicity we will take each $a_j = 0.1a$ to be constant.

As in Chapter 4, braid words should be read from right to left so that, for example, the word $\sigma_1\sigma_2$ means that the component σ_2 runs from $Z/a = 0$ to $Z/a = 1$ and

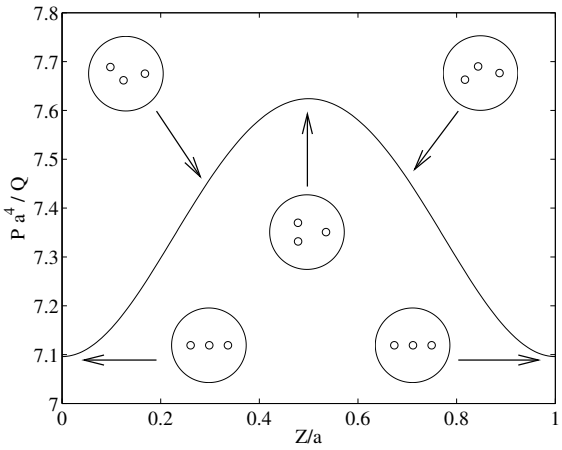

Figure 5.8: Plot of pressure gradient P against Z for each component unit described in Table 5.2. By symmetry, the pressure gradient function is the same for each braid letter. The cross-sectional location of the inner pipes is illustrated for $Z/a = 0,\ 0.25,\ 0.5,\ 0.75,\ 1$.

it is followed by σ_1 from $Z/a = 1$ to $Z/a = 2$. The sequence is then repeated, or, equivalently, fluid leaving the mixer at $Z/a = 2$ is reinjected in the same cross-sectional location at $Z/a = 0$.

To determine the flow completely we need to specify the flow rate Q. However, since Q has no influence beyond scaling all velocity components and the pressure gradient, we arbitrarily choose a unit flow rate for numerical simulations.

We note immediately, by symmetry, that the pressure gradient $P(Z)$ and, consequently, the total pressure drops across each braid component are exactly the same. This feature allows us to compare fairly the mixing performance of different braid topologies based on equal energy inputs and similar geometrical parameters. In his complementary study, Vikhansky [141] has shown that stretching in his braided pipe mixer is greater than in other commercially available mixers (based on equal driving pressure gradients), but has not performed a comparison of different braid topologies.

In Figure 5.8 we plot the variation of pressure gradient P with Z along a component unit, noting that $P(Z)$ is the same for all four braid letters. The plot reveals a maximum pressure gradient half-way along the pipe. Roughly speaking, this is

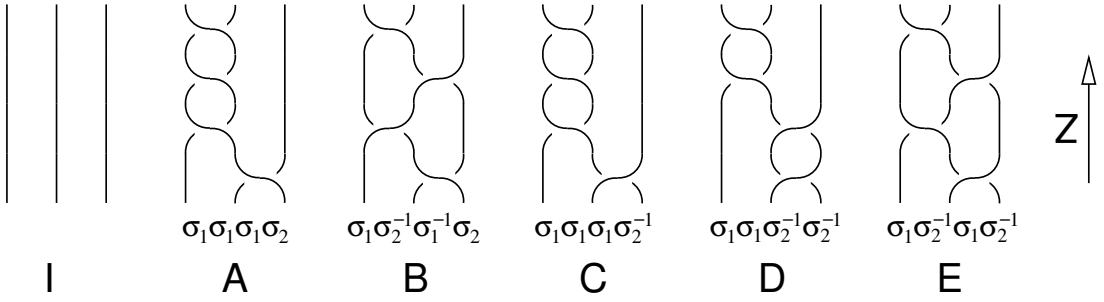

Figure 5.9: Braid diagrams and braid words for six arrangements of the BPM component units. The identity braid I corresponds to flow down a device of constant annular cross-section. The braids A–E correspond to those in Chapter 4.

where the inner-pipe locations are most widely spread over the domain, causing impedance to the flow over the greatest area. The minimum pressure gradient is found at the ends of the component, where the pipe centres are aligned along a diameter. In this configuration, regions above and below the pipes allow the bulk of fluid to flow relatively unhindered.

5.3.2 Effects of braid topology

We describe here numerical flow simulations performed to compare the effectiveness of braided pipe mixers using different braid words. In addition to the five braid words A–E considered in Chapter 4, we include, for reference, the identity braid I which is a device of constant annular cross-section with the p_j taking the constant values $p_1/a = -\frac{1}{2}$, $p_2/a = 0$, $p_3/a = \frac{1}{2}$. The braid diagrams for the six braid words are shown in Figure 5.9. We recall that the braids I and A have the same trivial topology, as we have discussed in Chapter 4.

An initial qualitative assessment of mixing performance for each braid word can be made by inspecting results of passive dye advection, where the position (x, y, Z) of a particle of dye satisfies the ordinary differential equation $(\dot{x}, \dot{y}, \dot{Z}) = \epsilon(u, v, w)$. As an initial condition we use the line of dye shown in Figure 5.10, which is represented numerically as a large number of particles. A plot showing the location of the dye after time $t = 10a^3/Q/\epsilon$ is shown in Figure 5.11 for the six different braid words I and A–E. This particular instant illustrates the dye after a noticeable

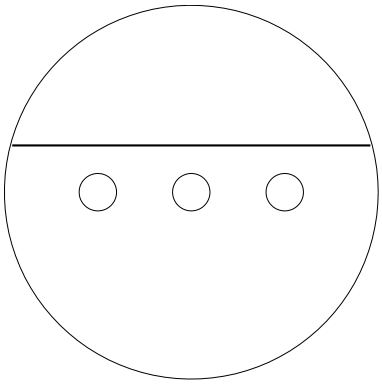

Figure 5.10: Plot of the initial line of dye used for simulations in this chapter. The line lies in the plane $Z = 0$ between the points $(x, y) = (-0.96, 0.25)$ and $(x, y) = (0.96, 0.25)$.

amount of mixing, but before the dye has become well-mixed.

We see immediately that the braids A–E all exhibit a qualitatively good mixing performance. It is difficult to distinguish by eye between the effects of the braid topologies B–E. This is in contrast to the corresponding two-dimensional BSD_3 dye advection plots in Figure 4.10 (Chapter 4) where there is a *visible* difference between stirring qualities. The dye in braid I is unmixed.

To probe mixing performance further, we compute some quantitative measures. As a comparison with the two dimensional batch stirrer from Chapter 4 we can compute the material stretch rates for each braid word. (Note, however, that we now have stretching in the axial direction, as well as in the cross-section.) In Figure 5.12 we show the variation of the stretched length $l(t)/l(0)$ of the line of dye with time. As in Chapter 4, we use the particle insertion algorithm of Krasnopolskaya *et al.* [83] to ensure that the stretched line remains well resolved.

Figure 5.12 shows that the identity braid I produces an algebraic stretch rate (in fact a linear one), whilst braids A–E provide exponential growth, as we would expect for a chaotic flow. Unlike in the two-dimensional BSD_3 from Chapter 4, however, it is difficult to differentiate between the stretch rates; if the line lengths grow according to $l(t)/l(0) \sim C\lambda^{tQ\epsilon/a^3}$ then we find that $\lambda \approx 1.22$ for each braid A–E, although the constant C is different for each of these braids. To explain why there is no clear distinction between the stretch rates we must consider more

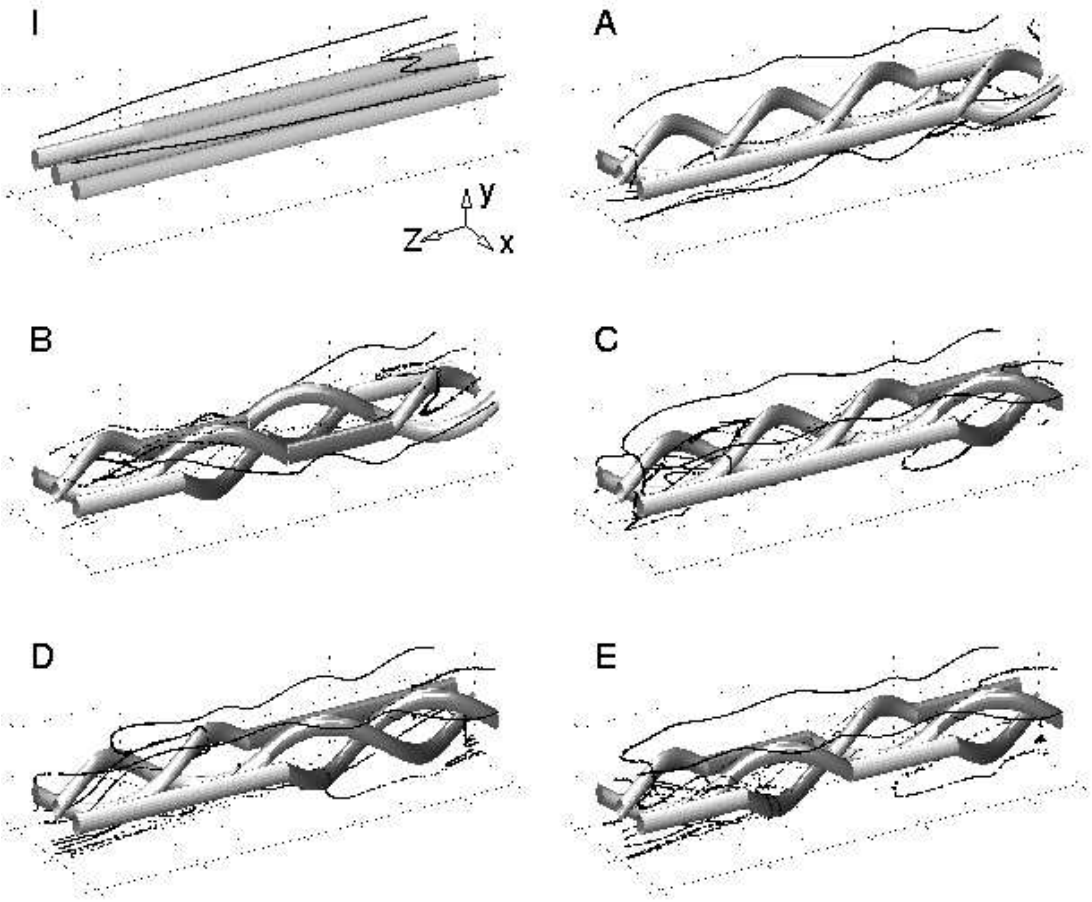

Figure 5.11: Plots of dye-streak locations after time $t = 10a^3/Q/\epsilon$ for the six braid words I, A–E. The outer cylinder is omitted from each plot for clarity. The axial locations Z of each particle are plotted modulo $4a$, so material leaving the end of the mixing unit is reinjected at the beginning, with the same cross-sectional coordinates.

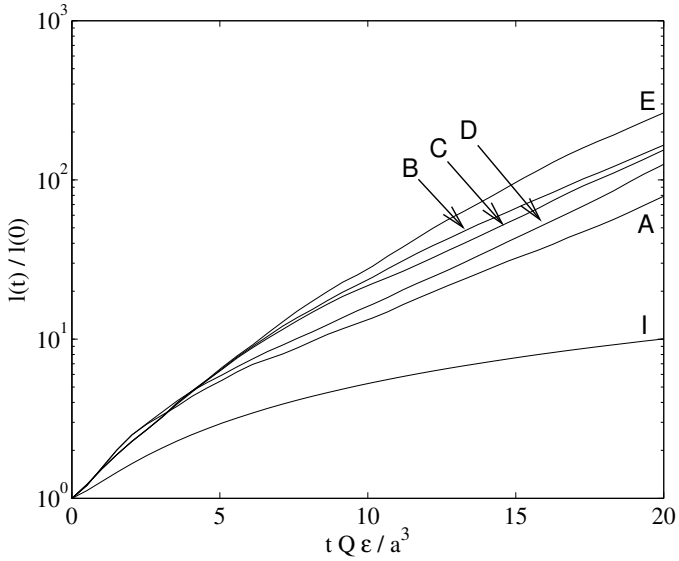

Figure 5.12: Plot of stretched line length $l(t)/l(0)$ with time t for the six braids I and A–E.

carefully how topological chaos is applied to flow in three dimensions; we discuss this in Section 5.4.1.

In a continuous-throughput mixer it is instructive to study the rate of axial transport along the pipe. It is straightforward to compute statistics for the axial distribution of material in the pipe, and these figures may be used to assess mixing quality. In Figure 5.13 we plot the axial coordinate Z_n of the nth particle ($1 \leq n \leq 10000$) at time $t = 30a^3/Q/\epsilon$, using the initial condition in Figure 5.10. This time was selected because the mean axial particle displacement is around $Z/a = 10.5$, or just under three complete braided pipe units. This corresponds to our plots given in Chapter 4 with mixing in the BSD_m for up to three complete braids. At time $t = 30a^3/Q/\epsilon$, the dye particles furthest along the pipe have progressed approximately as far as $Z/a = 16$, corresponding to the length of four complete braided pipe units.

The plots in Figure 5.13 demonstrate that the exact distribution of Z_n depends on which braid is used. However, the braid topology has only a small influence on the mean axial displacements, given in Table 5.3, in each case around $Z/a = 10$. The sharp spikes in Figure 5.13 correspond to where dye particles have been significantly retarded by close encounters with the inner pipe walls, and where

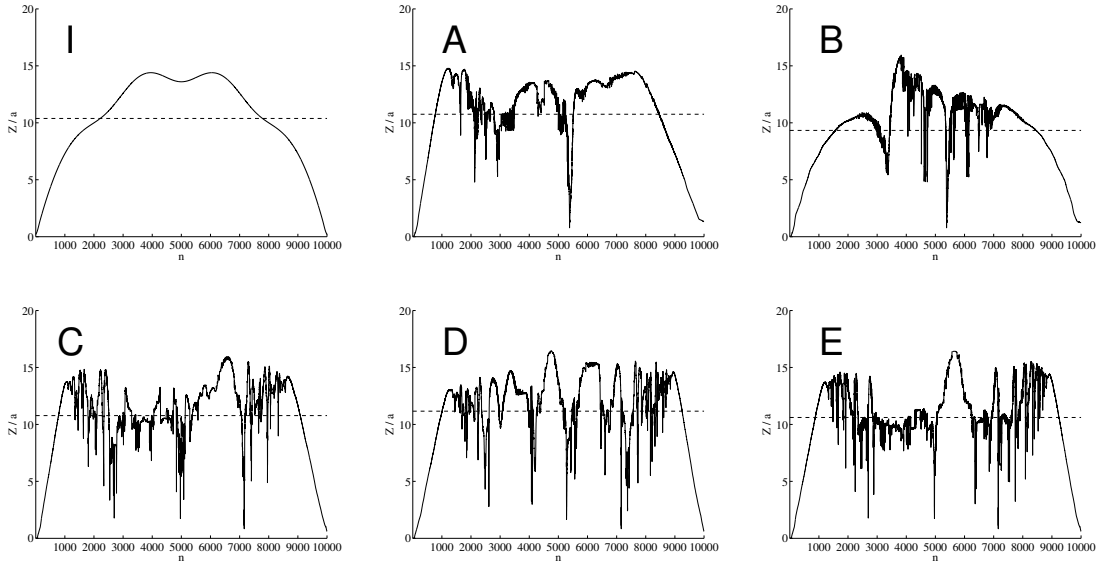

Figure 5.13: Axial distribution of dye for each of the six test braids. In each plot, the abscissa shows the particle number n ($1 \leq n \leq 10000$) and the ordinate axis records the corresponding Z-coordinate (i.e. the distance the dye has progressed in the axial direction) at time $t = 30a^3/Q/\epsilon$. The dashed line shows the mean displacement in the axial direction at that time.

Braid	mean$(Z)/a$	var$(Z)/a^2$	mean$(\|\delta Z\|)/a$	Spikes
I	10.4	14.6	0.00	0
A	10.7	14.0	0.08	33
B	9.3	11.2	0.11	53
C	10.8	11.8	0.09	37
D	11.2	13.4	0.10	48
E	10.6	11.7	0.14	70

Table 5.3: Axial distribution statistics for the six braids I,A–E.

dye has been stretched around a pipe.

Stretching of dye around the pipes is a possible mechanism that leads to good mixing and so we are interested in determining the number of such encounters. Where the spikes are found, there is a corresponding large difference between the final axial location of neighbouring particles from the initial line (Figure 5.10), and so it is worthwhile to study $\delta Z_n = |Z_{n+1} - Z_n|$ $(1 \leq n \leq 9999)$. The distributions of δZ_n for the six test braids are plotted in Figure 5.14. The numbers of 'large' spikes (which, rather arbitrarily, we will define here as those for which $\delta Z/a > 2$) in each plot gives an indication of mixing quality, and are recorded in Table 5.3. According to this crude 'mixing measure' the protocols are ordered (worst to best) as I, A, C, D, B, E. This is almost the same as the ordering according to line-stretch in Figure 5.12, except that C and D are swapped.

There are, of course, many other mixing measures that we could compute for the BPM flow. For example, most of the measures described in Chapter 3 for the two-dimensional BSD$_m$ could be used, provided appropriate modifications were made to account for the extra spatial dimension (e.g. iterated mappings could be replaced by Poincaré sections at fixed axial locations).

There are also further common measures that apply specifically to continuous-throughput mixers. Associated with the axial distribution plots in Figures 5.13 and 5.14, one common quantitative diagnostic of static mixers is the axial dispersion rate [71, 72, 73], which is the growth rate of the variance of the axial coordinates Z_n of the dye particles. The variance of the axial distribution of

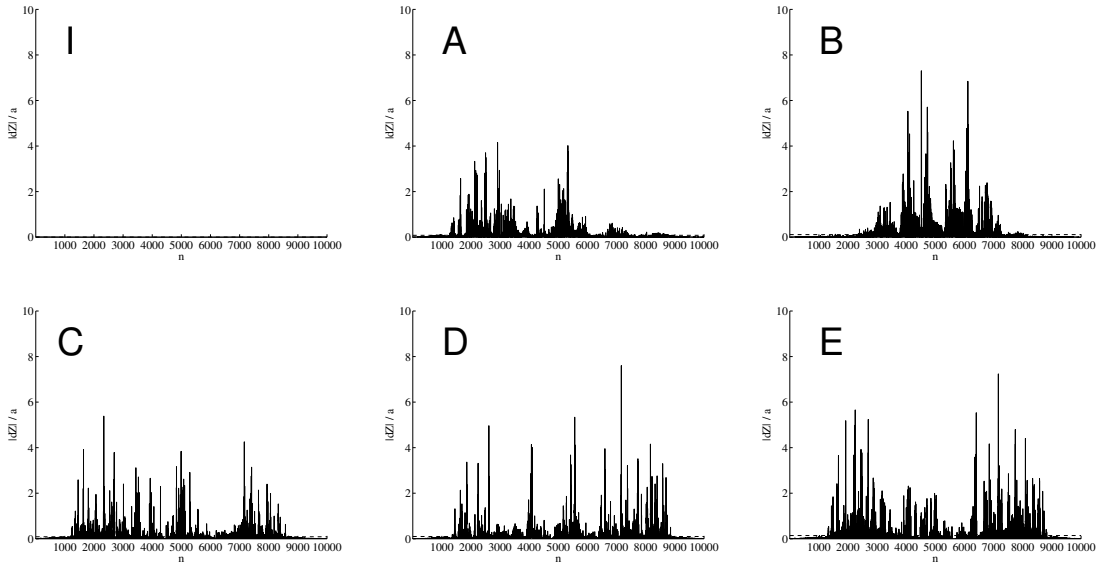

Figure 5.14: Distribution of $|\delta Z|$ for each of the six test braids. In each plot, the abscissa shows the particle number n ($1 \leq n \leq 9999$) and the ordinate axis records the value of $|Z_{n+1} - Z_n|/a$ at time $t = 30a^3/Q/\epsilon$.

perfect tracer in a regular laminar flow grows quadratically in time t [139]. With chaotic mixing in the cross-section, an averaging of the axial velocity for the different fluid particles occurs, and thus reduces the axial dispersion rate. With chaotic advection the variance growth is proportional t^ν with $1 < \nu < 2$ [71, 72, 73]. For example, in the twisted pipe flow studied by Jones [71], it is found that $\nu \approx 1.8$. In the BPM, we find that the variance of Z_n is similar for each of the braids A–E (see Table 5.3), and grows with $\nu \approx 2$, even though advection is chaotic. A possible explanation for this is that the inner pipe walls retard a significant number of particles, while others progress unhindered. This effect counters the averaging effect of the cross-sectional mixing and thus the exponent of axial dispersion is not significantly reduced below that in a regular flow.

Another continuous-throughput diagnostic of interest, related to axial dispersion, is the exit age distribution, or residence time distribution (RTD), which is the distribution of times taken for dye particles to pass a given axial location [22, 23, 87] (which could be the end of the mixer, for example). The integral with respect to time of the RTD gives the fraction of fluid that has passed a given axial location at a given time. In Figure 5.15 we plot, for braids I, A-E and the initial

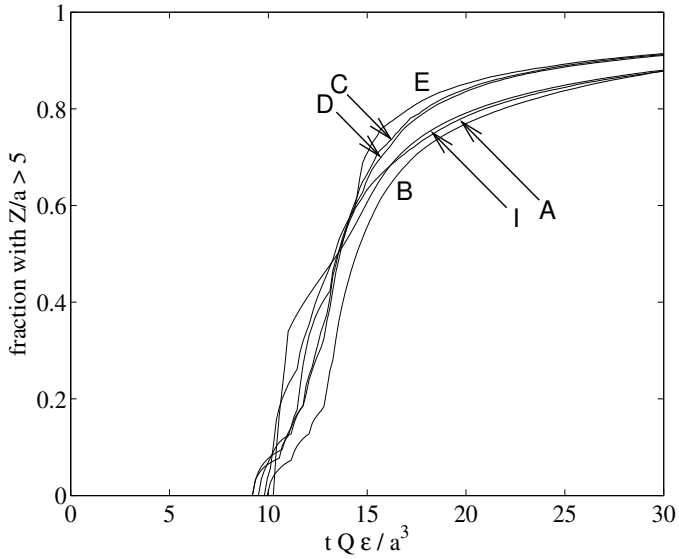

Figure 5.15: Residence time plots for the six braids I and A–E.

condition shown in Figure 5.10, the fraction of dye that has passed the axial location $Z/a = 5$ as a function of time. It can be seen that until approximately $t = 10a^3/Q/\epsilon$, no dye passes $Z/a = 5$ for any of the mixers, then the bulk of the dye particles flow past before $t = 20a^3/Q/\epsilon$. There is then a late tail for each curve representing the particles that flow more slowly in the axial direction due to their proximity to the pipe walls. As one might expect, since the mean and variance of the axial distributions of Z for the protocols I, A–E are similar (see Table 5.3), the curves in Figure 5.15 are also similar.

According to the diagnostics we have calculated in this section, there is not as clear a difference between the mixing quality of the different braids as we found in Chapter 4 for the corresponding braids in the two-dimensional batch stirring device. If it is possible to enhance static mixer performance by creating topologically nontrivial braids, as suggested by Boyland *et al.* [20] and Mackay [92], then we must consider more carefully the influence of the twisted internal pipes on the flow. We discuss topological chaos in three dimensions further in Section 5.4.1.

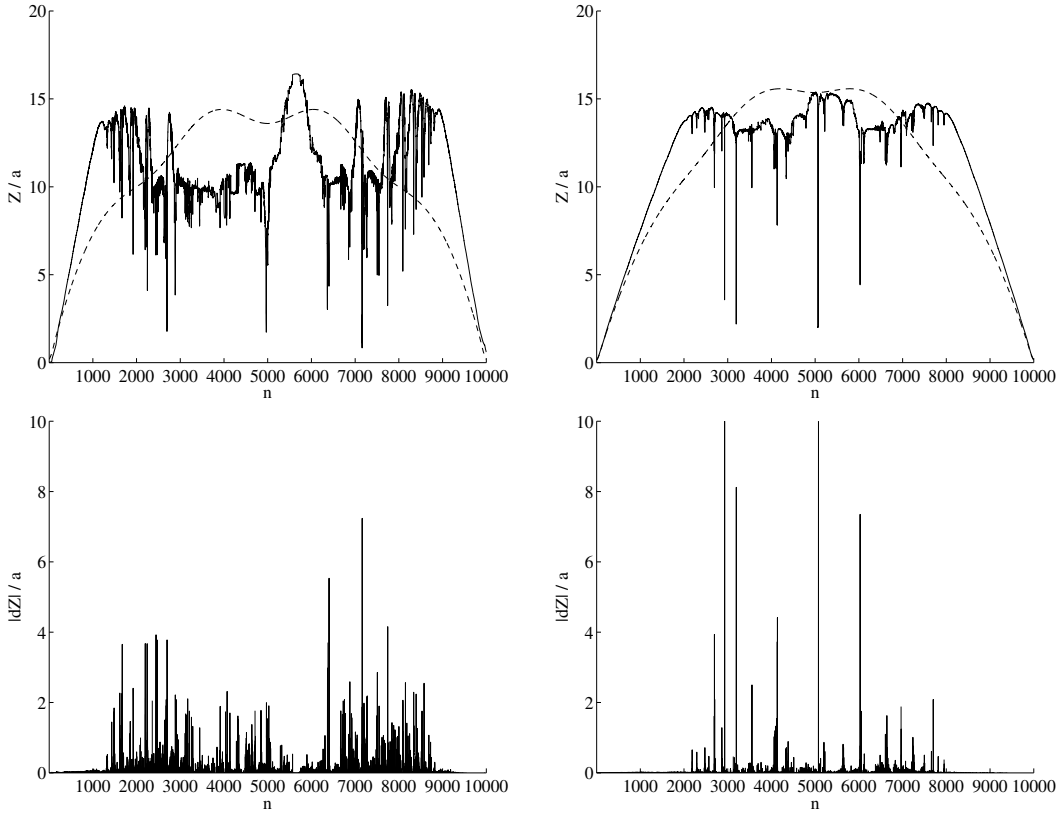

Figure 5.16: Simulations in BPM braid E with different radii of inner pipe. The plots on the left correspond to $a_j/a = 0.1$, and the right-hand plots to $a_j/a = 0.01$. The distributions of Z_n (top row) and δZ_n (bottom row) are plotted for the initial condition in Figure 5.10 at time $t = 30a^3/Q/\epsilon$. The dashed lines in the top plots show the corresponding results for braid I, for comparison.

5.3.3 Effects of pipe thickness

In Chapter 4, in the two-dimensional BSD_3, we examined the effect on the material stretch rate of the size of the inner stirring rods. We found that the stretch rates predicted by the topological theory are still achieved and that mixing still occurs over a wide area even using narrow stirring rods. In this section we study the effect of varying the radii of the inner pipes in the BPM. These radii have no influence on the topology of the braids, and so Boyland *et al.* [20] have speculated that as the pipe radii become small, effectively becoming wires, mixing quality might be preserved.

In Figure 5.16 we show axial distribution plots to examine the effect of shrinking

the radii a_j in the BPM braid E (see Section 5.3.2) from $a_j/a = 0.1$ to $a_j/a = 0.01$. It can be seen that as the inner pipes become thinner the distribution of the axial coordinates Z_n for braids E tends toward that of braid I. To be precise, the mean M of the absolute difference $|Z_n^E - Z_n^I|/a$ is $M = 3.1$ when $a_j/a = 0.1$ but reduces to $M = 1.6$ when $a_j/a = 0.01$. In the distributions of δZ (defined above), the number of large spikes (those for which $\delta Z/a > 2$) is also much lower, decreasing from 70 to 21 (some of the lines in Figure 5.16 actually conceal several spikes). This decrease occurs because the inner pipes have less of an influence on the velocity field, and consequently there is much slower transverse mixing and it takes much longer for new fluid particles to encounter the pipe walls. The largest spikes that remain for $a_j/a = 0.01$ correspond to the dye particles that are already 'lined up' with the pipes.

The suggestion by Boyland *et al.* [20] that shrinking the inner pipes should preserve mixing quality because it preserves flow topology appears not to apply in three dimensions as it did for the stirring rod size in the two dimensional BSD_m (see Chapter 4).

5.4 Discussion

Having numerically simulated the braided pipe mixer for various regimes of twisted inner pipe, we now discuss some general features of the BPM flow that have implications on mixing quality. We suggest in Section 5.4.1 why topological chaos does not appear to have a pronounced effect on stretching, as it did in the two-dimensional BSD_m. In Section 5.4.2 we discuss the features and limitations of our BPM model, in particular with regard to the assumption that the inner pipe twistedness parameter ϵ is small.

5.4.1 Topological chaos in three dimensions

The dye advection simulations we have seen in Section 5.3 show that there is little to distinguish between the mixing quality of different braid topologies. This is in contrast with the BSD_3 results from Chapter 4 where the mixing stretch

rates for braids A–E were distinct and in tight accordance with the theoretical predictions given by Boyland *et al.* [20]. That there is no clear ordering of the stretch rates for the different three-dimensional braids is perhaps unsurprising, since in a continuous-throughput mixer, particles travel at different axial speeds, and thus spend different amounts of time in each component unit: there is no clearly defined 'stirring period', in terms of which the stretch rate constant λ in the BSD_m was defined in Chapter 4.

If there is topological chaos present in the BPM, then we must reformulate the dynamics as a two-dimensional system in order to be able to apply the topological theory. The key to doing so is the analogy between time t for the BSD_m, and axial coordinate Z for the BPM: if the BPM dynamics are rewritten in terms of the independent variable Z then we can study motion in the xy-plane using Z effectively as 'time'. To this end, the passive dye-advection equations $(\dot{x}, \dot{y}, \dot{Z}) = (u, v, w)$ may be rewritten as

$$\left(\frac{\mathrm{d}x}{\mathrm{d}Z}, \frac{\mathrm{d}y}{\mathrm{d}Z}\right) = \left(\frac{u}{w}, \frac{v}{w}\right). \tag{5.50}$$

We identify $(u/w, v/w)$ as a 'pseudo-velocity field'. An example plot of this pseudo-velocity field is given in Figure 5.17 for the cross-section of the braid letter σ_1 at $Z/a = 0.4$. Note that even though $u, v, w \to 0$ as the inner pipes are approached, the ratios u/w and v/w give a finite pseudo-velocity on the pipe boundaries. By inspecting Figure 5.17 carefully it can be seen that the pseudo-velocity field resembles a hybrid of the inviscid and viscous regimes for the BSD_m (see Figures 4.2 and 4.3). Close to the inner pipes there is fluid slip, as in the inviscid BSD_m. However, the flow is not localised around the pipes, as with inviscid flow, and resembles more closely the viscous BSD_m.

Using Z as a 'time' coordinate, the inner pipes now appear to 'move' in exactly the same way as in the BSD_m; therefore we are guaranteed that the flow according to (5.50) is topologically chaotic. Why, then, is this topological chaos not manifest in the stretch rate measures in Section 5.3? To answer this question, we consider performing advection simulations according to (5.50): the greatest contribution to material stretching occurs for the fluid that becomes wrapped around the 'moving' inner pipes. Although this stretching occurs rapidly in 'time' Z, it is precisely the fluid near the inner pipes that moves slowly down the BPM in real

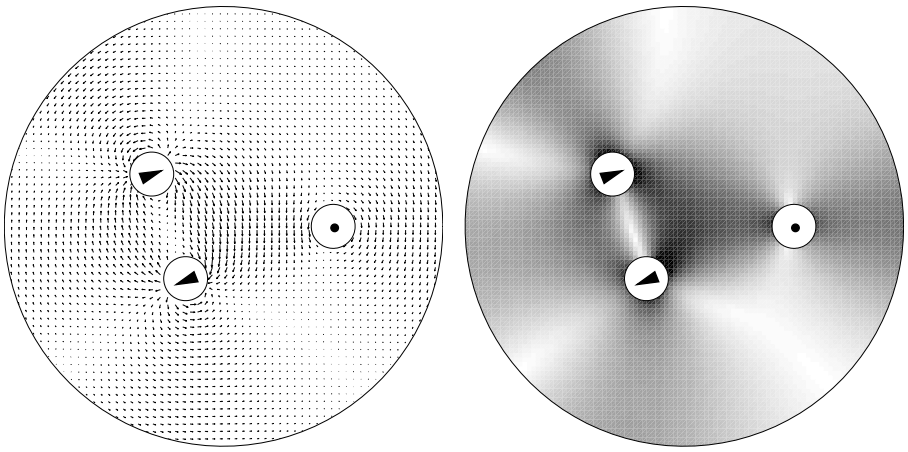

Figure 5.17: Plot of the pseudo-velocity field $(u/w, v/w)$. On the left is shown a vector plot of $(u/w, v/w)$ at the cross-section $Z/a = 0.4$ in for the braid letter σ_1. On the right is a grey-scale plot showing the magnitude of the pseudo-velocity field, with darker tones indicating faster flow. In the vicinity of the 'moving' cylinders the flow resembles the inviscid flow patterns of the BSD_3 in Chapter 4 (i.e. there is slip).

time t, and consequently this extreme stretching is not observed on a practical time scale. Since the topological theory could be applied to predict material stretch rates in the 'time' variable Z, using the pseudo-velocity, it would be a worthwhile future study to determine whether the topological chaos stretch-rate theory could then be modified in order to predict temporal material stretch rates for BPM flows.

5.4.2 Effects of pipe twistedness

So far in this chapter we have studied braided pipe mixers containing internal pipes that bend with a slow axial variation ($\epsilon \ll 1$). In this section we consider the possibility of increasing mixing quality and reducing the length of the BPM by coiling the inner pipes more tightly, so that $\epsilon = \mathcal{O}(1)$. This regime is not accessible to our present analysis because our derivation of the BPM velocity field in Section 5.2 assumed $\epsilon \ll 1$, so that flow was predominantly axial, with small transverse velocity components.

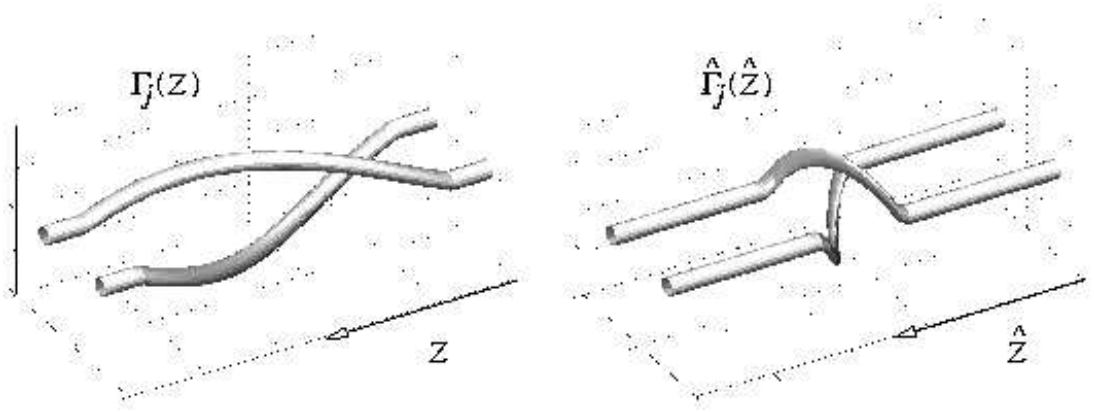

Figure 5.18: Stretching of a BPM component unit.

A simple thought experiment shows that the twistedness parameter ϵ must have a nontrivial influence on mixing quality. Consider the effect of squashing the braided inserts axially (see Figure 5.18), without altering their topology. As the inserts are coiled more tightly, fluid encountering the pipes traverses them like a bluff body, rather than as a guide to generate mixing in the cross-section. For example, in the BPM shown in Figure 5.18, as '$\epsilon \to \infty$', the inner pipes come to resemble two straight pipes with an annular ring placed at a given Z cross-section — the cross-sectional mixing is lost all together. In contrast, in two dimensions, the analogue to twistedness is the speed with which the stirrers are moved, and this has no effect on the mixing quality (in a Stokes flow regime).

With our model as currently developed, we cannot even capture the effects described in the above thought experiment by ignoring our assumption and using an artificially large value of ϵ (as we did in presenting the foregoing numerical results): the mixing performance depends only on the braid, not on the parameter ϵ. To see why this is so, consider two sections of a BPM (perhaps one of the 'component units' discussed in Section 5.3.1), one a non-uniformly stretched version of the other. As illustrated in Figure 5.18, we refer to the two component units with hatted and non-hatted variables. The stretching of one component with respect to the other is then characterised by the relations $Z = \Lambda(\hat{Z})$ and $\Gamma_j(Z) = \hat{\Gamma}_j(\hat{Z})$. The lengths of the sections are the same so that $\Lambda(0) = 0$ and $\Lambda(a) = a$ and we also insist that $\Lambda(\hat{Z})$ is strictly increasing (so $\Lambda'(\hat{Z}) > 0$ for all $\hat{Z}$).

The axial velocities in the two component units satisfy

$$w(x, y, Z) = \hat{w}(\hat{x}, \hat{y}, \hat{Z}) \tag{5.51}$$

(where $\hat{x} = x$ and $\hat{y} = y$), since w is determined by the location at Z of the boundaries and by the value of the flow rate Q. It follows from the chain rule that

$$\frac{\partial w}{\partial Z} = \frac{1}{\Lambda'} \frac{\partial \hat{w}}{\partial \hat{Z}} \tag{5.52}$$

and hence, by continuity, that

$$u(x, y, Z) = \frac{1}{\Lambda'} \hat{u}(\hat{x}, \hat{y}, \hat{Z}), \quad v(x, y, Z) = \frac{1}{\Lambda'} \hat{v}(\hat{x}, \hat{y}, \hat{Z}). \tag{5.53}$$

If particles are advected along the two components according to

$$\frac{\mathrm{d}x}{\mathrm{d}t} = u(x, y, Z), \quad \frac{\mathrm{d}y}{\mathrm{d}t} = v(x, y, Z), \quad \frac{\mathrm{d}Z}{\mathrm{d}t} = w(x, y, Z), \tag{5.54}$$

$$\frac{\mathrm{d}\hat{x}}{\mathrm{d}\hat{t}} = \hat{u}(\hat{x}, \hat{y}, \hat{Z}), \quad \frac{\mathrm{d}\hat{y}}{\mathrm{d}\hat{t}} = \hat{v}(\hat{x}, \hat{y}, \hat{Z}), \quad \frac{\mathrm{d}\hat{Z}}{\mathrm{d}\hat{t}} = \hat{w}(\hat{x}, \hat{y}, \hat{Z}), \tag{5.55}$$

and we have the initial condition $(x(0), y(0), Z(0)) = (\hat{x}(0), \hat{y}(0), \hat{Z}(0))$ for a given particle, then

$$x(t) = \hat{x}(\hat{t}), \quad y(t) = \hat{y}(\hat{t}), \quad Z(t) = \hat{Z}(\hat{t}) \tag{5.56}$$

provided we define $\hat{t}$ so that

$$\frac{\mathrm{d}\hat{t}}{\mathrm{d}t} = \frac{1}{\Lambda'}, \tag{5.57}$$

and $\hat{t} = 0$ when $t = 0$. Thus, according to our $\epsilon \ll 1$ model, particle trajectories are simply stretched in the same way as the component unit. In particular, the mapping of particle positions in the cross-section from one end of the component unit to the other is unchanged. The times of passage T for a particle across the components may differ according to how the components are stretched, however. This can be seen by considering

$$T = \int_0^a \frac{\mathrm{d}Z}{w}, \quad \hat{T} = \int_0^a \frac{\mathrm{d}\hat{Z}}{\hat{w}}, \tag{5.58}$$

then according to (5.51) we have

$$\hat{T} - T = \int_0^a \frac{1}{w} \left(\frac{1}{\Lambda'} - 1 \right) \mathrm{d}Z. \tag{5.59}$$

In general this quantity will be non-zero when $Z \not\equiv \hat{Z}$.

Since particles trajectories stretch in the same way as the component units as ϵ is increased, the cross-sectional velocity components (according to our model) become comparible in magnitude to the axial velocity component. The assumption that flow is predominantly in the axial direction is violated as ϵ becomes large. When $\epsilon = \mathcal{O}(1)$, equation (5.51) is no longer valid, because w depends not only on the locations of the pipes in the cross-section, but also on their axial variation. We might begin to capture such effects of increasing ϵ by including $\mathcal{O}(\epsilon^2)$ corrections to $\boldsymbol{u}$ in (5.6), which would depend implicitly on the axial derivatives of the p_j. It would then be found that particles follow different trajectories according to exactly how a component is stretched, and it would be possible to study the effect on mixing quality of using more compressed BPM component units.

5.5 Conclusions

Our aim in this chapter was to determine whether the ideas of topological chaos discussed in Chapter 4 could be applied to three-dimensional static mixers, such as the BPM. Although there is no formal topological chaos theory for flows in three dimensions, it was suggested [20, 92], based on an analogy with the two-dimensional batch stirring device studied in Chapter 4, that a static mixer containing a braided inner shape would mix well, provided the braid had an appropriate topology.

To investigate topological chaos we have developed an accurate series solution for the flow in the braided pipe mixer (BPM), a continuous-throughput device containing a number of twisted pipes that may be made into various types of braid. The mixing performances of different braids were evaluated by performing numerical dye advection simulations. We found that there was no clear influence of the braid topology on the mixing performance. This was in contrast to the results of Chapter 4, where calculated material stretch rates were in agreement with theoretical topological chaos predictions.

By considering a reformulation of dye particle advection using the axial coordinate Z as the independent variable we concluded that topological chaos is present in

the BPM, but that the regions where high material stretch rate occurs correspond precisely to where the flow is very slow — in the vicinity of the inner pipes. Hence over a practical time scale the effects of topological chaos are not evident, and so topological chaos may be of limited use in aiding continuous-throughput mixer design. In three dimensions the most intuitive way to obtain effective mixing seems to involve 'cutting and folding' of the fluid, such as in the Kenics static mixer (Figure 1.5). Of course questions of energy efficiency still remain, since devices like the Kenics mixer require a greater driving pressure gradient than for flow down a simple circular cylindrical pipe [91, 137].

It is worth noting that our analysis for the BPM may easily be extended to accommodate twisted pipes of non-circular cross-section by making local conformal transformations that map non-circular boundaries into circles [41, 44], as we have outlined in Chapter 4 for the batch stirring device. For example, flat ribbons may provide effective stirring whilst providing a lower flow resistance, as in the Kenics static mixer (see Figure 1.5). Although altering the stirrer cross-sections would not change the braid topologies in the current study, they offer a great deal of flexibility for tuning the BPM to increase mixing quality, and warrant future investigation. A future investigation should also be made into the flow solution for the BPM when the axial variation of the inner pipes is not slow. Tighter coiling of the inner pipes would enable the BPM to be shortened and may also enhance mixing.

Chapter 6

Conclusions

In this thesis we have studied various problems of fluid mixing and transport with applicability in biological and industrial settings. Specifically we have: developed a spatially-structured model for transport in an artificial kidney dialyser (Chapter 2); explored a range of fluid mixing measures and the empirical relationships between them for two-dimensional, laminar, chaotic flows (Chapter 3) [49]; examined the concept of topological chaos in two spatial dimensions and investigated whether it may be practical in batch mixer design (Chapter 4) [50]; and developed a model static mixer to investigate whether topological ideas may be extended to three-dimensional static mixer design (Chapter 5) [48].

6.1 Artificial kidney dialysis

In Chapter 2 we began by examining the problem of solute and water transport in an artificial kidney dialyser. By formulating a simple parallel-channel model of an artificial kidney we were able to predict the physiologically important effects on the solute and water removal rates of the membrane permeability properties, the solute and fluid input conditions and the dialyser geometry. We believe our model is the first to allow investigation of the simultaneous impact of ultrafiltration, osmosis and diffusion upon solute and water fluxes in a clinical dialyser regime.

The profiles of fluid velocity and solute concentration depend primarily on the

flow Reynolds number (Re), Péclet number (Pe) and the aspect ratio (δ) of the fluid channels. Parameter estimates based on measurements from a modern dialysis machine indicate that the system operates in a regime for which $Re \sim 1$, $Pe \sim 10^3$ and $\delta \sim 10^{-4}$. Steady-state asymptotic solutions for the fluid velocity and solute concentration were constructed as expansions in the small parameter $\epsilon = \delta Pe \sim 10^{-1}$. An excellent agreement was shown between this analysis and direct numerical solutions of the governing fluid flow and advection–diffusion equations.

We found that spatially averaged models (such as in [15, 16]) significantly overestimate diffusive and osmotic fluxes; these estimates were improved upon in our spatially-structured asymptotic model. Our model also shows how competition between the transport processes of ultrafiltration, diffusion and osmosis can either attenuate or enhance solute and water fluxes, compared with their values when these transport processes occur independently. Such behaviour is not evident in existing lumped-parameter, spatially-averaged models [15, 16]. For the cases of diffusion and osmosis, we find that there is an optimum ratio between blood and dialysate channel widths at which solute and solvent transport rates are maximised. We also found that by varying the channel inlet conditions it is possible, to an extent, to select independently the solute and water fluxes, and thereby perform a wide range of dialysis session profiles.

Although the mixing in our model dialyser is very simple, the mixing problem is non-trivial because of the interaction between the motions of the water (driven by ultrafiltration and osmosis) and solute (by diffusion and advection). In a more realistic model we would replace the parallel channels carrying blood and dialysate with a set of parallel pipes. Of course, as a careful examination of Figure 1.1 indicates, the hollow fibres in a clinical dialyser are not exactly parallel, but gently twist, mimicking the BPM geometry studied in Chapter 5. In view of this, chaotic advection may have a small influence on the cross-membrane transport rates.

Our model may also be applied directly to a human kidney nephron by considering a twisted pipe geometry and incorporating active transport of solutes by cells in the membrane walls. As we have discussed at the end of Chapter 2, our study

could also be generalised to investigate, in real or artificial kidneys, the combined effects of the multiple chemical species in blood, such as albumin, calcium, creatinine, glucose, phosphate, potassium, sodium and urea. This problem would be more analytically involved, due to different membrane permeabilities for each solute (caused by the different molecular sizes), and would also be complicated by the interplay between solvent and solute motions that arises because it is the total osmotic permeability of the solutes that drives the fluid motion.

Ultimately, the relationships between solute and water transport rates, flow rates, geometries and permeabilities could be employed in a dialyser control system, enabling the accurate regulation of a kidney dialysis session, where patient solute and water levels are brought to desired levels over a specified period, whilst causing minimum trauma.

6.2 Mixing measures for Stokes flow

A great deal of analytical, numerical and experimental research is directed towards improving mixer design, since, at an industrial level, mixing can be extremely costly and so there are tremendous economic benefits to be realised if significant improvements can be made. The fundamental basis of every mixer design process is to produce a mixer which maximises a given mixing measure. Though the measure used will naturally depend on the application (e.g. the yield of a chemical reaction [28, 30, 96, 104, 105, 113, 130]), an important question arises: how are the various mixing measures related, and to what extent can one measure be substituted for another when analysing a mixer? We have addressed this question in Chapter 3 as it has, to date, received very little attention in the literature.

In that chapter we studied the effectiveness of a large number of stirring protocols in a batch stirring device, the BSD_1 [20, 47] (see Figure 3.1), a model two-dimensional, laminar flow device which mimics the action of many of the planetary mixers used in industry (see Figure 1.3). The BSD_1 was assessed according to a variety of measures commonly found in the literature. Our focus was

whether the method of *sieving*, whereby candidate protocols are successively eliminated according to increasingly discerning mixing measures, could be effective in practice at determining globally optimal mixer designs. We also studied the extent to which different mixing measures were correlated, to determine whether computationally expensive measures could be replaced by cheaper ones, thereby accelerating the selection of good mixer designs.

It was found that, while no single protocol simultaneously optimised all measures, a small subset of the protocols performed well against most measures. However, it was difficult to elicit any general rule for selecting effective protocols: for example, superficially similar protocols were found to exhibit considerably different mixing capabilities. Also, some mixing protocols performed very poorly against some measures, but well against others. Our results indicate that the selection of effective protocols by sieving may therefore be ineffective in practice. The correlation coefficients between the ten measures studied were tabulated. Some pairs of measures were very poorly correlated; for example the correlation coefficient between lamellar width variance and Liapunov exponent was $|C| = 0.05$. This suggests that there is no simple relationship between the two measures, and therefore it might be possible to find a protocol that has a desirable value for both measures (e.g. low lamellar width variance and high Liapunov exponent). Other pairs of measures were well correlated. The correlation coefficients of mixing variance against Liapunov exponent or iterated mapping chaoticity were $|C| = 0.93$ and $|C| = 0.90$, respectively. It was noted that mixing variance was significantly simpler to compute, and so could serve as a useful substitute for Liapunov exponent or iterated mapping calculations.

It would be worthwhile to extend the investigation begun in Chapter 3 to three-dimensional mixing flows, either in a batch mixer or in a continuous-throughput device, such as the BPM from Chapter 5. Some of the mixing measures we have discussed may be applied directly to either case (for example, Liapunov exponent and energy usage measures – although in the continuous-flow case the appropriate measure would be the pressure drop); others can be extended by making straightforward changes to reflect the greater geometrical complexity of three spatial dimensions (for example, intermaterial density and interface stretch

measures apply to surfaces rather than curves). There is a somewhat subtle difference between a time-periodically operated two-dimensional mixer, such as the BSD_m from Chapters 3 and 4, and a continuous-flow device, such as the BPM from Chapter 5, in which the fluid becomes mixed as it passes down a pipe containing some spatial structure of axial wavelength L. In three dimensions, the analogue of the period-one map $\mathcal{M}$, used in Chapter 3, is the map taking particle positions in the cross-section at $\zeta = 0$ to their corresponding cross-sectional locations at $\zeta = L$.

Finally, the study in Chapter 3 was made possible because an exact expression was available for the velocity field, making accurate and fast advection simulations possible for a large number of stirring protocols; a future study should be focused on the problem of evaluating mixing devices where no exact solution is known for the velocity field, either by simplifying stirring dynamics without influencing the calculated values of the mixing measures, or by developing mixing measures that can be calculated without the need for computationally intensive Lagrangian particle tracking (e.g. measures that may be computed directly from the Eulerian velocity field [75, 151]).

6.3 Topological chaos in two dimensions

Mixer optimisation is usually achieved by parameter tuning. Multi-dimensional parameter spaces for mixer operation and a sensitive dependence on flow properties can make optimisation extremely difficult, as we have seen in Chapter 3. However, as we have studied in Chapter 4, the concept of Topological Chaos described by Boyland, Aref and Stremler [20] offers a way of enhancing mixer design in a robust way: a theoretical material stretch rate prediction may be made in a mixer, provided the flow has appropriate topological properties, irrespective of any details of the flow, such as compressibility or viscosity. However, this theoretical approach, while widely applicable, cannot predict the size of the region in which this stretch rate is achieved.

In Chapter 4 we studied the BSD_m, a generalisation of the BSD_1 from Chapter 3,

containing $m \geq 3$ stirring rods, and capable of generating topologically chaotic flow. We have provided numerical and experimental evidence to support the observation of Boyland *et al.* that the region of high stretch is comparable with that through which the stirring elements move during operation of the device. The key to our precise numerical simulations of dye advection was developing a fast, spectrally accurate series solution for computing the velocity field. Our technique was validated by comparing numerical simulations with experiments performed in a simple, inexpensive rig.

Calculated material stretch rates were found to be in close agreement with the corresponding theoretical predictions, irrespective of whether we modelled the fluid as inviscid or viscous, even though there are significant differences between the flow fields generated in the two cases. The theoretical stretch rates were achieved in the BSD_m even as the radii of the stirring rods were reduced, although the stretching became more localised around the trajectories of the rods.

At the end of Chapter 4 we discussed how to construct practical topologically chaotic mixers, and outlined how energy efficiency considerations could be addressed using the results of our analysis. It was shown how effective flow topologies could be created in batch mixers using simple gearing, by the strategic insertion of fixed rods into a flow with an otherwise trivial topology. However, it was found that the insertion of additional stirring rods may lead to a significant increase in energy input to the mixer, depending on the exact details of the geometry, so in some instances a non-topologically chaotic mixer may be deemed more economical. A more detailed investigation would be worthwhile to determine whether the flow topology with the greatest stretch rate can be achieved by the insertion of stationary stirring rods without a large increase in energy cost.

Another possible avenue for further investigation is the use of random stirring protocols [53]; since the stretch rate in a topologically chaotic mixer may be derived from a braid matrix, it may be possible to examine the effect of random stirring protocols in the BSD_m using the theory of random matrix products [45]. A practical implementation of random mixing protocols could be achieved by inserting and removing extra stationary stirring rods, in a similar way to the baffled cavity mixer [67].

The effect of stirrer shape on mixing quality is also a subject that deserves further investigation, because stirrers of circular cross-section are rather artificial and a 'paddle' cross-section might seem more natural. We have demonstrated at the end of Chapter 4 that other stirrers may be considered by using appropriate conformal transformations in the series solution for the velocity field (a flow with moving paddles is shown in Figure 4.20). Although the change in stirrer cross-section has no influence on the flow topology, it may have an impact on energy usage, or be used to enhance stretching rate above that predicted by the topological theory, allowing a more economical, or a higher performance mixer, respectively.

The numerical method used to compute the BSD_m flow also should receive further exploration, as it could be generalised for other non-fluid dynamical problems involving Laplace's equation or the biharmonic equation. Our method requires comparatively little computation and storage, and it may be possible to develop it so that it can be applied to more general problems as easily as finite-difference or finite-element methods, but with much greater accuracy.

6.4 Topological chaos in three dimensions

While batch-stirring flows (often modelled as two-dimensional) are commonly used in industry, for example in food production and brewing, they are rarely found in biological scenarios. Furthermore, many industrial applications, such as polymer blending or crude oil processing, demand continuous-throughput mixing in three-dimensional flows.

In Chapter 5 we investigated the braided pipe mixer (BPM) in which fluid flowing along a cylindrical pipe is stirred as it passes around smaller inner pipes judiciously twisted inside the mixer. It was recently suggested that the obstacles in the pipe could be braided in the ways described in Chapter 4 to generate effective chaotic mixing in three dimensions [20, 92]. Furthermore, it was conjectured that these pipes could be reduced in thickness and still generate good mixing. We addressed these hypotheses in Chapter 5.

We found that it was difficult to distinguish between the overall mixing perfor-

mance of different braids according to material stretch rate and measures based on the axial distribution of tracer. The results were in contrast to those of Chapter 4, where braid performances were clearly ordered, and in agreement with the two-dimensional topological theory. Although one might intuitively expect static mixer design could be improved using topological considerations, there is no formal topological chaos theory for three-dimensional flows. By considering particle advection in an artificial two-dimensional flow using the axial coordinate Z at 'time', we saw that topological chaos *was* present in the BPM; however, in this system, the regions of high material stretch rate corresponded exactly with those regions where particle advection down the device was very slow using the real time t. On a practical time scale the fast stretching is not seen, and therefore the topological chaos is of no significant benefit to the mixer design.

Our other important finding concerns the influence of geometrical parameters on the mixing quality. In the two-dimensional BSD_3 we found that effective mixing could be achieved even with narrow stirring rods. However, in the static mixer we find that decreasing the radii of the inner pipes reduces the mixing quality greatly, although it does lead to an energy saving due to the lower pressure gradient required to drive the flow.

Although topological considerations do not appear to offer significant help in enhancing the BPM, we have seen, nevertheless, that braided internal pipes do offer qualitatively good mixing (see Figure 5.11) and so it would be worth investigating how the pipe geometries could be altered to optimise the BPM. A future investigation should be made into the effects on mixing of increasing the twistedness ϵ of the internal pipes. By increasing ϵ, a more vigorous transverse flow could be created, and separation of fluid around the pipes could be made more pronounced, leading to more rapid mixing.

The role of flow separation around obstacles has been studied by Mizuno and Funakoshi [100] in the partitioned pipe mixer, where fluid flows around either side of baffles. A 'separation manifold' may be defined as the surface found by advecting backwards in time a streak of dye placed infinitesimally close to the ends of the baffles. Any fluid straddling such a manifold eventually becomes stretched around both sides of the corresponding baffle, and is subsequently well

mixed. Separation manifolds may also be studied for the BPM, to locate regions of effective mixing. The shape of these manifolds will depend quite sensitively on the exact geometry, not just the topology of the inner pipes. To compute the separation manifolds in the BPM it would be necessary to study the pseudo-velocity field in Section 5.4.1 more closely, to determine where separation occurs on the internal pipes.

In addition to simulating chaotic advection in the BPM, we could attempt to include additional processes. For instance, we could include multiple chemical species, in order to model industrial chemical reaction processes (e.g. polymerisation). We could also model biological processes, such as transport in a dialysis fibre bundle, relating back to Chapter 2, or describe blood flow in a catheterised artery [3, 35, 68]. As with the two-dimensional BSD_m, it would be relevant to applications to determine the effects of changing the inner pipe cross-section in the BPM by using conformal mapping techniques. For instance, flat ribbon-like inserts could be investigated (using the Joukowski transformation given in Chapter 4) similar to those used in the Kenics static mixer (Figure 1.5).

Bibliography

[1] M. Abbas and V. P. Tyagi. Analysis of a hollow-fibre artificial kidney performing simultaneous dialysis and ultrafiltration. *Chem. Eng. Sci.*, 42:133–142, 1987.

[2] D. J. Acheson. *Elementary Fluid Dynamics.* Clarendon Press, Oxford, 1996.

[3] R. R. Adabala. Peristaltic pumping in a circular tube in the presence of an eccentric catheter. *Trans. of the ASME*, 117:448–454, 1995.

[4] A. Adrover, M. Giona, F. J. Muzzio, S. Cerbelli, and M. M. Alvarez. Analytic expression for the short-time rate of growth of the intermaterial contact perimeter in two-dimensional chaotic flows and hamiltonian systems. *Phys. Rev. E*, 58:447–458, 1998.

[5] G. K. Aldis. The osmotic permeability of a tubule wall. *Bull. Math. Biol.*, 50:547–558, 1988.

[6] G. K. Aldis. The unstirred layer during osmotic flow into a tubule. *Bull. Math. Biol.*, 50:531–545, 1988.

[7] M. M. Alvarez, F. J. Muzzio, S. Cerbelli, A. Adrover, and M. Giona. Self-similar spatiotemporal structure of intermaterial boundaries in chaotic flows. *Phys. Rev. Lett.*, 81:3395–3398, 1998.

[8] M. M. Alvarez-Hernández, T. Shinbrot, J. Zalc, and F. J. Muzzio. Practical chaotic mixing. *Chem. Eng. Sci.*, 57:3749–3753, 2002.

[9] H. Aref. Stirring by chaotic advection. *J. Fluid Mech.*, 143:1–21, 1984.

[10] H. Aref. The development of chaotic advection. *Phys. Fluids*, 14:1315–1325, 2002.

[11] H. Aref and S. Balachandar. Chaotic advection in a Stokes flow. *Phys. Fluids*, 29:3515–3521, 1986.

[12] H. Aref and S. W. Jones. Enhanced separation of diffusing particles by chaotic advection. *Phys. Fluids A*, 3:34–38, 1989.

[13] T. Atobe, M. Funakoshi, and S. Inoue. Orbital instability and chaos in the Stokes flow between two eccentric cylinders. *Fluid Dyn. Res.*, 16:115–129, 1995.

[14] A. Avudainayagam and B. Jothiram. A circle theorem for plane Stokes flows. *Q. J. Mech. Appl. Math.*, 41:383–393, 1988.

[15] S. Baigent, R. Unwin, and C. C. Yeng. Mathematical modelling of profiled haemodialysis: a simplified approach. *J. Theo. Med.*, 3:143–160, 2001.

[16] S. Baigent and M. Williams. Mathematical modelling of profiled haemodialysis. In H. M. Byrne and J. R. King, editors, *Proceedings of the first mathematics in medicine study group*, pages 58–67, University of Nottingham, 2000.

[17] J.-M. Bourot. Sur l'application d'une méthode de moindre carrés à la résolution approchée du problème aux limites, pour certaines catégories d'écoulements. *J. Méc.*, 8:301–322, 1969.

[18] P. Boyland. Fluid mechanics and mathematical structures. In R. L. Ricca, editor, *An introduction to the geometry and topology of fluid flows*, pages 105–134. Kluwer, Dordrecht, 2002.

[19] P. Boyland, M. Stremler, and H. Aref. Topological fluid mechanics of point vortex motions. *Physica D*, 175:69–95, 2003.

[20] P. L. Boyland, H. Aref, and M. A. Stremler. Topological fluid mechanics of stirring. *J. Fluid Mech.*, 403:277–304, 2000.

[21] J. P. Butler and A. Tsuda. Effect of convective stretching and folding on aerosol mixing deep in the lung, assessed by approximate entropy. *J. Appl. Phys.*, 83:800–809, 1997.

[22] C. Castelain, D. Berger, P. Legentilhomme, A. Mokrani, and H. Peerhossaini. Experimental and numerical characterisation of mixing in a steady spatially chaotic flow by means of residence time distribution measurements. *Int. J. Heat Mass Transfer*, 43:3687–3700, 2000.

[23] C. Castelain, A. Mokrani, Y. L. Guer, and H. Peerhossaini. Experimental study of chaotic advection regime in a twisted duct flow. *Eur. J. Mech. B/Fluids*, 20:205–232, 2001.

[24] J. Chaiken, R. Chevray, M. Tabor, and Q. M. Tan. Experimental study of Lagrangian turbulence in a Stokes flow. *Proc. R. Soc. Lond. A*, 408:165–174, 1986.

[25] H. Chaté, E. Villermaux, and J.-M. Chomaz. *Mixing; chaos and turbulence.* Kluwer Academic / Plenum Publishers, New York, 1999.

[26] A. T. Chwang and T. Y. Wu. Hydromechanics of low-Reynolds-number flow. Part 2. Singularity method for Stokes flows. *J. Fluid Mech.*, 67:787–815, 1975.

[27] M. J. Clifford, S. M. Cox, and M. D. Finn. Reynolds number effects in a laminar chaotic mixer with changing geometry. Submitted to *Chem. Eng. Sci.*

[28] M. J. Clifford, S. M. Cox, and E. P. L. Roberts. Lamellar modelling of reaction diffusion and mixing in a two-dimensional flow. *Chem. Eng. J.*, 71:49–56, 1998.

[29] M. J. Clifford, S. M. Cox, and E. P. L. Roberts. Measuring striation widths. *Phys. Lett. A*, 260:209–217, 1999.

[30] M. J. Clifford, S. M. Cox, and E. P. L. Roberts. The influence of a lamellar structure upon the yield of a chemical reaction. *Inst. Chem. Eng.*, 78:371–377, 2000.

[31] C. K. Colton, K. A. Smith, P. Stroeve, and E. W. Merrill. Laminar flow mass transfer in a flat duct with permeable walls. *AIChEJ.*, 17:773–780, 1971.

[32] D. O. Cooney, S. Kim, and E. J. Davis. Analyses of mass transfer in hemodialyzers for laminar blood flow and homogeneous dialysate. *Chem. Eng. Sci.*, 29:1731–1738, 1974.

[33] D. A. Covey and N. O. Ewoldsen. Porosity in manually and machine-mixed resin-modified glass ionomer cements. *Oper. Dent.*, 26:617–623, 2001.

[34] D. D'Alessandro, M. Dahleh, and I. Mezić. Control of mixing in fluid flow: a maximum entropy approach. *IEEE Trans. on automatic control*, 44:1852–1863, 1999.

[35] P. Daripa and R. K. Dash. A numerical study of pulsatile blood flow in an eccentric catheterized artery using a fast algorithm. *J. Eng. Math.*, 42:1–22, 2002.

[36] J. T. Daugirdas, P. G. Blake, and T. S. Ing. *Handbook of dialysis.* Lippincott Williams and Wilkins, Philadelphia, 2001.

[37] R. L. Davidchack and Y. Lai. Efficient algorithm for detecting unstable periodic orbits in chaotic systems. *Phys. Rev. E*, 60:6172–6175, 1999.

[38] W. M. Deen, C. R. Robinson, and B. M. Brenner. Concentration polarization in an ultrafiltering capillary. *Biophys. J.*, 14:412–431, 1974.

[39] UK Department of Trade and Industry. Chemicals sector overview, May 2003. http://www.dti.gov.uk/sectors_chemicals.html.

[40] E. Dresselhaus and M. Tabor. The persistence of strain in dynamical systems. *J. Phys. A*, 22:971–984, 1989.

[41] T. A. Driscoll and L. N. Trefethen. *Schwarz–Christoffel mapping.* Cambridge University Press, Cambridge, 2002.

[42] C. Eckart. An analysis of stirring and mixing processes in incompressible fluids. *J. Mar. Res.*, 7:265–275, 1948.

[43] M. R. El-Saden. Heat conduction in an eccentrically hollow, infinitely long cylinder with internal heat generation. *J. Heat Transfer*, 83:510–512, 1961.

[44] M. Embree and L. N. Trefethen. Green's functions for multiply connected domains via conformal mapping. *SIAM Review*, 41:745–761, 1999.

[45] M. Embree and L. N. Trefethen. Growth and decay of random Fibonacci sequences. *Proc. R. Soc. Lond.*, 455:2471–2485, 1999.

[46] S. D. Fields and J. M. Ottino. Effect of stretching path on the course of polymerizations: applications to idealised unpremixed reactors. *Chem. Eng. Sci.*, 42:467–477, 1987.

[47] M. D. Finn and S. M. Cox. Stokes flow in a mixer with changing geometry. *J. Eng. Math.*, 41:75–99, 2001.

[48] M. D. Finn, S. M. Cox, and H. M. Byrne. Chaotic advection in a braided pipe mixer. Submitted to *Phys. Fluids*.

[49] M. D. Finn, S. M. Cox, and H. M. Byrne. Mixing measures for two-dimensional chaotic Stokes flows. Submitted to *J. Eng. Math.*

[50] M. D. Finn, S. M. Cox, and H. M. Byrne. Topological chaos in inviscid and viscous mixers. *J. Fluid Mech.*, 493:345–361, 2003.

[51] V. Fitt, J. R. Ockendon, and M. Shillor. Counter-current mass transfer. *Int. J. Heat Mass Transfer*, 28:753–759, 1985.

[52] J. G. Franjione, C. Leong, and J. M. Ottino. Symmetries within chaos: a route to effective mixing. *Phys. Fluids*, 1:1772–1783, 1989.

[53] J. G. Franjione and J. M. Ottino. Symmetry concepts for geometric analysis of mixing flows. *Phil. Trans. R. Soc. Lond. A*, 338:301–323, 1992.

[54] R. A. Frazer. On the motion of circular cylinders in a viscous fluid. *Phil. Trans. R. Soc. Lond. A*, 225:93–130, 1925.

[55] Y. C. Fung. *Biomechanics: Circulation.* Springer–Verlag, New York, 1997.

[56] O. S. Galaktionov, P. D. Anderson, and G. W. M. Peters. Symmetry of periodic structures in a 3d mixing cavity flow. *Phys. Fluids*, 12:469–471, 2000.

[57] A. E. Gill. *Atmosphere-Ocean Dynamics.* Academic Press Inc., London, 1982.

[58] M. Giona and A. Adrover. Global geometry and coarse-grained formulation of the evolution of pointwise intermaterial interface measure in chaotic flows. *Chem. Eng. Sci.*, 56:3387–3399, 2001.

[59] J. B. Grotberg. Pulmonary flow and transport phenomena. *Annu. Rev. Fluid Mech.*, 26:529–571, 1994.

[60] D. M. Hobbs and F. J. Muzzio. The kenics static mixer: a three dimensional chaotic flow. *Chem. Eng. J.*, 67:153–166, 1997.

[61] D. M. Hobbs and F. J. Muzzio. Optimization of a static mixer using dynamical systems techniques. *Chem. Eng. Sci.*, 53:3199–3213, 1998.

[62] D. M. Hobbs and F. J. Muzzio. Reynolds number effects on laminar mixing in the kenics static mixer. *Chem. Eng. J.*, 70:93–104, 1998.

[63] F. C. Hoppensteadt and C. S. Peskin. *Mathematics in medicine and the life sciences.* Springer–Verlag, New York, 1992.

[64] T. Hwu. Stretches of fluid materials for Stokes flow in circular cavity. *J. Eng. Mech.*, 126:554–557, 2000.

[65] T. Hwu, D. Young, and Y. Chen. Chaotic advections for Stokes flows in circular cavity. *J. Eng. Mech.*, 123:774–782, 1997.

[66] S. C. Jana, G. Metcalfe, and J. M. Ottino. Experimental and computational studies of mixing in complex Stokes flows: the vortex mixing flow and multicellular cavity flows. *J. Fluid Mech.*, 269:199–246, 1994.

[67] S. C. Jana, M. Tjahjadi, and J. M. Ottino. Chaotic mixing of viscous fluids by periodic changes in geometry: baffled cavity flow. *AIChEJ.*, 40:1769–1781, 1994.

[68] G. Jayaraman and R. K. Dash. Numerical study of flow in a constricted curved annulus: An application to flow in a catheterised artery. *J. Eng. Math.*, 40:355–376, 2001.

[69] G. B. Jeffery. Plane stress and plane strain in bipolar co-ordinates. *Phil. Trans. R. Soc. Lond. A*, 221:265–293, 1920.

[70] G. B. Jeffery. The rotation of two circular cylinders in a viscous fluid. *Proc. R. Soc. Lond. A*, 101:169–174, 1922.

[71] S. W. Jones. Chaotic advection and dispersion. *Physica D*, 76:55–69, 1994.

[72] S. W. Jones, O. M. Thomas, and H. Aref. Chaotic advection by laminar flow in a twisted pipe. *J. Fluid Mech.*, 209:335–357, 1989.

[73] S. W. Jones and W. R. Young. Shear dispersion and anomalous diffusion by chaotic advection. *J. Fluid Mech.*, 280:149–172, 1994.

[74] J. Keener and J. Sneyd. *Mathematical physiology.* Springer–Verlag, New York, 1998.

[75] G. P. King, G. Rowlands, M. Rudman, and A. N. Yannacopoulos. Predicting chaotic dispersion with eulerian symmetry measures: wavy Taylor-vortex flow. *Phys. Fluids*, 13:2522–2528, 2001.

[76] M. E. Klonowska and W. J. Prosnak. On a method of computing the Stokes flow in plane multiply connected domains. *Bulletin of the Polish Academy of Sciences Technical Sciences*, 34:505–518, 1986.

[77] M. A. Knepper and F. C. Rector Jr. Urinary concentration and dilution. In B. M. Brenner and F. C. Rector Jr., editors, *The kidney (volume 1)*, pages 445–482. W. B. Saunders Company, Philadelphia, 4th edition, 1991.

[78] S. D. Kolev and W. E. van der Linden. Laminar dispersion in parallel plate sections of flow systems used in analytical chemistry and chemical engineering. *Anal. Chim. Acta*, 247:51–60, 1991.

[79] S. D. Kolev and W. E. van der Linden. Analysis of transient laminar mass transfer in a parallel-plate dialyser. *Anal. Chim. Acta*, 257:331–342, 1992.

[80] S. D. Kolev and W. E. van der Linden. Influence of the main parameters of a parallel-plate dialyser under laminar flow conditions. *Anal. Chim. Acta*, 257:317–329, 1992.

[81] S. D. Kolev and W. E. van der Linden. Limits of mass transfer in parallel-plate dialysers. *Anal. Chim. Acta*, 256:301–305, 1992.

[82] S. D. Kolev and W. E. van der Linden. Application of Laplace transforms for the solution of transient mass- and heat-transfer problems in flow systems. *Int. J. Heat Mass Transfer*, 36:135–139, 1993.

[83] T. S. Krasnopolskaya, V. V. Meleshko, G. W. M. Peters, and H. E. H. Meijer. Mixing in Stokes flow in an annular wedge cavity. *Eur. J. Mech. B/Fluids*, 18:793–822, 1999.

[84] P. G. M. Kruijt. *Analysis and optimization of laminar mixing.* CIP-Data Library Technische Universiteit, Eindhoven, 2000.

[85] H. A. Kusch and J. M. Ottino. Experiments on mixing in continuous chaotic flows. *J. Fluid Mech.*, 236:319–348, 1992.

[86] H. Lamb. *Hydrodynamics.* Cambridge University Press, Cambridge, 6th edition, 1932.

[87] O. Levenspiel. *Chemical reaction engineering.* John Wiley and Sons, Inc., New York, 1999.

[88] F. H. Ling. Chaotic mixing in a spatially periodic continuous mixer. *Phys. Fluids A*, 5:2147–2160, 1993.

[89] F. H. Ling. The effect of mixing protocol on mixing in discontinuous cavity flows. *Phys. Lett. A*, 177:331–337, 1993.

[90] E. O. Macagno and J. Christensen. Fluid mechanics of the duodenum. *Ann. Rev. Fluid Mech.*, 12:139–158, 1980.

[91] D. A. MacDonald. Fully developed incompressible flow between non-coaxial circular cylinders. *J. Appl. Math. Phys. (ZAMP)*, 33:737–751, 1982.

[92] R. S. MacKay. Complicated dynamics from simple topological hypotheses. *Phil. Trans. R. Soc. Lond. A*, 359:1479–1496, 2001.

[93] V. V. Meleshko and H. Aref. A blinking rotlet model for chaotic advection. *Phys. Fluids*, 8:3215–3217 (Errata in *Phys. Fluids*, 10:1543, 1998), 1996.

[94] V. V. Meleshko, O. S. Galaktionov, G. W. M. Peters, and H. E. H. Meijer. Three-dimensional mixing in Stokes flow: the partitioned pipe mixer problem revisited. *Eur. J. Mech. B/Fluids*, 18:783–792, 1999.

[95] V. V. Meleshko and G. W. M. Peters. Periodic points for two-dimensional Stokes flow in a rectangular cavity. *Phys. Lett. A.*, 216:87–96, 1996.

[96] G. Metcalfe and J. M. Ottino. Autocatalytic processes in mixing flows. *Phys. Rev. Lett.*, 72:2875–2878, 1994.

[97] J. H. Michell. The inversion of plane stress. *Proc. Lond. Math. Soc.*, 34:134–142, 1901.

[98] L. M. Milne-Thomson. *Plane elastic systems*. Springer–Verlag, Berlin, 1960.

[99] L. M. Milne-Thomson. *Theoretical hydrodynamics*. Macmillan and Co., New York, 1968.

[100] Y. Mizuno and M. Funakoshi. Chaotic mixing due to a spatially periodic three-dimensional flow. *Fluid Dyn. Res.*, 31:129–149, 2002.

[101] F. J. Muzzio, M. M. Alvarez, S. Cerbelli, M. Giona, and A. Adrover. The intermaterial area density generated by time- and spatially periodic 2D chaotic flows. *Chem. Eng. Sci.*, 55:1497–1508, 1999.

[102] F. J. Muzzio and M. Liu. Chemical reactions in chaotic flows. *Chem. Eng. J.*, 64:117–127, 1996.

[103] F. J. Muzzio, C. Meneveau, P. D. Swanson, and J. M. Ottino. Scaling and multifractal properties of mixing in chaotic flows. *Phys. Fluids A*, 4:1439–1456, 1992.

[104] F. J. Muzzio and J. M. Ottino. Dynamics of a lamellar system with diffusion and reaction: Scaling analysis and global kinetics. *Phys. Rev. A*, 40:7182–7192, 1989.

[105] F. J. Muzzio and J. M. Ottino. Diffusion and reaction in a lamellar system: Self-similarity with finite rates of reaction. *Phys. Rev. A*, 42:5873–5884, 1990.

[106] F. J. Muzzio and P. D. Swanson. The statistics of stretching and stirring in chaotic flows. *Phys. Fluids A*, 3:822–834, 1991.

[107] E. B. Nauman. *Mixing in continuous flow systems.* John Wiley and Sons, Inc., New York, 1983.

[108] H. R. Neave and P. L. B. Worthington. *Distribution Free Tests*, pages 169–177. Unwin Hyman, Boston, 1988.

[109] F. S. Neto and R. M. Cotta. Improved hybrid lumped-differential formulation for double-pipe heat exchanger analysis. *J. Heat Transfer*, 115:921–927, 1993.

[110] R. J. Nunge and W. N. Gill. Analysis of heat or mass transfer in some countercurrent flows. *Int. J. Heat Mass Transfer*, 8:873–886, 1965.

[111] J. M. Ottino. *The kinematics of mixing: stretching, chaos and transport.* Cambridge University Press, Cambridge, 1989.

[112] J. M. Ottino. The mixing of fluids. *Sci. Amer.*, Jan:40–49, 1989.

[113] J. M. Ottino. Mixing and chemical reactions: a tutorial. *Chem. Eng. Sci.*, 49:4005–4027, 1994.

[114] T. J. Pedley. The interaction between stirring and osmosis. Part 1. *J. Fluid Mech.*, 101:843–861, 1980.

[115] T. J. Pedley. The interaction between stirring and osmosis. Part 2. *J. Fluid Mech.*, 107:281–296, 1980.

[116] T. J. Pedley. Calculation of unstirred layer thickness in membrane transport experiments: a survey. *Q. Rev. Biophys.*, 16:115–150, 1983.

[117] T. J. Pedley, P. Corieri, R. D. Kamm, J. B. Grotberg, P. E. Hydon, and R. C. Schroter. Gas flow and mixing in the airways. *Critical Care Medicine*, 22:S24–S36, 1994.

[118] T. J. Pedley and J. Fischbarg. The development of osmotic flow through an unstirred layer. *J. Theo. Biol.*, 70:427–447, 1978.

[119] C. Pozrikidis. *Introduction to theoretical and computational fluid dynamics*, pages 599–631. Oxford University Press, Oxford, 1997.

[120] S. Prakash and J. L. Kokini. Estimation and prediction of shear rate distribution as a model mixer. *J. Food Eng.*, 44:135–148, 2000.

[121] T. J. Price, T. Mullin, and J. J. Kobine. Numerical and experimental characterisation of a family of two roll mill flows. *Proc. R. Soc. Lond. A*, 459:117–135, 2003.

[122] F. Raynal and J. Gence. Energy saving in chaotic laminar mixing. *Int. J. Heat Mass Transfer*, 40:3267–3273, 1997.

[123] R. L. Ricca and M. A. Berger. Topological ideas and fluid mechanics. *Physics Today*, December:28–34, 1996.

[124] E. P. L. Roberts and M. R. Mackley. The simulation of stretch rates for the quantitative prediction and mapping of mixing within a channel flow. *Chem. Eng. Sci.*, 50:3727–3746, 1995.

[125] A. J. S. Rodrigo and J. P. B. Mota. On the optimization of mixing protocol in a certain class of three-dimensional Stokes flows. *Phys. Fluids*, 15:1505–1516, 2003.

[126] D. Rothstein, E. Henry, and J. P. Gollub. Persistent patterns in transient chaotic fluid mixing. *Nature*, 401:770–772, 1999.

[127] M. Rudman, G. Metcalfe, A. Brydon, and L. Graham. Numerical mixing experiments in the rotated arc mixer (RAM). In *Proceedings of the 6th World Congress of Chemical Engineering, Melbourne, Australia*, 2001.

[128] E. Saatdjian and N. Midoux. Chaotic mixing and heat transfer between confocal ellipses: experimental and numerical results. *Phys. Fluids*, 8:677–691, 1996.

[129] E. Saatdjian, N. Midoux, and J. C. André. On the solution of Stokes' equations between confocal ellipses. *Phys. Fluids*, 6:3833–3846, 1994.

[130] D. R. Sawyers, M. Sen, and H. Chang. Effect of chaotic interfacial stretching on bimolecular chemical reaction in helical-coil reactors. *Chem. Eng. J.*, 64:129–139, 1996.

[131] P. Schoenfeld and M. G. Cogan. *Introduction to dialysis. 2nd ed.* Livingstone, New York, 1991.

[132] S. K. Sen. Circle theorems for steady Stokes flow. *J. Appl. Math. Phys. (ZAMP)*, 40:139–145, 1989.

[133] P. J. Sheskey, R. T. Robb, R. D. Moore, and B. M. Boyce. Effects of lubricant level, method of mixing, and duration of mixing on a controlled-release matrix tablet containing hydroxypropyl methylcellulose. *Drug Dev. Ind. Pharm.*, 21:2151–2165, 1995.

[134] P. N. Shivakumar and C. Ji. On the Poisson's equation for doubly connected regions. *Can. Appl. Math. Q.*, 1:555–567, 1999.

[135] D. U. Silverthorn. *Human physiology. An integrated approach.* Prentice Hall, New Jersey, 1998.

[136] R. Smith. Diffusion in shear flows made easy: the Taylor limit. *J. Fluid Mech.*, 175:201–214, 1987.

[137] W. T. Snyder and G. A. Goldstein. An analysis of fully developed laminar flow in an eccentric annulus. *AIChEJ.*, 11:462–467, 1965.

[138] S. A. Suslov and A. J. Roberts. Advection–dispersion in symmetric field-flow fractionation channels. *J. Math. Chem.*, 26:27–46, 1999.

[139] G. Taylor. Dispersion of soluble matter in solvent flowing slowly through a tube. *Proc. R. Soc. Lond.*, 219A:186–203, 1953.

[140] N. B. Tufillaro, T. Abbot, and J. Reilly. *An experimental approach to nonlinear dynamics and chaos.* Addison Wesley, New York, 1992.

[141] A. Vikhansky. On the applicability of topological chaos to mixer design. Preprint, University of the Negev.

[142] A. Vikhansky. Simulation of topological chaos by immersed boundary method. Preprint, University of the Negev.

[143] A. Vikhansky. Enhancement of laminar mixing by optimal control methods. *Chem. Eng. Sci.*, 57:2719–2725, 2002.

[144] G. Walker and T. Davies. Mass transfer in laminar flow between parallel permeable plates. *AIChEJ.*, 20:881–889, 1974.

[145] G. H. Wannier. A contribution to the hydrodynamics of lubrication. *Q. Appl. Math.*, 8:1–32, 1950.

[146] H. E. De Wardener. *The kidney: an outline of normal and abnormal function.* Churchill, Livingstone, New York, 1985.

[147] E. J. Watson. The rotation of two circular cylinders in a viscous fluid. *Mathematika*, 42:105–126, 1995.

[148] A. A. Woering. *The optimization of mixing of viscous fluids.* Fedodruk B. V., Enschede, 1997.

[149] A. A. Woering, W. C. M. Gorissen, and A. Biesheuvel. Optimization of viscous mixing in a two-dimensional cavity transfer mixer. *Flow, Turbulence and Combustion*, 60:377–407, 1999.

[150] T. Yamanaka. The asymptotic analysis of mutual boundary value problems in parallel plate transport phenomena. *J. Chem. Eng. Japan*, 14:98–104, 1981.

[151] A. N. Yannacopoulos, I. Mezić, G. Rowlands, and G. P. King. Eulerian diagnostics for lagrangian chaos in three-dimensional Navier–Stokes flows. *Phys. Rev. E*, 57:482–490, 1998.

[152] H. M. Yeh, T. W. Chang, and Y. J. Chen. Analysis of dialysis coupled with ultrafiltration in cross-flow membrane modules. *J. Mem. Sci.*, 134:151–162, 1997.

[153] J. M. Zalc and F. J. Muzzio. Parallel-competitive reactions in a two-dimensional chaotic flow. *Chem. Eng. Sci.*, 54:1053–1069, 1999.

[154] D. A. Zumbrunnen and S. Inamdar. Novel sub-micron highly multi-layered polymer films formed by continuous flow chaotic mixing. *Chem. Eng. Sci.*, 56:3893–3897, 2001.