

Durham E-Theses

Dynamics of Surfactant Adsorption at Solid–Liquid Interfaces

DAVID ALEXANDER WOODS

How to cite:

WOODS, DAVID ALEXANDER (2011) Dynamics of Surfactant Adsorption at Solid–Liquid Interfaces. Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/3277/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Dynamics of Surfactant Adsorption at Solid–Liquid Interfaces

*A thesis submitted in partial fulfilment of the requirements for the
degree of Doctor of Philosophy in the University of Durham by,*

David A. Woods



Department of Chemistry
University of Durham

November 2011

Abstract

The adsorption kinetics of surfactants at the solid–liquid interface is of fundamental interest to a wide variety of process including detergency, wetting of solid surfaces, agricultural sprays and paper processing. Accordingly, a significant body of work has been carried out to understand this field. Much of this work has used the optical techniques of ellipsometry and optical reflectometry or mass measurements from the quartz crystal microbalance. These methods have the time resolution to measure surfactant adsorption kinetics but are insensitive to chemical composition and thus produce limited information on the adsorption of surfactant mixtures.

The technique I adopt here, total internal reflection (TIR) Raman spectroscopy, provides detailed information about the chemical composition of the surface with a time resolution of 2s. The short penetration depth of the probe laser into solution (~ 100 nm) provides surface sensitivity. The different components of the adsorbed film are distinguished by their vibrational Raman spectra. The Raman signal from a component in the adsorbed layer is linearly proportional to the amount of that component present, allowing straightforward interpretation of the acquired data. I use principal component analysis to deconvolute the recorded spectra.

First I look at the equilibrium and kinetic aspects of the adsorption of two model surfactants to a flat silica surface as single component systems: the cationic surfactant cetyltrimethylammonium bromide (CTAB) and the non-ionic surfactant Triton X-100. Use of the well-defined wall jet geometry provides known hydrodynamics allowing the mass transport to the surface to be modelled. The mass transport model is coupled with a kinetic model consistent with the Frumkin isotherm allowing the whole adsorption process to be captured. The fit between the model and the experimental results helps to understand interactions on the surface.

Secondly I look at the two model surfactants adsorbing to silica as a mixed system. The adsorption isotherm shows strong synergistic behaviour with the addition of small amounts of CTAB ($\sim 2\%$ of the 2 mM total surfactant concentration) doubling the adsorbed amount of Triton X-100. This synergism has a marked influence on the kinetics: for example, when Triton X-100 replaces CTAB the Triton X-100 surface excess overshoots its equilibrium value and returns only very slowly to equilibrium. For systems above the cmc, the repartitioning of surfactant between micelles and monomers results a local increase in the monomer concentration of Triton X-100 resulting in a temporary spike in the Triton X-100 surface excess during the rinsing of a mixed layer.

Finally I study alternative model surfaces to silica. The adsorption to CTAB and Triton X-100 to a cellulose surface is studied, and detailed equilibrium isotherms obtained by slow variation of the bulk concentration controlled with a continuous stirred tank mixer. The preparation of the model cellulose surface is also followed spectroscopically. Spectra are also acquired from mica surfaces in optical contact with silica hemispheres; it is unfortunately not yet possible to acquire useful data on adsorption at the mica–water interface.

Acknowledgements

Thanks are due to

- my supervisor, Prof. Colin Bain for his support, guidance and (I imagine) patience over the course of the project; he is still wrong about split infinitives though. . . ;
- Unilever, for funding me through a CASE award, and specifically Jordan Petkov and his colleagues at Port Sunlight for their help and support during the project;
- Eric Tyrode for his help when I was getting started, and for taking me to Sweden to do sum-frequency spectroscopy (results to appear elsewhere!);
- the staff of the glassblowing and mechanical workshops in the Chemistry Department at Durham, who made much of the apparatus I used;
- all my colleagues over the years who helped me with miscellaneous lab work and distracted me—or helped me distract myself—with crosswords, tea breaks and barbecues: without their help I would have done much more work over the years!
- Jenny.

Additional thanks to anyone who deserves thanks but has been missed from the list above.

Publications List

The following publications have arisen from work presented in this thesis:

1. “Surfactant adsorption kinetics by total internal reflection Raman spectroscopy. 1. Pure surfactants on silica”, D. A. Woods and J. Petkov and C. D. Bain, *J. Phys. Chem. B*, 2011, **115**, 7341–7352
2. “Surfactant adsorption kinetics by total internal reflection Raman spectroscopy. 2. CTAB and Triton X-100 mixtures on silica”, D. A. Woods and J. Petkov and C. D. Bain, *J. Phys. Chem. B*, 2011, **115**, 7353–7363
3. “Surfactant adsorption by total internal reflection Raman spectroscopy. Part III: adsorption onto cellulose”, D. A. Woods and J. Petkov and C. D. Bain, *Colloids and Surfaces A*, 2011, **391**, 10–18
4. “Total internal reflection Raman spectroscopy”, D. A. Woods and C. D. Bain, *Analyst*, 2011, accepted, DOI: 10.1039/C1AN15722A (review article)

In the interests of avoiding unnecessary reading I will describe how work is duplicated from these papers. Chapter 2 (Experimental) contains work from papers 1–3 above, but goes into more detail than the papers. Chapter 3 (Modelling) duplicated some of the modelling section in paper 1, but goes into considerably more detail with greater validation. Chapter 4 (single component surfactant systems) is substantially taken from paper 1 above. Chapter 5 (multicomponent surfactant systems) is substantially taken from paper 2 above, but contains some additional modelling work. Chapter 6 (on cellulose surfaces) is largely taken from the third of these papers, with a few extra experimental results. The literature review on total internal reflection Raman spectroscopy in Chapter 1 is based on paper 4 above.

Copyright notice

The copyright of this thesis rests with the author. The author is happy to permit short quotations or reproduction of all small number of figures provided that the source is acknowledged. The author is also happy to permit physical or electronic reproduction of the thesis, unmodified and in its entirety, provided that no charge is made for the reproduction. No other quotation from it should be published without prior written consent. Information derived from it should be acknowledged.

Contents overview

1	Introduction	5
2	Experimental Details	71
3	Numerical Modelling	88
4	Single component systems	119
5	CTAB/Triton X-100 multicomponent systems	134
6	Cellulose	160
7	Mica	178
A	Temperature variation of water	193
B	Anionic surfactants	197
C	Additional data for CTAB/TX-100 mixtures	201
D	Verification of mass transport	206
E	Computer Code	211
	Bibliography	237

Contents

1	Introduction	5
1.1	Motivation and aims	5
1.2	Thesis outline	6
1.3	Surfactants and Interfaces	7
1.3.1	Interfaces	7
1.3.2	Surfactants	8
1.3.3	Aggregation	10
1.3.4	Adsorption and interfacial energy	13
1.3.5	Adsorption at silica surfaces	15
1.3.6	Mixed surfactant systems	21
1.3.7	Diffusion	29
1.4	Adsorption kinetics and hydrodynamic control	33
1.4.1	Air–water interface	35
1.4.2	Adsorption onto colloidal gels and fibrous materials	35
1.4.3	Stirred cuvettes	35
1.4.4	Channel flow cells	36
1.4.5	Wall-jet flow cell	36
1.4.6	Dual-inlet channel flow cell	38
1.5	Optics	38
1.5.1	Reflection and refraction at interfaces	40
1.5.2	Total internal reflection	44
1.6	Total internal reflection Raman spectroscopy	46
1.6.1	Vibrational spectroscopy	46
1.6.2	Raman scattering	48
1.6.3	Review of TIR Raman	54
2	Experimental Details	71
2.1	Raman spectroscopy	71
2.1.1	Apparatus description	71
2.1.2	Operating parameters and spectral assignment	74
2.1.3	Polarisation selection	75
2.1.4	Target factor analysis	78
2.2	Wall-jet cell	82
2.3	Continuous stirred mixer	83
2.3.1	Introduction	83

2.3.2	Validation	86
3	Numerical Modelling	88
3.1	Introduction	88
3.2	Convection-diffusion	88
3.2.1	Flow profile	89
3.2.2	1-dimensional solution	95
3.2.3	Analytical solution	95
3.3	Finite difference model	97
3.3.1	Definitions	97
3.3.2	Convection-diffusion implementation	98
3.3.3	Simple boundary conditions	98
3.3.4	Langmuir Isotherm	99
3.3.5	Frumkin Isotherm	100
3.3.6	Micelles	106
3.3.7	Mixed systems	108
3.3.8	Grid parameters	109
3.3.9	Initial concentration profile	114
3.4	Comparison to other models	114
4	Single component systems	119
4.1	Introduction	119
4.1.1	Diffusion coefficients	120
4.2	Equilibrium Isotherms	120
4.2.1	Errors in fitting	121
4.3	Kinetics	122
4.3.1	CTAB	122
4.3.2	TX-100	125
4.3.3	Comparison with other work	130
4.4	Conclusions	131
5	CTAB/Triton X-100 multicomponent systems	134
5.1	Introduction	134
5.2	Experimental	137
5.2.1	Modelling	138
5.3	Results	139
5.3.1	Equilibrium	139
5.3.2	Kinetics of adsorption and desorption	141
5.3.3	Kinetics of displacement of surfactant layers	145
5.4	Conclusions	156
6	Cellulose	160
6.1	Introduction	160
6.2	Experimental	163
6.3	Removal of $-\text{Si}(\text{CH}_3)_3$ from cellulose	165
6.3.1	600–1400 cm^{-1} region	168

6.4	CTAB	170
6.5	TX-100	172
6.6	Mixed systems	174
6.7	Conclusions	176
7	Mica	178
7.1	Introduction	178
7.2	Experimental	179
7.3	CTAB	184
7.3.1	Isotherm	184
7.3.2	Kinetics	186
7.3.3	Conclusions	190
7.4	Other issues	191
7.5	Continuation	191
A	Temperature variation of water	193
B	Anionic surfactants	197
B.1	SDS–TX-100	197
B.2	SDS–lauryl sulfobetaine	197
C	Additional data for CTAB/TX-100 mixtures	201
D	Verification of mass transport	206
D.1	Model	206
D.2	Results	207
D.2.1	Transport to the surface	207
D.2.2	Transport from the surface	207
E	Computer Code	211
E.1	Data analysis	212
E.1.1	Target factor analysis	212
E.1.2	Other data analysis	215
E.2	Modelling	218
E.2.1	Basic model	218
E.2.2	Frumkin isotherm	226
E.2.3	Regular solution theory	231
E.2.4	Visualisation routines	234
	Bibliography	237

Chapter 1

Introduction

1.1 Motivation and aims

The central aim of this project is to understand the kinetics of adsorption at the solid–water interface, with a particular focus on mixed surfactant systems. To do so requires an experiment fulfilling two key criteria: a means of following the composition and concentration of multiple components at the interface and a cell design that allows the transport of surfactant to the interface to be understood. To fulfil the first requirement I will use total internal reflection (TIR) Raman spectroscopy. The total internal reflection aspect provides a means of selecting only those molecules falling within a thin evanescent wave close to the interface. Raman spectroscopy allows different chemical species in a mixed system to be distinguished based on their vibrational spectra. To ensure the second requirement is met, experiments were conducted in the hydrodynamically well-defined wall-jet flow cell, allowing the results obtained to be fitted to a quantitative model.

The kinetics of surfactant adsorption at the solid–liquid interface plays a central role in a wide range of practical applications of surfactants. Examples include the wetting of fabrics in detergency,¹ the spreading of agricultural sprays on leaves,² the adsorption of collectors on ore particles³ and the deposition of conditioners on hair and textiles.⁴ While pure surfactants provide important model systems for the study of the thermodynamics and kinetics of adsorption, practical applications of surfactants invariably involve mixtures for several reasons. First, commercial surfactants are not pure compounds but contain a range of chain lengths or degrees of polymerization, unreacted precursors or hydrolysis products. As an example, small amounts of dodecanol present in sodium dodecyl sulfate drastically change the properties of adsorbed films, especially at concentrations below the critical micelle concentration.^{5,6} Second, mixtures of surfactants often produced enhanced performance over the pure components.⁷ For example, hand-dishwashing detergents are typically mixtures of three or more surfactants. Third, surfactants may be added to formulations for distinct purposes⁸—as wetting agents, dispersants, detergents, foam stabilizers⁹ or friction modifiers.¹⁰ Understanding and controlling the interaction of different surfactants is an essential aid to effective formulation.

Methods for measuring surfactant adsorption largely fall into one of two classes:

those that can measure surface composition, and those that can take measurements on a short timescale; very few techniques can accomplish both these aims. Examples of techniques that provide compositional information are depletion studies, in which a powder is suspended in a surfactant solution and the surfactant remaining in the solution is measured either *in situ* or following removal of an aliquot, and neutron reflection, in which selective deuteration allows surfactants to be distinguished. Techniques that respond rapidly to sub-monolayer changes in the mass of an adsorbed film but lack the chemical selectivity to determine the composition of the adsorbed layer include ellipsometry, reflectometry, surface plasmon resonance and quartz crystal microbalance measurements. Ellipsometry and reflectometry in particular have been extensively used to study adsorption of single surfactants onto hydrophilic and hydrophobic silica surfaces.

Therefore there is a need for techniques which can follow the adsorption of the individual components of mixed surfactant systems on a timescale of seconds. Spectroscopy is a standard way to distinguish chemical species but most common surfactants do not absorb light in the visible or near-UV, nor do they fluoresce, so UV-vis absorption spectroscopy and fluorescence detection are not generally applicable. All surfactants have molecular vibrations, however IR absorption spectroscopy is experimentally challenging due to the large IR cross-section of water compared to the cross-section of surfactant hydrocarbon chains. Although Raman cross-sections are typically around 10 orders of magnitude weaker than IR cross sections, the greater sensitivity to hydrocarbons compared with water makes it viable from studying nanometre-thick layers.

1.2 Thesis outline

This thesis is structured as follows. This chapter—the introduction—will provide the necessary background to understand the remainder of the thesis. The first topic covered is the behaviour of surfactants in solution and at the surface, specifically at the silica surface since much of the work in this thesis concerns the silica–water interface. I also briefly summarise how adsorption kinetics have been measured in the past including the design of flow cells with controlled hydrodynamics. I present an introduction to optics sufficient to explain the total internal reflection aspect of TIR Raman spectroscopy. The final aspect of the introductory chapter explains Raman spectroscopy, includes a review of previous work using TIR Raman spectroscopy.

Chapter 2 details the experimental methods common to the majority of work in the thesis (issues specific to a single chapter only are still described in that chapter). Chapter 3 describes the numerical model used to understand the results acquired in subsequent chapters.

The fourth to seventh chapters present experimental results and associated discussion. Chapter 4 contains work on the adsorption of two surfactants—CTAB and Triton X-100—onto silica as single components together with detailed fitting to a numerical model; chapter 5 looks at mixtures of the same two surfactants adsorbing to silica with some qualitative numerical modelling explaining the results. Chapter 6 looks at the adsorption of the two surfactants to a model cellulose surface instead of a silica surface. Finally chapter 7 records my attempts—largely unsuccessful—to follow

adsorption to a mica–water interface using TIR Raman spectroscopy. It is included in the hope that it could present a starting point for future, more successful, work on mica interfaces.

Concluding the thesis there are five appendices. The first shows the temperature variation of water spectra, recorded to understand a side-project not presented in the thesis but which may be useful for any future studies using temperature variation. The second reports the non-adsorption of mixtures containing anionic surfactants onto silica. The third is a compilation of extra data from chapter 5, omitted from the main text for readability. The fourth attempts to verify the mass transport in the wall-jet flow cell using acetonitrile as a marker of concentration. The final appendix is a compilation of computer code written for the project.

1.3 Surfactants and Interfaces

1.3.1 Interfaces

An interface is a boundary between two immiscible phases. Typically, the term “surface” is used only to refer to a gas–condensed-phase interface while the term “interface” can be used generally for any interface. There is an energy penalty for forming a surface, due to unfulfilled interactions with neighbouring molecules at the edge of the condensed phase. The surface energy (or surface tension) is given by,¹¹

$$\gamma_1 = \frac{1}{2}W_{11} \quad (1.1)$$

where W_{11} is the work of cohesion: the energy to move two surfaces of material 1 from contact to infinite separation. The formation of an interface between two substances (1 and 2) also has a change in free energy, the interfacial energy, defined as,

$$\gamma_{12} = \gamma_1 + \gamma_2 - W_{12}, \quad (1.2)$$

where W_{12} is the work of adhesion, the energy required to move two surfaces of the different materials from contact to being infinitely separated (in a vacuum). The interfacial energy is a thermodynamic property and therefore is defined at chemical and mechanical equilibrium. Both surface and interfacial energies are always positive, otherwise the two substances are miscible. The interfacial energy can be interpreted as the minimum work required to generate a unit area of interface:¹²

$$W = \gamma d\sigma, \quad (1.3)$$

where σ represents the interfacial area. Alternatively, it can be seen as the force required to stretch expand an interface perpendicular to a unit length. For liquid–gas and liquid–liquid interfaces the interfacial energy can be measured directly by measuring the force needed to deform the interface (using a tensiometer). For solid–liquid interfaces the interfacial energy is not measurable directly.

Name	cmc / mM	Structure
Hexadecyltrimethylammonium bromide (CTAB)	0.92	
Sodium dodecyl sulphate (SDS)	8.9	
Triton X-100 (TX-100)	0.27	
N-Dodecyl-N,N-dimethyl-3-ammonio-1-propanesulfonate	3.5	
Alkyl polyethylene oxides (C _n E _m)		

Table 1.1: List of surfactants used in this work, together with critical micelle concentrations (cmc) and structures. The C_nE_m series of surfactants are not used in this work but are shown since they are discussed extensively in the text. TX-100 is a commercial mixture of surfactants similar to the structure illustrated: in practice the C₈H₁₇ alkyl chain could consist of multiple isomers and there is a distribution of chain lengths for the hydrophilic chain (9.5 is the mean number of ethylene glycol groups).

1.3.2 Surfactants

Surfactants are molecules that exist preferentially at interfaces. Aqueous surfactants have two components: a hydrophilic headgroup and a hydrophobic tailgroup (some surfactants may have multiple headgroups or tailgroups). The tailgroup is often a hydrocarbon chain, although other possibilities exist such as fluorocarbon or siloxane chains.¹³ Head-groups are usually classified by their ionic nature. Nonionic headgroups are uncharged (although do contain dipoles or higher order multipoles allowing them to interact with the aqueous solution). Typically examples are polyether chains (–CH₂–CH₂–O–) or chained glucose units. Anionic headgroups have a negative charge, with one of the most common examples being the sulphate group (–OSO₃[–]). Anionic headgroup surfactants are the main surface active component of commercial detergents. Cationic headgroups have a positive charge, with the most commonly used group being the quaternary ammonium salts (e.g. –N(CH₃)₃⁺). Finally zwitterionic headgroups have both a positive and negative charge with betaines and sulfobetaines being the most common examples. The surfactants used in this work are shown in table 1.1.

Surfactants exhibit a number of properties in solution as a consequence of the hydrophobicity of their tailgroup. First, they form aggregates in solution, termed micelles. The thermodynamics of micellisation are discussed later. Second, they adsorb




Figure 1.1: Behaviour of surfactants in solution and at interfaces. 1. monomers in solution, 2. micelles in solution, 3. adsorbed bilayer, 4. adsorbed admicelle, 5. adsorbed hemi-micelle, 6. adsorbed monolayer, and 7. adsorbed monomers.

to interfaces or surfaces. The preferred orientation and structure of the surfactant at the interface is determined by the nature of the interface: at hydrophobic interfaces they attach with the hydrophobic chain facing the interface, either as a monolayer or as a hemi-micelle aggregate; at hydrophilic interfaces they adsorb either as bilayers or as separate aggregated structures (admicelles) with the surfactants immediately adjacent to the interface orientated with their headgroup facing the interface. The air–water surface acts like a hydrophobic interface. The different behaviours of surfactants are illustrated in figure 1.1.

The primary driving force for removing the tailgroup from an aqueous solution is the “hydrophobic effect”. The hydrogen bonding structure of bulk water is disrupted by the non-polar tail. To reduce the loss of hydrogen bonding the water molecules form structures surrounding the tail allowing each water molecule to remain tetrahedrally coordinated. However, this organisation comes at a large entropy cost, which is roughly proportional to the surface area of the dissolved molecule.¹¹ The energetic contribution to dissolving the hydrophobic tails is small (and often favourable), since there is little difference in the dispersion interactions of water–water,

water–hydrocarbon and hydrocarbon–hydrocarbon.

1.3.3 Aggregation

For a system that can form aggregates of different sizes, the chemical potential of an aggregate of size N is,

$$\mu_N = \mu_N^\circ + \frac{kT}{N} \ln \frac{X_N}{N}, \quad (1.4)$$

where μ_N° is the standard chemical potential of monomers in aggregates of size N , X_N is the activity (approximately equal to the concentration, expressed as a mole or volume fraction of the system) of monomers in aggregates of size N and μ_N the chemical potential of a monomer in an aggregate of size N . At equilibrium, μ_N is equal for all N . When $N = 1$, the equation describes a monomer in solution. Rearranging equation 1.4, and using the monomer concentration as a reference, gives,¹¹

$$X_N = N \left[X_1 \exp \left(\frac{\mu_1^\circ - \mu_N^\circ}{kT} \right) \right]^N \quad (1.5)$$

Within a spherical aggregate, the aggregation number, N , is proportional to the volume of the aggregate while the number of monomers at the surface of the aggregate is proportional to $N^{2/3}$. If adjacent monomers within an aggregate interact with an energy of $-\alpha kT$ —where α is a parameter defining the strength of the interaction—and surface monomers have unfulfilled interactions then,

$$\mu_N^\circ = \mu_\infty^\circ + \frac{\alpha kT}{N^{1/3}}, \quad (1.6)$$

where μ_∞° is the chemical potential of a monomer in bulk solution of other monomers of the same compound. Similar expressions exist for alternatively shaped aggregates. Equation 1.6 produces an aggregate size distribution that peaks at infinity, corresponding to phase separation rather than aggregation in solution. Aggregation in solution is governed by the amphiphilic nature of surfactant molecule: the hydrophilic headgroups can be put on the edge of the micelle without a significant energy penalty. The expression for μ_N° arises from multiple different contributions:^{14,15}

- the free energy of chain transfer: the energy of removing a surfactant tail from solution into the aggregate, which is independent of aggregate size and equivalent to the first term in equation 1.6;
- the free energy of surfactant tail deformation: the cost of constraining the hydrocarbon chain so one end is as held at the edge of the micelle by the headgroup;
- the free energy which arises from headgroup–solvent and headgroup–headgroup interactions (the latter can include repulsion where the headgroups are oppositely charged but also specific interactions such as ionic bridging);
- the free energy of headgroup deformation (for “chain” headgroups such as in C_nE_m surfactants) arising from the same constraints as the surfactant tail deformation;




Figure 1.2: Schematic illustration of a typical aggregate size distribution, showing that the majority of the monomers are either part of aggregates in a narrow range around the aggregation number, M , or exist as free monomers ($N = 1$).

- a small contribution from steric repulsions between the headgroups and the chains;
- the interfacial energy from the interface between the hydrophobic chains and water molecules (which can often penetrate into the headgroups) which is closely related to the second term in equation 1.6.

Once all of the contributions to μ_N° have been accounted for the chemical potential of a monomer shows a minimum with respect to aggregate size, giving a size distribution similar to that illustrated in figure 1.2. The size distribution shows two peaks: one at $N = 1$ (monomers), and one at $N = M$ where M is a characteristic micelle size (termed the *aggregation number* and typically of the order 20–200). The latter peak has a width of $\sim\sqrt{M}$ for spherical micelles indicating a polydispersity in the micelle size.

Inserting equation 1.6 into equation 1.5 gives,¹¹

$$X_N = N \left[X_1 \exp \left(\alpha \left[1 - \frac{1}{N^{1/3}} \right] \right) \right]^N \approx N [X_1 \exp(\alpha)]^N. \quad (1.7)$$

The above expression works on the assumption that equation 1.6 is approximately correct, and that knowledge of the actual size distribution is only necessary to be able to set N to an appropriate value. Since X_N cannot exceed 1 (X_N is a fractional concentration), X_1 is inherently limited to a value of $e^{-\alpha}$. This concentration is termed the *critical micelle concentration* (cmc), and marks the point where the formation of aggregates becomes significant. The cmc, X_1^* , is given by,

$$X_1^* \approx \exp \left(-\frac{\mu_1^\circ - \mu_N^\circ}{kT} \right) \approx e^{-\alpha}. \quad (1.8)$$





Figure 1.3: Concentration of monomers in solution (blue lines) and monomers in aggregates (green lines) with respect to monomer concentration. Calculated using equation 1.7 and the requirement that $C_{\text{total}} = X_1 + X_M$. The calculation here assumes that aggregates are monodisperse, with size $M = 10$ (solid lines) or $M = 100$ (dashed lines).

Above the cmc any increase in total concentration will primarily result in an increase in micelle concentration rather than an increase in monomer concentration. The aggregation number of the micelles determines the sharpness of the transition. Figure 1.3 shows the concentration of monomers in solution and monomers in aggregated with respect to concentration, illustrating the effect of aggregation number.

At the cmc sharp changes in slope are seen in many properties of the solution, for example the diffusion coefficients (which affects the conductivity) and the interfacial tension.⁷ These provide a means of measuring the cmc.

In my modelling of surfactant systems (Chapter 3) I have assumed that the kinetics of the formation and breakdown of micelles are fast. This assumption is useful in that it simplifies the modelling, but is not necessarily a complete description. The generally accepted model for the kinetics of micelle growth and breakdown was derived by Aniansson and Wall in the 1970s.^{16–18} They stated that the size of an aggregate changes by one monomer at a time and that—in ionic surfactants—counterions are able to adjust more quickly than surfactant molecules and so do not influence the rate. The kinetics of the formation of micelles is described by a series of equilibria for every aggregation number (indicated here by the subscript g),



The system as a whole is described by a set of partial differential equations known as the Becker-Döring equations,

$$\frac{d[X_1]}{dt} = \sum_{g=1} k_d^g [X_{g+1}] - k_a^g [X_1][X_g] \quad (1.9a)$$

$$\frac{d[X_g]}{dt} = k_{g-1}^a [X_{g-1}][X_1] + k_g^d [X_{g+1}] - k_g^a [X_g][X_1] - k_{g-1}^d [X_g] \quad \text{for } g = 2, 3 \dots \quad (1.9b)$$

When a small perturbation is applied to the system the size distribution responds through two steps. Taking the example of a rapid dilution of the solution, first there is a rapid decrease in the average micelle size as all the micelles lose a few monomers, and second is a slower approach to the true equilibrium condition as some of the micelles break down completely, losing all their monomers. Under some conditions the calculated rate of the second step is unrealistically slow (\sim years), suggesting that the model above is not a complete description.¹⁴

1.3.4 Adsorption and interfacial energy

In general the change in Gibbs free energy (dG) for a system is defined by,

$$dG(T, P, A, n_i) = -SdT + VdP + \gamma dA + \sum_i \mu_i dn_i, \quad (1.10)$$

where A is the interfacial area. At constant temperature and pressure the change in Gibbs free energy is,

$$dG(A, n_i)|_{T,P} = \gamma dA + \sum_i \mu_i dn_i. \quad (1.11)$$

Integrating dG at constant temperature and pressure gives,

$$G = \gamma A + \sum_i \mu_i n_i, \quad (1.12)$$

which has the total differential,

$$dG(\gamma, A, \mu_i, n_i) = Ad\gamma + \gamma dA + \sum_i \mu_i dn_i + \sum_i n_i d\mu_i. \quad (1.13)$$

Comparing equations 1.11 and 1.13 gives the Gibbs-Duhem equation,

$$\sum_i n_i d\mu_i + Ad\gamma = 0. \quad (1.14)$$

The number of moles of each component can be divided onto two parts: n_i^0 , the amount of material that would be present if the concentration was equal to the bulk concentration, changing abruptly at the interface to the bulk concentration on the other side of the interface; and n_i^σ , the surface excess, which is the difference between




Figure 1.4: Schematic of the concentration profiles of solvent (blue) and surfactant (red) at a solid–liquid interface. Solid lines represent actual concentration profiles; Dashed lines represent the assumed concentration profile from n_i^0 , where the bulk concentration continues up to the dividing plane; shaded areas represents the surface excess (the difference between the actual concentration and n_i^0). The Gibbs dividing plane—chosen to give a solvent surface excess of 0—is shown as a thin black line. It is not the same as the physical solid–liquid interface, shown as a thin dotted line.

n_i^0 and the amount of material actually present:

$$n_i = n_i^0 + n_i^\sigma. \quad (1.15)$$

The actual position of the interface—the Gibbs dividing plane—can be chosen arbitrarily but is usually set so the surface excess of solvent is zero: the deficit of solvent on one side of the dividing plane cancels the excess of solvent on the other side. Figure 1.4 illustrates the Gibbs dividing plane and the surface excess. For surfactants the concentration at the surface is usually so much greater than in the bulk that the surface excess is approximately equal to the amount of surfactant at the surface.

Combining the Gibbs-Duhem equation, 1.14, with equation 1.15 gives,

$$\begin{aligned} \sum_i n_i^0 d\mu_i + \sum_i n_i^\sigma d\mu_i + Ad\gamma &= 0 \\ \sum_i n_i^\sigma d\mu_i + Ad\gamma &= 0. \end{aligned} \quad (1.16)$$

The second line uses the Gibbs-Duhem equation in the absence of an interface: $\sum_i n_i d\mu_i = 0$. Defining the surface excess per surface area, $\Gamma_i = n_i^\sigma/A$,

$$d\gamma = - \sum_i \Gamma_i d\mu_i, \quad (1.17)$$

At equilibrium the chemical potential of component i at the surface must be equal to

the chemical potential of i in the bulk, and hence

$$d\mu_i = -RT d \ln a_i, \quad (1.18)$$

where a_i is the activity of component i . It is worth looking at some common specialisations for the Gibbs adsorption isotherm. For a single non-ionic surfactant (given the index 1) in a solvent (given the index 0),

$$d\gamma = -RT(\Gamma_0 d \ln a_0 + \Gamma_1 d \ln a_1) \quad (1.19)$$

whereas for a surfactant that forms a completely dissociated 1:1 electrolyte (components labelled + and -),

$$d\gamma = -RT(\Gamma_0 d \ln a_0 + \Gamma_+ d \ln a_+ + \Gamma_- d \ln a_-). \quad (1.20)$$

The solvent term can be removed for two reasons: firstly the surface excess of solvent is 0 (set by the positioning of the dividing plane). Secondly, in a dilute solution the activity of the solvent is approximately constant, since $\sum_i n_i d\mu_i = 0$ and so when $n_0 \gg n_1, n_2, \dots$ then $|d\mu_0| \ll |d\mu_1|, |d\mu_2|, \dots$. In an ideal dilute solution the activity of the surfactant can be replaced by the concentration relative to 1 molal, c_1/c^\ominus . With these approximations, equation 1.19 becomes,

$$d\gamma = -RT\Gamma_1 d \ln c_1^\ominus \quad \text{or} \quad \Gamma_1 = \frac{-1}{RT} \left(\frac{d\gamma}{d \ln c_1^\ominus} \right). \quad (1.21)$$

For the 1:1 electrolyte we can also say that $\Gamma_+ = \Gamma_- = \Gamma_1$ (to preserve electroneutrality) and that $c_+ = c_- = c_1$, and hence equation 1.20 becomes,

$$d\gamma = -2RT\Gamma_1 d \ln c_1^\ominus \quad \text{or} \quad \Gamma_1 = -\frac{1}{2RT} \left(\frac{d\gamma}{d \ln c_1^\ominus} \right). \quad (1.22)$$

For a completely dissociated 1:1 electrolyte in the presence of an excess of a salt containing the same counterion as the surfactant, the activity of the counterion is constant and the expression for the nonionic, equation 1.21, is obtained.

Equations 1.21 and 1.22 are especially useful at the air–water interface where measurement of the surface energy is experimentally straightforward (and often easier than determining the surface excess). At the solid–liquid interface the equations still apply, but measurement of the surface energy is more difficult, and so usually the surface excess is measured directly.

1.3.5 Adsorption at silica surfaces

A number of recent reviews^{19–23} have covered adsorption of surfactants at the silica–water interface—either as their main focus or as part of a wider review—so I do not plan to present a full overview here. However, a large portion of the results in this thesis are looking at adsorption onto silica, so I will summarise the key features and areas of specific interest to this project.

Silica has a chemical formula SiO_2 , and its surface properties depend on the way






Figure 1.5: Functional groups on a silica surface: (a) silanol groups, including i) geminal, ii) paired and iii) isolated groups; and (b) siloxane groups. The density of silanol groups is greater on hydroxylated silica and less on pyrogenic silica.

that the groups are capped. The nature of the surface chemistry depends on the pH of the adjacent solution, according to the equilibria,



and



The isoelectric point for silica²⁴ is at approximately pH 2, and the silica becomes increasingly negatively charged between pH 6 and 11. Because of the negative charge on silica at neutral pH, cationic surfactant adsorb whereas anionic surfactants tend not to adsorb.

Two types of functional groups exist on the silica surface (see figure 1.5): siloxane groups predominate on the freshly oxidised surfaces (termed “pyrogenic” silica) and silanol groups predominate when the wafer has been soaked in water or basic solution (termed “hydroxylated” silica). Silanol groups are further classified according to their proximity to other silanol groups: where a silicon atom has two silanol groups attached it is termed geminal, silanol groups on adjacent Si atoms can form hydrogen bonds and are termed paired while silanol groups without other groups on adjacent Si atoms are isolated.²¹ Hydroxylated silica has lower surface charge than pyrogenic, since the attached silanol groups on hydroxylated silica are stabilised by hydrogen bonding with adjacent groups whereas the few silanol groups on pyrogenic silica tend not to have neighbouring groups to hydrogen bond with.¹⁹ The nature of the surface groups and any water involved in hydrogen bonding to them has been extensively characterised by IR spectroscopy.²¹ Flat silica surfaces are usually prepared by oxidising silicon wafers (heating the wafer in an oxygen atmosphere). An alternative to using flat silica surfaces is to use colloidal silica—often in conjunction with techniques that measure depletion from solution—which has both the advantage and disadvantage that the surface area to volume is significant: enough is adsorbed to make calorimetric measurements of the heat of adsorption possible but additionally there is enough surface that noticeable changes pH and ionic strength can occur as more surfactant is adsorbed.




Figure 1.6: Sketches of the (a) two-step and (b) four-step models of adsorption onto hydrophilic surfaces. The regions used in the interpretation are numbered, and the cmc and csac (critical surface aggregate concentration) are marked. Based on diagrams in refs. 26 and 28.

Ionic surfactants

Two models have been proposed to explain ionic surfactant adsorption onto silica and other hydrophilic surfaces: the two-step model^{25,26} and the four-region model.²⁷ In the two-step model the adsorption isotherm (surface excess vs. concentration) is plotted using linear scales, revealing two plateaus—the titular steps. Four distinct regions are assigned: the two slopes leading up to the plateaus, and the plateaus themselves (see figure 1.6(a)). The four-region model plots the adsorption isotherm on a log–log scale—which has the advantage of emphasising small features at low surface excess—to reveal three distinct slopes and a final plateau above the cmc (see figure 1.6(b)). Much of the behaviour of surfactants on silica also applies for other metal oxide surfaces, so conclusions can be—and are—drawn from studies on those surfaces too.

The most important interactions promoting surfactant adsorption to silica are electrostatic attraction to the surface and the hydrophobic effect promoting aggregation to already-adsorbed surfactant. At low concentrations the electrostatic attraction dominates with very little interaction between adsorbed molecules. Initially, adsorption is approximately linear with concentration.²⁹ Throughout this region the surface charge follows adsorption closely, indicating that every surfactant molecule that adsorbs displaces an ion.³⁰ At very low surface coverage charge is greater than adsorption indicating that the surfactant is inducing the loss of ions from neighbouring silanols group as well as the group to which it is adsorbed.³⁰ Above a certain point (known as the charge compensation point) adsorbing more surfactant has little influence on the surface charge suggesting the surfactant is not adsorbing directly to the surface. This point is often close to the common intersection point—illustrated in figure 1.7—where addition of electrolyte does not change the surface coverage,^{31,32} above the common intersection point adsorption is increased by electrolyte (closely packed surfactant charges are shielded from each other) while below it adsorption is




Figure 1.7: Isotherms of dodecylpyridinium chloride on a colloidal silica at pH 7 at two different KCl concentrations illustrating the common intersection point, where electrolyte has no effect on adsorption. The two graphs show the same data plotted on log–log and linear–log scales. Reprinted with permission from ref. 32. Copyright 1997 American Chemical Society.

decreased by electrolyte (competitive adsorption of electrolyte ions).

Aggregation on the surface is driven by very similar considerations as aggregation in the bulk, discussed earlier in this chapter. The critical concentration for aggregation at the surface (the c_{scac}) occurs at a lower concentration than in the bulk since there is less entropic penalty: movement is already restricted at the surface. Two distinct forms for aggregation have been proposed: a uniform bilayer or distinct aggregates (which can form a number of different shapes). Evidence from soft contact AFM imaging suggests that cationic surfactants generally form distinct spherical aggregates on silica.^{33,34} Addition of “soft” counterions, such as Br^- , $\text{S}_2\text{O}_3^{2-}$, CS_3^{2-} or HS^- , causes a transition to elongated cylindrical micelles,³⁵ due to stronger counterion association with the surfactant and hence less headgroup–headgroup repulsion allowing less curved structures to form. High surfactant concentrations also promoted the formation of cylindrical aggregates.³⁴ Examples of adsorbed aggregate shapes are shown in figure 1.8. Fluorescence quenching experiments—looking at pyrene contained within aggregated structures—confirms the view that cationic surfactants adsorb to silica as discrete aggregates, with aggregation numbers between 40 and 150.³⁶ Some double chained surfactants form featureless bilayers on silica (and other surfaces),³⁷ in keeping with their formation of large vesicles in solution; however, this behaviour is atypical for surfactant adsorption on silica in general. It is difficult to perform AFM on partially formed aggregates, partly because the AFM experiment relies on repulsion from the headgroup facing into solution¹⁹ and partly because adsorbed aggregates on partially covered interface may be too mobile to be imaged on the time scale of AFM. Therefore, Understanding of behaviour in this concentration range is poorer.

At intermediate concentrations—when the surface aggregation is only partially complete—Somasundaran and coworkers have concluded that surfactants are adsorbed to oppositely charged surfaces in micelle like environments³⁸ with their chain facing into solution in this second region,³⁹ causing the surface to become hydrophobic.²⁸ Recently, Tyrode *et al.* showed that no sum-frequency signal is detectable for CTAB adsorbing to silica at approximately neutral pH.⁴⁰ They concluded that




Figure 1.8: 300 nm AFM images of 10 mM CTAB adsorbed to silica in (a) the absence of added electrolyte and (b) the presence of 300 mM NaBr. The inserts are Fourier transforms of the main images. Reprinted with permission from ref. 35. Copyright 2000 American Chemical Society.

the adsorbed CTAB aggregates must be centrosymmetric at all concentrations. Subsequent work has shown sum-frequency signal,⁴¹ although that was attributed to a strongly charged surface at pH 11 rather than asymmetric aggregates. At low ionic strength (10 mM NaCl) no change in CTAC (the analogue of CTAB with a chloride in place of a bromide counterion) aggregate structure could be inferred as surface excess increased; at high ionic strength (300 mM NaCl) the adsorbed layer became more ordered with increased surface excess.⁴²

The four-region and two-step models can be explained in terms of the effects described in the paragraphs above. The regions in the two models do not overlap perfectly. Atkin *et al.* classified the first region as “the electrostatic concentration span”,¹⁹ and adsorption in the region can be thought of almost entirely as an electrostatic interaction. The second region of the four-region model represents the beginning of hydrophobic interactions. Atkin *et al.* classifies it as “the electrostatic and hydrophobic concentration span”,¹⁹ where both electrostatic and hydrophobic forces are important. The charge compensation point usually falls within this region. Aggregates begin to form within this region, and the number of aggregates increases as concentration increases.²⁹ The third region in the four-region model has been classified as “the hydrophobic concentration span”,¹⁹ where electrostatic forces are either unimportant or acting as a hindrance to adsorption. Within this region the number of surface aggregates no longer increases—since all charged sites on the surface have been filled—but the size of each aggregate increases with concentration.²⁹ The third region in the four-region model corresponds to the third region in the two-step model, with the first plateau of the two-step model happening only when charge neutralisation occurs well before the csac. For surfactants that fit the two-step model there is no point where electrostatic and hydrophobic forces act cooperatively. Above the cmc—the fourth region in both models—no further adsorption takes place, since the




Figure 1.9: Surface excess of CTAB on pyrogenic silica adsorbing from an initial concentration of 0.6 mM, illustrating the “slow adsorption region”. Reprinted with permission from ref. 43. Copyright 2000 American Chemical Society.

chemical potential of the monomers no longer changes significantly with concentration. For mixed systems there can be a drop in surface excess above the cmc as less surface active components are solubilised in micelles.

The kinetics of cationic surfactant adsorption onto silica have also been studied extensively.^{34,43–52} The exact rate of adsorption depends on transport to the surface, but in general adsorption kinetics for the C_n TAB series of surfactants (where n is the chain length) are quick: of the order of seconds to minutes. Atkin and coworkers reported that the “sticking ratio”—the ratio of the initial rate of adsorption compared to the theoretical limiting rate of adsorption—increases slightly with concentration and increases more sharply at the cmc, suggesting cooperative adsorption.¹⁹ Addition of electrolyte increases the rate of adsorption. One notable effect is the “slow adsorption region”, typically at around 60% of the cmc depending on the surfactant. In this region there is a rapid initial adsorption to about 1/3 of the final coverage and then a slow approach to the equilibrium surface coverage taking up to 24 h (see figure 1.9).^{43,44} The “slow adsorption region” can be elegantly demonstrated through a stepwise increase and decrease in surfactant concentration showing hysteresis in the values reached.^{19,53} The slow adsorption is attributed to slow equilibration of adsorbed aggregates, possibly due to their immobility on the surface.⁵³ The effect may be related to the incomplete desorption of CTAB from silica for adsorbed layers equilibrated for a long time, compared with complete desorption for adsorbed layers that are rinsed only a short time after formation.⁵⁴ Another kinetic effect relates to the exchange of surfactant monomers in adsorbed layers: Ducker and coworkers showed that exchange of a CTAB layer for a deuterated analogue was rapid and complete,⁴⁷ but exchange of a similar double chained surfactant was partial with only the top layer of the adsorbed bilayer being exchanged.⁵¹

Non-ionic surfactants

Non-ionic surfactants have been studied in less detail than ionic surfactants, since their properties are much more dependent on the exact method of preparation of the silica surface and so results obtained are less generally useful. Often there is




Figure 1.10: Kinetics of adsorption and desorption of the nonionic surfactant $C_{16}E_6$ to silica, illustrating the slow desorption kinetics. Reprinted with permission from ref. 63. Copyright 1998 American Chemical Society.

a small amount of adsorption at low concentrations driven by hydrogen bonding between the surfactant headgroup and the surface.²³ The majority of adsorption is hydrophobically driven and isotherms are typically quite sharp. Imaging using AFM reveals that, above the csac, the surfactants form spherical or elongated aggregates similar to those formed by ionic surfactants.^{55,56} The presence of surface aggregates is supported by analysis of the fluorescence spectra of pyrene embedded in the surface layer.⁵⁷

The kinetics of adsorption of non-ionic surfactants has also been studied, primarily by Tiberg and coworkers,⁵⁸⁻⁶¹ but also by other groups.^{48,56,62} An example of the kinetics of nonionic surfactant adsorption and desorption is shown in figure 1.10. Both the adsorption and (especially) the desorption kinetics can be slow (~ 100 s),^{56,62} particularly for surfactants with longer hydrophobic chains. Postmus *et al.* explained the slow desorption kinetics in terms of the low csac of many nonionic surfactants: during desorption the subsurface concentration is less than or equal to the csac and hence the concentration gradient is shallow meaning diffusion away from the surface is slow.⁶² Brinck *et al.* prepared a thorough model of the adsorption kinetics of single surfactant⁶³ and surfactant mixtures⁶⁴ adsorbing to silica. The single surfactant mixtures are primarily limited by transport to the surface: partly because adsorption does not start until just below the cmc so there is only a small concentration gradient of monomers (there is a larger concentration gradient of micelles, but micelle diffusion is slower), and partly because the experimental set-up used produces a thick diffusion layer. The surfactant mixture kinetics display interesting effects relating to transport and micelle composition of mixed systems but they are not specific to silica surfaces and are best discussed in chapter 5.

1.3.6 Mixed surfactant systems

This section does not aim to present a full review of work on mixed surfactant systems since the literature is extensive and not all relevant. The literature relevant to the specific systems I have looked at is presented in the introduction to chapter 5. In-

Symbol	Meaning
i, j or k	Component in a mixed system
x_i	Mole fraction of i in the mixed micelles
μ_i	Chemical potential of i
μ_i^\ominus	Standard chemical potential i ; the standard states are described in the text
f_i	Activity coefficient of i
f'_i	Activity coefficient of i , specific to liquid-like mixing with the micelle
c_i	Monomer concentration of i
c^\ominus	The standard concentration (1 molal = 1 mol Kg ⁻¹)
a_i	Activity coefficient of i
N_i	Number of monomers of component i in a micelle
N	Total number of monomers in a micelle

Table 1.2: Symbols used in thermodynamics derivations. This table only shows symbols that are common to more than one thermodynamic model. Those symbols that are specific to an individual model are presented in a shorter table at the start of that section. Subscripts of “mon” and “inmic” and “mic” refer to monomers in solution and monomers in micelles, respectively.

stead this section presents an overview of the basic thermodynamic models for mixed systems, some of the results of which are used in chapter 3.

This section incorporates derivations from a variety of different sources, many of which use different symbols to indicate the same property. Therefore, I have adopted a single set of symbols, which are detailed in table 1.2. Symbols specific to the phase separation model and regular solution theory are introduced later.

The chemical potential of a monomer in solution is given by,

$$\begin{aligned}\mu_{\text{mon},i} &= \mu_{\text{mon},i}^\ominus + RT \ln a_{\text{mon},i} \\ &= \mu_{\text{mon},i}^\ominus + RT \ln \left(f_{\text{mon},i} \frac{c_{\text{mon},i}}{c^\ominus} \right).\end{aligned}\tag{1.23}$$

Expressing activity as related to the monomer concentration makes the assumption that the concentration of monomer is much less than the concentration of solvent, and so the volume fraction of monomer is approximately proportional to the concentration of monomer. According to Henry’s law, $f_i \rightarrow 0$ as $c_{\text{mon},i} \rightarrow 0$. In an ideal dilute solution, $a_{\text{mon},i} = c_{\text{mon},i}/c^\ominus$, and so

$$\mu_{\text{mon},i} = \mu_{\text{mon},i}^\ominus + RT \ln \frac{c_{\text{mon},i}}{c^\ominus}.\tag{1.24}$$

The chemical potential of a monomer of i in a micelle with N monomers in total and a fixed composition of $\mathbf{x} = (x_i, x_j, x_k, \dots)$ is,

$$\begin{aligned}\mu_{\text{inmic},i}(\mathbf{x}) &= \mu_{\text{inmic},i}^\ominus(\mathbf{x}) + RT \ln a_{\text{inmic},i}(\mathbf{x}) \\ &= \mu_{\text{inmic},i}^\ominus(\mathbf{x}) + \frac{RT}{N} \ln \left(f_{\text{inmic},i}(\mathbf{x}) \frac{c_{\text{inmic},i}}{c^\ominus} \right)\end{aligned}\tag{1.25}$$

where the chemical potentials and activity (and thus the activity coefficients) are functions of the composition. The activity coefficient being a function of composition makes sense: as a micelle incorporates more of an ionic component and becomes more

charged it will repel other micelles with similar charge more strongly (in this case $f_{\text{inmic},i} > 1$). In a dilute solution $f_{\text{inmic},i} = 1$, and so,

$$\mu_{\text{inmic},i}(\mathbf{x}) = \mu_{\text{inmic},i}^{\ominus}(\mathbf{x}) + \frac{RT}{N} \ln \frac{c_{\text{inmic},i}}{c^{\ominus}}. \quad (1.26)$$

For non-ionic surfactants the micelles are not expected to be strongly interacting, so there is no reason why the Henry's law approximation should not apply. For ionic surfactants the micelles will be charged and repel each other, thus having a larger effective volume than the concentration dependence should imply. Surface tension measurements of cationic surfactants above the cmc show little variation over an order of magnitude of concentration range (ref. 65 illustrates this for various C_n TABs), showing that the concentration dependence of chemical potential is small and Henry's law can be applied.

The next issue to be addressed is the composition variation of the chemical potential of monomers in micelles. The centre of a micelle is similar to a liquid mixture made up of the various surfactants involved, and hence the form of the chemical potential can be assumed to be similar to that obtained from liquid-liquid mixing. Using $\mu_{\text{inmic},i}^{\ominus}$ as shorthand for the pure micelle case, $\mu_{\text{inmic},i}^{\ominus}(x_i = 1)$,

$$\mu_{\text{inmic},i}^{\ominus}(\mathbf{x}) = \mu_{\text{inmic},i}^{\ominus} + RT \ln f'_{\text{inmic},i} x_i, \quad (1.27)$$

with f' used instead of f to indicate that the activity coefficient is different from the one in equation 1.25. Therefore, equation 1.26 becomes,

$$\mu_{\text{inmic},i} = \mu_{\text{inmic},i}^{\ominus} + RT \ln f'_{\text{inmic},i} x_i + \frac{RT}{N} \ln \frac{c_{\text{inmic},i}}{c^{\ominus}}, \quad (1.28)$$

where $\mu_{\text{inmic},i} \equiv \mu_{\text{inmic},i}(\mathbf{x})$ (I am not longer explicitly writing the dependence on the composition since it is now contained within the equation).

The definition of the standard states for both monomers and micelles is hypothetical, since most surfactants form micelles long before $c_{\text{mon},i} = c^{\ominus}$ and $c_{\text{inmic},i} = c^{\ominus}$ represents a very concentrated micellar solution which would usually be accompanied by significant changes in structure compared to the micelles formed just above the cmc. Therefore, both standard chemical potentials are defined in terms of a solution approaching zero concentration:

$$\mu_{\text{mon},i}^{\ominus} = \lim_{c_{\text{mon},i} \rightarrow 0} \mu_{\text{mon},i} - RT \ln \frac{c_{\text{mon},i}}{c^{\ominus}} \quad (1.29a)$$

$$\mu_{\text{inmic},i}^{\ominus} = \lim_{c_{\text{inmic},i} \rightarrow 0} \mu_{\text{inmic},i} - \frac{RT}{N} \ln \frac{c_{\text{inmic},i}}{c^{\ominus}} \quad \text{when } x_i = 1. \quad (1.29b)$$

The above definitions are using the Henry's law approximation of $f_{\text{mon},i} = f_{\text{inmic},i} = 1$. Note also that the equation 1.29b only applies apply for a specific aggregation number, N ; there will be a different standard chemical potential for every aggregation number (in principle the symbol $\mu_{\text{inmic},i}^{\ominus}$ should also be labelled with the aggregation number to denote this, but I have not done so since I feel it is overly labelled already!).

Having established reasonably general expressions for the chemical potentials of mixed micelles, I will now look at the simplified expressions that are often used.

Phase separation model

Symbol	Meaning
α_i	Mole fraction of i in the whole mixed solution, excluding water
c_{total}	Total concentration of all surfactants in the mixed solution (both in monomers and micelles)
c^*	Cmc of mixed solution
c_i^*	Cmc of pure i

Table 1.3: Symbols used in the phase separation model. See also table 1.2.

The phase separation model is one of the simplest methods of modelling the effect of mixed micelles, so named because it treats the micellisation process as a separation into two bulk phases. The following derivation is largely based on work by J. Clint⁶⁶ and Holland and Rubingh;⁶⁷ the same result was also obtained by Lange and Beck,⁶⁸ but their work is not readily available in English. Additional symbols are defined in table 1.3.

Using a rearranged form of equation 1.8 defines the standard chemical potential of a monomer in a pure micelle relative to the cmc of that component,

$$\mu_{\text{inmic},i}^{\ominus} = \mu_{\text{mon},i}^{\ominus} + RT \ln \frac{c_i^*}{c^{\ominus}}. \quad (1.30)$$

Implicit in the treatment of micelles as a bulk phase is the approximation that $N = \infty$, and so the translational entropy term ($RT/N \ln c_{\text{inmic}}/c^{\ominus}$) of the micelle chemical potentials is neglected. Because of this equation 1.29b cannot be used to define the standard chemical potential of monomers in micelles, so equation 1.30 defines it using the cmc instead.

Monomers in solution are assumed to behave as an ideal dilute solution, so equation 1.23 applies. Mixing within the mixed micelle is assumed to be ideal ($f'_{\text{inmic},i} = 1$), so equation 1.28 becomes,

$$\mu_{\text{inmic},i} = \mu_{\text{inmic},i}^{\ominus} + RT \ln x_i, \quad (1.31)$$

with the translational entropy removed since the $1/N$ term is zero with $N = \infty$.

By definition, x_i is,

$$x_i = \frac{\alpha_i c_{\text{total}} - c_{\text{mon},i}}{c - \sum_j c_{\text{mon},j}}. \quad (1.32)$$

At equilibrium $\mu_{M,i} = \mu_i$. Using this equality to combine equations 1.23, 1.31 and 1.30 gives

$$\ln x_i = \ln \left(\frac{c_{\text{mon},i}}{c_i^*} \right), \quad (1.33)$$

and hence

$$c_{\text{mon},i} = x_i c_i^*. \quad (1.34)$$

At the mixed cmc, c^* ,

$$c_{\text{mon},i} = \alpha_i c^*. \quad (1.35)$$

Since $\sum_j x_j = 1$, equations 1.34 and 1.35 can be combined to give

$$\frac{1}{c^*} = \sum_{i=1}^N \frac{\alpha_i}{c_i^*} \quad (1.36)$$

Because this method assumes ideal mixing within the micelle, it will only produce sensible results for systems of very similar surfactants, for example mixtures of C_nE_m surfactants. Ionic surfactants which can dissociate into ions and counterions complicate the description.⁶⁹ In an ionic surfactant $c_{\text{mon},i}$ should be replaced by $c_{\text{mon},i+}c_{\text{mon},i-}$ where $c_{\text{mon},i+}$ and $c_{\text{mon},i-}$ are the concentrations as monomers of the positive and negative ions of the surfactant, respectively. In the absence of electrolyte $c_{\text{mon},i} = (c_{\text{mon},i+})^2$; in the presence of electrolyte the counterion concentration will depend on both the concentration of the ionic surfactant and the concentration of electrolyte. The overall expression depends on the degree of ionic dissociation within the micelle: if each component generates r_i particles (separate ions or molecules) within the micelle then

$$x_i = \left(\frac{c_{\text{mon},i}}{c_i^*} \right)^{2/r_i}. \quad (1.37)$$

In this case x_i is the mole fraction of particles generated by component i within the micelle. For a surfactant with two monovalent ions r_i can vary between 1 and 2; there is no reason to presuppose it will be one of the limiting values.

Regular solution theory

Symbol	Meaning
x_i^σ	Mole fraction of i in a mixed monolayer
β_{ij}	Interaction parameter between components i and j (for nonionic micelles)
β_{ij}^σ	Interaction parameter between components i and j (for a monolayer)
c^s	Concentration of mixed surfactant necessary to produce a given surface tension
c_i^s	Concentration of i necessary to produce a given surface tension
Z	The number of nearest neighbours of a surfactant molecule in a micelle
W_{ij}	Pairwise interaction energy between surfactant molecules in a micelle (including $i = j$)

Table 1.4: Table of symbols used for regular solution theory. See also tables 1.2 and 1.3.

Regular solution theory extends the phase separation model to account for non-ideal mixing. Although the use of one parameter, termed β , does not satisfactorily describe all systems, the theory has been heavily used to explain experimental studies of mixed surfactant systems, and has also formed the basis for many of the more complex theoretical extensions.⁷⁰ The derivations within this section are largely based on the work of Holland and Rubingh^{67,71}. Table 1.4 defines extra symbols used in this section.

The key change from the phase separation model is that ideal mixing within the mixed micelle is no longer assumed. An activity coefficient is reintroduced into

equation 1.31 to give

$$\mu_{\text{inmic},i} = \mu_{\text{inmic},i}^{\ominus} + RT \ln f'_{\text{inmic},i} x_i. \quad (1.38)$$

This leads to equation 1.36 becoming

$$\frac{1}{c^*} = \sum_{i=1}^N \frac{\alpha_i}{f'_{\text{inmic},i} c_i^*}. \quad (1.39)$$

The activity coefficients can be approximated using a regular solution approximation, as

$$\ln f'_{\text{inmic},i} = \sum_{\substack{j=1 \\ (j \neq i)}}^n \beta_{ij} x_j^2 + \sum_{\substack{j=1 \\ (i \neq j \neq k)}}^n \sum_{k=1}^{j-1} (\beta_{ij} + \beta_{ik} - \beta_{jk}) x_j x_k. \quad (1.40)$$

This is the leading order non-zero term of a power series expansion for $\ln f'_{\text{inmic},i}$.

In order to attribute physical meaning to the interaction parameter, β_{ij} is usually thought of in terms of the interaction enthalpies of the two species,⁷²

$$\beta_{ij} = \frac{Z(W_{ii} + W_{jj} - 2W_{ij})}{RT}. \quad (1.41)$$

This expression is purely enthalpic, meaning that micelle composition does not change with temperature. It was recognised early on that in the many cases this does not hold, and that calorimetric measurements suggest an entropic contribution.^{73,74} However—provided the system is held at constant temperature—it is often still possible to fit results to regular solution theory even when the interaction is not purely enthalpic in nature.⁷⁵

Figure 1.11 shows how some properties of the system are affected by the value of β_{ij} . The sign of β_{ij} provides information about the nature of the interaction between two surfactants. A negative value of β_{ij} implies an attractive interaction, meaning a more stable system. The vast majority of β values are negative; however there are some counterexamples, such as some fluorinated surfactants interacting with conventional hydrocarbon surfactants.⁷⁶

For a two component system, the formulae for $f'_{\text{inmic},1}$ and $f'_{\text{inmic},2}$ can be simplified to

$$\begin{aligned} \ln f'_{\text{inmic},1} &= \beta_{12}(1 - x_1)^2, \\ \ln f'_{\text{inmic},2} &= \beta_{12}x_1^2. \end{aligned} \quad (1.42)$$


In order to obtain a value of β_{12} for a binary system, the equation

$$x_1^2 \ln \left(\frac{\alpha_1 c^*}{x_1 c_1^*} \right) = (1 - x_1)^2 \ln \left(\frac{(1 - \alpha_1) c^*}{(1 - x_1) c_2^*} \right) \quad (1.43)$$


must be solved iteratively for x_1 , which is then used in the equation

$$\beta_{12} = \frac{\ln \left(\frac{\alpha_1 c^*}{x_1 c_1^*} \right)}{(1 - x_1)^2}. \quad (1.44)$$

Equation 1.44 comes from the relation $c_{\text{mon},i} = x_i f'_{\text{inmic},i} c_i^*$ (the regular solution



(a)



(b)

Figure 1.11: (a) The micelle composition at the cmc and (b) the cmc, for an ideal mixed surfactant system and one fitting regular solution theory with a parameter β of -2 .

theory equivalent of equation 1.34), equation 1.35, and equation 1.42. It only applies at the cmc, where as further above the cmc the micelle composition tends increasingly towards the bulk composition.

Rosen and Hua derived an analogous set of equations for adsorption onto a monolayer.^{76,77} For a two component system, x_1^σ is obtained from the equation

$$(x_1^\sigma)^2 \ln \left(\frac{\alpha_1 c^s}{x_1^\sigma c_1^s} \right) = (1 - x_1^\sigma)^2 \ln \left(\frac{(1 - \alpha_1) c^s}{(1 - x_1^\sigma) c_2^s} \right), \quad (1.45)$$

which is then used to calculate β_{ij}^σ from


$$\beta_{ij}^\sigma = \frac{\ln \left(\frac{\alpha_1 c^s}{x_1^\sigma c_1^s} \right)}{(1 - x_1^\sigma)^2}. \quad (1.46)$$

When using these equations, c^s , c_1^s and c_2^s are the concentrations of mixed surfactant, pure component 1 and pure component 2 necessary to produce a given surface tension. The actual surface tension chosen should be taken at a point when the surfactant molecules in the monolayer are closely packed, hence with as low a surface tension as possible.⁷⁶


There are a number of caveats when using regular solution theory. Most importantly, for ionic surfactants it needs adjustment in the same way as was done for the ideal mixing case.⁶⁹ Much of the work on with ionic surfactants does not make this adjustment and uses β_{12} as a fitting parameter, even though the parameter can no longer be directly related to a thermodynamic model. I also take the approach of using the non-adjusted equations (1.43 and 1.44) in my model since the value of β_{12} I am using is from a fit to those equations, and I am only trying to model an attractive interaction rather than an exact functional form. A secondary consideration is that the theory assumes that the micelles and monolayers contain no water (although water could be included as a third component). Therefore, it is most useful to obtain β_{ij}^σ at a high surface excess, where the surface is likely to be entirely covered with surfactant.⁷⁶

A more general failing of the regular solution theory approach is that it cannot account for asymmetry in excess free energy with respect to micelle composition that often occurs in real strongly interacting systems.⁷⁰ Causes of the asymmetry include changes in micelle structure with composition (for example rod-like to spherical), non-random packing within mixed micelles (RST assumes a mean field approach)⁷⁵ and differing sizes of headgroups (so that as the composition changes the number of nearest neighbour interactions varies).⁷⁸ I do not intend to describe more advanced models in any detail, since there are a great variety of them and the modelling of mixed systems presented later in the thesis is only intended to be qualitative so greater detail is not necessary.

Regular solution theory is useful in the sense that it can often predict behaviour within an entire range of mixture compositions (see, for example, reference 79) from the cmcs of the pure components and one mixture cmc. However, it is not a predictive theory in the sense that it is based on fitting parameters derived from data. Additionally, for strongly interacting solutions (such as SDS and *n*-dodecyl-



(a) Self diffusion; “tagged” molecules spread throughout a homogeneous solution.



(b) Mutual diffusion; an inhomogeneous solution becomes more homogeneous with time.

Figure 1.12: The two forms of diffusion.

N,N-dimethylamino betaine and many other examples) a single value of β_{12} is often not sufficient to describe the whole concentration range.⁸⁰

1.3.7 Diffusion

In order to study surfactant adsorption kinetics it is very important to understand the mass transport of surfactants, since without knowing how fast surfactants are able to reach a surface any information on how fast they adsorb to the surface is largely useless. Therefore, when studying the kinetics of surfactant adsorption, it is useful to understand the diffusion within those systems.

Self- and mutual-diffusion

Two different diffusion coefficients can be defined: the self diffusion coefficient and the mutual diffusion coefficient.⁸¹ These two possibilities are shown in figure 1.12.

The self diffusion coefficient describes the rate at which an individual object moves through a homogeneous solution. For a sphere of radius r , the self diffusion coefficient, D_s , is given by the Stokes-Einstein equation,

$$D_s = \frac{k_B T}{6\pi\eta r} \quad (1.47)$$

where η is the viscosity of the medium. For a spherical micelle it is often not possible to use the actual micelle radius in the Stokes-Einstein equation since electrostatic

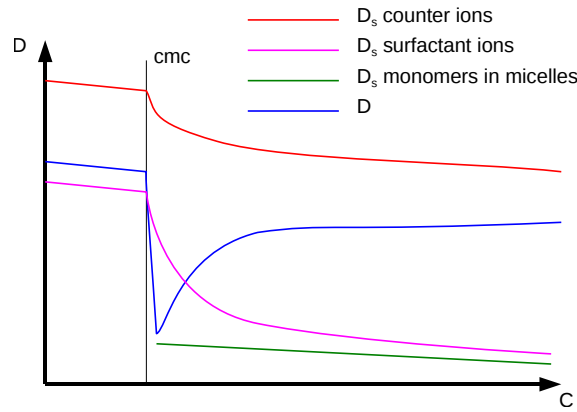


Figure 1.13: Self diffusion coefficients, D_s , for different components of an ionic surfactant and mutual diffusion coefficient, D , for the surfactant as a whole.

interactions with other micelles or the surrounding solution may cause drag, giving an effective radius that is larger than the physical radius. For the case where exchange between monomers and micelles is rapid, the observed self diffusion coefficient for a surfactant is an average of the self diffusion coefficients for monomers in micelles and monomers in solution;⁸¹

$$D_s = (1 - \beta)D_s^{\text{mon}} + \beta D_s^{\text{inmic}}, \quad (1.48)$$

where β is the fraction of the surfactant monomers that exist as micelles.

The mutual diffusion coefficient describes how diffusion acts to reduce a chemical potential gradient in the system, usually by moving molecules from a region with high concentration to a region with low concentration. The overall flux of molecules due to diffusion is

$$J = -\frac{Dc}{RT} \frac{\partial \mu}{\partial x}. \quad (1.49)$$

At low concentrations—where activity is equal to concentration—this reduces to $J = -D \frac{\partial c}{\partial x}$. Because of the requirement for the solution to have an overall neutral change, the mutual diffusion coefficient of the surfactant is heavily coupled to the self diffusion coefficient of the surfactant counterions (in an ionic surfactant). Counterions usually diffuse significantly faster than micelles—due to their smaller size—and so can “pull along” the oppositely charged micelles with them, accelerating the rate of micelle diffusion. The coupling of self diffusion coefficients for monomers, micelles and counterions into a mutual diffusion coefficient for a surfactant can be quite complicated; for example Leaist obtained the follow equation for the mutual diffusion coefficient of an ionic surfactant in a solution with no added salt:^{81,82}

$$D = \frac{c^{\text{mon}} c^{\text{ion}} D_s^{\text{mon}} D_s^{\text{ion}} + (n - q)^2 c^{\text{mon}} c^{\text{inmic}} D_s^{\text{mon}} D_s^{\text{inmic}} + n^2 c^{\text{ion}} c^{\text{inmic}} D_s^{\text{ion}} D_s^{\text{inmic}}}{c^{\text{mon}} D_s^{\text{mon}} + c^{\text{ion}} D_s^{\text{ion}} + q^2 c^{\text{inmic}} D_s^{\text{inmic}}} \times \frac{c^{\text{mon}} + c^{\text{ion}} + q^2 c^{\text{inmic}}}{c^{\text{mon}} c^{\text{ion}} + (n - q)^2 c^{\text{mon}} c^{\text{inmic}} + n^2 c^{\text{ion}} c^{\text{inmic}}}, \quad (1.50)$$

where c and D_s refer to the concentration and self diffusion coefficient of the appropriate component, q to the charge of the micelle and n to the aggregation number. In solution with sufficiently high ionic strength the effect of the counterions is masked, and so the mutual diffusion coefficient becomes equal to the self diffusion coefficient.

Just above the cmc there is a sharp drop in the mutual diffusion coefficient of the surfactant. When the concentration is above the cmc, diffusion happens primarily through micelles. At the point just above the cmc the monomer concentration is still relatively high compared to the total concentration and tends to outweigh the counterion concentration that would normally accelerate micellar diffusion. As the surfactant concentration increases the monomer concentration becomes smaller relative to the counterion concentration, so the counterions are able to couple with the micelles to a much greater effect, leading to an increase in mutual diffusion coefficient.

An alternative way of looking at the sharp drop in the mutual diffusion coefficient of a surfactant near the cmc is in terms of the chemical potential gradient—the driving force for diffusion (equation 1.49).⁸³ As binary solutions approach the point of phase separation the chemical potentials of the two components become equal and hence the chemical potential gradients reduce to zero, which in turn reduces diffusion to zero.⁸⁴ Although the formation of micelles can be modelled as a phase separation, micelles are not a true separate phase, so the diffusion coefficient is not reduced all the way to zero.⁸³

Diffusion in mixed systems

For a solution containing N components, the flux within the solution is given by the coupled Fick equations⁸⁵

$$J_i = - \sum_{k=1}^N D_{ik} \nabla c_k \quad (i = 1 \dots N), \quad (1.51)$$

where c_k is the concentration of component k . Where $i = k$, D_{ik} is the main diffusion coefficient whereas, for $i \neq k$, D_{ik} is a cross-term diffusion coefficient.⁸¹ Since the Fick equations deal with uncharged species there is no contribution from electric field terms and so they are accurate at all concentrations.⁸⁶

The overall mass transport is described by⁸⁵

$$\frac{\partial c_i}{\partial t} = \sum_{k=1}^N \nabla (D_{ik} \nabla c_k) \quad (i = 1 \dots N), \quad (1.52)$$

or, in situations where the diffusion coefficients can be treated as constants,

$$\frac{\partial c_i}{\partial t} = \sum_{k=1}^N D_{ik} \nabla^2 c_k \quad (i = 1 \dots N). \quad (1.53)$$

Successful use of these equations requires either that the various diffusion coefficients have been measured across the target concentration and composition range, or some method of predicting the diffusion coefficients.

Even for mixed systems of ionic surfactants where the only difference between the

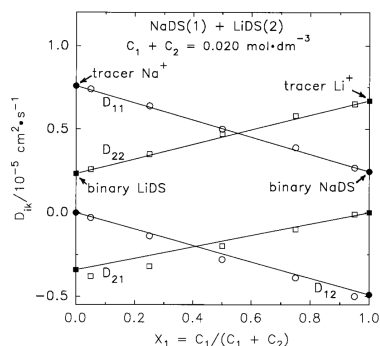


Figure 1.14: Mutual diffusion coefficients of SDS (component 1) and LiDS (component 2) at 25°C and 0.020 mol·dm⁻³ total concentration. Reprinted with permission from ref. 87. Copyright 2001 American Chemical Society.




Figure 1.15: Diffusion coefficients for the SDS-DSB system (see text) at 20 mM total surfactant concentrations, where SDS is component 1 and DSB component 2, measured using the Taylor dispersion method. Reprinted with permission from ref. 83. Copyright 2002 American Chemical Society.

two surfactants is the counter-ion—for example sodium and lithium dodecyl sulphate (SDS and LiDS)⁸⁷—behaviour is not always simple. As shown in figure 1.14, the total main diffusion coefficients (D_{11} and D_{22}) decrease as the mole fraction of the surfactant with which they are associated increases. At zero mole fraction the main diffusion is simply that of a tiny quantity of counterion diffusing in a solution of the other surfactant. In contrast, the coupled diffusion coefficients (D_{12} and D_{21}) increase (become less negative) as the mole fraction of surfactant with which they are associated increases, becoming zero as the mole fraction becomes zero, since it does not make sense for a gradient of SDS to drive diffusion of LiDS where there is no LiDS present (for example).⁸⁷

For mixed systems of surfactants that differ in the surface active ion the results can be more complicated. For example, in a mixture of SDS and dodecylsulphobetaine (DSB) above the cmc—summarised in figure 1.15—the cross term diffusion coefficients are larger than the main diffusion coefficients at several different compositions.⁸³ Therefore, for some compositions, an SDS concentration gradient can cause a large flux of DSB towards an area with a low concentration of SDS. At other compositions, the D_{12} coefficient (where SDS is component 1 and DSB is component 2) is negative, and hence SDS tends to flow towards areas rich in DSB. Additionally, between at

mole fractions of between 55% and 75% SDS the main diffusion coefficient of D_{22} is actually negative causing DSB to flow from lower to higher concentrations.⁸³

The two examples of mixed systems shown illustrate that modelling the diffusion of mixed systems is a very challenging problem. These effects are compounded by changes in the size and the shape of the mixed micelles with both concentration and mole fraction⁸⁸ (see equation 1.47). While the diffusion of mixed systems is important in the adsorption behaviour of mixed surfactant systems (as will be seen in Chapter 5) a full quantitative model is a serious undertaking, even without considering the surface behaviour. I have therefore taken the approach that it is best to model the diffusion of mixed surfactants as simply as possible, but to be aware that the reality is much more complicated.

1.4 Adsorption kinetics and hydrodynamic control

There are a number of processes that take part during the overall surfactant adsorption process, summarised in figure 1.16:

- Firstly the surfactant molecules must arrive close to the surface. This can happen by a mixture of convection (the transport of material due to a current in the surrounding medium) and diffusion (the transport of material due to a chemical potential gradient). The surfactant can exist in solution as monomers and micelles and can be exchanged between the two forms.
- Secondly surfactant molecules in the *subsurface*—the region immediately adjacent to the surface—can adsorb and stick to the surface. This is usually thought to happen only to monomers but there is evidence that micelles can adsorb at the air–water interface⁸⁹ and some models have assumed micellar adsorption at the solid–liquid interface.⁶³
- Finally, surfactant adsorbed to the surface may rearrange, for example exchanging between adsorbed monomers and adsorbed aggregates, changing in structure to pack more efficiently or phase-separating on the surface.

The latter two processes are what this work is primarily concerned with. However, in order to understand adsorption as a whole the transport process must also be understood; if transport to the subsurface is not understood then it is impossible to know what is limiting the overall rate of adsorption. For this reason surfactant adsorption kinetics should be conducted with well-defined hydrodynamics. The remainder of this section discusses this problem. It is worth pointing out that the issue of well-defined hydrodynamics is not unique to surfactant adsorption: it is also a consideration in other fields including colloidal deposition and electrochemistry. Electrochemistry can be viewed as the adsorption of ions onto an electrode surface, and the majority of relevant hydrodynamical modelling has been done to help understand electrochemical experiments.




Figure 1.16: Processes involved in surfactant adsorption, illustrated for a hydrophilic solid–liquid interface. The numbered processes are: 1) exchange between monomers and micelles in solution, 2) transport of monomers to the subsurface, 3) transport of micelles to the subsurface, 4) adsorption of monomers from the subsurface, 5) adsorption of micelles from the subsurface (which may not be possible) and 6) rearrangement of adsorbed molecules.

1.4.1 Air–water interface

Adsorption at the air–water interface is clearly related to adsorption at the solid–liquid interface, but has one key difference: it is possible to create expanding surfaces so that the surface excess is being diluted by the formation of new surface to which bulk surfactant can adsorb. Most commonly this is done by inflating liquid droplets or air bubbles from the end of a nozzle while the surface tension is probed (either by the drop shape or bubble pressure).⁹⁰ Diffusion controlled adsorption can be modelled by the Ward-Tordai equation,

$$\Gamma(t) = 2\sqrt{\frac{D}{\pi}} \left[c_0\sqrt{t} - \int_0^{\sqrt{t}} c(0, t - \tau) d\sqrt{\tau} \right], \quad (1.54)$$

and related approximations for the $c(0, t)$, the subsurface concentration at time t as well as an isotherm linking Γ and the subsurface concentration (D is the surfactant diffusion coefficient and c_0 the bulk concentration). An alternative approach is to use a continuously expanding surface such as an overflowing cylinder⁹¹ or liquid jet.⁸⁹ Here a non-equilibrium steady-state with a defined surface age can be established allowing the use of techniques requiring long measurement times, such as neutron reflectometry. It is difficult to see how such measurements could be performed at the solid–liquid interface.

1.4.2 Adsorption onto colloidal gels and fibrous materials

Many experiments have been performed on adsorption kinetics at colloidal silica or fibrous materials such as cotton and filter paper. A number of such studies are reviewed in ref. 20. These studies usually show adsorption timescales of hundreds of minutes, leading to the conclusion that they are primarily measuring transport of surfactant to the surface and not the intrinsic adsorption process. Since the adsorption process is usually followed by measuring depletion of surfactant from the bulk solution then adsorption must be accompanied by an appreciable reduction in the surfactant concentration around the interface or the experiment cannot be followed. Therefore, the rate at which that lost surfactant is replenished from the remainder of the solution becomes critical to the observed adsorption rate.

1.4.3 Stirred cuvettes

A number of studies of adsorption of non-ionic surfactants^{58–61} and surface active proteins^{92,93} used a flat sample placed in cuvette, stirred using a magnetic stirrer at the bottom of the cuvette. The solution adjacent to the sample was assumed to be stationary, due to a no-slip boundary condition parallel to the interface and the incompressibility of the fluid perpendicular to the interface. Therefore the authors made the assumption that the surfactant was transported across the stationary solution (the *stagnation layer*) entirely by diffusion, while the remainder of the solution was equalised to uniform concentration by the stirring. The thickness of the stagnation layer was around 100 μm ,⁵⁸ although other similar experiments had reported layer thickness of 2.5–20 μm .⁹³ The thickness was not measured exactly, but estim-

ated from the rates of desorption, and therefore enters the model as a further fitting parameter.

This method has a number of issues that reduce its usefulness. Firstly, the need to use the stagnation layer thickness as a fitting parameter is unsatisfactory since it reduces the amount of information that can be extracted from the model. Secondly, the use of a single value for the stagnation layer thickness is unrealistic: in reality the influence of convection will tend to zero approaching the interface, and the point at which diffusion becomes dominant will depend on the diffusion coefficient of the adsorbant, which can vary from surfactant to surfactant and is different for monomers and micelles. Thirdly, the assumed starting conditions of a constant concentration right up to the diffusion layer are unrealistic meaning that the experimental start of adsorption will be less abrupt than the modelled start of adsorption. This final point is addressed for a different geometry in chapter 3, and may not have as big an influence on the results as might initially be expected.

1.4.4 Channel flow cells

The channel flow cell provides a means of exchanging one solution for another in a controlled manner. It consists of a long, wide channel with narrow depth (in the z -axis). Solution is flushed in along the x -axis, and forms a parabolic flow in the xz -plane (shown in figure 1.17). The cell is wide enough in the y -axis that the flow profile can be treated as uniform through the majority of this axis. Close to the edge of the cell—where the velocity is slow—transport occurs primarily through diffusion. If the surface acts as a perfect sink then a steady-state diffusion layer is set up, with a limiting flux to the surface of,⁹⁴

$$J = \frac{1}{\Gamma(4/3)9^{1/3}} \left(\frac{v_0}{Dhx} \right)^{1/3} Dc_0, \quad (1.55)$$

where $\Gamma(4/3)9^{1/3} \approx 0.67$ with Γ representing the Gamma function, v_0 is the velocity at the centre of the channel and h is the depth of the channel in the z direction. Note that the flux is *non-uniform*; it depends on the position along the x -axis! The point $x = 0$ is defined as the point where the perfect-sink boundary condition begins and is well-defined in electrochemistry as the edge of the electrode but less well defined in surfactant adsorption kinetics because surfactant can often adsorb to surfaces upstream of the sample. One approach to avoiding the issues involved in modelling the cell is to increase the flow rate until the adsorption curve no longer changes:⁹⁵ at this point the rate must be wholly controlled by the final step of transfer to the surface. Despite these limitations channel flow cells have been used extensively for applications including electrochemistry,⁹⁶ surfactant adsorption^{34,47,48,51,95,97,98} and protein adsorption.⁹⁹

1.4.5 Wall-jet flow cell

The wall-jet flow cell is the design used in this work and is discussed in more detail in chapter 3 (figure 3.1 shows the flow profile in the cell). In brief, a jet of liquid is directed into the flat sample from perpendicular to the sample, and so the jet spreads




Figure 1.17: Channel flow cell, showing the parabolic profile of the newly injected surfactant solution. The y -axis is perpendicular to the page.

out as it hits the sample. At the surface there is a no slip boundary condition, meaning that solution is completely stationary and transport is by diffusion alone. Directly above the centre of the jet there is a *stagnation point*, and near this point there is almost no radial convection since it is close to the axis of symmetry of the jet. At the stagnation point the limiting flux to the interface (assuming the interface adsorbs all the surfactant that hits it) is given by the equation,¹⁰⁰

$$J = 0.77646D^{2/3}\alpha^{1/3}c_0, \quad (1.56)$$

where α is a constant that can be calculated from hydrodynamic modelling.¹⁰¹ As I will show later, the use of the above equation to describe surfactant adsorption is unsatisfactory, since significant adsorption is expected to take place well before the limiting flux is ever reached. Ignoring that issue, the key issue with the wall-jet flow cell is the uncertainty in the concentration profile being injected into the cell, which depends on the way the surfactant solution is switched for injection into the cell and the diffusion within the inlet tube that leads into the cell. I will attempt to address this issue in chapter 3 and show that the errors arising from it are small; however, it remains an inherent design problem in such cells. The wall-jet flow cell has been used extensively in electrochemistry,^{102,103} colloidal deposition^{101,104} and surfactant/polymer adsorption.^{43,45,50,100,105,105–109}

Note that when used in electrochemistry the term “wall-jet” usually refers to apparatus with a detector (the electrode) much larger than the inlet; the term “wall-tube” refers to apparatus with a detector much smaller than the inlet.¹¹⁰ Much of the recent work in electrochemistry has focussed on miniaturised flow cells, where the electrode is usually larger than the jet so the “wall-jet” situation applies. All the work on surfactant adsorption reviewed here is done with a detector (a laser spot) much smaller than the jet, so would be referred to as a “wall-tube”. Despite this I will use




Figure 1.18: Illustration of the operation of the dual-inlet channel flow cell. (a) A steady-state concentration profile is established within the cell; (b) flow is stopped and surfactant diffuses to the surface. The measurement is taken from the point where the laser beam reflects on the lower surface. Reprinted with permission from ref. 49. Copyright 2007 American Chemical Society.

the term “wall-jet” throughout this thesis.

1.4.6 Dual-inlet channel flow cell

The final cell design I will describe is the dual-inlet channel flow cell, used by Curwen *et al.* to study surfactant adsorption kinetics^{49,94,111} and Compton and coworkers previously to study electrochemical reactions.^{112,113} The operating principles are illustrated in figure 1.18. In brief, surfactant solution and water are injected continuously into two different inlets, setting up a steady-state concentration profile within the cell. Injection into both inlets is then stopped, allowing the surfactant already within the cell to adsorb to the surface. By establishing a well-defined profile at the start of the experiment the uncertainty in initial conditions that most other flow cells suffer from is elegantly avoided. The disadvantages to this cell design are that it requires large volumes of solution to set up the steady-state profile and that surfactant must diffuse across $\sim 500\ \mu\text{m}$ of cell (compared to $\sim 50\ \mu\text{m}$ for the wall-jet cell). The latter consideration limits the rates of adsorption that can be measured.

1.5 Optics

The information within this section is covered in most optics textbooks. References 114, 115 and 116 all cover the majority of the material, and so will not be cited individually in the following text.

Maxwell’s equations are

$$\nabla \times \mathbf{H} - \frac{\partial \mathbf{D}}{\partial t} = \mathbf{J}, \quad (1.57a)$$

$$\nabla \times \mathbf{E} + \frac{\partial \mathbf{B}}{\partial t} = 0, \quad (1.57b)$$

$$\nabla \cdot \mathbf{D} = \rho, \quad (1.57c)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (1.57d)$$

where \mathbf{H} is the magnetic intensity vector, \mathbf{D} the electric displacement vector, \mathbf{J} the

electric current density, \mathbf{E} the electric field vector, \mathbf{B} the magnetic induction vector and ρ the electric charge density. The pairs \mathbf{D} , \mathbf{E} and \mathbf{B} , \mathbf{H} can be related together by

$$\mathbf{D} = \varepsilon_0 \mathbf{E} + \mathbf{P}, \quad (1.58a)$$

$$\mathbf{H} = \mu_0^{-1} \mathbf{B} - \mathbf{M}, \quad (1.58b)$$

where \mathbf{P} is the electric polarisation vector and \mathbf{M} the magnetic polarization vector. Where the field is harmonic in time, the material is isotropic and at rest the following relations apply:

$$\mathbf{D} = \varepsilon \mathbf{E}, \quad (1.59a)$$

$$\mathbf{H} = \mu^{-1} \mathbf{B}. \quad (1.59b)$$

In a vacuum, ρ and \mathbf{J} both equal 0. Taking the curl of equations 1.57b and 1.57a and substituting in \mathbf{B} from equation 1.59b and \mathbf{D} from equation 1.59a respectively gives,

$$\nabla \times \nabla \times \mathbf{E} + \nabla \times \left(\mu \frac{\partial \mathbf{H}}{\partial t} \right) = 0, \quad (1.60a)$$

$$\nabla \times \nabla \times \mathbf{H} - \nabla \times \left(\varepsilon \frac{\partial \mathbf{E}}{\partial t} \right) = 0. \quad (1.60b)$$

In a homogeneous medium μ and ε are constant in space, so can appear before the curl. Differentiating equations 1.57a and 1.57b with respect to time and substituting in \mathbf{D} from equation 1.59a and \mathbf{B} from equation 1.59b respectively gives,

$$\nabla \times \frac{\partial \mathbf{H}}{\partial t} - \varepsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} = 0, \quad (1.61a)$$

$$\nabla \times \frac{\partial \mathbf{E}}{\partial t} + \mu \frac{\partial^2 \mathbf{H}}{\partial t^2} = 0, \quad (1.61b)$$

and so,

$$\nabla \times \nabla \times \mathbf{E} + \varepsilon \mu \frac{\partial^2 \mathbf{E}}{\partial t^2} = 0, \quad (1.62a)$$

$$\nabla \times \nabla \times \mathbf{H} + \varepsilon \mu \frac{\partial^2 \mathbf{H}}{\partial t^2} = 0. \quad (1.62b)$$

Using the identity $\nabla \times \nabla \times \mathbf{A} = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A}$ gives,

$$\nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} + \varepsilon \mu \frac{\partial^2 \mathbf{E}}{\partial t^2} = 0, \quad (1.63a)$$

$$\nabla(\nabla \cdot \mathbf{H}) - \nabla^2 \mathbf{H} + \varepsilon \mu \frac{\partial^2 \mathbf{H}}{\partial t^2} = 0, \quad (1.63b)$$

Using equation 1.57c with \mathbf{D} from equation 1.59a and equation 1.57d with \mathbf{B} equa-

tion 1.59b gives

$$\varepsilon \nabla \cdot \mathbf{E} = 0, \quad (1.64a)$$

$$\mu \nabla \cdot \mathbf{H} = 0. \quad (1.64b)$$

Under these conditions, equations 1.63 and equations 1.64 can be combined to give

$$\nabla^2 \mathbf{E} - \varepsilon \mu \frac{\partial^2 \mathbf{E}}{\partial t^2} = 0, \quad (1.65a)$$

$$\nabla^2 \mathbf{H} - \varepsilon \mu \frac{\partial^2 \mathbf{H}}{\partial t^2} = 0. \quad (1.65b)$$

These are equations describing wave motion and suggest that the electric and magnetic field vectors propagate with a speed of $\nu = 1/\sqrt{\varepsilon\mu}$. For a monochromatic plane wave equations 1.65 have the solution

$$\mathbf{E} = \mathbf{E}_0 e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t + \delta)}, \quad (1.66a)$$

$$\mathbf{H} = \mathbf{H}_0 e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t + \delta)}, \quad (1.66b)$$

where \mathbf{E}_0 and \mathbf{H}_0 are constant vectors with the directions giving linear polarisations, \mathbf{k} the propagation vector with magnitude $2\pi/\lambda$, ω the angular temporal frequency and δ is a phase shift. \mathbf{H}_0 is perpendicular to \mathbf{E}_0 ; the polarisation of the light as a whole is usually expressed in terms of \mathbf{E} . Any possible polarisation of light can be generated by taking linear combinations of two perpendicularly polarised light waves. Varying the phase shifts of the two polarisations can produce circularly and elliptically polarised light.

One further useful vector is the Poynting vector of a wave, \mathbf{S} representing the flux of energy crossing a unit area normal to the direction of the vector. It is defined by

$$\mathbf{S} = \mathbf{E} \times \mathbf{H}. \quad (1.67)$$

1.5.1 Reflection and refraction at interfaces

The equations for total internal reflection are easiest to express in terms of a planar, monochromatic incident light wave,

$$\mathbf{E}_i = \mathbf{E}_{0i} e^{-i(\mathbf{k}_i \cdot \mathbf{r} - \omega_i t + \delta_i)}. \quad (1.68)$$

The subscript i designates that this is the incident light vector. When the incident light hits an interface transmitted (subscript t) and reflected (subscript r) plane waves are also produced. These waves are given by the equations,

$$\mathbf{E}_r = \mathbf{E}_{0r} e^{-i(\mathbf{k}_r \cdot \mathbf{r} - \omega_r t + \delta_r)}, \quad (1.69a)$$

$$\mathbf{E}_t = \mathbf{E}_{0t} e^{-i(\mathbf{k}_t \cdot \mathbf{r} - \omega_t t + \delta_t)}. \quad (1.69b)$$

The coordinate system used for the remainder of this section is defined in fig-




Figure 1.19: Schematic illustrating the coordinate system used in the discussion of reflection and refraction at the interface, the direction of the parallel and perpendicular components of the electric field vectors, and the direction of the propagation vectors of the various plane waves. Additionally, the incident light is labelled as S and P polarised, to match the convention I use when discussing Raman spectroscopy later.

ure 1.19. At the interface—defined as the plane $z = 0$ —boundary conditions mean that the following must be continuous:

1. the component of the overall electric field, \mathbf{E} , tangential to the interface;
2. the component of \mathbf{D} , the electric displacement vector, normal to the interface;
3. the component of \mathbf{H} , the magnetic intensity vector, tangential to the interface;
4. the component of the magnetic induction vector, \mathbf{B} , normal to the interface.

From the first boundary condition, at the interface,

$$\hat{\mathbf{u}}_n \times \mathbf{E}_{0i} e^{-i(\mathbf{k}_i \cdot \mathbf{r} - \omega_i t + \delta_i)} + \hat{\mathbf{u}}_n \times \mathbf{E}_{0r} e^{-i(\mathbf{k}_r \cdot \mathbf{r} - \omega_r t + \delta_r)} = \hat{\mathbf{u}}_n \times \mathbf{E}_{0t} e^{-i(\mathbf{k}_t \cdot \mathbf{r} - \omega_t t + \delta_t)}, \quad (1.70)$$

where $\hat{\mathbf{u}}_n$ is the unit vector normal to the interface. In order for that condition to apply, the time-dependent terms all need to behave identically, constraining,

$$\omega_i = \omega_r = \omega_t. \quad (1.71)$$

The spatial variation of the wave must be the same at all points on the interface, giving,

$$(\mathbf{k}_i \cdot \mathbf{r} + \delta_i) = (\mathbf{k}_r \cdot \mathbf{r} + \delta_r) = (\mathbf{k}_t \cdot \mathbf{r} + \delta_t), \quad (1.72)$$

and so,

$$(\mathbf{k}_i - \mathbf{k}_r) \cdot \mathbf{r} = \delta_r - \delta_i, \quad (1.73a)$$

$$(\mathbf{k}_i - \mathbf{k}_t) \cdot \mathbf{r} = \delta_t - \delta_i. \quad (1.73b)$$

These equations state that the plane of the interface—defined by the values of vector \mathbf{r} when it points to the interface—must be perpendicular to the vectors $\mathbf{k}_i - \mathbf{k}_r$ and $\mathbf{k}_i - \mathbf{k}_t$. This also means that $\mathbf{k}_i - \mathbf{k}_r$ and $\mathbf{k}_i - \mathbf{k}_t$ must both be parallel to $\hat{\mathbf{u}}_n$ and thus,

$$(\mathbf{k}_i - \mathbf{k}_r) \times \hat{\mathbf{u}}_n = (\mathbf{k}_i - \mathbf{k}_t) \times \hat{\mathbf{u}}_n = 0. \quad (1.74)$$

Therefore $k_i \sin \theta_i = k_r \sin \theta_r = k_t \sin \theta_t$, giving $\theta_i = \theta_r$ (since the magnitudes of k_i and k_r are equal as they are in the same medium) and $n_i \sin \theta_i = n_t \sin \theta_t$ (using $k = \omega n/c$). When $n_t \geq n_i$, there is a real value of θ_t for every possible incident angle θ_i ; when $n_t < n_i$ there are only real values of θ_t when $\sin \theta_i \leq n_t/n_i$. The angle $\theta_i = \arcsin n_t/n_i$ is termed the critical angle, θ_c , and above this angle there is no net transmission of energy across the interface.

Since any polarisation of incident wave can be represented as a linear combination of two orthogonal linear polarisations, the incident wave is split into two linear combinations, defined with respect to the plane of incidence. For the case where \mathbf{E} is perpendicular to the plane of incidence, the electric field lies in the plane of the interface. Hence boundary condition 1 applies, and so

$$\mathbf{E}_i + \mathbf{E}_r = \mathbf{E}_t. \quad (1.75)$$

Boundary condition 3 also applies: taking the components of $\mu^{-1}\mathbf{B}$ tangential to the interface gives

$$-\frac{B_i}{\mu_i} \cos \theta_i + \frac{B_r}{\mu_r} \cos \theta_r = -\frac{B_t}{\mu_t} \cos \theta_t. \quad (1.76)$$

Boundary conditions 2 and 4—relating to the normal components—are automatically fulfilled with the other two boundary conditions. Using $\mathbf{k} \times \mathbf{E} = \nu \mathbf{B}$, where $\nu = c/n$ is the speed of light in the medium, equation 1.76 becomes,

$$\frac{1}{\mu_i \nu_i} (E_i - E_r) \cos \theta_i = \frac{1}{\mu_t \nu_t} E_t \cos \theta_t, \quad (1.77)$$

which can be further simplified—with the knowledge that E_i , E_r and E_t must all have the same time dependence and that, for dielectric materials, $\mu_i \approx \mu_t \approx \mu_0$ —to give,

$$\frac{n_i}{\nu_i} (E_{0i} - E_{0r}) \cos \theta_i = \frac{n_t}{\nu_t} E_{0t} \cos \theta_t. \quad (1.78)$$

Following some rearrangement the amplitude reflection and amplitude transmission coefficients for perpendicularly polarised light are obtained:

$$\begin{aligned} r_{\perp} &= \left(\frac{E_{0r}}{E_{0i}} \right)_{\perp} = -\frac{\sin(\theta_i - \theta_t)}{\sin(\theta_i + \theta_t)}, \\ t_{\perp} &= \left(\frac{E_{0t}}{E_{0i}} \right)_{\perp} = \frac{2 \sin \theta_t \cos \theta_i}{\sin(\theta_i + \theta_t)}. \end{aligned} \quad (1.79)$$

Similarly, when \mathbf{E} is parallel to the plane of incidence, then boundary condition 1 gives

$$E_i \cos \theta_i - E_r \cos \theta_r = E_t \cos \theta_t, \quad (1.80)$$




Figure 1.20: Amplitude coefficients of reflection for parallel (solid blue line) and perpendicular (dashed green line) polarised light as a function of incident angle for internal reflection at the silica-water interface ($n_i = 1.461$, $n_t = 1.336$). Above the critical angle (66°) all the light is reflected. At the Brewster angle, θ_b , no parallel polarised light is reflected.

while boundary condition 3 gives

$$\frac{\mathbf{B}_i + \mathbf{B}_r}{\mu_i} = \frac{\mathbf{B}_t}{\mu_t}. \quad (1.81)$$

The amplitude reflection and amplitude transmission coefficients for parallel polarised light are,

$$r_{\parallel} = \left(\frac{E_{0r}}{E_{0i}} \right)_{\parallel} = \frac{\tan(\theta_i - \theta_t)}{\tan(\theta_i + \theta_t)}, \quad (1.82)$$

$$t_{\parallel} = \left(\frac{E_{0t}}{E_{0i}} \right)_{\parallel} = \frac{2 \sin \theta_t \cos \theta_i}{\sin(\theta_i + \theta_t) \cos(\theta_i - \theta_t)}.$$

Below the critical angle, all the angles in equations 1.79 and 1.82 are real and hence the phase change of the transmitted and reflected light can only be 0 or π . None of the terms in t_{\parallel} or t_{\perp} allow a change of sign with incident angles between 0 and 90° , hence the phase of the transmitted light must be the same as that of the incident light. r_{\perp} is negative when $\theta_i > \theta_t$ ($n_i < n_t$), and so the phase differs by π . The parallel polarised reflected light is in phase with the incident light when $\theta_i + \theta_t < \pi/2$ and $n_i < n_t$ but out of phase under the same conditions when $n_t < n_i$. The angle at which the reflected light first changes phase, $\theta_i + \theta_t = \pi/2$ is known as Brewster's angle. At this angle no parallel polarised light is reflected, linearly polarising the reflected light.

Reflectivity coefficients for internal reflection at the silica-water interface are shown in figure 1.20, together with the critical angle and Brewster's angle.




Figure 1.21: Penetration depth with respect to angle for total internal reflection at the silica-water interface ($n_i = 1.461$, $n_t = 1.336$). Both the penetration depth (solid blue line) and the penetration depth of the square of the electric field (dashed green line) are shown, together with values for 73° , the angle used for most of my experiments.

1.5.2 Total internal reflection

Above the critical angle all the incident light is reflected back, giving total internal reflection. However, for the boundary conditions to be fulfilled, there must be an electric field on the transmitted side of the interface. The $\mathbf{r} \cdot \mathbf{k}_t$ term in equation 1.69b factorises to become $k_t(x \sin \theta_t + y \cos \theta_t)$, and together with

$$\begin{aligned} \sin \theta_t &= \frac{n_i}{n_t} \sin \theta_i, \\ \cos \theta_t &= \pm i\beta = \pm i \sqrt{\left(\frac{n_i}{n_t}\right)^2 \sin^2 \theta_i - 1}, \end{aligned} \quad (1.83)$$

where β is defined for brevity later, equation 1.69b becomes

$$\mathbf{E}_t = \mathbf{E}_{0t} e^{-i\left(k_t x \frac{n_i \sin \theta_i}{n_t}\right)} e^{\pm k_t \beta z}. \quad (1.84)$$

In the latter term—describing the z dependence—only the negative sign makes physical sense: the positive sign would result in an ever increasing wave further from the interface. The transmitted wave propagates along the plane of incidence, but decays exponentially as it moves further from the interface. The electric field has a characteristic decay length in the z direction of $(k_t \beta)^{-1}$, and the square of the electric field—which defines the intensity of the light—has a characteristic decay length of $(k_t \beta)^{-1}/2$. These decay lengths are of the order of one wavelength of the light. The component of the time-averaged Poynting vector in the direction normal to the interface vanishes, showing that there is no overall flow of energy across the interface. The penetration depth for the silica-water interface is plotted in figure 1.21.

Following some rearrangement and substitution the Fresnel coefficients for the

reflected light then become

$$r_{\parallel} = \frac{n^2 \cos \theta_i - i\sqrt{\sin^2 \theta_i - n^2}}{n^2 \cos \theta_i + i\sqrt{\sin^2 \theta_i - n^2}}, \quad (1.85a)$$

$$r_{\perp} = \frac{\cos \theta_i - i\sqrt{\sin^2 \theta_i - n^2}}{\cos \theta_i + i\sqrt{\sin^2 \theta_i - n^2}}, \quad (1.85b)$$

where $n = n_r/n_i$. The Fresnel factors for the transmitted beam are enhanced by the overlap between the incident and reflected beam. At this point it is worth splitting the factors into separate factors for light in each of the three axes (rather than simply parallel and perpendicular to the plane of incidence), since the presence of the surface means that molecules at the surface may be ordered with respect to the z axis and so any spectroscopy conducted with the evanescent wave could have a different dependency on the intensity in the z axis than on the intensity in the x and y axes. So that the notation matches the notation used to describe the Raman experiments conducted later I have added a second notation, using subscript P to indicate parallel and subscript S to indicate perpendicular polarisation.

Perpendicularly polarised (S) incident light has no component in the x and z axes, so $t_{Sx} = t_{Sz} = 0$ (or $t_{\perp x} = t_{\perp z} = 0$); similarly, parallel polarised (P) incident light has no component in the y axis, so $t_{Py} = 0$ (or $t_{\parallel y} = 0$). The non-zero Fresnel factors are¹¹⁷

$$t_{Px} = t_{\parallel x} = \frac{2 \cos \theta_i (\sin^2 \theta_i - n^2) + 2in^2 \cos^2 \theta_i \sqrt{\sin^2 \theta_i - n^2}}{n^4 \cos^2 \theta_i + \sin^2 \theta_i - n^2}, \quad (1.86a)$$

$$t_{Sy} = t_{\perp y} = \frac{2 \cos^2 \theta_i - 2i \cos \theta_i \sqrt{\sin^2 \theta_i - n^2}}{1 - n^2}, \quad (1.86b)$$

$$t_{Pz} = t_{\parallel z} = \frac{2n^2 \cos^2 \theta_i \sin \theta_i - 2i \cos \theta_i \sin \theta_i \sqrt{\sin^2 \theta_i - n^2}}{n^4 \cos^2 \theta_i + \sin^2 \theta_i - n^2}. \quad (1.86c)$$

Note that equations 1.86 only apply above the critical angle, since they make an assumption regarding the sign of equation 1.83 that is only true above the critical angle. The absolute values of the Fresnel factors with respect to angle are shown in figure 1.22, using the refractive indices of a silica–water interface. At the critical angle there is an enhancement in t_{Sy} and t_{Pz} while t_{Px} is zero, making analysis of spectra arising from this angle simpler. However, aligning an experiment to be exactly at the critical angle is challenging. At 73° —the angle I normally use for experiments— t_{Pz} is close to its local maximum.

Where there is a thin film at the surface with refractive index n' the Fresnel factors change slightly.¹¹⁸ Assuming the film is thin enough compared to the wavelength of light, reflection and refraction within the film can be neglected. For light polarised with an electric field parallel to the interface, E_{\parallel} must be continuous across the interface (as before). E is not affected by the film and thus the factors t_{Px} and t_{Sy} do not change. For light polarised perpendicular to the interface, $D_{\perp} = \epsilon E_{\perp}$ must be continuous across the interface. To meet this criteria t_{Pz} must be multiplied by $(n_i/n')^2$.




Figure 1.22: Absolute values of the Fresnel factors as a function of incident angle, at the silica-water interface ($n_i = 1.461$, $n_t = 1.336$). t_{Px} is the solid blue line, t_{Sy} the dashed green line and t_{Pz} the dotted red line.

The phase shift, δ of the reflected light is given by

$$\begin{aligned} \tan \frac{\delta_{\parallel}}{2} &= -\frac{\sqrt{\sin^2 \theta_i - n^2}}{n^2 \cos \theta_i} \\ \tan \frac{\delta_{\perp}}{2} &= -\frac{\sqrt{\sin^2 \theta_i - n^2}}{\cos \theta_i}. \end{aligned} \quad (1.87)$$

Derivation of the above equations is based on equations 1.85 being of a form $z(z^*)^{-1}$, where z^* designates a complex conjugate: $e^{i\delta} = z(z^*)^{-1} = e^{2i\alpha}$ where α is the argument of z expressed in Euler form.

1.6 Total internal reflection Raman spectroscopy

1.6.1 Vibrational spectroscopy

For a nonlinear molecule containing N atoms, there are $3N$ coordinates needed to specify the position of every atom (one per axis per atom). Three of those correspond to translations of the centre of mass of the molecule, and a further three to rotation about that centre of mass. Therefore, there are $3N - 6$ coordinates remaining which correspond to vibrations of the molecule. For a linear molecule there are $3N - 5$ coordinates, since only two rotations can be specified.¹¹⁹

Using a Taylor series expansion of the potential energy, V , of the molecule

$$V = V(0) + \sum_i \left. \frac{\partial V}{\partial x_i} \right|_0 x_i + \frac{1}{2} \sum_{i,j} \left. \frac{\partial^2 V}{\partial x_i \partial x_j} \right|_{0,0} x_i x_j + \dots, \quad (1.88)$$

where x is the displacement of an atom along a coordinate. At the potential energy minimum the first derivatives, $\frac{\partial V}{\partial x}$, are all zero. By taking linear combinations of the different coordinates—weighted by the mass of the atom so that $x'_i = \sqrt{m_i} x_i$ —a

replacement expression for equation 1.88 can be found so that there are no cross terms ($i \neq j$):

$$V = \frac{1}{2} \sum_i K_i \mathbf{q}_i^2, \quad (1.89)$$

where \mathbf{q}_i is such a linear combination of atomic coordinates, termed a *normal coordinate*, and K_i is a constant arising from the second derivative of potential energy and the masses of the atoms involved.¹¹⁹ $V(0)$ is dropped from the equation above since it only affects the absolute energy of the molecule and does not alter the vibrations of the molecule. The total vibrational wavefunction of the molecule, ψ is the product of the wavefunctions for each normal mode:

$$\psi = \prod_i \psi_{v_i}(\mathbf{q}_i). \quad (1.90)$$

If the potential energy were actually harmonic, as implied by equation 1.89, there would be an infinite number of discrete energy levels available for each normal mode given by,

$$E = \left(v_i + \frac{1}{2} \right) \hbar \omega_i, \quad (1.91)$$

where v_i is the vibrational quantum number (an integer ≥ 0) and $\omega_i = \sqrt{K_i}$. The total vibrational energy of the molecule is the sum of the energies of each normal mode. In the lowest vibrational quantum number, $E = \frac{1}{2} \hbar \omega_i$, the zero-point energy. Because the potential energy is not actually harmonic (we have neglected higher order terms in the Taylor series expansion early), equation 1.91 does not apply exactly but serves as a good approximation at low quantum numbers. At higher quantum numbers the molecule has enough energy for the atoms to reach infinite separation, breaking the bonds. Therefore, there are actually a finite number of available vibrational energy states. Figure 1.23 shows the difference in shape and energy level spacing between a harmonic and anharmonic potential. The anharmonic potential is modelled using a Morse potential, which is often used as an approximation for the potential energy of a diatomic molecule. For most vibrational transitions only the lowest energy level is significantly occupied at room temperature ($\Delta E \ll k_B T$), and hence only transitions from this first level have a significant impact on the spectrum.

The vibrational modes of a molecule can be probed by infrared and Raman spectroscopy. Raman spectroscopy will be discussed in detail in the following section. Infrared spectroscopy is an absorption spectroscopy: a sample is irradiated with a range of infrared frequencies and the light passing through the sample is detected. Where the energy of the incident light matches the energy gap between two vibrational levels (and other selection rules are met) the light can be absorbed by the molecule.

The 3 (or 2 for linear molecules) rotational modes also have energy levels associated with, with smaller energy gaps than for the vibrational levels. In the gas phase the rotational spectrum can be observed simultaneously with the vibrational spectrum (since a vibrational transition can be accompanied by a rotational transition). In the liquid phase the lifetime of rotational states is shortened due to collisions with adjacent molecules. The linewidths of the bands are blurred by $\delta E \approx \hbar/\tau$, where τ is




Figure 1.23: Harmonic (solid blue) and anharmonic (dotted green) potentials with respect to atomic displacement, with quantised energy levels indicated using horizontal lines. The constants used have been chosen to be reasonably close to a 3000 cm^{-1} C–H stretch, although the displacements are unphysically large.

the average lifetime of the state,¹²⁰ and hence it is not possible to resolve rotational energy levels in liquid phase spectra.*

1.6.2 Raman scattering

When light interacts with a molecule there are number of possibilities which can happen: if the energy of the light exactly matches a transition between the current energy state and a higher allowed energy state within the molecule then the light can be absorbed. If the energy of the light matches that of a transition to a lower energy state then two photons can be emitted—the principle by which laser radiation is produced. If the energy of the light does not match any transitions within the molecule, the light will be scattered.

The majority of the scattered radiation will be at the same frequency to the original radiation; this is referred to as Rayleigh scattering. A portion of the radiation will have a shifted frequency, and this is referred to as Raman scattering. Raman scattering was first observed by C. V. Raman in 1928,^{121,122} although it had been predicted in 1923 by A. Smekal.¹²³ Further possibilities exist, based on three or more photon processes which allow the frequency to be doubled, tripled and so forth, but these are unimportant to the work here and so will not be discussed further.

Radiation at lower frequency than that of the incident light is referred to as a Stokes band, whereas radiation at a higher frequency to the incident light is referred to as an anti-Stokes band. Both of these possibilities are illustrated in figure 1.24. In each case the system proceeds through a virtual energy state, shown in the figure by a dotted line. The difference between the frequencies of the scattered light is determined by the difference in energy between the initial and final states of the molecule involved.

*As Heisenberg’s wife complained, “when you’ve got the time you don’t have the energy and when you’ve got the energy you don’t have the time.”




Figure 1.24: Different types of Raman Scattering. ν_s is the frequency of the scattered light, and is given by $\nu_s = \nu_i - \nu_{fi}$ or $\nu_s = \nu_i + \nu_{if}$ depending on the type of scattering. The height of a level represents the amount of energy held by the scattering molecule when in that state.

For my work I am only looking at vibrational energy levels, and so spectra generated are very similar to the infra-red spectra of the molecules in question—although with different selection rules governing which modes are seen. At room temperature only the lowest vibrational energy levels are significantly occupied, meaning that only Stokes scattering will be observed. Therefore, only Stokes scattering is used in these experiments.

Neglecting derivatives of the electric field (the electric dipole approximation), \mathbf{p} , the dipole moment of a molecule interacting with an electric field is

$$\mathbf{p} = \mu^0 + \alpha \mathbf{E}' + \beta \mathbf{E}\mathbf{E} + \gamma \mathbf{E}\mathbf{E}\mathbf{E} + \dots, \quad (1.92)$$

where μ^0 is the static dipole. α is the polarisability, \mathbf{E}' is the applied electric field local to the molecule. β and γ are the first and second order hyperpolarisabilities, respectively. In weak electric fields β and γ are unimportant.

\mathbf{E}' is not equal to the electric field applied to the material, \mathbf{E} , but is affected by electric fields from surrounding dipoles. Therefore, it is often easier to use the a related series expressed in terms of the electrical susceptibility, a bulk property:

$$\mathbf{P} = \mathbf{P}^{(0)} + \varepsilon_0 \chi^{(1)} \mathbf{E} + \varepsilon_0 \chi^{(2)} \mathbf{E}\mathbf{E} + \dots, \quad (1.93)$$

where \mathbf{P} is the bulk polarisation (see equation 1.58a), $\chi^{(1)}$ and $\chi^{(2)}$ are the first order (linear) and second order electrical susceptibilities and $\mathbf{P}^{(0)}$ is the static polarisation (zero for most materials). As with equation 1.92, the second order and higher terms can usually be neglected.

In an isotropic medium, the local field is greater than the mean incident field by

$$\mathbf{E}' = \mathbf{E} + \frac{\mathbf{P}}{3\varepsilon_0} \quad (1.94)$$

Therefore, the molecular polarisability is related to its bulk material property, the

first order susceptibility, by¹¹⁶

$$\chi^{(1)} = \frac{N\langle\alpha\rangle}{1 - \frac{1}{3\epsilon_0}N\langle\alpha\rangle}, \quad (1.95)$$

where N is the number of dipoles, and $\langle\alpha\rangle$ is an orientational average of the true polarisability tensor of the molecule. The molecular polarisability can also be related to the permittivity and refractive index of the material by the Clausius-Mossotti equation (also known as the Lorentz-Lorenz equation):

$$\frac{N\langle\alpha\rangle}{3\epsilon_0} = \frac{\epsilon - 1}{\epsilon + 2} = \frac{n^2 - 1}{n^2 + 2}. \quad (1.96)$$

Since Raman spectroscopy acts on a molecular level all the fields discussed will be local fields. For simplicity I will switch to using \mathbf{E} to denote the local field throughout the remainder of this section. Raman scattering occurs when an electric dipole is induced in the scattering molecule using the electric field of the light: the linear term in equation 1.92:

$$\mathbf{p} = \boldsymbol{\alpha}\mathbf{E}. \quad (1.97)$$

The polarisability is a 3×3 tensor; the meaning of its elements will be discussed later in this section.

The electric field of light with a frequency ν_0 can be expressed as,

$$\mathbf{E} = \mathbf{E}_0 \cos 2\pi\nu_0 t, \quad (1.98)$$

and hence equation 1.97 can be rewritten as,

$$\mathbf{p} = \boldsymbol{\alpha}\mathbf{E}_0 \cos 2\pi\nu_0 t, \quad (1.99)$$

representing the induction of Rayleigh scattered light at the frequency ν_0 . The polarisability of a molecule varies with movement along a vibrational coordinate, \mathbf{q} , of that molecule. Using a Taylor series expansion for the polarisability,

$$\boldsymbol{\alpha} = \boldsymbol{\alpha}_0 + \left. \frac{\partial \boldsymbol{\alpha}}{\partial \mathbf{q}} \right|_0 \mathbf{q} + \dots \quad (1.100)$$

Assuming that the vibrational mode behaves as a harmonic oscillator, with frequency ν_v , then,

$$\mathbf{q} = \mathbf{q}_0 \cos 2\pi\nu_v t. \quad (1.101)$$

Substituting the previous two equations into equation 1.99 gives,¹²⁴

$$\mathbf{p} = \boldsymbol{\alpha}_0 \mathbf{E}_0 \cos 2\pi\nu_0 t + \frac{1}{2} \mathbf{E}_0 \mathbf{q} \left. \frac{\partial \boldsymbol{\alpha}}{\partial \mathbf{q}} \right|_0 [\cos 2\pi(\nu_0 + \nu_v)t + \cos 2\pi(\nu_0 - \nu_v)t]. \quad (1.102)$$

This equation describes Raman scattering in a classical manner (as opposed to the quantum mechanical description using the wavefunctions of the initial and final states of the molecule). The three terms on the right-hand side of the above equation are Rayleigh, anti-Stokes and Stokes scattering respectively. Stokes and anti-Stokes

scattering can only occur when the $\frac{\partial \alpha}{\partial \mathbf{q}} \neq 0$: when the polarisability changes during the vibration. The frequency of the Stokes and anti-Stokes scattering can be seen as beat frequencies between the probe light (frequency ν_0) and the molecular vibration (frequency ν_v).¹²⁵

As mentioned before, the polarisation is a 3×3 tensor. Expressed in lab coordinates its elements are

$$\boldsymbol{\alpha} = \begin{pmatrix} \alpha_{xx} & \alpha_{xy} & \alpha_{xz} \\ \alpha_{yx} & \alpha_{yy} & \alpha_{yz} \\ \alpha_{zx} & \alpha_{zy} & \alpha_{zz} \end{pmatrix}, \quad (1.103)$$

where the first element of the subscript denotes the polarisation of the emitted radiation and the second element denotes the polarisation of the incident field. Using time-dependent perturbation theory, an element of the Raman tensor can be given by,¹²⁵

$$\alpha_{\rho\sigma} = \frac{1}{\hbar} \sum_{ev} \left\{ \frac{\langle \psi_b | \hat{\mu}_\rho | \psi_{ev} \rangle \langle \psi_{ev} | \hat{\mu}_\sigma | \psi_a \rangle}{\omega_{ev} - \omega_0 - i\Gamma_{ev}} + \frac{\langle \psi_b | \hat{\mu}_\sigma | \psi_{ev} \rangle \langle \psi_{ev} | \hat{\mu}_\rho | \psi_a \rangle}{\omega_{ev} + \omega_0 - i\Gamma_{ev}} \right\}, \quad (1.104)$$

where $|\psi_b\rangle$ and $|\psi_a\rangle$ are the initial and final vibrational wavefunctions in the ground electronic state and $|\psi_{ev}\rangle$ is the wavefunction of an excited electronic state. ω_0 is the angular frequency of the incident light, ω_{ev} the angular frequency of light required to reach the excited state and Γ_{ev} is the homogenous line width of the transition. In non-resonant Raman spectroscopy, where $\omega_0 \ll \omega_{ev}$, the denominators of the two terms in equation 1.104 are essentially equal. Therefore, $\alpha_{\rho\sigma} = \alpha_{\sigma\rho}$, and the number of independent components of $\boldsymbol{\alpha}$ is reduced. In resonance Raman spectroscopy—where the energy of the incident light is close to an excited electronic state—the Raman signal is greatly enhanced, and $\alpha_{\rho\sigma} \neq \alpha_{\sigma\rho}$. The elements of the Raman tensor that my experiment can probe are discussed in the experimental chapter, section 2.1.3.

The frequency dependence of Raman scattering can be deduced by considering the energy transmitted (the Poynting vector) by an oscillating electric dipole, p with angular frequency $\omega = 2\pi\nu$. At a distance, $R \gg c/\omega$, far from the dipole,¹¹⁶

$$\varepsilon_0 E_\theta \approx \frac{H_\phi}{c} \approx \frac{d^2 p(t - R/c)}{dt^2} \frac{1}{4\pi c^2 R} \sin \theta, \quad (1.105)$$

where θ is the angle from the direction the dipole is vibrating in, and ϕ the angle in the plane perpendicular to the dipole. The $\frac{d^2 p}{dt^2}$ term arises from the need to induce an oscillating electromagnetic field: a stationary dipole induces a constant electric field, a moving dipole induces a constant magnetic field and an accelerating dipole induces a changing magnetic field.¹¹⁵ The dipole term is retarded, in that it has a time offset equivalent to the light travelling the distance R . Putting the values for E_θ and H_ϕ into equation 1.67 gives a Poynting vector of,

$$S = \left(\frac{\partial^2 p(t - R/c)}{\partial t^2} \right)^2 \frac{1}{16\pi^2 \varepsilon_0 c^3 R^2} \sin^2 \theta. \quad (1.106)$$

Using $p(t) = p_0 \exp(-i\omega t)$ gives

$$\varepsilon_0 E_\theta = \frac{H_\phi}{c} \approx - \left(\frac{\omega}{c}\right)^2 p_0 \sin \theta \frac{\exp(-i\omega(t - R/c))}{R}. \quad (1.107)$$

The time averaged Poynting vector is therefore^{116,125}

$$\langle S \rangle = \frac{\omega^4 p_0^2}{32\pi^2 \varepsilon_0 c^3 R^2} \sin^2 \theta. \quad (1.108)$$

Ignoring any resonance effects, the intensity of Raman scattering scales with the fourth power of the frequency of the emitted light (and so has a similar dependency on the frequency of the exciting light). Experimentally, Raman scattering is detected as the number of photons hitting the detector. Given that a photon has an energy of $E = \hbar\omega$, equation 1.108 expressed in terms of a rate of photos, \mathcal{N} , becomes,

$$\mathcal{N} = \frac{\omega^3 p_0^2}{16\pi h \varepsilon_0 c^3 R^2} \sin^2 \theta. \quad (1.109)$$

There is a practical trade-off to be made with the choice of the incident light use to excite Raman scattering since higher frequencies of light often induce fluorescence which can overwhelm the Raman signal.

Additionally, equation 1.108 predicts a strong angular dependence for the transmitted light, mostly in the plane perpendicular to the dipole and not in the direction of the dipole itself. The angular dependence has important consequences for the experimental ability to detect certain polarisations of scattered light, as will be discussed in Chapter 2.

Selection rules

One of the selection rules for Raman scattering was shown earlier in this section: there must be a change in polarisability of the molecule during the vibration excited. The general requirement for incident light inducing a transition between quantum states is that the transition dipole moment, $\boldsymbol{\mu}_{fi}$, is non-zero. In the general case, the transition dipole moment is given by,

$$\boldsymbol{\mu}_{fi} = \langle \psi_f | \hat{\boldsymbol{\mu}} | \psi_i \rangle, \quad (1.110)$$

where ψ_f and ψ_i are the wavefunctions of the final and initial states respectively, and $\hat{\boldsymbol{\mu}}$ is the transition dipole moment operator. Equation 1.110 can be used directly for absorption spectroscopy (such as infrared vibrational spectroscopy), where incident light directly excites a vibration within the molecule. In Raman spectroscopy, the electric dipole is induced by the incident light, and so equation 1.110 can be rewritten as,

$$\boldsymbol{\mu}_{fi} = \langle \psi_f | \hat{\boldsymbol{\alpha}} | \psi_i \rangle \mathbf{E}, \quad (1.111)$$

where $\hat{\boldsymbol{\alpha}}$ is the transition polarisability tensor operator.

The integral $\langle \psi_f | \hat{\boldsymbol{\alpha}} | \psi_i \rangle = \int \psi_f^* \hat{\boldsymbol{\alpha}} \psi_i$ will be zero unless the product $\psi_f^* \hat{\boldsymbol{\alpha}} \psi_i$ contains the total symmetric irreducible representation.¹²⁶ The symmetry properties for an element $\alpha_{\rho\sigma}$ are the same as for the function $\rho\sigma$, hence the elements of the Ra-

man tensor have symmetry properties of the form x^2 , xy , xz , y^2 etc.. The ground state wavefunction is always totally symmetric; the symmetry of the excited vibrational state must be known to work out if the vibration is Raman active. For large molecules, vibrations are often localised on specific groups, and they can be treated in isolation from the molecule as a whole. For example, the CH₂ methylene group (with a C_{2v} pointgroup) has a symmetric stretch (A₁ symmetry), and antisymmetric stretch (B₂ symmetry) and a bending mode (A₁ symmetry). The symmetric stretch and bending mode can appear in xx , yy or zz polarisations, while the antisymmetric stretch can only appear in the yz polarisation¹²⁶ (remembering that the three axes are orientated to the molecule, rather than the lab). Experimentally these are best viewed as polarised (incident polarisation = emitted polarisation) and depolarised (incident polarisation \neq emitted polarisation).

Expanding the transition polarisability tensor operator using a Taylor series—as with equation 1.100 above—gives,¹²⁵

$$\boldsymbol{\mu}_{fi} = \left(\boldsymbol{\alpha}_0 \langle \psi_f | \psi_i \rangle + \frac{\partial \boldsymbol{\alpha}}{\partial \mathbf{q}_j} \Big|_0 \langle \psi_f | \mathbf{q}_j | \psi_i \rangle + \frac{1}{2} \frac{\partial^2 \boldsymbol{\alpha}}{\partial \mathbf{q}_j \partial \mathbf{q}_k} \Big|_{0,0} \langle \psi_f | \mathbf{q}_j \mathbf{q}_k | \psi_i \rangle + \dots \right) \mathbf{E}, \quad (1.112)$$

where q_j and q_k are coordinates along vibrational modes j and k . The zeroth order first term is zero since the wavefunctions are orthonormal so $\langle \psi_f | \psi_i \rangle = 0$. The second order (and higher) terms can normally be neglected. Assuming harmonic vibrational modes, the wavefunctions for the initial and final states can be expressed as Hermite polynomials, and¹²⁰

$$\langle \psi_f | q | \psi_i \rangle = A^2 N_{v_f} N_{v_i} \left[v_i \int_{-\infty}^{\infty} H_{v_f} H_{v_i-1} e^{-y^2} dy + \frac{1}{2} \int_{-\infty}^{\infty} H_{v_f} H_{v_i+1} e^{-y^2} dy \right], \quad (1.113)$$

where A is a constant depending on stiffness of the harmonic oscillator and relating $y = q/A$, v_i and v_f are the vibrational quantum numbers of the initial and final states, N_{v_i} and N_{v_f} are scaling constants and the H_v terms are the Hermite polynomials for quantum number v . The first integral is zero unless $v_f = v_i - 1$ and the second is zero unless $v_f = v_i + 1$, giving a selection rule of

$$\Delta v = \pm 1. \quad (1.114)$$

This selection rule arises because polarisability changes only once per oscillation (for rotational transitions the selection rule is $\Delta J = 0, \pm 2$ because the polarisability changes twice per transition).¹¹⁹ Anharmonicity—where either the polarisability is not proportional to displacement from equilibrium or the potential energy well is not harmonic—can relax the selection rule. The second and higher order terms of equation 1.112 represent multiple vibrational modes excited simultaneously, also a consequence of anharmonicity.

1.6.3 Review of TIR Raman

As discussed in more detail at the start of this chapter, information about the behaviour of surfaces is of great relevance to a wide range of physical problems. A wide range of analytical techniques exist to probe surfaces of which many—such as tensiometry, ellipsometry, quartz crystal microscopy, neutron reflectivity and atomic force microscopy—are not chemically specific meaning that the different species cannot easily be distinguished. For this reason the use of spectroscopic techniques is desirable. Some spectroscopic techniques, for example sum-frequency spectroscopy and second harmonic generation, are innately surface sensitive due to their selection rules. Others can be made surface sensitive by careful control of the region probed. This can sometimes be achieved by only collecting emitted light from certain regions (as in confocal microscopy) but is often better achieved by only passing the probe light through the region of interest. A common way of doing this is to use the evanescent wave from totally internally reflected (TIR) light, for example in attenuated total reflection infrared spectroscopy (ATR-IR),^{127–129} TIR-fluorescence microscopy¹³⁰ (although fluorescence spectroscopy usually provides little chemical detail without specific labels) and TIR-Raman spectroscopy, which this review deals with. A further surface sensitive technique is surface enhanced raman spectroscopy (SERS) where a large enhancement of Raman signal occurs at a limited range of metallic substrates.^{131,132} SERS can be combined with total internal reflection illumination. I exclude SERS from this review since we feel the analysis of the spectra is considerably more complicated than plain TIR-Raman since it is difficult to predict which spatial regions will be enhanced in SERS, whereas an evanescent wave is well understood.

TIR-Raman spectroscopy was first demonstrated by Ikeshoji and coworkers in 1973,¹³³ from carbon disulphide beneath a glass substrate (spectra are shown in figure 1.25). Subsequent early work looked at dyes adsorbed to a transparent electrode, showing a shift in Raman spectrum with potential;¹³⁴ dyes at the quartz–water interface;¹³⁵ dyes at the air–water interface;¹³⁶ and dyes at the CCl₄–water interface (including the determination of orientation from multiple polarisations).^{137,138} However, the signal was still weak—even with resonance enhancement from the dye—due to large amounts of Raman scattering from the substrate. Very little work has subsequently been done using TIR Raman at electrode interfaces, probably due to the prevalence of SERS for spectroscopy of metallic surfaces. Iwamoto *et al.* demonstrated the power of unenhanced TIR Raman for looking at adsorbed surface layers, showing that it was possible to record spectra from a 1.1 μm thick polystyrene film without seeing peaks from a thicker polyethylene film that it was mounted on, out of the range of the evanescent wave.^{139,140} Additionally, they demonstrated that it was possible to look at biological samples, and reported enhancement to the spectra due to the total internal reflection.¹⁴¹ Their key improvement was their careful choice of substrate: they picked a substrate that did not fluoresce and so reduced the background signal to manageable levels. Around the same time Hölzer and co-workers¹⁴² published a series of papers using TIR-Raman spectroscopy to study polymers attached to surfaces, including polyethylene (to determine the degree of crystallinity) and partially




Figure 1.25: The first reported TIR-Raman spectra, of CS_2 at the glass-liquid interface for a variety of angles of incidence. Reproduced from ref. 133 with permission.

polymerised styrene/polystyrene mixtures^{143,144} (to determine the degree of polymerisation). Pagannone and Suëtaka both studied the angular dependence of the Raman scattered radiation^{145,146}. The angular distribution of scattered light is a function of the polarisation of the emitted light,¹⁴⁷ the polarisation of incident radiation and the orientation of molecules in the probed layer, but the greatest signal is always measured at the critical angle (see figure 1.26 for an example).

In order to improve the sensitivity some groups used guided wave enhancement from a thin dielectric layer between the bulk substrate and the sample layer. In these cases the evanescent wave actually usually emerges from below the sample layer and the predominant enhancement comes from the multiple reflections within the sample layer,¹⁴⁸ although for thinner films the guiding layer is placed next to the sample and the evanescent field in the sample is probed instead.^{149,150} The advantage of waveguide Raman compared to TIR-Raman is that a long section of the interface can be probed, as opposed to a single reflection. Although waveguide Raman is a related technique, I consider it outside the scope of the current review; such applications are reviewed in references 151 and 128.

This review is structured as follows. First, I set out some of the general experimental considerations necessary when using TIR-Raman spectroscopy, especially with regards to the substrate. Second I compare TIR-Raman to competing techniques: principally ATR-IR, but also sum-frequency generation and non-chemically-specific techniques such as ellipsometry. Third, I look at the use of TIR-Raman at a variety of




Figure 1.26: Angular dependence of scattered Raman radiation from the 1530 cm^{-1} band of a 5 nm thick copper phthalocyanine film phthalocyanine film on glass. The solid line is a calculated dependence. The critical angle is labelled as α_c . Reprinted from reference 146, Copyright (1984), with permission from Elsevier.

different interfaces: the solid–liquid interface, where we discuss surfactant adsorption, lipid bilayers, liquid crystals, adsorption onto particles and water structure at the interface; the liquid–liquid interface, covering adsorption of fluorescent dyes; and the solid–solid interface, discussing tribological contacts, strained silicon surfaces, polymer thin films and biological surface. Finally I discuss depth profiling, a challenging but potentially useful extension of TIR-Raman spectroscopy, and compare it with confocal Raman spectroscopy which remains the current state of the art in this area.

Experiment

The TIR-Raman spectrometer at Durham University is detailed extensively in Chapter 2 of this thesis. Here I only look at general experimental considerations, and also briefly discuss an alternative set-up of TIR-Raman spectroscopy that has recently seen some use.

A range of substrates have been used successfully in TIR-Raman including fused silica, CaF_2 , SrTiO_3 , sapphire and a variety of high purity glasses. Ideally, the substrate should not have a strong Raman spectrum in the region of interest, since it will tend to dominate the recorded spectra (the surface selectivity does not apply to the material the probe light is incident though). For example, glasses including fused silica are clear in the C–H stretching region around 2900 cm^{-1} , but have strong peaks up to 1600 cm^{-1} , while sapphire is clear from around 800 cm^{-1} .¹⁴⁰ Where there is spectral overlap, crystalline compounds (such as sapphire) should often be preferred, since they have sharper peaks than amorphous substrates like glasses. A bigger problem with TIR-Raman spectroscopy is fluorescence from the substrate, which can often dominate over the Raman scattering. Many optical glasses are prone to such fluorescence due to impurities in the glass, as is sapphire at certain excitation wavelengths.

Use of a longer wavelength excitation can reduce fluorescence, at the cost of reduced Raman signal (due to the ν^3 dependence of Raman scattering). Pulsed lasers and time-gated excitation can also help reduce fluorescence,¹⁵² although this approach has not yet been applied to TIR-Raman. Additionally, use of a birefringent material is undesirable since it will give two critical angles depending on polarisation. The substrate must be transparent from the wavelength of the probe laser to the wavelength of emitted light (for a 532 nm probe laser the 2900 cm^{-1} C–H bands are emitted at around 630 nm). A small amount of absorption at the excitation wavelength in the substrate can lead to damage due to heating. For some substrates it is not possible to produce a flat, optically transparent prism for use in TIR Raman. In these cases an alternative option is to coat a transparent substrate with a thin film of particles; this is discussed further in the “solid–liquid” section.

An alternative approach is to use annular illumination,^{153,154} typically through a hemispherical lens (often termed a solid immersion lens or SIL). The beam passes along the standard beam path of the Raman microscope, and is delivered to the sample through the same objective lens that is used to collect Raman scattered light. If the numerical aperture of the objective lens is high enough then some of the beam will be incident above the critical angle. Placing a circular block in the beam path gives annular illumination, removing light that would reach the sample at less than the critical angle. The advantage of this technique is that the insertion of the circular mask into the sample is straightforward and needs few changes to the default beam path. The main disadvantages are that the polarisation of the incident beam is scrambled and that the angle of incidence is typically less well defined than for a custom beam path. Loss of much of the incident laser power at the mask is a further disadvantage, although is typically unimportant since the laser is often attenuated anyway to reduce damage to the sample. A simpler variant of this approach is to use a full angular range of illumination and collect the Raman spectrum arising from both the evanescent wave and the transmitted light,¹⁵⁵ providing the enhancement of the signal from the surface but no surface selectivity.

Comparison with other techniques

The most similar technique to TIR-Raman spectroscopy is ATR-IR spectroscopy, which measures the adsorption of infrared radiation within an evanescent wave. TIR-Raman has better surface selectivity, since the penetration depth is proportional to the wavelength of light in the prism, which is in the visible or near-IR for Raman scattering and in the mid-IR for ATR-IR. TIR-Raman spectra are also simpler to interpret for two reasons: first, in ATR-IR the penetration depth varies with wavelength, and second, strong bands in ATR-IR are distorted by the rapid variations in the refractive index of the sample. Both TIR-Raman and ATR-IR have the sensitivity to detect monolayers of organic molecules in a single spot. TIR-Raman offers two significant simplifications in the analysis of the recorded spectra: first, absorption of the light in the evanescent wave can usually be neglected in Raman, while it is necessary to record an ATR-IR spectrum (when it is small a multiple bounce arrangement can be used); second the penetration depth of TIR-Raman is a single value, since a monochromatic

	Raman cross-section / $\text{cm}^{-2} \text{ molecule}^{-1} \text{ sr}^{-1}$	IR cross-section / $\text{cm}^{-2} \text{ molecule}^{-1}$
water	4×10^{-30} (550 nm excitation) 4×10^{-28} (250 nm excitation)	1×10^{-18}
C–H	6×10^{-29} (532 nm excitation) 3×10^{-24} (239 nm excitation)	3×10^{-19}

Table 1.5: Representative absorption/scattering cross-sections for Raman and IR spectroscopy. Raman cross-sections of water are for the 3400 cm^{-1} band,¹⁵⁷; Raman cross sections for the C–H stretch¹⁵⁸ are for cyclohexane between $2850\text{--}2950 \text{ cm}^{-1}$; the IR cross-section of water¹⁵⁹ is measured at 3450 cm^{-1} ; the IR cross section for the C–H stretch is for *n*-heptane¹⁶⁰ at 2900 cm^{-1} .

probe is used, whereas for ATR-IR it varies across the spectral range.

For applications at the solid–water interface, one of the most important considerations is the intensity of the water background, and this is one of the areas where Raman and IR differ most strongly. Table 1.5 compares Raman and IR cross-sections of the O–H stretch modes of water and C–H stretching modes in hydrocarbons. For IR, C–H bands have a $3\times$ lower cross-section than the O–H band, whereas for Raman the cross-section is around $10\times$ greater (with excitation in the visible). Therefore, although Raman scattering is a comparatively weak effect compared to infrared absorption, it has much better sensitivity for hydrocarbons relative to water than infrared, making it preferable when studying aqueous solutions. Using a deuterated solvent or sample can also be an option to reduce overlap with solvent peaks in both Raman and ATR-IR, although it does add to the cost of the experiment.^{47,51} A further issue is the small signal from a single bounce of the infrared beam; when studying rapidly changing processes multiple bounces are usually necessary to obtain sufficient signal.^{47,156} Increasing the sampling area often increases the amount of material that must be used and—for kinetic measurements—increases volume of solution that must be rapidly and controllably exchanged if the results are to be meaningful.

As mentioned earlier, Raman spectroscopy is very sensitive to fluorescence, which can dominate the spectra if present. Often fluorescence can be avoided by careful choice of materials or the wavelength of probe light. While the average fluorescent background can be subtracted from Raman spectra, the noise in the background cannot; it is this noise that degrades the quality of the Raman spectrum. ATR-IR has the advantage that it does not suffer from fluorescence. The problem of fluorescence does not just apply to the substrate: Beattie *et al.* found the surfactant 2-mercaptobenzothiazole fluoresced strongly when adsorbed as a dimer,³ whereas such problems were not observed in ATR-IR.¹⁶¹ IR does have the requirement that the material the light is incident through must be transparent in the infrared, which is often more limiting than the requirement of Raman that the material is transparent close to the wavelength of the probe light (usually visible).

For studying the thickness of single component layers, ellipsometry, optical reflectometry and the quartz crystal microbalance (QCM) can provide a similar quality of data to TIR-Raman, often with higher sampling rates where kinetic processes are important. However, when more than one component is present none of these techniques offers any way to distinguish the components, and so a spectroscopic technique

becomes essential. Additionally, QCM is sensitive to water trapped in the adsorbed layer, which complicates analysis and reduces the accuracy of measured surface excesses.

A further group of spectroscopic techniques are often used for studying surface behaviour: infrared–visible sum-frequency generation^{162–167} (SFG) and second harmonic generation. Both of these are genuinely surface selective in that their selection rules ensure that no signal is obtained from the bulk sample since most bulk materials are SF-inactive. Both methods are sensitive to the symmetry of the surface layer, and will yield no signal from a centrosymmetric layer, which can make the determination of an adsorbed amount impossible. In some of the examples I will give, TIR-Raman and sum-frequency spectroscopy will be used together to provide complimentary information, since Raman is readily converted to a surface excess while sum-frequency is more sensitive to structure

SFG is also the preferred tool for probing structure at the air–water interface:¹⁶⁵ the relative intensity of the peaks from ‘ice-like’ and ‘water-like’ O–H stretches provides local structural information and non-hydrogen bonded O–H bonds at the surface give a characteristic sharp peak at high wavenumbers; the tilt of O–H oscillators to the surface can be inferred from the polarisation-dependence of the SF signal; surface charge can be inferred from the phase of the water signal in the electrical double layers; there is no competing signal from bulk water. While Nickolov *et al.* used TIR Raman to compare water structure at hydrophilic (sapphire) and hydrophobic (an ω -tricosenoic acid monolayer) interfaces,^{168,169} based on the intensity and position of the ice-like and water-like peaks around 3200 cm^{-1} and 3420 cm^{-1} , the information from SFG is both more detailed and more surface sensitive.

Third-order nonlinear effects, such as coherent anti-Stokes Raman scattering (CARS) and stimulated Raman scattering (SRS), have been developed into chemically sensitive microscopies with sub-micron spatial resolution and with enhanced sensitivity compared to linear Raman scattering. Unlike SFG, CARS and SRS have no intrinsic surface sensitivity and any advantages as surface spectroscopies over TIR Raman would arise from improved signal levels or better spatial resolution (arising from the nonlinear mixing processes involved in the signal generation). SRS also been combined with TIR illumination to look at the water structure at the air-water interface.¹⁷⁰ The enhancement occurs transiently before the formation of a plasma from the liquid induced by the excitation laser; these unusual experimental conditions may not represent the properties of room temperature water. Both SRS¹⁷¹ and CARS^{172,173} have been combined with waveguide excitation to achieve surface sensitivity.

Solid–liquid and solid–air interfaces

Tyrode *et al.* used a combination of TIR-Raman and sum frequency spectroscopy to study the adsorption of a cationic surfactant onto silica.⁴⁰ TIR-Raman provided a convenient means of determining the surface excess, to a precision of 0.01 mg m^{-2} (1% of the saturated surface layer). The spectra were largely insensitive to surface concentration (there is a small change in one part of the spectrum), and the relative intensities of peaks in different polarisations did not vary with concentration, showing

that the orientation and local structure of the molecules in the adsorbed layer were independent of surface coverage. SFG provided complementary information about the conformation of CTAB at the silica interface, with a lack of C–H signal indicating that the adsorbed molecules were in a centrosymmetric environment. The combination of TIR-Raman and sum-frequency generation was particularly powerful, since SFG alone provides no information about the amount of surfactant adsorbed, while TIR-Raman does not provide the same level of detail about conformation at the interface as sum-frequency. For work on surfactants adsorbed at the solid–liquid interface, TIR-Raman often provides little information of the order or tilt of the surfactant chains, since the adsorbed layer is always quite disordered; in this regard sum-frequency spectroscopy has a clear advantage.

TIR-Raman has also been used to investigate planar supported lipid bilayers. Okamura *et al.* used TIR-Raman to study a saturated lipid (DPPC) film (9 monolayers thick) at the quartz–air interface.¹⁷⁴ The ratio of intensities of various C–H bands allowed the study of conformational order (the ratio of gauche and trans conformers) with respect to the surface pressure at which the lipid was deposited onto the substrate. Subsequently, Lee *et al.* studied phase transitions of lipids with temperature¹⁷⁵ as well as the formation of lipid bilayers from mixed lipid surfactant systems.¹⁷⁶ In the latter case, deuteration of the lipid was used to distinguish the two components. Lipids are structurally similar to organic surfactants so the experimental sensitivity is somewhat better than that obtained for surfactant adsorption. Lipid bilayers are more ordered than adsorbed films of simple detergents studied and TIR-Raman spectroscopy provides conformational markers with which to probe this order. The ratio of the intensities of the methylene antisymmetric stretch (d^-) and symmetric stretch (d^+) bands increases with increased chain order, while the ratio of the intensity of the d^- band and the Fermi-resonance of d^- and the methylene scissoring mode changes in the opposite manner. The ratio of the intensities of the S- and P-polarised spectra for both the CH₂ stretching (2800–3000 cm⁻¹) scissoring (1450 cm⁻¹) modes provides an indication of chain tilt, with a greater intensity in the S- polarised spectra indicating that the chains are orientated closer to the surface normal. The C–H twisting mode at 1300 cm⁻¹ shows the opposite dependence.¹⁷⁵ Shifts in vibrational frequency can also provide information about chain order. Such spectral markers are not unique to TIR-Raman and are also available in other vibrational spectroscopic techniques such ATR-IR¹⁷⁷ and sum-frequency spectroscopy¹¹⁷ (the markers in SFG are different, but convey similar information). The use of S- and P-polarisations is a unique feature of TIR/ATR-based techniques since the surface provides a reference plane for the orientation of molecules; the tilt of molecules cannot be obtained from polarised spectra of lipid vesicles in the bulk.

Morikawa and coworkers used TIR-Raman to study the behaviour of a nematic liquid crystal at a solid surface, including the time dependence of the structure following a change in electric field.¹⁷⁸ The spectra were taken through a glass substrate coated with a transparent 50 nm indium tin oxide electrode then covered with a 60 nm silicon oxide layer. They determined the director angle in the liquid crystal molecule from the intensities of different Raman polarisations (see figure 1.27). The kinetics of reorientation following a brief application of an electric field could be measured




Figure 1.27: (a) Time resolved TIR Raman spectra of the liquid crystal 4-*n*-hexyl-4'-cyanobiphenyl measured using Px polarisation at a variety of time delays (shown on the graph) after application of an electric field. Spectra were accumulated (shown on the graph) over many pulses (up to 15 mins at 5 or 10 Hz) with the same time delay after the application of the electric field. (b) Director rotation angles for the same liquid crystal derived from the intensity ratios of different Raman polarisations. The three different symbols represent different amplitudes of the applied electric field. Taken from ref. 178. Reprinted from “Time-resolved Total Internal Reflection Raman Scattering Study on Electric-Field-Induced Reorientation Dynamics of Nematic Liquid Crystal of 4-Hexyl-4'-Cyanobiphenyl”, ref. 178, with permission from Taylor and Francis, Ltd..

on a 1 ms timescale: the electric field was pulsed periodically, and Raman scattering was excited by a 15 ns laser pulse with a variable delay. Comparison with transmission IR spectroscopy showed a difference between the behaviour of the surface layer (probed with TIR Raman) and the bulk (probed with transmission IR). TIR Raman worked well for this application because the depth of penetration of the evanescent wave (~ 50 nm) was comparable to the thickness of the liquid crystal perturbed by the interface. The authors noted that they were still seeing averaged behaviour within the evanescent wave, and suggested that more detail might be resolved by varying the angle of incidence (see section 1.6.3).

Willems and co-workers looked at individual multi-walled carbon nanotubes held perpendicular to the surface within the evanescent field by affixing them to an AFM tip.¹⁷⁹ By varying the angle to incidence they were able to vary the depth sensitivity of the Raman scattering from the nanotube. The spectra showed two bands: a G band intrinsic to multi-walled nanotubes, and a D band associated with defects in the structure. The intensity of the D band increased at shorter penetration depths showing that there were more defects closer to the end of the nanotube. The spectra were also strongly sensitive to the orientation of the nanotube to the surface, with no signal usually seen for S polarised incident light and a nanotube perpendicular to the surface. The combination of AFM and TIR Raman provides a novel means of probing a single microscopic object on length scales below the diffraction limit.

When looking at the solid-liquid interface (or the solid-solid interface) only a limited range of substrates can be used; the constraints are discussed in the experimental




Figure 1.28: TIR-Raman spectra of a 5 mM 1:1 mixture of 2-mercaptobenzothiazole and dibutyl dithiophosphate adsorbing onto sphalerite. Spectra are shown at 5 min (light grey line), 15 min (dark grey line) and 25 min (black line). (a) the CH stretching region and (b) the fingerprint region. Reprinted from ref. 3, Copyright (2006), with permission from Elsevier.

section. The use of thin layers of non-transparent particles can relax these constraint (contact with the aqueous sample solution limits the laser-induced heating that takes place). Beattie and co-workers investigated the adsorption of a variety of surfactants onto sphalerite,³ a form of ZnS with Fe impurities. Figure 1.28 shows the adsorption of a mixture of 2-mercaptobenzothiazole and dibutyl dithiophosphate; the rate of adsorption was relatively slow, which is often the case for adsorption onto particulate surfaces. This method of studying adsorption onto non-transparent particles has been used extensively with ATR-IR,¹⁸⁰ but only rarely with TIR-Raman. In contrast to TIR Raman on planar substrates it has the disadvantages that the penetration depth of the evanescent wave is less well defined, due to interaction of the evanescent field with the particulate film; and no orientational information can be obtained, since the adsorbed layer is adsorbed on all sides of the particles rather than orientated relative to the flat substrate, which were randomly orientated to the excitation light.

Liquid–liquid interface

The liquid–liquid interface is extremely challenging to study using TIR-Raman spectroscopy, and to my knowledge only two groups have published work demonstrating it thus far. One of the key problems is the strong spectra from the incident medium: it is usually possible to select solid substrates with non-overlapping spectra, but most organic solvents have strong spectra that swamp those of organic surfactants. One approach to maximising the signal from the surfactants is to use resonance Raman, which often requires the use of aromatic organic dyes. All studies at the liquid–liquid interface used this method. Fujiwara and Watarai looked at the adsorption of the dye *meso*-tetrakis(N-methylpyridyl)porphyrinatomanganese(III) adsorbing to the toluene–aqueous interface in the presence of an anionic surfactant.¹⁸¹ Figure 1.29 shows the design of the optical cell used at the liquid–liquid interface. Only the resonantly enhanced manganese complex and not the surfactant was seen in the Raman spectra. Yamamoto and Watarai then went on to study tetraphenylprophyrin at the dodecane–water interface using a variety of evanescent wave scattering techniques including TIR-Raman,¹⁸² observing peak shifts depending on the counterions present.




Figure 1.29: (a) Optics configuration and (b) optical cell design for resonant enhanced TIR Raman spectroscopy at the liquid–liquid interface. Reprinted with permission from ref. 181. Copyright 2003 American Chemical Society.

However the study was mainly concerned with domain formation in the adsorbed layer which was imaged using TIR Rayleigh scattering. An additional approach is to select a non-overlapping incident medium, such as CCl_4 . Takenaka and Nakanaga recorded the adsorption of a mixture of an number of dyes onto the CCl_4 –aqueous interface.^{137,138,183} For the dye methyl orange co-adsorbing with CTAB,¹³⁷ the tilt of the dye molecules in the monolayer was estimated at 50–60° based on the relative intensities of polarised spectra. The absolute adsorbed amount was obtained from surface tension measurements and not from the Raman intensity. It was only possible to study the resonantly enhanced dye, and not the CTAB. An interesting experimental enhancement used in these papers was an extra mirror to increase the laser power by directing the totally reflected light back at the sample, and another extra mirror to capture Raman scattering emitted away from the objective.

TIR-Raman spectroscopy is not currently the preferred technique at the liquid–liquid interface. Instead, genuinely surface selective techniques such as second harmonic generation or sum-frequency spectroscopy allow the interface to be probed without any contribution from the incident medium. Alternatively total internal reflection fluorescence spectroscopy can be used when the surfactant contains a fluorophore, since most pure solvents do not fluoresce. Perera and Stevens recently reviewed spectroscopic techniques for looking at the liquid–liquid interface,¹⁸⁴ and a wide range of references can be found there. Finally, non-spectroscopic techniques such as ellipsometry are viable for probing the liquid–liquid interface.¹⁸⁵

Liquid–air interface

The liquid–air interface has barely been studied using TIR-Raman spectroscopy since the surface selectivity of the evanescent wave is working in the wrong direction: the liquid contributes strongly to the Raman signal while the evanescent wave decays into air. External reflection infrared spectroscopy^{186,187} and SFG^{164,188} are more suitable techniques, as both show surface selectivity. Takenaka and coworkers used TIR-Raman to look at monolayers on water of deuterated stearic acid and deuterated cadmium stearate under various amounts of compression.¹⁸⁹ Acquisition times were between 1 and 2 h and it was necessary to use deuterated layers to avoid spectral overlap between the O–H and C–H stretching bands. Greene conducted a preliminary study of the adsorption of CTAB at the air–water interface, and found the C–H peaks of CTAB barely visible above the background from water.¹¹⁸

Solid–solid interface

TIR Raman spectroscopy can be used to look at thin films at a solid–solid contact point. With monolayer films this provides a model of boundary layer lubrication. In order to get spectroscopic information about films of less than 10 nm the a TIR-based technique is essential. SFG has also been used at such interfaces.¹⁹⁰ Reflection IR has been used to measure films with a minimum thickness of around 20 nm,¹⁹¹ but it is not yet possible to study monolayer films. Laser based spectroscopies have advantages when studying small areas and both SFG^{190,192} and Raman scattering^{117,193} have been used in TIR mode to study lubricated contacts under pressure and shear. An important limitation with all spectroscopic techniques is the need to have transparent substrates (either in the IR or visible region), whereas metal substrates are most relevant to real applications. A typical experiment at a solid–solid contact uses a high refractive index solid (for example sapphire or SF10 glass) in contact with a curved surface of lower refractive index solid (e.g. MgF₂ or silica). Either or both of the two solids can be coated with a thin film, often using Langmuir-Blodgett deposition. The Raman spectrum can be collected through either of the solids. Information about the ordering can be acquired based on the ratio of the symmetric and antisymmetric methylene stretch,¹¹⁷ as with the lipid systems discussed earlier. In addition to Langmuir-Blodgett deposited films, Beattie *et al.* looked at hydrocarbon layers as they were squeezed out from between two solid surfaces.¹⁹³ They were able to detect small pockets of hydrocarbon remaining in gaps in the rough surface at a quantity well below a monolayer.

A number of groups have looked at a strained silicon thin film under partially total internal reflected conditions using illumination through a high numerical aperture objective.^{155,194} Lermen and co-workers reported that for their optical alignment, approximately half the Raman signal came from the evanescent wave and half from transmitted light.¹⁵⁵ The shift in the characteristic silicon peak at 520 cm⁻¹ allowed stressed and unstressed silicon to be identified. Samples include the area around a scratch in a silicon wafer,¹⁹⁴ and a thin film of strained silicon separated from the bulk silicon substrate by a 150 nm SiO₂ layer.¹⁵⁵ In the latter case the peak from the bulk silicon was not completely eliminated, but was less prominent than when

illuminated without any total internal reflection.

Grausem *et al.* used light collected from a small optical fibre probe brought into an evanescent wave to perform high resolution Raman imaging of a patterned sample.¹⁹⁵ In the region around the collection probe there is a frustrated evanescent wave (energy is transmitted into the tip). This experiment is primarily taking advantage of the superior lateral resolution available using near-field illumination, rather than the surface sensitivity arising from total internal reflection. Damage to the sample probe caused by the excitation laser is a problem in such measurements.

Mattei, Fornari and Pagannone predicted¹⁴⁵ and measured¹⁹⁶ Raman spectra from surface polaritons at the sapphire–NaClO₃ interface, showing a shift in frequency with azimuthal angle of scattering. Similar work looked at surface polaritons at the CaF₂–air interface (illuminated by an evanescent wave from the glass–air interface across a thin gap).¹⁹⁷

The final use of TIR Raman spectroscopy at the solid–solid interface is to look at thin films of one substance on top of a second thicker layer; the optical element is in contact with the thin film. Such experiments can be seen as a precursor to the depth profiling experiments described in the next section, however most of these experiments were conducted simply to demonstrate the efficacy of TIR Raman in detecting only the thin layer. Iwamoto and coworkers were the first to demonstrate such selectivity using a thin film of polystyrene on a polyethylene base.^{139,140} Ishizaki and Kim looked at two different polymer systems: a thin films of polystyrene on poly(vinyl methyl ether) and a fluorinated polyimide.¹⁹⁸ In the latter case near infrared excitation light was need to avoid fluorescence from the sample. Greene looked at polyethylene naphthalate on polyethylene terephthalate (PET)¹⁹⁹ and Michaels looked at a thin poly(3,4-ethylenedioxythiophene)/poly(styrenesulfonate) (PEDOT:PSS) film on PET.¹⁵⁴ Michaels acquired Raman images (figure 1.30) by scanning the hemispherical lens and taking spectra at each point, revealing a partially covered surface. The choice of samples is not limited to synthetic polymers: Greene and Bain looked at epicuticular leaf waxes *in vivo*.²⁰⁰ They used the ratio of the symmetric and anti-symmetric methylene stretches to distinguish crystalline waxes in the top layer from amorphous waxes when the top layer had been removed with a cellulose acetate (figure 1.31). The limited penetration depth (40 nm) achieved with a cubic zirconia reflection element avoided fluorescence from the underlying pigments, which is often a problem in Raman spectroscopy of plant material.

Depth profiling

The variation in the penetration depth with angle of incidence raises the prospect of using TIR Raman spectroscopy for depth profiling thin film on the 100-nm level. The Raman signal, I , from a specific angle of incidence, θ_i , and Stokes shift, $\bar{\nu}$, is given by²⁰¹

$$I(\theta_i, \bar{\nu}) = \int_0^\infty C(z, \bar{\nu}) |E(z, \theta_i)|^2 dz, \quad (1.115)$$

where E is the electric field from the incident light and C depends on the Raman cross-section and the collection efficiency. Both E and C are dependent on the material profile with z , which affects the cross-section and variation in E with depth. The




Figure 1.30: TIR Raman image of a micropatterned PEDOT:PSS film on PET, generated from the peak areas of (a) the PEDOT:PSS marker band at 440 cm^{-1} and (b) the PET marker band at 1726 cm^{-1} . The scale bar is $1\text{ }\mu\text{m}$. The image was acquired by scanning the position of the excitation/collection optics across the sample. Reproduced with permission from ref. 154.




Figure 1.31: TIR-Raman spectra from the epicuticular wax layer of a leaf. (a) Crystalline wax of the top layer and (b) amorphous wax below the crystalline wax. Reprinted from ref. 200, Copyright (2005), with permission from Elsevier.

equation above is a Fredholm integral of the first kind, and in general is extremely difficult to solve.

A limiting case that simplifies the analysis greatly is when the refractive index is constant in the low-index medium. This case is realistic for many situations: for example, many polymers have refractive indices around 1.5. In this case the intensity of Raman scattering is

$$\frac{I(\theta_i, \bar{\nu})}{t(\theta_i)^2 E_0^2} = \int_0^\infty C(z, \bar{\nu}) e^{-2k_t \beta(\theta_i) z} dz, \quad (1.116)$$

where β is defined in equation 1.83, and t is one of the expressions given in equation 1.86 and depends on the polarisation used. This equation describes the Laplace transform of the function $C(z, \bar{\nu})$. One way of solving problems such as equation 1.116 is to divide the sample into a series of layers, j , of equal thickness, Δz . Then

$$\frac{I(\theta_i, \bar{\nu})}{t(\theta_i)^2 E_0^2} = \sum_{j=0} e^{-2k_t \beta(\theta_i) z_j} C(z_j, \bar{\nu}) \Delta z, \quad (1.117)$$

with the sum truncated at an appropriate value of z , where $e^{-2k_t \beta z}$ is negligible. In Raman depth profiling, spectra are taken at a range of incident angles in order to vary the transform variable, β . Eq. 1.117 can then be expressed in matrix form:²⁰²

$$\mathbf{a} = \mathbf{G} \mathbf{c}, \quad (1.118)$$

where \mathbf{a} is a column vector representing the experimentally measured Raman intensity at different incident angles (with elements $a_k = I(\theta_i|_k, \bar{\nu}) / (\theta_i|_k^2 E_0^2)$), \mathbf{c} is a column vector representing the Raman cross-section at different values of z and \mathbf{G} is a matrix with each row corresponding to a different incident angle, and each column corresponding to a different value of z (elements are calculated as $e^{-2k_t \beta z_j}$). Equation 1.118 applies for a specific wavelength of scattered light; however for simplicity the dependence of $\bar{\nu}$ is no longer written. Superficially, the problem of calculating the material dependence with z requires a straightforward matrix inversion:

$$\mathbf{c} = \mathbf{G}^{-1} \mathbf{a}, \quad (1.119)$$

however, the matrix \mathbf{G} is ill-conditioned and so the calculation is extremely sensitive to tiny experimental error in \mathbf{a} . The ill-conditioning arises because each row of \mathbf{G} has a very similar weighting of experimental response with respect to depth.

The solution is to regularise \mathbf{G} (Tikhonov regularisation is commonly used and illustrated here, but other approaches exist): a pseudoinverse for \mathbf{G} is created using²⁰²

$$\mathbf{G}^+ = (\mathbf{G}^T \mathbf{G} + L^2 \mathbf{I})^{-1} \mathbf{G}^T, \quad (1.120)$$

so that,

$$\mathbf{c}^+ = \mathbf{G}^+ \mathbf{a}, \quad (1.121)$$

where L is a scalar parameter and \mathbf{c}^+ is an approximation to the actual experimental response with respect to depth, \mathbf{c} . Larger values of L decrease the sensitivity of




Figure 1.32: Modelled intensity dependence with respect to angle of incidence for a range of different film thicknesses. (a) Uniform films adjacent to the surface with a thickness, t ; (b) Uniform films stretching from t_{start} to $z = \infty$. $n_i = 1.72$ (SF10 glass), $n_r = 1.5$ (a typical polymer). The dashed vertical line shows the critical angle.

\mathbf{c}^+ to experimental errors in \mathbf{a} , at the cost of a less detailed reproduction of \mathbf{c}^+ . Details of the choice of value of L and the regularisation process are given in refs. 203 and 204. The effect of the regularisation process is similar to representing \mathbf{c} with a truncated Fourier series. For the similar problem of depth profiling using in ATR-IR, the problem is sufficiently ill-defined that only very smooth concentration profiles can be modelled.²⁰²

To understand why the errors are amplified, \mathbf{G} is split using a singular value decomposition (or equivalently a singular value expansion for continuous functional forms such as that given in equation 1.116):

$$\mathbf{U}\mathbf{S}\mathbf{V}^T = \mathbf{G}. \quad (1.122)$$

The inverse of \mathbf{G} is,²⁰²

$$\mathbf{G}^{-1} = \mathbf{V}\mathbf{S}^{-1}\mathbf{U}^T. \quad (1.123)$$

In \mathbf{G}^{-1} , the weighting of single column of \mathbf{V} , \mathbf{v}_i , is given by μ_i^{-1} , where μ_i is the i th singular value (taken from the diagonal of \mathbf{S}). For the Laplace transform described above, μ_i decreases very rapidly with i , causing a massive amplification of any experimental error.

Modelling the experimental response to varying depth profiles illustrates the limitations of TIR Raman for depth profiling. Figure 1.32 shows the intensity dependence with respect to angle of incidence for a variety of film thicknesses (assuming a sharp step in the concentration profile). The dependence is strong enough that it should be possible to fit a characteristic thickness to a layer, especially if the signal from an infinitely thick layer is known and can be used as a reference. The dependence with respect to the functional form of the concentration profile is much weaker; figure 1.33 shows that there is almost no sensitivity to a blurring of the interface (modelled using $c = \frac{1}{2}(-\text{erf}((z - 50 \text{ nm})/30 \text{ nm}) + 1)$), and only a slight difference compared to an exponential concentration profile.

Due to the difficulty in deriving unique depth profiles from TIR Raman, very little




Figure 1.33: (a) Concentration profiles with respect to depth, for a step function (solid blue), an exponential decay (dashed green) and a profile derived from the error function (dotted red). (b) Modelled Raman intensity with respect to angle of incidence arising from those concentration profiles. In all cases the characteristic thickness is 50 nm.




Figure 1.34: The angle dependence of the Raman intensity from two thick polymer layers around the critical angle (marked as a vertical dotted line). The inset focusses on small variations from the critical angle. The dashed line shows the expected signal from a single angle of incidence while the solid line shows the expected signal calculated using a spread of incident angle which cause non-evanescent fields in the sample. Reproduced from ref. 201 with permission.

work has been published on the subject. Initial work using variable angles looked at homogeneous samples,^{140,201} and aimed to recover the correct dependence on Raman signal with respect to angle of incidence. However, Fontaine and Furtak found that even in this simplest case there was a wide range of experimental difficulties that affect the measurement.²⁰¹ For example, back reflections from the exit side of a symmetric optical element can strongly affect the intensity of the electric field in the sample. Additionally, when very close to the critical angle the small spread of angles in the incident beam can lead to some light being transmitted, which has a large impact on the spectrum since it samples to an infinite depth (as shown in figure 1.34); strategies to mitigate this effect are discussed in ref. 201.

Rather than attempting the matrix inversion as required by equation 1.121, the more common approach is to model the sample as a small number of well-defined layers, changing the layer thicknesses to fit the data. Iwamoto deliberately picked

two-layer samples with identical refractive indices, and obtained reasonable estimates of the thickness of the first layer.¹⁴⁰ However, there was a considerable difference between the predicted and measured values for the angle dependence of the spectra, attributed to the spread of angles.

Fontaine and Furtak subsequently looked at two-layer polymer films and were able to identify the position of the break between the layers with a precision of ± 2 nm.²⁰⁵ It should be noted that this latter study mostly exploited the position of waveguide modes below the critical angle, rather than the changing penetration depth above the critical angle. A recent development is the production of a scanning angle Raman spectrometer,²⁰⁶ allowing greater automation of the measurement of Raman spectra at a range of angles. For a homogeneous benzonitrile sample the measured Raman signal agreed well with the expected angular variation of signal, providing the angular spread of the converging incident light was accounted for.

Other total internal reflection techniques could also be used for depth profiling. The use of ATR-IR is slightly more mature than for TIR Raman, however the majority of work within the last decade has looked only at qualitative changes in spectra with angle of incidence and has not attempted a full depth profile.^{207,208} Analysis is more complicated than for TIR Raman since the refractive index (and hence the penetration depth) varies across the spectral range, whereas for Raman there is only one incident wavelength. Additionally when absorption of incident light by the sample is appreciable the Laplace transform (eqn. 1.116) no longer applies. A range of attempts have been made to extract quantitative depth profiles from ATR-IR,²⁰⁹ often by assuming set functional forms for the concentration profiles with respect to depth, for example exponential²¹⁰ and step-function forms^{211,212} (both the exponential and the step function have an analytical solution for the Laplace transform).

For depth profiling, the most important alternative to TIR Raman spectroscopy is confocal Raman microscopy. A recent review by Everall discussed the limitations of confocal Raman microscopy in some detail:^{213,214} the depth resolution is approximately 1–2 μm . Much better depth resolution (see Fig. 1.32) can potentially be achieved with TIR Raman; however a quantitative analysis is only feasible for simple cases where the functional form of the concentration profile is well-known. At one level, TIR Raman and confocal Raman are complementary: TIR Raman can only probe structure near a planar interface ($z \leq \lambda$), but with a surface sensitivity of 100 nm or less; confocal microscopy has poorer depth resolution but can study much thicker samples, limited only by the working distance of the objective.

Chapter 2

Experimental Details

This section details the experimental set-up used throughout this thesis. Where there are specific deviations from the scheme outlined here then they are detailed in the chapter to which they relate.

2.1 Raman spectroscopy

2.1.1 Apparatus description

The history behind total internal reflection (TIR) Raman spectroscopy and the general principles of operation have been discussed in the introductory chapter. The TIR-Raman system used here has been described in ref. 40. However, it will be detailed in full within this chapter. The optical set-up is shown in figure 2.1.

The pump laser is a continuous-wave, frequency-doubled solid-state laser (Opus 532, Laser Quantum, Manchester, UK) that produces horizontally polarised light with a wavelength of 532 nm. For measuring multicomponent kinetics, the laser was typically operated at 1.5 W, yielding ~ 1.0 W at sample. The high laser power was necessary to obtain good enough signal to resolve multiple components within short acquisition times. For single component experiments the laser was often operated at




Figure 2.1: The optical set-up, viewed from above. This diagram is not to scale.

0.7 W, since the extra signal was not necessary to analyse the results. In order to ensure that the high laser powers were not affecting the adsorbed surfactant experiments were conducted with the surfactant tetradecyltrimethyl ammonium bromide ($C_{14}TAB$)* to see if increases laser power caused changes in the observed spectrum. Increasing the laser power to 1.0 W, 1.3 W and finally 1.5 W while continuously acquiring spectra revealed no change in the observed spectra except for the expected linear relationship between laser power and Raman signal. The final power of 1.5 W was maintained on the sample for 30 min with no sign of degradation of the surfactant. None of the surfactants used in this thesis are expected to be damaged by 532 nm light and none showed any sensitivity to laser power. Mica substrates were damaged by the laser; this is covered in chapter 7.

The beam passes through a $\lambda/2$ -plate, which can be set to either change or preserve the polarisation of the laser, allowing the incident light to be either S or P polarised. Following the $\lambda/2$ -plate, the beam is expanded using a telescope consisting of a -25 mm and a $+125$ mm focal length lens. Following the telescope is an iris which is left open during normal operation of the experiment, but which can be contracted to aid with alignment. The beam is then deflected up using a periscope, and gently focussed using a gradient index lens ($f = 120$ mm) to a diameter of ~ 10 μm on the sample surface. The Raman light is collected with commercial Raman microscope (Ramascop 1000, Renishaw, Wootton-under-edge, UK).

Graph paper placed under the microscope provides a marker so that a vertical ruler (with a cubic support at its base allowing it to stand upright) can be placed at a known position on either side of the microscope (typically 15 or 16 cm), allowing the height of the laser—and thus the angle—to be set accurately. The height of the focus of the microscope is 20.5 cm above the plate it is mounted on. Throughout the alignment process the iris in the beam-path is contracted, producing a ~ 1 mm diameter spot for greater consistency in positioning. Once the correct angle is selected the mirrors controlling the angle are no longer changed. Once the incident angle is set, the next stage is to insert the gradient index lens into the beam path using translation movements to control its position and rotation to ensure it is normal to the beam (determined from the back reflection). It is also important to ensure that the laser focus roughly coincides with the plane of the microscope focus: the lens mount allows the focus to be fine-tuned over a range of roughly 1 cm. Once the lens is in place, the laser spot is centred on the hemisphere by changing the angle of the beam perpendicular to the plane of incidence and by using a translation stage which adjusts the height of both the top mirror of the periscope and the gradient index lens simultaneously; the angle of the beam within the plane of incidence is not changed. Finally, the laser spot is focused at the surface using the fine adjustment of the lens holder.

A schematic of the Raman spectrometer is shown in figure 2.2. In brief²¹⁵

1. the light passes through a pair of long-pass edge filters,[†] which remove the

*This surfactant was used because the test coincided with an experiment involving the surfactant (results of which will appear elsewhere but not in this thesis); there was no other specific reason for the choice of surfactant.

[†]The spectrometer manual describes holographic notch filters, but these have subsequently been replaced with edge filters.




Figure 2.2: Schematic of the beam path of the collection optics and within the Raman spectrometer.

Rayleigh scattered component;

2. the light passes through a half-wave plate and a polariser, which can be moved in and out of the beam-path to remove specific polarisations of the scattered light;
3. the light is focused through a slit (width set to $200\ \mu\text{m}^{\ddagger}$) by a lens, thereby ensuring that divergent light from out-of-focus areas of the sample is removed;
4. the light is reflected by a prism onto a diffraction grating ($1800\ \text{lines}\ \text{mm}^{-1}$), which disperses light according to its wavelength;
5. finally the light passes through a focussing lens then strikes a charge coupled device (CCD) camera, which records the intensity of light across it. The lens ensures that all light with the same angle (i.e. the same frequency) is focussed to the same point on the CCD irrespective of where on the grating it originated from.

The solid substrates were 10-mm diameter fused silica hemispheres (scratch/dig ratio 40/20), obtained from Global Optics (Bournemouth, UK). Hemispheres were chosen to minimise optical aberrations. Both UV- and IR-grade hemispheres were used and produced identical results (the grade refers to regions of the spectrum where the hemisphere is transparent and is accomplished by a slight variation of the silica composition). Silica was selected as the substrate since it is often used as a model surface for surfactant adsorption, it is optically transparent, non-fluorescent and has a Raman spectrum that does not overlap with C–H stretching bands around $2900\ \text{cm}^{-1}$. Below $1300\ \text{cm}^{-1}$ silica has a Raman spectrum with intense, broad peaks and so is not appropriate when this spectral region is of interest. The angle of incidence at the silica-water interface was 73.0° giving an illuminated region of $30 \times 10\ \mu\text{m}$, a penetration depth for the electric field $d_p = 206\ \text{nm}$ and a sampling depth for Raman scattering of $d_p/2 = 103\ \text{nm}$. The incident laser was S polarised (perpendicular to the plane of incidence), since this polarisation gave the highest signal levels. The Raman

[‡]I do not believe the slit width actually corresponds to the values marked on the knob used to control it, but I think $200\ \mu\text{m}$ is reasonably close to the actual value.

scattered radiation was collected through the fused silica prism with a $50\times$ ULWD, 0.55 NA objective (Olympus) and directed into the spectrometer. Data were acquired over a fixed wavenumber range encompassing the CH stretching region (from 2600 to 3200 cm^{-1}) of the Stokes scattering.

The sample cell was mounted in a commercial upright microscope (Leica DM LM). The position of the cell could be adjusted using three perpendicularly orientated computer-controlled translation stages (Newport). The sample could be imaged optically using a camera mounted at the top of the microscope when illuminated in white light directed using a partially reflective mirror in a rotating filter magazine between the objective and the camera. A second rotating filter magazine contains a mirror to divert the collected light into the spectrometer. This mirror must be rotated out of the light-path for the sample to be viewed using the camera; it is not possible to simultaneously image the sample and record spectra of the sample.

The centre of the hemisphere can be found using a combination of two methods: firstly, by lowering the curved surface of the hemisphere below the plane of focus, then attempting to refocus the hemisphere by moving it parallel to the plane of focus—iteration of this method eventually finds the top of the hemisphere. Secondly, near to the centre of the flat surface of the hemisphere there is a floating spot of white light, caused by the hemisphere focussing the microscope illumination source. The spot is centred in the image of the sample, and then the height adjusted slightly to bring the hemisphere surface into focus. To ensure that the hemisphere was flat, the height of the laser spot is compared on both the incident and reflected beams at equal distances from the centre. The cell was mounted in a large optics mount allowing the angle to be adjusted.

2.1.2 Operating parameters and spectral assignment

The spectrometer has two modes of data acquisition. In the “extended scan” mode the grating is turned during the acquisition and data is read out continuously. This has the advantage that the spectral range is not limited by the range of wavelengths that can be fit on the CCD, but the disadvantage that rotation of the grating is a slow process so the rate of acquisition is limited. In the “single scan” mode the grating is held at a fixed position and the entire width of the CCD is read-out at once. The advantage is the entire width of the CCD can be recorded instantly (a range of Raman shifts of 660 cm^{-1}) but the disadvantage that the acquisition is limited to that range. In this work I have exclusively used “single scan” mode. In “single scan” mode the CCD functions as follows: a specific horizontal position on the CCD camera corresponds to a specific wavelength of light. Each column of pixels corresponds to Raman shifts with a wavenumber range 1.1 cm^{-1} wide. Signal from 10–15 rows of pixels is summed; these rows are chosen to encompass the height on the camera where the Raman scattered light arrives but to omit all other pixels which will only contribute to the noise (figure 2.3 shows this diagrammatically). The CCD is front illuminated; front illuminated CCDs typically have quantum efficiencies of around 35–40% in the visible region, meaning that only 35–40% of photons will be detected. The spectrometer records 1 detected photon as 1 count, but presents the data in terms

of counts per wavenumber so that spectra recorded using different gratings appear on the same scale.²¹⁶ Given that each pixel is 1.1 cm^{-1} wide, each “count” in the spectrometer software will on average account for 2.75 photons hitting the detector.[§] The spectrometer software displays two gain settings: “high” and “low”. Spectra recorded with both settings look close to identical; “high” was used since it is the default setting.

For measurements on equilibrium systems, a typical acquisition consisted of ten co-added scans of 30 s each, with two extra scans added by the spectrometer software’s automatic cosmic ray removal function (which removes sharp spikes that occasionally instantaneously appear in the spectrum). I did not use 1 scan of 300 s for two reasons: first, because the two extra scans for cosmic ray removal would then take 300 s each and second, because the water band would saturate the detector after 300 s. Using 30 s acquisitions was a good compromise since the detector would not be saturated at any laser power used or in any polarisation used, but the extra readout time and noise from multiple scans was kept sensibly small. Figure 2.4a shows the raw spectra of the two surfactants and is dominated by the O-H stretching mode of the water in the evanescent wave: even though the Raman scattering cross-section of the O-H bond is much smaller than the C-H bond, there are many more water molecules than surfactant molecules within the evanescent field. Figure 2.4b shows the surfactant spectra after subtraction of the water background. The main spectral features in the CTAB spectra (red, dashed line) are the symmetric and antisymmetric methylene stretches at 2852 and 2890 cm^{-1} respectively, the Fermi resonance of the symmetric stretch with the CH_2 scissoring mode at 2928 cm^{-1} ,²¹⁷ the antisymmetric methyl stretch at 2960 cm^{-1} , the symmetric methyl stretch from the trimethylammonium headgroups at 2985 cm^{-1} and the corresponding antisymmetric stretch at 3040 cm^{-1} .²¹⁸ The small peak at 2700 cm^{-1} is the overtone of a C-H bending mode. TX-100 is a nonylphenyl ethoxylate with 9–10 ethylene oxide (EO) groups on average. The Raman spectrum of TX-100 shows a broad and relatively featureless band between 2800 and 3000 cm^{-1} arising from the CH stretches of the alkyl chain and the EO groups. The distinct peak at 3080 cm^{-1} is the symmetric C-H stretch of the aromatic ring.

For the kinetics measurements a much shorter acquisition time of 1 s was used, with a 1-s readout time between each measurement limited by the spectrometer software. Figure 2.5 shows a set of 10 sequential spectra from a single component and a multicomponent kinetics experiment, illustrating both the signal to noise achieved in a 1-s acquisition and the amount of change in spectra observed over these experiments. A typical kinetics run would contain 300 spectra starting from before the injection of the surfactant until after the surface had reached equilibrium.

2.1.3 Polarisation selection

The coordinate system used in my experiment matches that illustrated earlier in the chapter with figure 1.19. S polarised light has an electric field vector perpendicular to the plane of incidence and P polarised light has an electric field vector parallel to the plane of incidence. S polarised light is pure y polarised when considered in

[§]In ref. 40 this was quoted as 5 photons per count; I believe this figure was calculated with an incorrect pixel width.




Figure 2.3: An image of the output from the CCD camera. Each column corresponds to one point on the spectrum (a range of 1.1 cm^{-1}). Only the area between the two red lines is actual read in any acquisition, since the remaining area contains no useful information so will only add to the noise while slowing down the readout time. The bottom part of the figure shows the resulting spectrum (note the inverted wavenumber-axis). 30 s acquisition time, 800 mW laser power, Sy polarisation. The spectrum is of a dirty hydrophobically coated (with 1,1,1,6,6,6-hexamethyldisilazane) silica-water interface.




Figure 2.4: Spectra for CTAB, TX-100 and water. a) Raw spectra; b) Spectra after subtraction of water background. S polarisation. Acquisition time = 360 s.

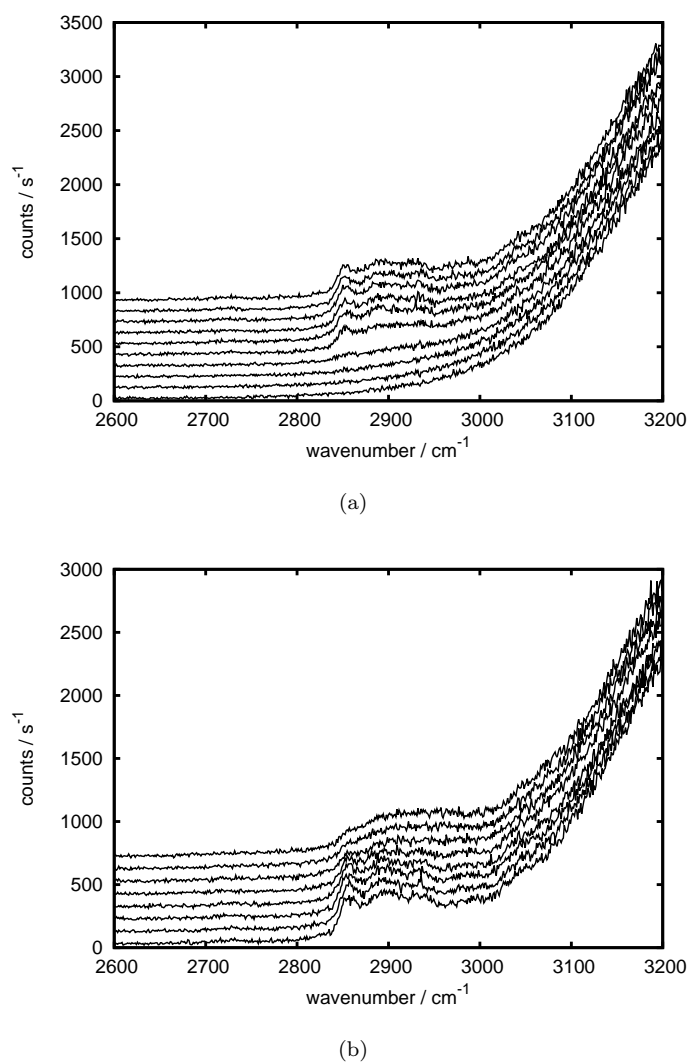


Figure 2.5: Sets of sequential spectra from (a) a kinetics experiment with 3 mM CTAB. 1-s acquisition time, S polarisation and (b) the replacement of 2 mM CTAB by 2 mM TX-100. The spectra are offset on the y-axis for clarity.

Experimental polarisation	Detected tensor elements	Equivalent tensor elements
Sy	α_{yy}	α_{xx}
Sx	α_{xy}	α_{yx}
Py	α_{yx} and α_{yz}	α_{xy} , α_{zy} , α_{xz} and α_{zx}
Px	α_{xx} and α_{xz}	α_{yy} , α_{zx} , α_{yz} and α_{zy}
—	α_{zz}	

Table 2.1: Elements of the Raman tensor probed by each of the four polarisation combinations used available in TIR Raman spectroscopy.

lab coordinates while P polarised light is a mixture of x and z polarisations, the weightings of which vary with angle of incidence (discussed in section 1.5.2). The polariser and half-wave plate in the spectrometer can collect two polarisations of scattered light: y , which is in the same plane as both the surface and the S-polarised light; and x , which is in the same plane as the surface but perpendicular to y . A dipole oscillating perpendicular to the surface will produce z -polarised light, but that is scattered primarily within the plane of the surface and so only negligible amounts are collected with an objective lens positioned normal to the sample. Therefore, I have access to four possible combinations of polarisations: Sy, Sx, Py and Px.

Ordering of molecules within the system is induced by the presence of the flat surface in the xy plane. The silica surface is isotropic, and hence there will be no overall difference between the alignment of molecules in the x axis and the y axis; individual molecules—or even groups of molecules—may be aligned with respect to these axes but when averaged over a sufficient area molecules at the surface will not be aligned. Molecules can, however, be aligned in the z axis: as an example in an adsorbed bilayer molecules might stand upright on the surface in preference to lying flat. Considering the elements of the Raman polarisability tensor given in equation 1.103 and discussed in the surrounding text, in lab coordinates, $\alpha_{xx} = \alpha_{yy}$, $\alpha_{xy} = \alpha_{yx}$ and $\alpha_{xz} = \alpha_{yz} = \alpha_{zx} = \alpha_{zy}$. α_{zz} remains unique but is scattered very weakly in the direction normal to the surface and so cannot be probed effectively in my experiment. The elements of the Raman tensor probed by each experimental polarisation combination are given in table 2.1.

2.1.4 Target factor analysis

Both the equilibrium and kinetic measurements were analysed using a chemometric method known as target factor analysis (TFA).²¹⁹ Target factor analysis is closely related to principal component analysis (PCA), with an additional step where TFA attempts to represent the results using physically realistic spectra rather than abstract mathematical components. The first step in both TFA and PCA is to decompose the data set, \mathbf{D} , (which consists of c spectra containing r points each, giving a size of $r \times c$) into its principal components, so that

$$\mathbf{D} = \mathbf{RC}, \quad (2.1)$$

where \mathbf{R} (size $r \times s$) and \mathbf{C} (size $s \times c$) are abstract row and column matrices (abstract meaning that the matrices do not correspond to physically meaningful spectra). s is

the number of eigenvectors needed to describe the data, and is equal to the smallest of r or c . A number of possible methods are available to generate \mathbf{R} and \mathbf{C} ; the technique we use here is singular value decomposition (SVD) since it is well known and fast implementations exist for a wide range of programming languages (including MATLAB). SVD generates three matrices

$$\mathbf{D} = \mathbf{USV}', \quad (2.2)$$

where \mathbf{U} and \mathbf{V} are orthonormal matrices (multiplication of the matrices by their transpose gives an identity matrix). \mathbf{S} is a diagonal matrix with elements that are the square root of the eigenvalues. $\mathbf{R} = \mathbf{US}$ while $\mathbf{V}' = \mathbf{C}$.

The next stage is to represent the system in terms of n significant factors. These factors represent the data, while the remaining $s - n$ factors are noise. Noise in the direction of the vectors defining the n significant factors is retained. For pure surfactants on silica, $n = 2$ and the first two components contain orthogonal linear combinations of the water and surfactant spectra. The compressed versions of \mathbf{R} and \mathbf{C} are designated $\bar{\mathbf{R}}$ and $\bar{\mathbf{C}}$ and have sizes $r \times n$ and $n \times c$ respectively. They can be multiplied to regenerate a version of the data, designated $\bar{\mathbf{D}}$, without the factors from the noise.

The second step is a coordinate rotation to extract the refined spectra corresponding to the water and the surfactant and their component weights in each of the kinetic spectra. A transformation of the form

$$\hat{\mathbf{X}} = \bar{\mathbf{R}}\mathbf{T} \quad (2.3a)$$

$$\hat{\mathbf{Y}} = \mathbf{T}^{-1}\bar{\mathbf{C}} \quad (2.3b)$$

where $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ are matrices representing a set of physically realistic spectra and component weights for the contribution of each spectrum respectively. \mathbf{T} is a transformation matrix of size $n \times n$. $\bar{\mathbf{D}}$ can be produced by multiplying $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$. The target transformation approach using in TFA looks at each component individually, taking a predicted component, \mathbf{x}_l —usually the measured Raman spectrum of a single component—and attempting to generate a column of the transformation matrix. Looking only at the l th column of $\hat{\mathbf{X}}$, and the l th column of \mathbf{T} ,

$$\hat{\mathbf{x}}_l = \bar{\mathbf{R}}\mathbf{t}_l. \quad (2.4)$$

To attempt to find \mathbf{t}_l from \mathbf{x}_l (the predicted component) I use a least squares approach:

$$\mathbf{a}_l = \mathbf{B}\mathbf{t}_l, \quad (2.5)$$

where $\mathbf{a}_l = \bar{\mathbf{R}}\mathbf{x}_l$ and $\mathbf{B} = \bar{\mathbf{R}}'\bar{\mathbf{R}}$. Hence,

$$\mathbf{t}_l = (\bar{\mathbf{R}}'\bar{\mathbf{R}})^{-1}\bar{\mathbf{R}}'\mathbf{x}_l. \quad (2.6)$$

Alternatively,

$$\mathbf{t}_l = \bar{\mathbf{R}}^+\mathbf{x}_l, \quad (2.7)$$




Figure 2.6: Example of target and refined spectra for TX-100. Two refined spectra are shown: one from the analysis of an equilibrium TX-100 isotherm and one from the analysis of a single kinetic run.

where $\bar{\mathbf{R}}^+$ is the Moore-Penrose pseudoinverse of $\bar{\mathbf{R}}$, which is implemented in most programming languages. Hence,

$$\hat{\mathbf{x}}_l = \bar{\mathbf{R}}\bar{\mathbf{R}}^+ \mathbf{x}_l. \quad (2.8)$$

Equation 2.8 describes the process of taking a predicted spectrum and producing a refined spectrum that represents part of the information contained within the first n abstract factors.

To perform this rotation a set of target spectra (each spectrum comprising one of \mathbf{x}_l) are required that approximate to the refined water and refined surfactant spectra. The target spectrum for water was acquired at the beginning of each experiment with pure water in the cell. The target spectrum for each surfactant was obtained by manual subtraction of the water target spectrum from the spectrum of the highest concentration surfactant solution. Figure 2.6 compares the target spectrum for TX-100 with the refined spectrum obtained from a set of equilibrium spectra at different concentrations and a kinetic spectrum from a single experiment. The differences between the target and refined spectra are very small, though the noise on the refined kinetic spectrum is higher owing to the much shorter acquisition times. The component weight of the surfactant spectrum from the TFA was divided by the component weight of the water spectrum to compensate for any drift in the microscope focus or laser power and to account for differences in the acquisition time (so that the longer equilibrium measurements appeared on the same scale as the brief kinetic measurements).

There are a few systematic errors that can occur when using TFA. If the water (or other background) factor is taken from a surface that is not completely clean, the resulting surfactant component weights will often be offset, making a genuinely clean surface have negative surface excess. Apart from the offset, the resulting component weights will be unchanged. If the spectra analysed have a small, uniform offset on the intensity-axis, but the target factors do not, then the uniform offset will sometimes be incorporated into the surfactant factors, giving the component weight of clean surface a positive offset, but also reducing the quality of the fit in general. This can usually be corrected by removing the offset from the spectra so that the lowest point of each

spectrum is at 0 intensity. In systems that have multiple overlapping surfactant spectra, the component weights of the two surfactants usually show negatively correlated noise (while the total surfactant kinetic weight will typically be less noisy). This is a consequence of the difficulty in unambiguously separating the different components in real experimental data, and is impossible to avoid completely.

Conversion to surface excess

Target Factor Analysis yields component weights, \hat{Y} , that determine the relative contribution of the refined spectra to each raw spectrum in the data set. The absolute values of the component weights have no meaning. A calibration procedure is required to convert the component weight of the surfactant into a surface excess (in moles m^{-2} of surface). For single component systems, the surface excess was calibrated from the contribution to the equilibrium Raman spectra from surfactants molecules in bulk solution within the evanescent wave: the number of bulk molecules contributing to the signal is simply the bulk concentration multiplied by the sampling depth. Above the cmc of a single surfactant system, the surface excess remains constant and therefore the increase in Raman signal with increasing surfactant concentration may be ascribed to bulk surfactant. So above the cmc, the component weight of the surfactant should be linear in concentration, with a slope that yields a calibration factor to convert component weight into surface excess. This calibration relies on four assumptions: first, that the adsorbed and bulk surfactant have the same spectra; second, that the spread of angles around the incident angle is small enough to be neglected; third, that the bulk concentration is uniform; fourth, that the surfactant is pure. For disordered surfactant layers the first assumption holds well, as does the second provided that the angle of incidence is not too close to the critical angle. The third assumption holds for non-ionic surfactants and for ionic surfactants provided that the Debye length is small compared to the penetration depth. The fourth assumption does not hold for Triton X-100, which is a mixture of isomers and EO chain length; the slope of the isotherm well above the cmc gives an average calibration factor for the surfactant. For mixed component systems the surface excess was determined by comparison of each normalised surfactant component weight to that of the relevant pure system. The highest bulk concentration used to calibrate the surface excess should be at least 10 mM although ideally should be 30–50 mM (if the cmc is greater than these values then the bulk concentrations used should be higher still).

The error in the conversion from component weight to surface excess can be estimated. Using the TX-100 isotherm shown later in figure 4.2 as an example, the slope above the cmc arising from the bulk signal is $(22 \pm 3) \text{ mol}^{-1} \text{ dm}^3$ (with the error calculated from the linear regression). The angle of incidence is $(73 \pm 1.5)^\circ$ (based on an error in the position of the beam in both height and distance from the microscope focus of $\pm 1 \text{ mm}$). Assuming that the wavelength of incident light and refractive indices of the hemisphere and water are known exactly, the error in the penetration

depth of the squared electric field is

$$\begin{aligned}\Delta d_p &= \left| \frac{d[(k_t\beta)^{-1}]}{d\theta_i} \right| \Delta\theta_i \\ &= \frac{(n_i/n_r)^2 \sin\theta_i \cos\theta_i}{2k_t((n_i/n_r)^2 \sin^2\theta_i - 1)^{3/2}} \Delta\theta_i\end{aligned}\tag{2.9}$$

(β is defined earlier in equation 1.83). Using 532 nm incident light and the refractive indices for water and silica, $d_p = 103 \pm 20$ nm. Thus the error in the calibration factor, A (with $\Gamma = y/A$ where y is the component weight), arising from the slope in the data-points and the the error in the angle of incidence is

$$\frac{\Delta A}{A} = \pm \sqrt{\left(\frac{20}{103}\right)^2 + \left(\frac{3}{22}\right)^2} = \pm 0.24,\tag{2.10}$$

(units have been omitted for ease of reading, since all the important values are fractional). Therefore, there is approximately a 25% systematic error in the absolute surface excesses calculated. For the purposes of this analysis the spread of angles is ignored, but will tend to increase the effective experimental penetration depth, and so lead to a small underestimate in the surface excess.

2.2 Wall-jet cell

The sample cell is illustrated schematically and as a photograph in figure 2.7. The custom-made glass cell consists of an inner chamber with a volume of ~ 6 mL surrounded by an outer jacket through which temperature-controlled water is passed. A tube allows a thermocouple probe to be inserted into the outer jacket. The top of the sample chamber is capped with a 10-mm diameter silica hemisphere sealed to the chamber with a Viton O-ring. An inlet tube (1-mm inner radius) is positioned 1.8 mm below the hemisphere surface. Unless otherwise specified the surfactant solutions were injected at a rate of 0.5 mL min^{-1} corresponding to a mean flow rate in the inlet pipe of 0.27 mm s^{-1} . The inlet pipe is sufficiently long that parabolic flow is established in the nozzle.

The cell itself is connected to a syringe pump through PTFE tubing (1/16" inner diameter, 1/8" outer diameter, bought from Cole-Palmer). The tubing is linked using a variety of 2-way and 3-way connectors made by Omnifit (from their "1.5 mm" range). The Omnifit connectors also fits over the 4 mm outer diameter tubing used as inlets and outlets for glassware—including the cell. Figure 2.8 shows diagrams of how the cell is connected. A luer-lock connector links the tubing to the syringe. The fluid flows into a three-way connector and can either flow towards the cell, or into an outlet (which is used to remove air bubbles from the tubing after connecting the syringe). In one variant of the set-up, the fluid passes through the continuous stirred mixer (see section 2.3) before entering the cell; otherwise it goes directly to the cell. A valve on the cell inlet can be closed to prevent the cell from draining. After the cell the fluid passes through another valve—which is usually left open—before emerging from an open ended tube to be caught in a beaker.





Figure 2.7: (a) Diagram and (b) photograph of the flow cell. (a) is based on the original design drawing by Eric Tyrode.

The cell, the tubing connecting the cell, and all the other glassware used in the experiment was cleaned with a commercial alkaline cleaning agent (Borer 15PF concentrate), then rinsed with copious high-purity water. The fused silica hemispheres were cleaned by soaking in chromosulphuric acid for at least 4 hours, then rinsed with high purity water. Between each experiment the cell was flushed with at least 100 mL of high purity water to wash any residual surfactant off the surface. Typically, equilibrium measurements and kinetic measurements were taken back-to-back: the cell was rinsed with water, the surfactant solution was flushed into the cell while the adsorption kinetics were recorded, then once adsorption was complete the longer equilibrium spectra were taken. If desorption measurements were also required then the cell was emptied and refilled with the surfactant solution (to ensure that the concentration was correct) before pure water was flushed in while the desorption was measured.


2.3 Continuous stirred mixer

2.3.1 Introduction

Conventionally, adsorption isotherms are acquired by a stepwise increase in concentration, waiting for an equilibration time and then acquiring data at a fixed concentration before stepping to the next higher concentration. For some of the work in this thesis I used a different procedure in which the concentration in the cell was


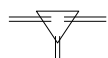
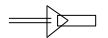


(a) Without mixer



(b) With mixer

Figure 2.8: Schematic of the two arrangements used to deliver sample to the cell.

-  On/off valve
-  Three-way junction
-  Luer-lock connection to syringe




Figure 2.9: Diagram of the continuously stirred mixer. The measurements are taken from the design drawing and will vary slightly from the final product. OD is an abbreviation for “outer diameter”.

slowly but continuously varied and Raman spectra were acquired continuously. This ‘quasi-equilibrium’ approach provides a much larger number of data points (hundreds as opposed to tens) in a shorter total time (1–2 h as opposed to 1–2 days). Provided that the surface equilibrates quickly on the time scale of the change in concentration, then an equilibrium isotherm will be obtained. To test whether equilibrium has been achieved, isotherms are acquired with increasing and decreasing concentration: hysteresis in the isotherm is a hallmark of slow kinetics.

To obtain a smoothly varying concentration profile at the surface of the hemisphere a continuous stirred mixer was installed in the inlet tube to the wall-jet cell. The mixer (figure 2.9) is a flat cylinder (volume = 8.3 mL) with an inlet and an outlet (leading to the cell) on opposite sides. A magnetic stirrer bar is sealed inside it and rotates rapidly (>200 rpm). The mixer is initially filled with one solution, and a second solution is then pumped in at a constant rate. I will refer to experiments where a surfactant is added to water as “in” experiments and the inverse—where a surfactant is diluted with water—as “out” experiments. The mixer can also be used to effect a change in the composition of mixtures of surfactants.

The continuous stirred mixer (or continuous stirred tank reactor) is a well established concept within chemical engineering.^{220–223} However, the principle has only seen limited use within the field of surface science. A similar mixing scheme has previously been used with drop-shape tensiometry measurements.²²⁴ Although this method has been used as a means of recording an isotherm,^{224,225} it is more typically used to look at the kinetics of slowly desorbing systems such as proteins.²²⁶ Because of the emphasis on kinetics it has not been used extensively since more rapid methods of solution exchange—such as double capillaries—are preferred.

The inlet concentration of the mixer is constant, at $[A]_{\text{in}}$. The outlet concentration, $[A]$, will be the concentration within the mixer as a whole, and vary according to

$$\frac{d[A]}{dt} = \frac{R}{V} ([A]_{\text{in}} - [A]), \quad (2.11)$$

where R is the pumping rate and V the volume of the mixer. Hence, the concentration within the mixer—and so also at the outlet—varies according to a simple first-order rate equation

$$[A] = [A]_{\text{in}} - ([A]_{\text{in}} - [A]_0) e^{-\frac{R}{V}t} \quad (2.12)$$

where $[A]_0$ is the initial concentrations and t the time since pumping started. Two simpler forms of equation 2.12 can be used when taking concentration isotherms: for an “in” experiment

$$[A] = [A]_{\text{in}} \left(1 - e^{-\frac{R}{V}t}\right), \quad (2.13)$$

while for an “out” experiment

$$[A] = [A]_0 e^{-\frac{R}{V}t}, \quad (2.14)$$

The time taken for the solution to travel from the outlet of the mixer to the sample surface was measured by observing the flow of liquid into an empty cell, and was determined to be 100 s at a pumping rate of 0.5 mL min^{-1} (and was scaled accordingly at other pumping rates). Assuming parabolic flow through the circular tubes the velocity at the centre of the tube will be double the effective flow rate (which is measured when the cell is filled from empty); neglecting diffusion the surfactant at the centre of the tube should reach the surface in about half that time. In practice the time offset is fitted to the data, and is usually around 75 s for the 0.5 mL min^{-1} flow rate, suggesting some amount of diffusion in the inlet tube. The uncertainty in the exact offset adds to the uncertainty in concentration, and is most significant near the start of an experimental run.

The acquisition time of each spectrum was 0.17–0.5 min, which is short compared to the characteristic time over which the concentration varies: $V/R = 17 \text{ min}$. The characteristic time for surfactant to cross the diffusion layer adjacent to the interface in our wall-jet flow cell is $\sim 4 \text{ s}$ (chapter 4), therefore we do not expect the final transport step of diffusion to the surface to be significant when using the mixer.

2.3.2 Validation

Although the derivation of equation 2.12 is simple and only relies on the assumption that the concentration within the cell is uniform (which can be controlled through the stirring rate) it is still useful to validate the concentration dependence. Acetonitrile is a useful indicator since we have found Raman signal is proportional to concentration. Therefore, we can measure the concentration within the cell by following acetonitrile signal with time.

Figure 2.10(b) shows the change in acetonitrile signal with time using a flow rate of 0.42 mL min^{-1} , together with the expected signal. Although the overlap is not quite perfect, it is generally very good, and almost certainly less than the error in concentration from making a solution through successive dilutions. The expected signal is calculated from the volume of the mixer and the flow rate, and is *not* a fit to the data. Figure 2.10(a) confirms that the acetonitrile signal is linearly proportional to its concentration (this might not be the case either if acetonitrile adsorbs or the refractive index of the solution varies with acetonitrile concentration thus changing the penet-




Figure 2.10: (a) Dependence of acetonitrile component weight on concentration, showing a linear dependence (the highest concentration represents 9% by mass). (b) Experimental (+) and theoretical (lines) variation of concentration of acetonitrile using the continuously stirred mixer with a flow rate of 0.42 mL min^{-1} , and an inlet concentration of acetonitrile of 1.5 M (6% by weight). The right hand axis converts component weight to a fraction of the inlet concentration.

ration depth). Changes in solution viscosity or acetonitrile diffusion coefficient with acetonitrile concentration should only affect the final transport across the diffusion layer, which is known to be fast compared to the rate at which concentration varies.

The mixer is used extensively in chapter 6. Additionally, the quality of data obtained can be seen by comparing figure 5.2 (manual solution preparation) with figure C.2 in the appendices (with mixer). *Note that figure 5.2 is the better dataset for reasons explained in the appendix.*

Chapter 3

Numerical Modelling

3.1 Introduction

In order to draw meaningful conclusions about the kinetics of surfactant adsorption, it is useful to be able to compare experimental data and a theoretical model. Much of the previous work on surfactant adsorption kinetics (using wall-jet flow cells as well as other geometries) has chosen to analyse results in terms of the “sticking ratio”, the ratio of measured flux compared to maximum possible flux. I believe this approach is unsatisfactory: firstly the surface reaches saturation coverage before the maximum flux can be reached and secondly—in a system controlled by the adsorption kinetics to the surface—the sticking ratio should vary strongly with pumping rate and hence is not really describing the properties of the surfactant. These points will be discussed in the final section of this chapter.

The first section of the chapter describes how the rather complicated problem of hydrodynamics in a wall-jet flow cell simplifies to a simpler one-dimensional problem near the axis of the wall-jet, and presents limiting analytical solutions for that case. The second section describes the formulation of a time-dependent model, based on the one-dimensional problem and adds boundary conditions that replicate realistic surface behaviour. Finally, the model is compared to previous models of surfactant adsorption in wall-jet flow cells. The application of the model and comparison to experimental data is primarily in chapter 4.

3.2 Convection-diffusion

The flux, \mathbf{J} , of a species with concentration c is given by the Nernst-Planck equation

$$\mathbf{J} = \mathbf{v}c - D\nabla c - \frac{Dze}{k_B T}c\nabla\phi, \quad (3.1)$$

where \mathbf{v} is the velocity of the surrounding solution, D is the diffusion coefficient of the species, z is the charge on the species, e is the elementary charge and $\nabla\phi$ is the electric field. The first term in the equation is the convection term, accounting for the movement of material due to the movement of the surrounding solution; the second term is the diffusion term, accounting for the spread of particles from random




Figure 3.1: Schematic diagram of streamlines in a wall jet. The region from which spectra are taken is circled. The wall jet is axisymmetric around the z -axis.

movements caused by thermal energy; the final term is the migration term, which describes movement induced by electric fields. The migration term can be neglected for neutral species (since $z = 0$) or where there is excess background electrolyte and no external field. For a solution with no background electrolyte, electric fields arise as a result of the different ions diffusing at different rates (excess electrolyte ensures that the potential is approximately equal everywhere). If the mutual-diffusion coefficient is used rather than the self-diffusion coefficient then the migration term can be neglected for systems with no background electrolyte and no external electric field. Mutual and self-diffusion coefficients are discussed in chapter 1.

Conservation of mass means that

$$\frac{\partial c}{\partial t} = -\nabla \cdot \mathbf{J}. \quad (3.2)$$

For an incompressible fluid $\nabla \cdot \mathbf{v} = 0$, and so,

$$\frac{\partial c}{\partial t} = D\nabla^2 c - \mathbf{v} \cdot \nabla c, \quad (3.3)$$

assuming D is constant. In cylindrical polar coordinates this equation becomes

$$\frac{\partial c}{\partial t} + v_z \frac{\partial c}{\partial z} + v_r \frac{\partial c}{\partial r} - D \frac{\partial^2 c}{\partial z^2} - \frac{D}{r} \frac{\partial}{\partial r} \left(r \frac{\partial c}{\partial r} \right) = 0. \quad (3.4)$$

The angular coordinate (usually designated as ϕ) is neglected in the above equation since in the systems used here there is no angular flow ($v_\phi = 0$) and the input concentration is rotationally isotropic ($\frac{\partial c}{\partial \phi} = 0$) so no transport occurs in that coordinate.

3.2.1 Flow profile

In section 1.4.5 of the introduction I briefly introduced the wall-jet flow cell and the wide range of applications it has been used for. The hydrodynamics within such cells have been modelled extensively. The wall-jet flow cell consists of a pipe perpendicular to a flat substrate. Solution enters the cell through the pipe and then diffuses outwards

into the remainder of the cell. A schematic illustration of flow within such a cell is shown in figure 3.1. The stagnation point is the point on the surface where the fluid velocity would be zero, *even without* the no-slip boundary condition imposed by the surface.

Ideal wall-jet

For an idealised version of the stagnation point geometry, it is possible to obtain an exact solution for the flow profile.²²⁷ The exact solution does not describe the influence of real experimental details—for example the presence of walls of the inlet tube—but does serve as an illustrative approximation. The geometry described is a flow of liquid (travelling from positive to negative z) impinging on a plane at $z = 0$. At the plane there is a no-slip boundary condition, and so at $z = 0$ the fluid cannot move parallel to the surface. The motion of fluids is described by the Navier-Stokes equations, which expressed in cylindrical polar coordinates are,

$$v_r \frac{\partial v_r}{\partial r} + v_z \frac{\partial v_r}{\partial z} = -\frac{1}{\rho} \frac{dp}{dr} + \nu \left(\frac{\partial^2 v_r}{\partial r^2} + \frac{1}{r} \frac{\partial v_r}{\partial r} - \frac{v_r}{r^2} + \frac{\partial^2 v_r}{\partial z^2} \right), \quad (3.5a)$$

$$v_r \frac{\partial v_z}{\partial r} + v_z \frac{\partial v_z}{\partial z} = -\frac{1}{\rho} \frac{dp}{dz} + \nu \left(\frac{\partial^2 v_z}{\partial r^2} + \frac{1}{r} \frac{\partial v_z}{\partial r} - \frac{v_z}{r^2} + \frac{\partial^2 v_z}{\partial z^2} \right), \quad (3.5b)$$

where ν is the viscosity. The fluid is assumed to be incompressible, meaning it must also satisfy the continuity equation,

$$\frac{\partial v_r}{\partial r} + \frac{v_r}{r} = \frac{\partial v_z}{\partial z} = 0, \quad (3.6)$$

which states that the rate of flow into a unit volume is equal to the flow out of that volume.

The model for an ideal wall-jet is divided into two regions. Close to the surface there is a *boundary layer*, where there is a stress induced by parallel streamlines moving at different rates which is proportional to the viscosity of the fluid. Outside the boundary layer all the parallel streamlines are moving at the same rate, and so the viscosity of the fluid has no influence on its behaviour. This situation is described as a “frictionless fluid”.²²⁷ Figure 3.2 illustrates the division.

The flow outside the boundary condition is considered first, since it is easier to derive and it represents the solution the boundary layer flow must reach at large values of z . The rate of flow towards the surface must reduce with z so that fluid does not enter the surface ($v_z(z=0) = 0$). The radial flow is axisymmetric about the centre of the jet. Since this flow is far from the surface the no-slip boundary condition at the surface is ignored and the rate of flow in the radial direction (v_r) is constant with respect to z . For a frictionless fluid, the pressure is described by Bernoulli’s equation:

$$p + \frac{1}{2} \rho |\mathbf{v}|^2 = p + \frac{1}{2} \rho (v_r^2 + v_z^2) = \text{constant}, \quad (3.7)$$

where ρ is the density of the liquid. The Navier-Stokes equations (3.5) and the




Figure 3.2: Schematic illustration of flow within an ideal wall-jet geometry. The horizontal dotted line indicates the division between the boundary layer and the bulk flow (where viscosity does not affect the flow). Based on ref. 227.

continuity equation (3.6) are satisfied by the flow rates

$$v_r^* = ar \quad (3.8a)$$

$$v_z^* = -2az, \quad (3.8b)$$

where a is a constant and the superscript ‘*’ denotes that the flow is outside the boundary layer.

Within the boundary layer there is a viscosity induced stress caused by the value of v_r changing from 0 to v_r^* across the boundary layer. It is reasonable to assume that $v_r \propto r$ within the boundary layer, so that it matches the behaviour of v_r^* outside the boundary layer ($z = \infty$). At this stage the assumption should be treated as a guess, which will be correct if the resulting equations for velocity are consistent with the Navier-Stokes equations and continuity equation. The behaviour of v_z is not yet known, but it is assumed to be independent of r and thus varies according to an arbitrary function, $f(z)$. By considering the continuity equation (3.6), flow within the hydrodynamic boundary layer is,²²⁷

$$v_r = rf'(z), \quad (3.9a)$$

$$v_z = -2f(z), \quad (3.9b)$$

where the prime represents differentiation with respect to z .

Substituting v_z and v_r into equations 3.5 gives,

$$rf'^2 - 2rf f'' = -\frac{1}{\rho} \frac{\partial p}{\partial r} + \nu r f''', \quad (3.10a)$$

$$4f f' = -\frac{1}{\rho} \frac{\partial p}{\partial z} - 2\nu f''. \quad (3.10b)$$

Therefore, from equation 3.10a,²²⁸

$$\begin{aligned} -\frac{1}{\rho} \int \frac{\partial p}{\partial r} dr &= \int r(f'^2 - 2ff'' - \nu f''') dr = \int rH(z) dr, \\ -\frac{1}{\rho} p &= \frac{r^2}{2} H(z) + C(z), \end{aligned} \quad (3.11)$$

where $C(z)$ is an arbitrary function of z (the form of $C(z)$ is only important to evaluate the pressure, p); the substitution

$$H(z) = f'^2 - 2ff'' - \nu f''' \quad (3.12)$$

is made for simplicity. Substitution of p into equation 3.10b gives,

$$4ff' = \frac{r^2}{2} H'(z) + C'(z) - 2\nu f'', \quad (3.13)$$

which can only be true for all r and z if $H' = 0$ and so H is a constant. Adding a requirement that the transition to the frictionless solution at $z = \infty$ is smooth means that $f''(\infty)$ and $f'''(\infty)$ are zero, and thus $H = a^2$. Therefore, equation 3.12 is²²⁸

$$a^2 = f'^2 - 2ff'' - \nu f''' ; \quad (3.14)$$

a numerical solution to this equation with the boundary conditions $f(0) = f'(0) = 0$ and $f'(\infty) = a$ provides the flow near the surface.

$f(z)$ can be non-dimensionalised as $\phi(\zeta) = \frac{f(z)}{\sqrt{a\nu}}$ where $\zeta = \sqrt{\frac{a}{\nu}}z$. The non-dimensional form of equation 3.14,

$$\phi''' + 2\phi\phi'' - \phi'^2 + 1 = 0. \quad (3.15)$$

was solved by Frössling,²²⁹ and the results are tabulated in ref. 227. They are presented in figure 3.3. Since the equation for ϕ can be solved to produce a consistent solution, equations 3.9 are valid.

Very close to the surface the velocity of the solution can be expressed using the leading terms of a power series as,²³⁰

$$v_r = 1.31a^{3/2}\nu^{-1/2}rz \quad (3.16a)$$

$$v_z = -1.31a^{3/2}\nu^{-1/2}z^2. \quad (3.16b)$$

The ideal case described above is a good approximation when the region of interest on the surface (usually a detector) is much smaller than the radius of the inlet tube. This applies to my cell, where the inlet radius is 1 mm and the radius of the laser spot probing the surface is tens of μm . As noted in the introductory chapter, this situation is referred to as a “wall-tube” in electrochemistry.

Realistic wall-jet

Flow within a real cell will deviate from the ideal solution due the use of an inlet tube of finite width and obstruction to the flow leaving the stagnation region by the cell




Figure 3.3: Dependence of the parameters ϕ (solid blue), $\sqrt{\phi}$ (dashed green) and ϕ' (dotted red) on ζ , a scaled coordinate along the z -axis. The data was taken from ref. 227, which in turn collated it from ref. 229.

walls. Flow within the inlet tube is a parabolic profile centred over the stagnation point rather than a plug flow. In order to model such cells it is necessary to solve the Navier-Stokes equation for the experimental cell geometry, typically using approximate methods such as the finite-difference, finite-element or finite-volume methods. Because of the difficulty in modelling the hydrodynamics of the whole of the cell, I am choosing to use the results of existing work focused on a similar cell design, primarily that of Dąbroś and van de Ven.¹⁰¹ With the knowledge of the full hydrodynamic modelling, it is possible to model the mass transport within a small region of the cell, close to the surface.

Dąbroś and van de Ven modelled flow in an extended region from the end of the inlet to where solution left the cell at the side. However, within the hydrodynamic boundary layer they expressed the flow as a second order Taylor series, to give the equations¹⁰¹

$$\begin{aligned} v_r &= \alpha z r, \\ v_z &= -\alpha z^2, \\ v_\phi &= \bar{\beta} r z = 0. \end{aligned} \quad (3.17)$$

This approximation is valid for large Schmidt numbers, where the diffusion boundary layer is much thicker than the hydrodynamic boundary layer. For a typical surfactant in water

$$Sc = \frac{\nu}{D} \approx \frac{1 \times 10^{-6} \text{ m}^2 \text{ s}^{-1}}{5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}} \approx 2000. \quad (3.18)$$

The approximation given in equations 3.17 also agrees with linear relationship between $\sqrt{f(z)}$, $f'(z)$ and z seen in the ideal case.

The vorticity of the flow at the surface, ω , is given by

$$\omega = \nabla \times \mathbf{v} = \frac{\partial v_r}{\partial z} - \frac{\partial v_z}{\partial r}. \quad (3.19)$$




Figure 3.4: Variation of the non-dimensional form for the hydrodynamic constant, $\bar{\alpha}$, with Reynolds number, for two different geometries: (1) $h/R = 1.7$ (close to the geometry of the cell I use in my experiments) and (2) $h/R = 1.0$. Reproduced with kind permission from Springer Science+Business Media: *Colloid and Polymer Science* “A direct method for studying particle deposition onto solid surfaces”, **261**, (1983), 694–707, T. Dąbrosz and T. G. M. van de Ven, figure 5 (ref. 101).

Therefore, α can be defined in terms of the vorticity at the surface,

$$\alpha = \frac{\omega}{r}, \quad (3.20)$$

allowing the approximate equations at the surface (3.17) to be readily related to the full hydrodynamic model of the cell. Values of α for different Reynolds numbers are calculated empirically in ref. 101 and shown in figure 3.4 in a non-dimensional form, $\bar{\alpha}$, where $\alpha = \frac{U}{R^2} \bar{\alpha}$ (U is the average inlet velocity and R the tube diameter).

Equation 3.4 can be combined with equation 3.17 to give the convection-diffusion equation describing flow close to the stagnation point in the wall-jet flow cell:

$$\frac{\partial c}{\partial t} - \alpha z^2 \frac{\partial c}{\partial z} + \alpha z r \frac{\partial c}{\partial r} - D \frac{\partial^2 c}{\partial z^2} - \frac{D}{r} \frac{\partial}{\partial r} \left(r \frac{\partial c}{\partial r} \right) = 0. \quad (3.21)$$

Validity

Rees *et al.* identified five issues that may limit the applicability of the simple model described in this section.¹⁰² These issues are described below (condensed into four points since I feel two of the points are very closely related):

1. The inlet flow must be laminar and not turbulent. This is controlled by selecting an appropriate Reynolds number (flow rate and tube dimensions).
2. The radial diffusion must be small. This is discussed in section 3.2.2.
3. Formation of vortices at the side of the jet may influence flow within the jet; thin nozzle walls promote vortices. Most modelling has been done on flow systems comprising a hole in a plate rather than a nozzle, in order to minimise the formation of vortices either side of the nozzle.²³¹ However, Melville *et al.*¹⁰³

showed that the flux at the surface was only $\sim 2\%$ greater for a submerged nozzle compared to a hole in a plate, and so we feel confident in using an submerged pipe and neglecting the difference between the two set-ups.

4. The behaviour is sensitive to the position of the probe (often an electrode, in our case the laser spot) relative to the centre of the nozzle; a good illustration is given in reference 232. This is more of an issue for the microjets used in electrochemistry than it is here.

3.2.2 1-dimensional solution

Equation 3.21 would be much easier to solve if it could be reduced to a 1-dimensional—rather than a 2-dimensional—problem. In order to show that it can be treated as 1-D it is useful to non-dimensionalise the equation, thereby making the characteristic magnitudes of the terms clearer. z is non-dimensionalised by $(D/\alpha)^{1/3}$, the thickness of the diffusion layer ($\sim 50 \mu\text{m}$); r by the nozzle radius, R ($=1 \text{ mm}$); and t by $(D\alpha^2)^{-1/3}$, the time taken to diffuse across the diffusion layer ($\sim 4 \text{ s}$). c is left dimensional at this stage, since the non-dimensionalisation of concentration only becomes relevant once adsorption to the surface is introduced (later). Therefore equation 3.21 becomes

$$\frac{\partial c}{\partial t'} - \frac{(D/\alpha)^{2/3}}{R^2 r'} \frac{\partial}{\partial r'} \left(r' \frac{\partial c}{\partial r'} \right) - \frac{\partial^2 c}{\partial z'^2} + z' r' \frac{\partial c}{\partial r'} - z'^2 \frac{\partial c}{\partial z'} = 0, \quad (3.22)$$

where r' , t' and z' are the non-dimensionalised coordinates. The second term can immediately be neglected since $(D/\alpha)^{2/3}/R^2 < 10^{-2}$. The fourth term is more complicated: as surfactant is transported into the stagnation region c is either flat (in which case $\frac{\partial c}{\partial r'} = 0$ and the term disappears, or—during the early stages of transport— c goes through a maximum at $r' = 0$. Using a Taylor series approximation near $r' = 0$, $c(r') \approx c_0(1 - Ar'^2)$ and hence $\frac{\partial c}{\partial r'} \approx -2c_0Ar'$. Thus the fourth term in equation 3.22 scales as r'^2 and can be neglected when r' is small provided A is small. When surfactant is being rinsed out of the stagnation region c is at a minimum at $r' = 0$, and the same argument applies.

Therefore equation 3.21 can be simplified to the one dimensional equation

$$\frac{\partial c}{\partial t} - \alpha z^2 \frac{\partial c}{\partial z} - D \frac{\partial^2 c}{\partial z^2} = 0. \quad (3.23)$$

Note that I mostly use the dimensional form of the equations since the computer programs written to solve equation 3.23 are written in this form.

3.2.3 Analytical solution

It is only possible to generate an analytical solution to equation 3.23 for a few simple sets of conditions—the majority of work will have to use numerical approximations. However, it is still worth looking at a few of the steady state solutions.

When the system has reached a steady state $\frac{\partial c}{\partial t} = 0$, and hence

$$-\alpha z^2 \frac{\partial c}{\partial z} - D \frac{\partial^2 c}{\partial z^2} = 0, \quad (3.24)$$

which can be integrated to give

$$\frac{\partial c}{\partial z} = C_1 e^{-\frac{\alpha z^3}{3D}}, \quad (3.25)$$

where C_1 is a constant of integration.

There are two physically sensible steady state system that can be applied to the 1-D convection-diffusion equation: the perfect sink and the no-flux boundary condition. They are both discussed more thoroughly later with respect to the numerical approximation applied. In brief, a perfect sink boundary condition fixes the concentration at the subsurface (typically to 0), while the no-flux boundary condition sets $\frac{\partial c}{\partial z} \Big|_{z=0} = 0$.

For the no-flux boundary condition, the constant of integration C_1 must be 0, and hence $\frac{\partial c}{\partial z}$ has no z dependence and is 0 everywhere. Thus the concentration must be the same everywhere, which is the logical outcome for the boundary condition.

Bijsterbosch *et al.*¹⁰⁵ performed the derivation for the perfect sink boundary condition. Defining x as

$$x = \left(\frac{\alpha}{3D} \right)^{1/3} z, \quad (3.26)$$

then

$$c = C_1 \left(\frac{3D}{\alpha} \right)^{1/3} \int_0^x e^{-x^3} .dx. \quad (3.27)$$

Using the identity of the gamma function, Γ (not to be confused with the alternative use of Γ for the surface excess), defined as²³³

$$\Gamma(y) = \int_0^\infty s^{y-1} e^{-s} ds, \quad (3.28)$$

then, if $s = x^3$,

$$\Gamma(1/3) = 3 \int_0^\infty e^{-x^3} .dx. \quad (3.29)$$

Hence,

$$c_{z=\infty} = C_1 \left(\frac{3D}{\alpha} \right)^{1/3} \frac{\Gamma(1/3)}{3} + C_2, \quad (3.30)$$

where C_2 is a second constant of integration, equal to the subsurface concentration, c_s . C_1 can then be determined based on the requirement that $c_{z=\infty} = c_{\text{in}}$ so that

$$C_1 = \left(\frac{\alpha}{3D} \right)^{1/3} \frac{3}{\Gamma(1/3)} (c_{\text{in}} - c_s). \quad (3.31)$$

The flux to the surface is then

$$J = D \frac{\partial c}{\partial z} \Big|_{z=0} = \frac{3^{2/3}}{\Gamma(1/3)} D^{2/3} \alpha^{1/3} (c_{\text{in}} - c_s) \quad (3.32)$$

This flux should be interpreted as a limiting flux: the largest value possible if every molecule that reached the surface stuck to the surface. Where c_s is fixed to zero the equation is equal to the form evaluated by Dijt *et al.*,¹⁰⁰

$$J = 0.77646 D^{2/3} \alpha^{1/3} c_{\text{in}}. \quad (3.33)$$

Equation 3.33 has been used extensively to model surfactant adsorption kinetics.^{43,45,100,106}

3.3 Finite difference model

This section describes the working of the finite difference model for surfactant adsorption kinetics in a wall-jet flow cell. The code is given in Appendix E; however this section is written so that the program could substantially be recreated without reference to the code. Reference 234 provides a good introduction to the formulation of finite difference models similar to the one presented here; at time of writing it is available to read at <http://compton.chem.ox.ac.uk/john/Thesis/index.html>.

The first part of this section will look at solving the one-dimensional convection-diffusion equation (equation 3.23). The latter parts will look at modelling the adsorption processes to the surface too.

3.3.1 Definitions

Throughout this section I use the symbol c_j^n to refer to the concentration at the grid-point j on timestep n . Values of j range from 0, the subsurface, to jj , the “inlet”. The step sizes in space and time are designated Δz and Δt respectively. The choice of appropriate values of Δz , Δt and the length of region to model are discussed later in this chapter. Where surface excess is used later on in the model it is given the symbol Γ^n (there is only one surface point so no index for the z axis is needed).

The model is based on backwards finite differences backwards in timestep and centred in space. Thus

$$\frac{\partial c}{\partial t} = \frac{c_j^n - c_j^{n-1}}{\Delta t}, \quad (3.34a)$$

$$\frac{\partial c}{\partial z} = \frac{c_{j+1}^n - c_{j-1}^n}{2\Delta z}, \quad (3.34b)$$

$$\frac{\partial^2 c}{\partial z^2} = \frac{c_{j+1}^n - 2c_j^n + c_{j-1}^n}{\Delta z^2}. \quad (3.34c)$$

The model uses a backwards-implicit calculation: a set of simultaneous equations are derived so that three points in the current timestep (c_{j-1}^n , c_j^n and c_{j+1}^n) all depend on a single point in the previous timestep (c_j^{n-1}). The alternative forwards-explicit calculation—where a future point (c_j^{n+1}) is calculated on the basis of three points in the current timestep (c_{j-1}^n , c_j^n and c_{j+1}^n)—is prone to develop large numerical instabilities for all but the smallest step sizes, and so fully explicit calculations are rarely used. As a technical note, the code implementing the model uses a degree of implicitness, θ , which can vary between 0 and 1, since values of 0.5 sometimes give more accurate results. For reasons discussed briefly later, I have exclusively used $\theta = 1$ in my calculations and so for simplicity all the equations here assume that value.

3.3.2 Convection-diffusion implementation

Equation 3.23 can be written as the finite difference form,

$$-\left(\frac{D}{\Delta z^2} - \frac{j^2 \alpha \Delta z}{2}\right) c_{j-1}^n + \left(\frac{2D}{\Delta z^2} + \frac{1}{\Delta t}\right) c_j^n - \left(\frac{D}{\Delta z^2} + \frac{j^2 \alpha \Delta z}{2}\right) c_{j+1}^n = \frac{c_j^{n-1}}{\Delta t} \quad \text{for } j = 1 \dots jj - 1. \quad (3.35)$$

A similar equation exists for $j = 0$, but it depends on the boundary conditions and will be given later. The complete system of equations can then be expressed in matrix form as

$$\mathbf{d}^{n-1} = \mathbf{T}\mathbf{u}^n \quad (3.36)$$

where \mathbf{d}^{n-1} is a vector with elements based on concentration at the previous timestep: $d_j^{n-1} = c_j^{n-1}/\Delta t'$, and \mathbf{u}^n is a vector of concentrations at the current time step. The Thomas (or double-sweep) algorithm²³⁵ is used to invert the matrix \mathbf{T} and thus calculate the concentrations at the current time step.*

Equation 3.36 has to be solved subject to boundary conditions at $z = 0$ describing adsorption to the surface. Section 3.3.3 describes the limiting cases of a perfect sink and no-flux boundary condition; sections 3.3.4 and 3.3.5 describe kinetic models consistent with the Langmuir and Frumkin isotherms.

Whichever boundary conditions are used at $z = 0$, the boundary condition far from the surface, c_{jj}^n , is simply the inlet concentration. This is simply substituted into the instance of equation 3.35 centred around point $jj - 1$. The boundary conditions at $t = 0$ depend on the experiment being modelled: for adsorption kinetics the concentration throughout the whole cell is set to 0; for desorption kinetics the concentration throughout the whole cell is set to the bulk concentration prior to rinsing.

3.3.3 Simple boundary conditions

The first limiting case is the perfect sink model, where all molecules that reach the surface adsorb irreversibly. The surfactant concentration in the cell adjacent to the surface is held at a specific concentration c_{edge} . c_{edge} is usually set to 0, but does not have to be: for example a “step-function” type isotherm could be modelled kinetically by setting c_{edge} to be the concentration at which the adsorption first takes place. The steady-state solution to the convective diffusion equation with the sink boundary condition has been discussed earlier in this chapter. Figure 3.5 shows the steady state concentration profile for a perfect sink boundary condition. There are two regions: the convection dominated region at higher z , where the concentration profile is flat, and the diffusion layer close to the surface where the concentration varies with z .

In our experiments, the surface excess saturates well before the flux to the surface reaches its limiting value. Consequently, equation 3.33 is of limited quantitative value. See the final section of this chapter for more details.

* *Technical detail:* although the Thomas algorithm is a quick means of inverting tridiagonal matrices such as \mathbf{T} , subsequent testing shows that it does not significantly outperform modern computer linear algebra systems and thus could be replaced with MATLAB code similar to “ $\mathbf{u} = \mathbf{d}/\mathbf{T}$;” without loss of speed.




Figure 3.5: Steady state concentration profile for a perfect sink boundary condition. $\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$, $D = 5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$ and $c_{\text{in}} = 1 \text{ mM}$. The dashed vertical line indicates the thickness of the diffusion layer, $(D/\alpha)^{1/3}$ (where the non-dimensional $z' = 1$).

Alternatively, if the surfactant does not adsorb at all, then the appropriate boundary condition is one of no flux:

$$\left. \frac{\partial c}{\partial z} \right|_{z=0} = 0 \quad (3.37)$$

This condition is not useful for monomeric surfactants (except where one wishes to ignore the effect of the surface), but it is useful in solutions above the critical micelle concentration if one makes the conventional (though not necessarily correct⁸⁹) assumption that micelles do not adsorb.

The two possibilities can be combined: if the sink model is used with $c_{\text{edge}} \geq 0$ then the no-flux boundary condition can be used until $c_0 > c_{\text{edge}}$. Similarly a “source” boundary condition can be devised to look at desorption processes by using the no-flux boundary condition until $c_0 < c_{\text{edge}}$ then switching to the sink calculation.

3.3.4 Langmuir Isotherm

The simplest monolayer adsorption isotherm is due to Langmuir:²³⁶


$$\frac{\Gamma}{\Gamma_{\infty}} = \frac{K_L c}{1 + K_L c} \quad (3.38)$$

where Γ_{∞} is the number of adsorption sites per unit area and K_L the Langmuir constant. The Langmuir isotherm is derived from the assumptions that

1. each site on the surface can have a single solute or solvent molecule adsorbed (alternatively this can be viewed as the solute and solvent molecules having equal surface area);
2. the surface is homogeneous (the adsorption sites are identical);
3. there is ideal behaviour on the surface (all interactions between adjacent molecules are the same regardless of the species involved);
4. the bulk phase is an ideal dilute solution;
5. the adsorbed film is monomolecular.


The last point is particularly problematic, since many surfactants adsorb as bilayers or surface aggregates.

Langmuir derived the isotherm on the basis of the consistent kinetic model.²³⁶ Adsorption of a surfactant molecule (A) is an exchange between a solvent molecule (S) and the surfactant molecule at an adsorption site (X):



The rate of reaction thus depends on the concentration of A (designated by c), and the number of solvent filled sites. The number of solvent filled sites is equal to the total number of sites, Γ_∞ , minus the number of solute-filled sites.

Desorption is the reverse process and so



The rate of reaction depends the number of solute-filled sites (given by Γ) and the concentration of solvent ($[S]$). A rate equation can then be written

$$\frac{d\Gamma}{dt} = k_a c (\Gamma_\infty - \Gamma) - k_{des} \Gamma [S]. \quad (3.39)$$

Since the surfactant solution is dilute, $[S]$ is constant and thus can be merged into the rate constant giving $k_d = k_{des}[S]$. At equilibrium $\frac{d\Gamma}{dt} = 0$. Thus,

$$0 = k_a c (\Gamma_\infty - \Gamma) - k_d \Gamma, \quad (3.40)$$

which can be rearranged to give equation 3.38 using $K_L = k_a/k_d$.

The consistent kinetic model based on the Langmuir isotherm is implemented in my code. It is not detailed further here since the majority of my modelling uses the Frumkin isotherm (discussed below). Given that the Langmuir isotherm is a special case for the Frumkin isotherm then the only advantage to having the Langmuir isotherm implemented separately is a small increase in speed of calculation and an internal consistency check that the results for the two isotherms are equal when the Frumkin interaction parameter is zero.

Both the Langmuir and Frumkin isotherms are thermodynamic and so only describe equilibrium behaviour. The consistent kinetic models only suggest a possible kinetic scheme that could fit with the isotherm, therefore the fact that an adsorption isotherm fits the Langmuir or Frumkin isotherm does not necessarily mean the kinetics of that surfactant will fit the kinetic model suggested in this section.

3.3.5 Frumkin Isotherm

An extension to the Langmuir isotherm due to Frumkin²³⁷ accounts for adsorbate-adsorbate interactions through an interaction parameter, ω , within the Bragg-Williams approximation (i.e. random mixing):

$$\frac{\Gamma}{\Gamma_\infty} = \frac{K_L e^{\omega\Gamma/\Gamma_\infty} c}{1 + K_L e^{\omega\Gamma/\Gamma_\infty} c} \quad (3.41)$$

It is useful to note that when $\omega = 0$, the Frumkin model reduces to the Langmuir model. The Frumkin isotherm is also known as the Fowler-Guggenheim isotherm, especially when applied to gas-phase adsorption.

Following the approach of Curwen *et al.*,⁴⁹ we adopt the kinetic scheme

$$\frac{d\Gamma}{dt} = k_a e^{\omega_a \Gamma / \Gamma_\infty} c(\Gamma_\infty - \Gamma) - k_d e^{\omega_d \Gamma / \Gamma_\infty} \Gamma \quad (3.42)$$

where $K_L = k_a/k_d$ and $\omega = \omega_a - \omega_d$. This scheme is consistent with the equilibrium Frumkin isotherm. It is worth mentioning an alternative scheme consistent with the Frumkin isotherm used by Ferri and Stebe,²³⁸ where adsorption is proportional to surface coverage, and the interaction energy is only important during desorption. This is equivalent to the scheme used by Curwen with ω_a fixed to 0.

We split this equation into the finite difference form

$$\Gamma^{n-1} - \Gamma^n + \Delta t k_a e^{\omega_a \Gamma / \Gamma_\infty} c_0^n (\Gamma^n - \Gamma_\infty) - \Delta t k_d e^{\omega_d \Gamma^n / \Gamma_\infty} \Gamma^n = 0 \quad (3.43)$$

and create a special case of equation 3.35 for $j = 0$, describing the solution immediately next to the solid surface:

$$\left(\frac{1}{\Delta t} + \frac{D}{\Delta z^2} + \frac{k_a e^{\omega_a \Gamma^n / \Gamma_\infty}}{\Delta z} (\Gamma^n - \Gamma_\infty) \right) c_0^n - \frac{D}{\Delta z^2} c_1^n = \frac{c_0^{n-1}}{\Delta t} + \frac{k_d e^{\omega_d \Gamma^n / \Gamma_\infty}}{\Delta z} \Gamma^n \quad (3.44)$$

Equation 3.44 is nonlinear in Γ therefore the set of partial differential equations cannot be written in the form of equation 3.36. The method used to circumvent the problem has undergone a number of iterations, of which only the latest variant is described. Reference 239 details the most important earlier method (both cases tend towards the same result; the improved version here only improves calculation speed).

To solve the equations we generate linear approximations and then iterate using the Newton-Raphson method. To generate a sensible starting point for the iteration an estimated value for Γ^n is calculated from eqn. 3.43 with $c_0^n = c_0^{n-1}$. This value of Γ^n is then substituted into eqn. 3.44 to give linear equations of the form of eqn. 3.36 which are solved using the Thomas algorithm. The value of c_0^n that arises is used in eqn. 3.43 to generate a better estimate of Γ^n . The Newton-Raphson iteration uses these values (c_0^n and the improved Γ^n) as its initial values.[†] The matrix form of Newton-Raphson method is written as

$$\mathbf{u}^{m+1} = \mathbf{u}^m - \frac{\mathbf{F}(\mathbf{u}^m)}{\mathbf{J}(\mathbf{u}^m)} \quad (3.45)$$

or

$$\mathbf{J}(\mathbf{u}^m) \mathbf{u}^{m+1} = \mathbf{J}(\mathbf{u}^m) \mathbf{u}^m - \mathbf{F}(\mathbf{u}^m) \quad (3.46)$$

where \mathbf{u}^m is a vector of concentrations at timestep n and iteration m (in this case the first element is Γ^n and the concentrations c_0^n to c_{jj}^n follow in the second and

[†] *Note:* the calculation of the starting points for the iteration in this manner is largely an artefact of how the program was developed; Γ^{n-1} and c_0^{n-1} would work equally well as starting points and give a simpler program, at the cost of a few extra iterations.

subsequent elements). \mathbf{J} is the jacobian matrix (calculated using the current best estimates for Γ^n and c_0^n where necessary), given by

$$\mathbf{J}(\mathbf{u}) = \begin{pmatrix} \frac{\partial f_1}{\partial u_1} & \frac{\partial f_1}{\partial u_2} & \cdots & \frac{\partial f_1}{\partial u_N} \\ \frac{\partial f_2}{\partial u_1} & \frac{\partial f_2}{\partial u_2} & \cdots & \frac{\partial f_2}{\partial u_N} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_N}{\partial u_1} & \frac{\partial f_N}{\partial u_2} & \cdots & \frac{\partial f_N}{\partial u_N} \end{pmatrix} \quad (3.47)$$

and \mathbf{F} —the error in the solution of each constituent equation—by

$$\mathbf{F}(\mathbf{u}) = \begin{pmatrix} f_1 \\ f_2 \\ \vdots \\ f_N \end{pmatrix} \quad (3.48)$$

The elements f_1 to f_N are the equations that make up the nonlinear system, so here f_1 , f_2 and $f_{3\dots N}$ are the left-hand sides minus the right-hand sides of equations 3.43, 3.44 and 3.35 respectively. Where the largest element in $|\mathbf{F}| < 10^{-5}$ or the mean value of $|\mathbf{F}| < 10^{-10}$ the answer is deemed acceptable and iteration stops. Typically between 0 (where the values obtained in estimating Γ^n and c_0^n are good enough to use as is) and 2 iterations are required at each time step, with the most iterations required when the surface excess is changing most rapidly.

The elements of \mathbf{J} are

$$\frac{\partial f_k}{\partial u_{k-1}} = \frac{\partial f_k}{\partial c_{j-1}^n} = \frac{-D}{\Delta z^2} + \frac{j^2 \alpha \Delta z}{2} \quad \text{for } k = 3 \dots N, \quad (3.49a)$$

$$\frac{\partial f_k}{\partial u_k} = \frac{\partial f_k}{\partial c_j^n} = \frac{2D}{\Delta z^2} + \frac{1}{\Delta t} \quad \text{for } k = 3 \dots N, \quad (3.49b)$$

$$\frac{\partial f_k}{\partial u_{k+1}} = \frac{\partial f_k}{\partial c_{j+1}^n} = \frac{-D}{\Delta z^2} - \frac{j^2 \alpha \Delta z}{2} \quad \text{for } k = 3 \dots N - 1, \quad (3.49c)$$

$$\frac{\partial f_1}{\partial u_1} = \frac{\partial f_1}{\partial \Gamma^n} = -1 + \left[k_a (\omega_a - 1 - \omega_a \Gamma^n / \Gamma_\infty) e^{\omega_a \Gamma^n / \Gamma_\infty} c_0^n - k_d (1 + \omega_d \Gamma / \Gamma_\infty) e^{\omega_d \Gamma / \Gamma_\infty} \right] \Delta t, \quad (3.49d)$$

$$\frac{\partial f_1}{\partial u_2} = \frac{\partial f_1}{\partial c_0^n} = k_a e^{\omega_a \Gamma^n / \Gamma_\infty} (\Gamma_\infty - \Gamma^n) \Delta t, \quad (3.49e)$$

$$\frac{\partial f_2}{\partial u_1} = \frac{\partial f_2}{\partial \Gamma^n} = \left[k_a (\omega_a - 1 - \omega_a \Gamma^n / \Gamma_\infty) e^{\omega_a \Gamma^n / \Gamma_\infty} c_0^n - k_d (1 + \omega_d \Gamma / \Gamma_\infty) e^{\omega_d \Gamma / \Gamma_\infty} \right] / \Delta z, \quad (3.49f)$$

$$\frac{\partial f_2}{\partial u_2} = \frac{\partial f_2}{\partial c_0^n} = \frac{1}{\Delta t} + \frac{D}{\Delta z^2} + \frac{k_a e^{\omega_a \Gamma^n / \Gamma_\infty}}{\Delta z} (\Gamma^n - \Gamma_\infty) \quad (3.49g)$$

$$\frac{\partial f_k}{\partial u_l} = 0 \quad \text{otherwise.} \quad (3.49h)$$

In the above equations, $j = k - 2$, due to the need to fit in Γ^n and c_0^n as indexes 1 and 2.

Phase separation

Surfactants such as TX-100—where the isotherm approaches a step function—are characterised by a large value of ω , representing a strong attractive interaction between neighbouring adsorbates. When $\omega > 4$ there is a phase separation within the adsorbed layer. We have based our treatment of phase separation on the work of Ferri and Stebe.²³⁸ The Gibbs free energy at constant temperature and pressure is given by combining equation 1.12 with equation 1.15 and eliminating the contributions from bulk surfactant on the basis that $\sum_i n_i d\mu_i = 0$. For a system with a single surfactant, and the solvent surface excess set to zero through the choice of the Gibbs dividing plane, and thus,

$$G = \gamma A + \mu n^\sigma, \quad (3.50)$$

where μ is the chemical potential of the surfactant. This equation also describes the free energy of a single phase on a surface with multiple phases. The surface energy, γ of the two phases must be equal, otherwise it would be favourable to expand one of the phases until it covered the whole surface. At equilibrium both surface phases are in equilibrium with surfactant in solution and therefore their chemical potentials must be equal and so,

$$\begin{aligned} \frac{G_1}{n_1^\sigma} - \frac{\gamma A_1}{n_1^\sigma} &= \frac{G_2}{n_2^\sigma} - \frac{\gamma A_2}{n_2^\sigma} \\ \frac{G_1}{n_1^\sigma} - \frac{G_2}{n_2^\sigma} &= \gamma \left(\frac{A_1}{n_1^\sigma} - \frac{A_2}{n_2^\sigma} \right), \end{aligned} \quad (3.51)$$

using the subscripts 1 and 2 to denote properties of the two separate phases. Since $\frac{\partial G}{\partial A} = \frac{\partial(G/n^\sigma)}{\partial(A/n^\sigma)} = \frac{\partial(G/n^\sigma)}{\partial a} = \gamma$ where $a = A/n^\sigma$ is the area per molecule,²³⁸

$$\int_{a_2}^{a_1} \gamma da = \gamma(a_1 - a_2), \quad (3.52)$$

where a_1 and a_2 are the area per molecule of the surfactant at the two binodal points ($1/\Gamma_1$ and $1/\Gamma_2$, respectively).

Therefore, the surface excesses of the two coexisting phases are found by solving equation 3.52, with the surface energy given by,

$$\gamma = \gamma_0 + RT\Gamma_\infty \left[\ln \left(1 - \frac{\Gamma}{\Gamma_\infty} \right) + \frac{\omega}{2} \left(\frac{\Gamma}{\Gamma_\infty} \right)^2 \right], \quad (3.53)$$

and plotted in figure 3.6. γ_0 is the surface energy at the pure solvent interface, but cancels out and need not be known for any calculation. $RT\Gamma_\infty$ is a constant factor on both sides of equation 3.52 so need not be known. We avoid the trivial solution where $\Gamma_1 = \Gamma_2$. The surface energies of the two phases must be equal and—when plotted against area per mole (see figure 3.7)—the area between the curve and a flat line connecting Γ_1 and Γ_2 must be equal above and below the flat line.

Equation 3.52 is solved by numerically minimising the function,

$$f(a_1, a_2) = \left(\int_{a_2}^{a_1} \gamma da - \gamma[a_1 - a_2] \right)^2 + ([\gamma(a_1) - \gamma(a_2)][a_1 - a_2])^2. \quad (3.54)$$




Figure 3.6: Variation of surface free energy with surface coverage, for values of ω of 0 (Langmuir isotherm), 4 (just critical) and 6 (greater than critical).




Figure 3.7: Variation of surface free energy with area per mole. $\omega = 6$ and $\Gamma_\infty = 5 \mu\text{mol m}^{-2}$ (a fairly typical value). A thin black line connects the surface excesses of the two phases.

The extra term of $([\gamma(a_1) - \gamma(a_2)][a_1 - a_2])^2$ is added to ensure that γ is equal at both points. When this condition is met the term disappears and so does not affect the solution. The differences are squared to ensure $f(a_1, a_2) \geq 0$. The integral is evaluated as,

$$\frac{\int_{a_2}^{a_1} \gamma da}{RT\Gamma_\infty} = \frac{\ln\left(1 - \frac{\Gamma}{\Gamma_\infty}\right)}{\Gamma} - \frac{\ln\left(\frac{\Gamma_\infty}{\Gamma} - 1\right)}{\Gamma_\infty} - \frac{\omega\Gamma}{2\Gamma_\infty^2}. \quad (3.55)$$

In order to ensure the calculation arrives at the correct answer bounds are set on acceptable values for the surface coverages of the two coexisting phases:

$$0 \leq \Gamma_1 < \Gamma_{\min} < \Gamma_{\max} < \Gamma_2 \leq \Gamma_\infty, \quad (3.56)$$

where Γ_{\min} is the local minimum in surface free energy (around $0.2\Gamma_\infty$ in the $\omega = 6$ curve on figure 3.6) and Γ_{\max} is the local maximum in surface free energy (around $0.85\Gamma_\infty$ for $\omega = 6$). I have found that the centre of the allowed region is a sensible initial value for the minimisation routine.

When the surface excess is below Γ_1 or above Γ_2 the calculation of concentrations and surface excesses proceeds as described above using equations 3.43 and 3.44; when the surface excess is between the two binodal points we use a lever rule, so that

$$\begin{aligned} \frac{\partial\Gamma}{\partial t} = (1-x) & \left[k_a e^{\omega_a \Gamma_1 / \Gamma_\infty} c(\Gamma_1 - \Gamma_\infty) - k_d e^{\omega_d \Gamma_1 / \Gamma_\infty} \Gamma_1 \right] + \\ & x \left[k_a e^{\omega_a \Gamma_2 / \Gamma_\infty} c(\Gamma_2 - \Gamma_\infty) - k_d e^{\omega_d \Gamma_2 / \Gamma_\infty} \Gamma_2 \right], \quad (3.57) \end{aligned}$$

where x is the fraction of adsorbed surfactant in the second state, given by

$$x = \frac{\Gamma - \Gamma_1}{\Gamma_2 - \Gamma_1} \quad (3.58)$$

The substitutions into equations 3.43 and 3.44 are reasonably straightforward, with any term involving Γ^n being replaced by the appropriately weighted terms for Γ_1 and Γ_2 if Γ^n is within the phase separated region. The modified values of $\frac{\partial f}{\partial u}$ used in the Newton-Rhapson iteration are as follows:

$$\begin{aligned} \frac{\partial f_1}{\partial u_1} = \frac{\partial f_1}{\partial \Gamma^n} = -1 + & \left[\frac{k_a}{\Gamma_2 - \Gamma_1} \left(e^{\omega_a \Gamma_2 / \Gamma_\infty} (\Gamma_\infty - \Gamma_2) - \right. \right. \\ & \left. \left. e^{\omega_a \Gamma_1 / \Gamma_\infty} (\Gamma_\infty - \Gamma_1) \right) c_0^n - \right. \quad (3.59a) \end{aligned}$$

$$\begin{aligned} & \left. \frac{k_d}{\Gamma_2 - \Gamma_1} \left(e^{\omega_d \Gamma_2 / \Gamma_\infty} \Gamma_2 - e^{\omega_d \Gamma_1 / \Gamma_\infty} \Gamma_1 \right) \right] \Delta t, \\ \frac{\partial f_2}{\partial u_1} = \frac{\partial f_2}{\partial \Gamma^n} = & \left[\frac{k_a}{\Gamma_2 - \Gamma_1} \left(e^{\omega_a \Gamma_2 / \Gamma_\infty} (\Gamma_\infty - \Gamma_2) - \right. \right. \\ & \left. \left. e^{\omega_a \Gamma_1 / \Gamma_\infty} (\Gamma_\infty - \Gamma_1) \right) c_0^n - \right. \quad (3.59b) \\ & \left. \frac{k_d}{\Gamma_2 - \Gamma_1} \left(e^{\omega_d \Gamma_2 / \Gamma_\infty} \Gamma_2 - e^{\omega_d \Gamma_1 / \Gamma_\infty} \Gamma_1 \right) \right] / \Delta z. \end{aligned}$$

$\frac{\partial f_1}{\partial u_2}$ and $\frac{\partial f_2}{\partial u_2}$ only require a simple substitution using a lever rule, and so are not given. The other terms are unchanged.

The rates of adsorption to the two phases is typically unequal. The model allows

surfactant to adsorb to either of the phases and assumes rapid equilibration within the layer to the composition given by the lever rule. However, if exchange between the two phases is slow compared to adsorption then an alternative possibility arises. The adsorbed molecules may be able to exist in the energetically unstable region of the energy–molecular-area curve (following the red line in figure 3.7 rather than the black tie-line).²⁴⁰ I have not attempted to model this possibility.

Non-dimensionalisation

The Frumkin kinetic scheme can be non-dimensionalised in the same manner that the convection–diffusion equation was. Γ is divided by Γ_∞ and c by $\Gamma_\infty(\alpha/D)^{1/3}$. The latter represents the bulk concentration for which the amount of surfactant in the diffusion layer is equal to the amount in a saturated surface layer. For a typical value of $\Gamma_\infty = 5 \mu\text{mol m}^{-2}$, $\Gamma_\infty(\alpha/D)^{1/3} \sim 0.1 \text{ mM}$. The non-dimensionalised form of equation 3.42 is

$$\frac{d\Gamma'}{dt'} = k'_a e^{\omega_a \Gamma'} c'(1 - \Gamma') - k'_d e^{\omega_d \Gamma'} \Gamma' \quad (3.60)$$

where $k'_a = \Gamma_\infty(D^2\alpha)^{-1/3}k_a$ and $k'_d = (D\alpha^2)^{-1/3}k_d$. Where we have fitted data to obtain constants for the Frumkin model, the dimensional forms are presented. The non-dimensional rate constants are fairly meaningless on their own, but the combined effect of a rate constant and the exponential contribution from the interaction parameter at a specific surface coverage (e.g. $k'_a e^{\omega_a \Gamma'}$ or $k'_d e^{\omega_d \Gamma'}$) can be interpreted usefully. The size of the coverage-dependent effective non-dimensional rate constants describes the relative influence of the adsorption kinetics compared to mass-transport: if they are much less than one then the observed rate is kinetically limited, while values much greater than one indicate a mass transport limitation.

3.3.6 Micelles

At concentrations above the cmc, the surfactant exists as both monomers and micelles. Micelles normally have a lower diffusion coefficient than monomers due to their larger size and therefore mass transport of the micelles must be modelled separately from that of monomers. Following the example of Curwen *et al.*,⁴⁹ we assume that micelles are monodisperse, the diffusion coefficients are concentration independent and the exchange between monomers and micelles is fast on the timescale of the experiment. The mass transport of the surfactant in micelles and as monomer is calculated separately, with a no-flux boundary condition for the micelles, and the concentrations summed at the end of each time step. The total concentration is then repartitioned between micelles and monomers so that the concentration of monomers is capped at the cmc and any remaining surfactant is assigned to the micelles. The micelle component (and the associated diffusion coefficient) is given as the concentration of monomers in micelles, rather than the concentration of whole micelles.

The presence of micelles has a significant effect on mass transport to the surface: using the perfect sink boundary condition—since it is easiest to visualise—figure 3.8 shows the concentration profile for a variety of different variations of mixtures of monomers and micelles. Where monomers and micelles are allowed to coexist, a

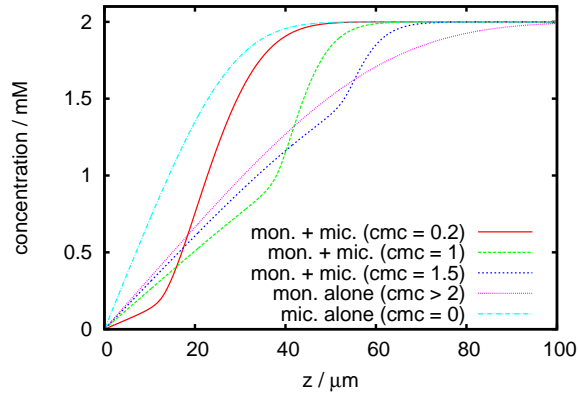


Figure 3.8: Steady-state concentration profiles for different combinations of monomers ($D = 5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$) and micelles ($D = 0.5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$) being transported to a surface ($\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$), assuming a perfect sink boundary condition. The total inlet concentration was 2 mM. The parameters are typical for the surfactants studied in this thesis. The concentration profiles were modelled to $z = 200 \mu\text{m}$ but are truncated to $100 \mu\text{m}$ for clarity.




Figure 3.9: Limiting flux to the surface (calculated using $J = D \frac{\partial c}{\partial z}$ at $z = 0$), for a variety of different cmcs, using the parameters given in the caption of figure 3.8.

change in slope can be observed as the concentration passes the cmc. Figure 3.9 shows the limiting flux to the surface as a function of cmc, illustrating the sensitivity of rate of adsorption to cmc.

3.3.7 Mixed systems

Bulk properties

As discussed in chapter 1, the simplest models for modelling mixed systems are ideal mixing and regular solution theory (RST). I have chosen to use regular solution theory, since it “downgrades” to ideal mixing given an interaction parameter of 0. My RST code can process a mixed system containing two surfactants (A and B). The RST model does not take into account the effect of counterions binding to micelles in mixed ionic–nonionic systems (see the brief discussion in the introduction, on pages 25 and 28, for some more details); this approach is suitable for a qualitative description of the mixed system—especially since experimental fits to RST often do not account for counterion binding either—but may lead to a poorer description of the real system.

The partitioning of surfactant into monomers and mixed micelles is treated in the same way as micelles are treated in a single-component system: at the end of each timestep the total concentration of each surfactant is summed, and then those concentrations are used to determine partitioning into monomers (of components A and B) and monomers in micelles (of components A and B). The components are then all allowed to diffuse separately during the next timestep, before being added together once again. This leads to the slightly counter-intuitive possibility of two components in a mixed micelle being able to diffuse in opposite directions, irrespective of the behaviour of the micelles as a whole. In reality the conclusion is quite reasonable: by assuming that exchange between monomers and micelles is fast, we are allowing the components to diffuse according to their own chemical potential gradients. This can mean, for example, monomers within two micelles of different composition could be rapidly exchanged, allowing the two surfactants to diffuse independently of overall micelle behaviour. As discussed in section 1.3.7 of the introductory chapter, diffusion in mixed systems can be very complex; I have not attempted to capture that complexity here.

The first step is to calculate the cmc of the mixed system from

$$c^* = \frac{x_1^* c_A^* e^{(\beta(1-x_A)^2)}}{\alpha_A}, \quad (3.61)$$

(except where $\alpha_A = 0$ when $c^* = c_B^*$) where c^* is the mixed cmc, α_A is the overall mole fraction of component A and c_A^* and c_B^* are the pure cmcs of A and B. x_A^* here is the mole fraction of A within micelles at the cmc, calculated by solving the equation

$$x_A^* - \frac{\alpha_A c_A^* e^{(\beta x_A^{*2})} (1 - x_A^*)}{(1 - \alpha_1) c_B^* e^{(\beta(1-x_A^*)^2)}} = 0. \quad (3.62)$$

Equation 3.62 is solved with the function minimisation solvers in MATLAB. For any grid-points with a total concentration less than the cmc the concentration of

monomers in micelles is set to 0.

x_A , the mole fraction of A within micelles at any concentration, can then be calculated using from

$$\frac{c\alpha_A}{c + c_A^*(1 - x_A)e^{\beta(1-x_A)^2} - c_B^*(1 - x_A)e^{\beta x_A^2}} - x_A = 0, \quad (3.63)$$

where c is the total concentration. Again, the equation is solved with the MATLAB minimisation solvers. Monomer concentrations for the two components are calculated from

$$c_A^m = x_A c_A^* e^{\beta(1-x_A)^2}, \quad (3.64a)$$

$$c_B^m = (1 - x_A) c_B^* e^{\beta x_A^2}. \quad (3.64b)$$

Concentrations of monomers in micelles are readily calculated from mass balance.

The solution of the two equations using the minimisation solvers can be reasonably slow, typically requiring around 15 iterations. Thus it would be expensive to do for every grid-point in at timestep. Since x_1 and c^* both vary smoothly with respect to α_A and (for x_1) c , they can be well-approximated using an interpolation. Therefore, for the actual calculation I precalculate x_1 and c^* for 60 concentrations and 50 compositions, then uses interpolated values between each timestep.

Interfacial properties

The modelling of interfacial properties for multicomponent systems is currently not implemented. In the current version of the model, each component could adsorb to the interface independently, but could not interact with—or even block adsorption of—other adsorbed components. Such a situation is clearly unrealistic. The main reason for not including mixed interface behaviour was the complexity of the system: for a mixed system without an interface one already needs to know the interaction parameter, β for the two surfactants. When an interface is involved then simply to describe the isotherm there would be three interaction parameters (ω_{AA} , ω_{BB} and ω_{AB}). When converted to a consistent kinetic model then each interaction parameter would have to be split into two components. Fitting such a wide range of parameters would seem unrealistic, especially given that the isotherms observed for the CTAB/TX-100 system adsorbing on silica (see chapter 5) could not be fitted to a regular solution model.

3.3.8 Grid parameters

As with all finite-difference based models, the answer obtained is not an exact solution to the differential equations it approximates. As the size of the grid (Δz and Δt) decreases then the quality of the approximation improves. In this section I determine the values necessary to obtain good results from the model. It is divided into three sections: the first looks at the effect of degree of implicitness. The second uses simple (perfect sink) boundary conditions and the third uses the Frumkin isotherm.

In addition to the grid parameters, the length of the region modelled is also im-




Figure 3.10: The effect of changing the length of cell modelled away from the wall. Modelled using perfect sink boundary conditions, $\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$, $D = 5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$, $\Delta z = 1 \times 10^{-7} \text{ m}$ and $\Delta t = 1 \times 10^{-2} \text{ s}$.




Figure 3.11: Flux to surface as a function of time for various lengths of modelled cell. Conditions are as for figure 3.10. For part (b) the times have been offset to show that the lines are identical in shape.

portant. The length must be great enough to contain the whole diffusion layer. Figure 3.10 shows concentration profiles for various lengths of cell; when the length is greater than $\sim 200\ \mu\text{m}$ the final concentration profile is insensitive to cell length. Figure 3.11 shows the flux to the surface as a function of time in both unchanged and offset forms. Once the length modelled is large enough for the diffusion layer, extending the length adds a small time offset (because the surfactant at the inlet at $t = 0$ is placed further from the surface), but the results are otherwise identical.

The absolute maximum distance from the surface that it is realistic to model is $\sim 1000\ \mu\text{m}$. At that point v_z is approximately equal to the inlet velocity at the centre of the parabolic flow in the inlet tube (using experimentally realistic values of $\alpha = 5000\ \text{m}^{-1}\ \text{s}^{-1}$, an inlet tube radius and a total flow rate within the inlet of $0.5\ \text{mL}\ \text{min}^{-1}$). In reality the transition from the z^2 dependence of the flow close to the surface to the flow in the tube—which is independent of z —will not be sharp and so $1000\ \mu\text{m}$ is too long to model as z^2 . Melville *et al.* provide a detailed breakdown of how far the z^2 dependence is valid, albeit for a much smaller cell with a different R/h ratio and much higher Reynolds numbers.²³⁰ They found an increasing range of validity for decreasing Reynolds number, with a range of $z = h/3$ for $Re \sim 5000$. It is unclear how that applies my experiment where $Re = 2.7$ for typical flow rates.

There is obviously a trade-off between the grid and length of calculation. Typical calculations involving simple boundary conditions on a 2000 point grid over 1000 time steps are completed in seconds (using 1.5–2 GHz computers). The length of calculation is close to linear with respect to the number of timesteps multiplied by the number of spatial grid points. The length of calculation is doubled by the addition of micelles, and increases further if regular solution theory is used ($4\times$ since there are 4 components, plus an extra $\sim 125\%$ from partitioning using RST, plus a fixed time ($\sim 5\ \text{s}$) at the start preparing lookup tables for the RST calculations. Using the Frumkin isotherm takes longer than the simple boundary condition calculations, because of the iterative nature of the procedure. Typically the length of calculation is tripled (since two extra iterations are usually required), although that can vary depending on the region of the isotherm probed.

Degree of implicitness (θ)

In section 3.3.1 I briefly introduced the concept “degree of implicitness”, θ . In general I have avoided using θ in the discussion of the model since (a) when using the model I have universally set $\theta = 1$ and (b) including the explicit term in the equations makes them more complicated for little added benefit. In this section I will briefly discuss the influence of the θ parameter.

An implicit calculation ensures that the calculation remains numerically stable—numerically unstable calculations are liable to show rapidly increasing waves building up as the calculation progresses—at the cost of introducing numerical diffusion (numerical diffusion means that errors in the calculation are spread out rather than accumulating at individual grid points). Introducing an explicit contribution can help make the calculation more accurate, and up to $\theta = 0.5$, the calculation should still remain numerically stable. However, using perfect sink boundary conditions, the




Figure 3.12: Concentration profiles at $t = 1$ s for different degrees of implicitness, θ . All other parameters are as specified in figure 3.10.

$\Delta z / \mu\text{m}$	$\Delta t / \text{ms}$				
	100	50	10	5	1
1.0	8.1980	8.2019	8.2047	8.2050	8.2053
0.5	8.2675	8.2712	8.2740	8.2743	8.2745
0.1	8.3239	8.3276	8.3302	8.3305	8.3308
0.05	8.3310	8.3347	8.3373	8.3376	8.3379
0.01	8.3367	8.3403	8.3430	8.3433	8.3435

a Limiting flux to the surface, in $\mu\text{mol m}^{-2} \text{s}^{-1}$. The exact limiting flux is $8.3641 \mu\text{mol m}^{-2} \text{s}^{-1}$.

$\Delta z / \mu\text{m}$	$\Delta t / \text{ms}$				
	100	50	10	5	1
1.0	-1.9866	-1.9403	-1.9065	-1.9024	-1.8992
0.5	-1.1557	-1.1110	-1.0782	-1.0743	-1.0712
0.1	-0.4810	-0.4374	-0.4056	-0.4018	-0.3988
0.05	-0.3960	-0.3526	-0.3209	-0.3171	-0.3141
0.01	-0.3279	-0.2846	-0.2530	-0.2493	-0.2463

b Percentage error compared to the exact limiting flux (the negative values indicates the fluxes are below the exact answer).

Table 3.1: Effect of grid size on limiting flux calculation, using perfect sink boundary conditions, $\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$ and $D = 5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$.

difference between $\theta = 0.5$ and $\theta = 1$ is less than 0.01% (the difference between that and the exact answer from equation 3.33 is around 0.4%).

Figure 3.12 shows the effect of changing θ on the simulation. Although the lines do not overlap perfectly, the most important feature is the sharp fluctuation at $z = 200 \mu\text{m}$ for the partially implicit calculation, caused by the fixed boundary condition at the high z edge. I do not believe this affects the overall calculation significantly—and the fluctuations get smaller as the calculation reaches a steady state—but I feel such fluctuations are undesirable, and hence prefer to use $\theta = 1$ to avoid such issues.

Simple boundary conditions

The perfect-sink boundary condition provides a simple test of the accuracy of a calculation, made more appealing since the exact answer is known from equation 3.33. In order to evaluate the effect of changing the grid spacing, I performed calculations at a variety of values of Δz and Δt . In all cases the total simulated time was 10 s

$\Delta z / \mu\text{m}$	$\Delta t / \text{ms}$				
	100	50	10	5	1
1.0	1.6932	1.7330	1.7627	1.7665	1.7695
0.5	1.6989	1.7362	1.7666	1.7704	1.7734
0.1	1.7029	1.7391	1.7697	1.7735	1.7765
0.05	1.7034	1.7396	1.7701	1.7739	1.7769

a Maximum fluxes to the surface, in $\mu\text{mol m}^{-2} \text{s}^{-1}$.

$\Delta z / \mu\text{m}$	$\Delta t / \text{ms}$				
	100	50	10	5	1
1.0	-4.7107	-2.4718	-0.8013	-0.5892	-0.4190
0.5	-4.3922	-2.2915	-0.5811	-0.3686	-0.1977
0.1	-4.1660	-2.1268	-0.4061	-0.1929	-0.0219
0.05	-4.1397	-2.0995	-0.3848	-0.1716	0.0000

b Percentage error compared to the highest calculated flux (smallest grid spacings).

Table 3.2: Effect of grid size on kinetics using a Frumkin model, $\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$, $D = 5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$, $k_a = 1 \text{ mol}^{-1} \text{ m}^3 \text{ s}^{-1}$, $k_d = 1 \text{ s}^{-1}$, $\omega_a = 3$ and $\omega_d = 0$. When using a grid-size of $0.01 \mu\text{m}$ none of the calculations converged within the 100 iterations during the Newton-Raphson iteration.

and total length modelled was $200 \mu\text{m}$. Table 3.1 shows that the flux is less sensitive to the length of the timestep than the spatial step-size. For the shorter timesteps, the answers can be improved by running the simulation for longer (the calculation is tending towards the steady-state answer, but has not reached it completely).

Frumkin isotherm

Although the perfect-sink boundary condition is easy to evaluate, since results can be compared with an analytical solution, it does not represent the typical use of the model developed here. The majority of the modelling throughout the remainder of this thesis uses the Frumkin isotherm and a consistent kinetic model. Therefore, I have evaluated the maximum rates of adsorption using the Frumkin isotherm for a variety of different grid spacings. In this case there is no known “exact” answer to compare with. Table 3.2 shows the maximum fluxes from the Frumkin model (other parameters were kept the same as for the perfect sink calculations). The values for the Frumkin model were chosen as being typical for the surfactants measured. For the Frumkin model the length of the timestep is more important than the grid spacing (this makes sense given that the Frumkin isotherm is applied at one grid point alone).

Conclusions on grid size

The optimum value of Δt is around 10ms and the optimum value of Δz is around $0.1 \mu\text{m}$. Beyond this point the accuracy of the calculations does not improve significantly, at the cost of a much longer calculation time. It is also worth noting that the calculation always under-estimates the flux to the surface compared to the analytical solution.




Figure 3.13: (a) Initial concentration profiles: the red line uses a step function, the other lines use an error function with widths, σ , of 10, 50 and 100 μm (lowest values are the steepest lines). (b) Fluxes calculated using the perfect-sink boundary condition starting with the concentration profiles from (a). All other parameters are as specified in figure 3.10. Line colours in (a) and (b) match. The lines in (b) overlap to the extent that they cannot be distinguished, indicating an insensitivity to the initial concentration profile.

3.3.9 Initial concentration profile

A potential concern when using the model is that the initial conditions for an adsorption run are specified as an empty cell with step change to the inlet concentration at the high- z boundary. In reality the concentration profile will become less sharp as it is transported down the inlet tube due to the influence of diffusion. Therefore, it is useful to check the influence of the sharpness of the initial concentration profiles. Figure 3.13(a) shows a variety of different initial concentration profiles, defined using

$$c(z) = \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{z - 400 \mu\text{m}}{\sigma\sqrt{2}} \right) \quad (3.65)$$

where σ is a characteristic width and $\operatorname{erf}(x)$ is the error function. The sharpest function is defined differently, as a step function centred at $z = 400 \mu\text{m}$. Figure 3.13(b) shows the flux to the surface from a calculation starting with the concentration profiles given. The time dependence of the flux is almost identical indicating that the model does not depend on the sharpness of the initial concentration profile. This is not too surprising: the concentration profile gets modified as it nears the surface to form a diffusion layer of characteristic thickness, $(D/\alpha)^{1/3}$.

3.4 Comparison to other models

The simplest model for adsorption kinetics together with mass transport is based on the limiting flux, given by equation 3.33. This model has been used widely,^{43,45,100,106} often in combination with a “sticking-ratio”, which compares the fraction of molecules hitting the surface with the fraction that actually adsorb. At low surface coverages there is likely to be negligible desorption, and hence the concept of a sticking-ratio is a reasonable way to view the adsorption coefficient. However, the greatest flaw in




Figure 3.14: Surface excess (red solid line, left axis) and flux to the surface (green dashed line, right axis) with a perfect sink boundary condition (a sticking-ratio of 1), as the solution first reaches the surface ($\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$, $D = 5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$). The horizontal red dashed line illustrates a typical saturated surface of $5 \text{ } \mu\text{mol m}^{-2}$.

the sticking-ratio approach is that, by the time the flux to the surface has actually reached its limiting value, the surface is completely saturated (using typical values of the sticking-ratio⁴⁵ of 0.1–1). Figure 3.14 models this effect using a perfect sink boundary condition (a sticking-ratio of 1). Hence, the actual mass transport to the surface at the time adsorption begins is not well described by equation 3.33, and any sticking ratios obtained will be sensitive to the flow rate.

Avena and Koopal used a more sophisticated approach to the idea of the sticking ratio, and introduced the boundary condition at the surface,²⁴¹

$$\frac{d\Gamma}{dt} = k_a c_s \phi - k_d f(\Gamma), \quad (3.66)$$

where k_a and k_d are rate constants, c_s is the subsurface concentration, ϕ is the available surface area (in the absence of competitive adsorbers $\phi = 1$) and $f(\Gamma)$ is a function relating to the adsorbed amount, set to 0 since they only attempt to model the initial rate of adsorption. The initial rate of adsorption

$$\left. \frac{d\Gamma}{dt} \right|_{t=0} = \frac{k_a \phi J_0 c_{\text{in}}}{k_a \phi c_{\text{in}} + J_0}, \quad (3.67)$$

where J_0 is the limiting rate of adsorption, given by equation 3.33. Equation 3.67 reduces to two simplified cases when either $k_a \phi \gg J_0$ or $k_a \phi \ll J_0$. They went on to merge the model with kinetic models consistent with a variety of isotherms, including the Langmuir isotherm, the Frumkin isotherm and a treatment of electrostatic interactions.²⁴² The Langmuir variant was used to model the adsorption of $C_n E_m$ surfactants onto cellulose.¹⁰⁶ However the concern still remains that the flux is not well known during the time when most adsorption takes place. I have attempted to model the effect of starting adsorption with $J = J_0$ in figure 3.15: it compares the effect of starting from an empty cell (realistic) with starting with the concentration profile obtained from using perfect sink boundary conditions for two different sets of




Figure 3.15: Effect assuming the initial flux is the limiting flux. Red solid lines represent starting from an empty cell, while green dashed lines represent starting with the concentration profile calculated using a perfect sink boundary condition (as in figure 3.5). The arbitrary delay before adsorption is not of interest here. Both (a) and (b) have the same Langmuir isotherm, but different kinetic parameters: (a) $k_a = 100$, $k_d = 1$ (the non-dimensional $k'_a \approx 46$, hence adsorption is much faster than mass transport) and (b) $k_a = 0.1$, $k_d = 0.001$ ($k'_a \approx 0.046$, hence adsorption is much slower than mass transport). In both cases $\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$, $D = 5 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$, $c_{\text{in}} = 1 \text{ mM}$ and $\Gamma_{\infty} = 5 \mu\text{mol m}^{-2}$.

Langmuir kinetic parameters (adsorption much faster than transport and adsorption much slower than transport). There is little difference when adsorption is slow, but when adsorption is fast the difference in rate (especially the initial rate) is significant. For the adsorption of C_{12}E_6 surfactants onto cellulose¹⁰⁶ transport was much faster than both adsorption or desorption, and so always setting the flux to be J_0 should not introduce large errors. Wertz and Santore also used the approximation that the flux to the surface can be set to the limiting flux to model enzyme adsorption kinetics.²⁴³ The hydrodynamics were not those of the wall-jet flow cell, but a similar expression for J_0 applied for their experimental set-up. In that case the main interest was an overshoot in the adsorbed amount attributed to a change in molecular configuration while adsorbed, and so the model primarily focused on that feature. Hence, inaccuracies in the value of J_0 at the start of adsorption were probably unimportant to the effect being studied. Tabor *et al.* modelled a double-Langmuir isotherm with a limiting transport rate.⁵² The double-Langmuir isotherm accounted for slower adsorption rates at high surface coverages, a shape that could also be modelled using a Frumkin isotherm with a positive value of ω_a . Which model of the surface is preferred would have to be based on a judgement about whether it was more realistic to have two different types of surface site (as in the double-Langmuir model), or a decrease in rate due to the difficulty of inserting into the adsorbed layer (as in the Frumkin model). A double-Langmuir isotherm could also be combined with the full mass-transport description used here.

The model used by Curwen *et al.*^{49,94} is very similar in approach to the one used here, although it describes different hydrodynamics to the model used here. The maximum rate of mass transport is lower than in the wall-jet flow cell, since the mass-transport of surfactant is controlled entirely by diffusion, and must travel

400 μm —approximately 4/5 of the width of the flow cell. For the wall-jet, only the last 50–100 μm is diffusion controlled, and transport up to that region comes from convection. The advantage of using the dual channel flow cell—as opposed to the wall-jet flow cell—is that the initial concentration profile at $t = 0$ is better defined: the cell is first set to that profile using a continuous flow to reach a steady-state condition, and then the flow is stopped, allowing transport to begin from the well-defined initial state. With the wall-jet flow cell there is always some uncertainty in how lateral and axial diffusion affects the concentration profile as it passes through the inlet pipe although, as I showed earlier in this chapter, the flux to the surface is insensitive to a spread initial concentration profile. Aside from the differences in hydrodynamics the model presented in this thesis introduces partial modelling of multicomponent systems (although not at the surface itself), and an improved method of dealing with the non-linear Frumkin isotherm at the surface.

Tiberg and co-workers used a substantially different model that viewed adsorption and desorption as mass transport controlled. They modelled a flat surface in contact with a stirred solution, for which a characteristic diffusion layer thickness could be experimentally determined (and was a fitting parameter in the model). The model went through several iterations. The original version,⁵⁸ published in 1994, models exchange between monomers and micelles as two extra terms on the transport equations so that breakdown of micelles is given by $-k_-c_{\text{mic}}$ and formation of micelles by $k_+c_{\text{mon}}^N$ where N is the aggregation number, and c_{mic} the concentration of monomers in micelles (here, I have deviated from the notation used in the original papers which is expressed in terms of the concentration of micelles). Using the idea that the establishment of a concentration profile across the diffusion layer is quick compared to transport to the surface allows the simplification of mass transport to a linear expression:

$$D_{\text{mon}}c_{\text{mon}} + D_{\text{mic}}c_{\text{mic}} = Ax + B, \quad (3.68)$$

where A and B are constants, and x is distance from the surface. Transport to the surface is then simply equal to the slope at the subsurface, A . This model predicts a linear increase in surface excess with time near the onset of both adsorption and desorption. Desorption at long times is explained by viewing the adsorbed layer as a thin layer of mostly surface aggregates, with the rate limiting step being the breakdown of the aggregates into monomers with an exponential decay. The next iteration of the model (1998),⁶³ replaced the single rate constants for micelle growth and breakdown with a more complex expression derived from the Aniansson and Wall model of micelle breakdown (but still splitting the solution into a monomer species and a micelle species rather than modelling every aggregation number as the full Aniansson and Wall model would require). The assumption of a linear concentration profile from the surface was dropped, and the full diffusion layer was modelled. The adsorbed layer was modelled as before: as a thin layer within which everything was assumed to be adsorbed. Adsorption could take place through three mechanisms: the transport of monomers into the adsorption layer followed by aggregation (only a very minor contribution), the transport of monomers into the adsorption layer to join an existing surface micelle (diffusion controlled, usually the dominant step), and the

transport of micelles directly onto the surface (diffusion controlled, with a Langmuir isotherm-type dependence on the free surface area). The lack of a sharp concentration cut-off for micelles in solution was important for the second of these steps, since a small initial adsorption of micelles allowed monomers to transfer much more quickly. The final iteration of the model (published concurrently),⁶⁴ allowed mixed surfactant systems, assuming ideal mixing in micelles and at the surface (reasonable for C_nE_m surfactants). The mechanism of adsorption where monomers diffused into the adsorbed layer then aggregated was neglected since it had been of limited importance in single component systems.

The treatment of micellisation used by Tiberg and coworkers is more sophisticated than the one presented in this chapter, however it is by no means a full picture. In a more complete model every possible micelle aggregation number (below a maximum cut-off) would have to be calculated together with exchange between them and transport. The largest limitation of relating kinetics of micelle formation to the Aniansson and Wall model is that complete micelle breakdown is unrealistically slow, although the behaviour following small perturbations to equilibrium are well described.¹⁴ The view of adsorption as primarily an aggregation process differs from my work, although the Frumkin ω parameter accounts for similar effects (especially when $\omega > 4$ and phase separation takes place). Direct adsorption of micelles was another difference compared to my model, although micellar adsorption varied between 7% of the total amount and “negligible”⁶³ (even though the majority of the adsorbed layer existed as aggregates).

Chapter 4

Single component systems

4.1 Introduction

For reasons described in the introduction, surfactant adsorption kinetics are of practical interest in a wide range of fields. However, classical methods for measuring adsorption isotherms on powders, based on the depletion of surfactant in solution, are not readily extendable to kinetic measurements other than on very long timescales (minutes). Penetration experiments into fibres or packed beds can be carried out on shorter timescales, but are of comparatively little benefit since it is difficult to relate the penetration rate to a microscopic kinetic model. Quantitative studies therefore tend to use model systems in which adsorption occurs at planar interface that can be interrogated with optical or electrochemical techniques. Cell designs with well-defined hydrodynamics (such as the channel-flow cell or wall jet) then permit quantitative modelling of adsorption kinetics. The closely related techniques of ellipsometry and reflectometry have been extensively used to study adsorption of single surfactants onto hydrophilic and hydrophobic silica surfaces. These optical techniques have high sensitivity (<1% of a monolayer) and rapid acquisition times (milliseconds) and are well-suited for studying adsorption at the solid-liquid interface, which in most studies takes place on timescales from seconds to minutes depending on whether mass transport or interfacial kinetics limit the adsorption rate. These techniques essentially measure the amount of adsorbed material and provide little or no information on the chemical nature of the adsorbed species.

Practical formulations invariably contain multiple ingredients that are added for a diverse range of purposes, for example to control rheology, disperse particulates, inhibit corrosion or oxidation and prevent bacterial growth. These additives are frequently surface-active. In addition, mixtures of surfactants are frequently employed to improve the effectiveness or efficiency of a formulation. To understand the mode of action of surfactants in mixtures one needs analytical techniques that have chemical as well as interfacial selectivity. ATR-IR is a spectroscopic technique that can be used to study surfactant adsorption kinetics,^{47,244} however, it does have significant limitations. The strong IR water band necessitates the use of either deuterated surfactant or solvent to be able to see the comparatively weak hydrocarbon signal. A large sampling area is required to be able to record spectra rapidly, requiring hydro-

dynamic control over a large region (Clark and Ducker solved this problem elegantly by using a small bubble between incoming and outgoing solutions to force a rapid exchange). Finally, the penetration depth of the IR radiation into the solution varies strongly with IR wavelength, complicating the data analysis. Neutron reflectivity (NR) has been widely used to study pure and mixed layers at solid-liquid interface and is especially valuable for providing insight into surface structure. Acquisitions times are typically tens of minutes to hours, which is too slow for most kinetic processes of surfactants at surfaces. Only a limited range of substrates are compatible with NR, of which silicon, quartz and sapphire are the most prevalent,²⁴⁵ although the range of surfaces can be increased by functionalising the basic substrate.^{246,247} Chemical specificity in NR normally requires selective deuteration of surfactants.

Recently, Tyrode *et al.* used TIR-Raman scattering to study the adsorption of the cationic surfactant CTAB (hexadecyltrimethylammonium bromide) at the silica-water interface⁴⁰ and demonstrated that a sensitivity of 1% of the saturation coverage was achievable ($\sim 3 \times 10^{-8} \text{ mol m}^{-2}$). We now turn our attention to adsorption kinetics. In this chapter, I report the adsorption kinetics of single surfactant systems at the silica-water interface in the controlled environment of a wall jet cell. CTAB and the non-ionic surfactant Triton X-100 (TX-100) are used as exemplars. I demonstrate a time resolution of 2 s with a typical precision of $1 \times 10^{-7} \text{ mol m}^{-2}$ in the adsorbed amount. In chapter 5, I address the competitive adsorption of CTAB and TX-100 and demonstrate that the same sensitivity and time resolution is achievable in a mixed surfactant system without the use of deuterated surfactants.

This chapter is structured as follows. Section 4.2 presents the equilibrium isotherms of the two surfactants and the best fits to the Frumkin adsorption model. Section 4.3 presents kinetics of adsorption and desorption of CTAB and TX-100 together with a quantitative analysis of these data. Section 4.4 concludes with a discussion of the strengths and limitations of the TIR-Raman technique and of the models employed to describe the adsorption kinetics of CTAB and TX-100. The experimental design has already been discussed in chapter 2 and the mass transport model in chapter 3.

4.1.1 Diffusion coefficients

For TX-100 we used values for the diffusion coefficient, $D_{\text{mon}} = 2.8 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$ for monomers²⁴⁸ and $D_{\text{mic}} = 0.53 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$ for micelles.²⁴⁹ As TX-100 is a non-ionic surfactant the mutual diffusion coefficient in dilute solutions is equal to the self-diffusion coefficient. For CTAB micelles we used a value of $D_{\text{mic}} = 0.53 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$ whereas for monomers we used the formula $D_{\text{mon}} = 2D_{\text{CTA}^+}D_{\text{Br}^-}/(D_{\text{CTA}^+} + D_{\text{Br}^-})$ (with $D_{\text{CTA}^+} = 5.6 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$ and $D_{\text{Br}^-} = 20 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$) to give a mutual diffusion coefficient of $8.8 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$ (all CTAB values from Lindman *et al.*²⁵⁰).

4.2 Equilibrium Isotherms

I used the CTAB isotherm from Tyrode *et al.*⁴⁰ and fitted a Frumkin isotherm (described by equation 3.41) to it. Both the original data and the fit are shown in




Figure 4.1: CTAB isotherm (points, taken from ref. 40) and best Frumkin model fit (line, $\Gamma_{\infty} = (5.49 \pm 0.1) \times 10^{-6} \text{ mol m}^{-2}$, $K_L = 0.33 \pm 0.02 \text{ mol}^{-1} \text{ m}^3$, $\omega = 3.1 \pm 0.1$). The experimental data shown were acquired with an S-polarised pump laser and y-polarised detection. Isotherms obtained with other polarisation combinations were very similar. Section 4.2.1 describes how the errors in the fits were calculated. There is an additional error of $\pm 25\%$ in the value of Γ_{∞} arising from the systematic error in the calibration of the surface excess from component weight.

figure 4.1. The fit captures the steep rise in the isotherm near the cmc, but the plateau at low surface coverage is not well represented; this plateau arises from electrostatic interactions with surface charges that are not included in the Frumkin model. Electrostatic interactions also have an effect above the plateau region. Here, the electrostatic interactions hinder adsorption due to repulsion between neighbouring CTAB molecules. The effect of the electrostatic interactions is primarily seen in reducing the value of Γ_{∞} , but may also reduce the interaction parameter, ω .

The TX-100 isotherm and best fit to the Frumkin model are shown in figure 4.2. The experimental isotherm approximates to a step function with a peak at a concentration just above the step. Such peaks are usually attributed to impurities in the surfactant:^{34,251} TX-100 is a mixture of different EO chain lengths and isomers of the hydrophobic moiety. We excluded all points with $\Gamma > 4.8 \mu\text{mol m}^{-2}$ in fitting an isotherm to the data, since no single-component isotherm can account for such a shape. A possible physical interpretation of the step in the isotherm is the coexistence of monomers and aggregates on the surface.

The parameters K_L and ω are strongly correlated and acceptable fits could be achieved for a range of values of $\omega > 4$. For fitting the kinetic data we used the best fit value of $\omega = 5.5$.

4.2.1 Errors in fitting

It is useful to have an estimate of the error in the fit to the equilibrium isotherms. Figure 4.3 shows the sensitivity of the quality of the fit to the parameters K_L and ω . In both cases, K_L and ω are strongly correlated, as evidenced by the long valley on the contour plot. The lack of sensitivity is worse for TX-100, whereas for CTAB there is a clear best fit (shown on the zoomed-in, logarithmic scale version of the plot, figure 4.3(c)). Figure 4.4 shows the sensitivity fit to Γ_{∞} , with K_L and ω held constant. In both cases there is a clear best value of Γ_{∞} .

The error in measured surface excess of each data-point can be estimated from the standard deviation in the surface excesses of the five highest concentrations of




Figure 4.2: a) Raw spectra used to generate the TX-100 isotherm, and b) the isotherm itself. Circles and squares are both experimental measurements, with the circles being excluded from the fitting process. The lines represent the best fit Frumkin isotherm ($\Gamma_{\infty} = (4.30 \pm 0.08) \times 10^{-6} \text{ mol m}^{-2}$, $K_L = 0.26 \pm 0.08 \text{ mol}^{-1} \text{ m}^3$, $\omega = 5.5 \pm 0.8$). The discontinuity in the line occurs where there is coexistence between two phases of different surface densities. Section 4.2.1 describes how the errors in the fits were calculated. There is an additional error of $\pm 25\%$ in the value of Γ_{∞} arising from the systematic error in the calibration of the surface excess from component weight.

TX-100, and is ± 0.2 . Propagating that error through the fitting procedure would be extremely challenging, especially since the value of Γ is solved from equation 3.41 using a minimisation procedure. Instead, I have re-fitting both the CTAB and TX-100 datasets 20 times, applying randomly-generated, normally-distributed errors to each data-point. The standard deviation of the fitting parameters generated gives the estimated error in the parameter. The errors in K_L , ω and Γ_{∞} are given in the captions for figures 4.1 and 4.2, along with the values themselves. Note that the errors in K_L and ω are 4–8 times greater for TX-100 than for CTAB, due to the correlation between them. The systematic error from the calibration of component weight to surface excess—discussed on page 82 is neglected here, but only affects the value of Γ_{∞} (and adds a much bigger error than the scatter of the data).

4.3 Kinetics

4.3.1 CTAB

We recorded the adsorption kinetics of CTAB onto a clean silica surface at concentrations of CTAB from 0.03 to 10 mM at a pumping rate of 0.5 mL min^{-1} , for which $\alpha = 5000 \text{ m}^{-1} \text{ s}^{-1}$ in equation 3.17. We also measured the adsorption kinetics of a 1.2 mM CTAB solution as a function of pumping rate between 0.12 and 1.6 mL min^{-1} . Figure 4.5 shows examples of adsorption profiles for 4 surfactant concentrations. Time $t = 0$ corresponds to the start of the adsorption process; the data are offset to remove the delay between the time when the surfactant was pumped into the flow system and when it reached the surface. The kinetic parameters that provided the best fit to the adsorption curves for the 6 concentrations measured above 0.4 mM are given in the caption. At the lowest concentrations of CTAB ($< 0.4 \text{ mM}$), the equilibrium isotherm underestimates the adsorbed amount and consequently the fit does not match the




Figure 4.3: Contour plots of the sum of the squared differences between the measured data and the fit for (a) CTAB and (b) TX-100 as a function of the fitting parameters K_L and ω . Γ_∞ is held constant across the plot. (c) shows a zoomed in section of (a), with a logarithmic scale on the height axis, to emphasise the position of the minimum. I believe the slightly “blocky” appearance of the plots is a consequence of the iterative procedure used to calculate the position of the tie-lines when $\omega > 4$, while the lines of dense contours arise in places where the concentration at which the surface phase-separates changes enough as to exclude a data-point from the fit.




Figure 4.4: Sensitivity of the sum of the squared differences between the measured data and the fit for (a) CTAB and (b) TX-100 as a function of the fitting parameter Γ_∞ . K_L and ω are held constant.




Figure 4.5: Experimentally measured kinetics (points) and Frumkin model fits (lines) for CTAB adsorbing onto a clean silica surface at concentrations of a) 0.03 mM b) 0.3 mM c) 1 mM d) 3 mM. The fitting parameters used were $k_a = 1.7 \text{ mol}^{-1} \text{ m}^3 \text{ s}^{-1}$, $k_d = 5 \text{ s}^{-1}$, $\omega_a = -1.9$ and $\omega_d = -5$.

data (Figure 4.5 (a) and (b)).

At higher concentrations the fit between experiment and the model is good. The negative values of ω_a and ω_d shows that both adsorption and desorption become slower at high surface coverage. The former is readily explicable in terms of steric hindrance at high surface coverage. The latter is best thought of as a thermodynamic consequence of the positive value of ω , which represents a favourable interaction between neighbouring surfactant molecules.

The non-dimensionalised kinetic constants show directly whether the adsorption/desorption rate is limited by mass transport or by kinetic barriers to adsorption or desorption, depending on whether the rate constant is greater or less than unity. For CTAB, the adsorption rate constant $k'_a e^{(\omega_a \Gamma')}$ varies from 0.6 at low coverage to 0.1 at high coverage, indicating a mixed kinetic / diffusion regime at low coverage and kinetically limited adsorption at high coverage. The desorption rate constant $k'_d e^{(\omega_d \Gamma')}$ varies from 0.1 at full coverage to 18 at low coverage, the large variation arising from the large negative value of ω_d . The initial desorption from the surface is slow and kinetically limited, but once the coverage begins to drop the desorption rate accelerates until it becomes mass transport limited.

The kinetics of desorption from surfaces with a high coverage into pure water (e.g. green crosses in figure 4.6) are close to diffusion-controlled and therefore rather insensitive to the kinetics at the surface. Nevertheless, it is worthwhile to check how




Figure 4.6: Examples of desorption curves for initial concentrations of 0.3 mM (data: +, model: solid line) and 3 mM (data: ×, model: dashed line). Model parameters taken from figures 4.1 & 4.5.

well the desorption kinetics are described by the kinetic parameters calculated from the equilibrium isotherm and the adsorption kinetics. This prediction is shown as the dashed line in figure 4.6. The model Frumkin parameters describe the initial desorption kinetics well, but do not fit the long tail in the desorption curve. This disagreement is entirely expected since the Frumkin isotherm does not capture the plateau in the adsorption isotherm of CTAB at low concentrations. For the same reason, the Frumkin parameters provide a very poor fit to the desorption kinetics at the lower initial surface excess.

The sensitivity of the adsorption and desorption curves to the kinetic parameters is illustrated in figure 4.7 for a final (adsorption) and initial (desorption) concentration of 3 mM. In each case, the ratio of k_a/k_d , and the difference $\omega_a - \omega_d$ were constrained so that K_L and ω were unchanged. The values shown in the caption of figure 4.7 provide a reasonable estimate as to how well the kinetic parameters can be determined from the fit. Additionally there is a contribution from the error in the fit of the isotherm (figure 4.1), however this is insignificant and can be neglected at this stage. The comparative insensitivity of the kinetics to the fitting parameters is a consequence of the mixed kinetic and mass-transport control. There is an additional—harder to quantify—error arising from the fact that the experimental data are not instantaneous readings at a single point in time but averaged values over the course of 1 s; for the fastest kinetics that effect could be significant.

Figure 4.8 shows the maximum adsorption rate in experiments where the concentration of CTAB was kept constant at 1.2 mM and the pumping speed varied. The predicted adsorption rates from the kinetic parameters in the caption of figure 4.5 are shown. Agreement is fair for flow rates of $<1 \text{ mL min}^{-1}$. At higher flow rates the measured rates are lower than predicted, but at these flow rates we are running into limitations imposed by the sampling frequency of the Raman spectra (0.5 Hz).

4.3.2 TX-100

The kinetics of adsorption of TX-100 was measured at a range of concentrations from 0.31 to 10 mM, shown in figure 4.9. The plateau value at 0.47 mM is elevated compared to the higher concentrations in agreement with the equilibrium isotherm presented in figure 4.2. The rise in the plateau at 10 mM arises from the contribution to the





Figure 4.7: Sensitivity of adsorption and desorption model to changing parameters, shown for 3 mM CTAB. The dashed lines represent a change in ω_a and ω_d (± 1 from the preferred value) whereas the dotted lines represent a change in k_a and k_d ($\times / \div 2$ of the preferred value).




Figure 4.8: Maximum adsorption rate with respect to surfactant solution (1.2 mM CTAB) injection rate. The solid line is the calculated values while the circles are experimental.




Figure 4.9: Adsorption of TX-100 at different concentrations.




Figure 4.10: Comparison of experimentally measured (red crosses) and simulated rates of adsorption. The parameters used in the Frumkin simulation are $k_a = 13 \text{ mol}^{-1} \text{ m}^3 \text{ s}^{-1}$, $k_d = 50 \text{ s}^{-1}$, $\omega_a = -3.5$ and $\omega_d = -9$.

signal from bulk surfactant within the evanescent wave. The adsorption kinetics of TX-100 were more difficult to model than those of CTAB, even with the addition of the two-phase region of the Frumkin isotherm. The adsorption curves all have a similar shape but with a maximum adsorption rate that increases with increasing bulk concentration.

Figure 4.10 compares the experimentally measured maximum rates and calculated rates according to two models. The simpler model uses a no-flux boundary condition for subsurface concentrations below the step in the isotherm at 0.24 mM and perfect sink boundary condition above that concentration. This diffusion-controlled model is in reasonable agreement with the experimental data for concentrations $< 0.5 \text{ mM}$. For higher bulk concentrations, the maximum rate of adsorption measured experimentally reaches a plateau indicating that the adsorption rate is limited by interfacial kinetics. Figure 4.11 shows fits using the Frumkin model on a representative adsorption and desorption curve.

To understand the interplay of kinetic and mass transport limitations, it is helpful to compare the experimental adsorption rates with the maximum adsorption rates from the kinetic model. Since micelles do not adsorb in our model, the maximum adsorption rate is reached when the subsurface concentration is equal to the cmc. The solid red line in figure 4.12a shows this kinetically limited flux as a function of surface excess, while the dashed green line shows the flux from the best fit Frumkin parameters for a bulk concentration of 10 mM. This comparison shows that at low coverages adsorption is mass-transport limited while at higher coverages it is kinetically limited.




Figure 4.11: a) Adsorption curve for 0.46 mM TX-100 together with best fit to the Frumkin model (see Fig. 15 for parameters). b) Desorption curve for 2 mM TX-100, and predicted behaviour from Frumkin model.

A similar set of curves is shown in figure 4.12b for the desorption of surfactant from the surface when 2 mM TX-100 is replaced by pure water. This comparison shows that the desorption rate is mass-transported limited throughout the whole desorption process. We note that this conclusion cannot be drawn directly from the desorption data in figure 4.11, where the measured desorption rate is limited by the sampling speed of the Raman spectrometer.

The modelled subsurface concentration shows interesting behaviour as a result of the slow initial desorption followed by rapid desorption once the surface reaches the two phase region. The delay before rapid desorption begins is sufficient that the subsurface concentration drops well below the step in the isotherm (around 0.25 mM). When the surface enters the two-phase region and desorption becomes rapid, the loss from the surface is great enough that the subsurface concentration temporarily increases. The process is shown in figure 4.13 for the desorption of a 0.27 mM TX-100 solution (the choice of concentration does not match any experimental measurements, but is chosen for ease of display since there are no micelles present). Figure 4.13b shows concentration profiles at various times during the process: note the increase in $\frac{dc}{dz}$ at the surface when desorption is fastest. There is no physical reason why an overshoot in concentration should not happen, since such overshoots are common in non-equilibrium systems, for example nucleation in supercritical solutions (the process seen here is analogous to the opposite of a nucleation process). However, at the moment it is unclear whether the overshoot modelled is a real process or an artefact of the model; confocal Raman spectroscopy focused around $\sim 10 \mu\text{m}$ from the surface might be able to detect the subsurface concentration if the time resolution was good enough, but would probably also be dominated by surfactant at the surface.

For TX-100, the Frumkin model gives very different values for the rate constants of adsorption and desorption in the low and high coverage phases. The non-dimensionalised adsorption rate constant decreases from ~ 7 in the low coverage phase to around 0.3 in the high coverage phase, while the desorption rate constant varies even more, from >200 in the low coverage phase to <0.1 in the high coverage phase. Thus adsorption to and desorption from the low coverage phase are fast and mass transport limited while in the high coverage phase both processes are kinetically lim-




Figure 4.12: a) Kinetically limited rates of adsorption from the Frumkin model using the preferred fitting parameters and a subsurface concentration equivalent to the cmc. The short-dashed curve (blue) does not account for the phase separation while the solid curve (red) does. The long-dashed curve (green) shows modelled flux during the adsorption of a 10 mM TX-100 solution under experimental conditions; initial rate is limited due to the need for surfactant to diffuse to the surface although later in the process the modelled rate is close to the maximum rate. b) As in a), but showing the kinetically limited rate of desorption when the subsurface concentration is 0. The long-dashed curve (green) shows the modelled flux for the desorption of 2 mM Triton under experimental conditions.




Figure 4.13: (a) Modelled desorption of a 0.27 mM (i.e. just at the cmc) TX-100 solution, showing how surface excess (solid red, left axis) and subsurface concentration (dashed green, right axis) change with time. (b) Concentration profiles with respect to distance from the surface at various times during the process shown in (a).

ited. Our model implicitly assumes that the two phases equilibrate rapidly on the surface so the kinetically slow high-coverage phase can desorb rapidly via the labile low-coverage phase once the surface enters the two-phase region. The observed behaviour is consistent with this interpretation: once desorption starts it is rapid. We note, however, that we have neglected micellar processes in our model, which may provide alternative mechanisms that explain the experimental data.

4.3.3 Comparison with other work

The kinetic parameters determined here can be compared with those in the literature for similar surfactants. Curwen *et al.* used the same kinetic model (though in a cell with different hydrodynamics) to look at the adsorption of the cationic surfactant cetyl pyridinium chloride (CPC) onto a silica surface.⁴⁹ In the absence of added electrolyte, the two parameters describing the isotherm on silica (K_L and ω) are comparable for CPC and CTAB. However, the kinetic parameters differ significantly: k_a and k_d are both $\sim 100\times$ lower for CPC than CTAB, while the values of ω_a and ω_d are roughly 5 greater for CPC than they are for CTAB. When the combined effect of the two terms is taken into account the difference in rates is less dramatic: at 50% coverage, $k_a e^{(0.5\omega_a)}$ gives an effective adsorption rate constant of $0.34 \text{ m}^3 \text{ mol}^{-1} \text{ s}^{-1}$ for CPC and $0.66 \text{ m}^3 \text{ mol}^{-1} \text{ s}^{-1}$ for CTAB. Similarly, $k_d e^{(0.5\omega_d)}$ gives an effective desorption rate constant of 0.20 s^{-1} for CPC and 0.41 s^{-1} for CTAB. The key difference between the two surfactants is not so much the rate as the shape of the adsorption and desorption curves with respect to surface coverage.

Atkin *et al.* also looked at the adsorption of CTAB onto silica (in the form of oxidised silicon).⁴⁵ They analysed their data in terms of “sticking ratio”: the ratio of the maximum measured flux to the surface compared to the maximum possible flux, as given by equation 3.33 (see chapter 3). In the absence of KCl, they recorded sticking ratios ranging from almost 0, at low concentrations, to 0.2, at or above the cmc (in the presence of KCl the ratios were much higher). For 1 mM CTAB (just above the cmc), the sticking ratio we record is 0.04; almost 5 times lower. This is consistent with the measured fluxes, which are also higher in their work.

Brinck *et al.* studied the adsorption and desorption of non-ionic surfactants C_nE_m on silica.⁶³ They used a rather different kinetic model which is not easily compared to the model we employed here. In their model, the majority of the material adsorbing is assumed to consist of aggregates, with transport of aggregates onto the surface assumed to be diffusion-controlled. Monomers could also adsorb to the surface through a process of merging with the adsorbed micelles with a rate constant identical to that of merging with micelles in the bulk. In their model, surfactant adsorption is dominated by micelles due to the very low cmcs of the C_nE_m surfactants studied (and hence low monomer concentrations); a situation that does not apply here. The principle advantage of the wall-jet geometry compared to the stirred cuvette used by Brinck *et al.* is that the thickness of the stagnant layer is well-defined and known for the wall jet, whereas for the stirred cuvette the thickness of the stagnant layer must be inferred from the desorption kinetics, which provides no way to deconvolute mass transport effects from the desorption process.

4.4 Conclusions

In this chapter I have shown that TIR Raman scattering is a viable method for the study of surfactant adsorption kinetics at the solid–liquid interface. Table 4.1 compares TIR Raman scattering to some of the other techniques that have been used to study surfactant adsorption kinetics. TIR Raman is competitive with optical reflectometry and ellipsometry in terms of sensitivity but not sampling rate. For the present TIR-Raman study the sampling rates achieved were sufficient under most conditions, though a faster acquisition rate would have been useful for some of the desorption experiments. The key advantage of TIR Raman over reflectometry is the chemical selectivity of the Raman spectroscopy, which is demonstrated in the following chapter. The related spectroscopic technique of ATR-IR has comparable performance in terms of the sampling frequency and noise, but it has additional complications that are not present for Raman: water is very IR-active, forcing the use of either D₂O or deuterated surfactant to avoid overlap with the water peak; much larger sampling areas are required, which presents problems controlling the hydrodynamics over the whole sampled area, and the interpretation of the data is more complex due to the variable penetration depth across the spectrum. The restriction of ATR-IR to infrared-transparent substrates is rather more restrictive than for TIR-Raman which only requires the substrate to be transparent near the wavelength of the excitation laser.

Technique	Sampling frequency / Hz	Typical point-to-point variation (standard deviation)	Comments	Reference
TIR-Raman	0.5	1-5%	Non-fluorescent materials	This work
Ellipsometry / optical reflectometry	Up to 1000	10% at 1kHz; 1% at 1Hz	No chemical selectivity	Curwen <i>et al.</i> ^{49,111}
ATR-IR	Up to 0.5	1% at 1 Hz	Large sampling areas are needed for short time resolution	Clark and Ducker ⁴⁷
QCM	Up to 200	0.3% at 0.2 Hz; 1% at 1 Hz	Water trapped in the adsorbed surfactant layer complicates interpretation of data; no chemical selectivity.	Q-Sense ²⁵²

Table 4.1: Comparison of the limitations of experimental techniques used to study surfactant adsorption kinetics.

For both CTAB and TX-100, fitting to the Frumkin isotherm yields a positive interaction parameter, ω , indicating a strong interaction between neighbouring molecules on the surface. For both surfactants the interaction parameters (ω_a and ω_d) in the kinetic expressions are negative, showing that the rate constants of adsorption and desorption decrease with increasing surface coverage. The negative value of ω_d is consistent with the favourable interchain interactions, while the negative value of ω_a may arise from the steric hindrance to adsorption at high surface coverage. The non-dimensionalised values of the rate constants indicate whether adsorption is under kinetic or diffusion control. The large and negative values of ω_a and ω_d result in a strong coverage dependence in the kinetic constants and a change from diffusion control at low surface coverages to a kinetically controlled or mixed regime at high coverages.

The methodology we present here for studying adsorption kinetics is general and robust. The kinetic parameters are determined from the equilibrium isotherm and the adsorption kinetics alone and the resulting desorption profiles are then predicted with no free parameters. The predicted and measured desorption profiles are in good agreement except in regions of low coverage of CTAB where the Frumkin isotherm fails to represent the adsorption isotherm of CTAB adequately.

Chapter 5

CTAB/Triton X-100 multicomponent systems

In the interests of readability, only an illustrative selection of the experiments relating to this work are shown here. The remainder are shown in appendix C.

5.1 Introduction

In chapter 4, I demonstrated the sensitivity and time resolution achievable through TIR-Raman spectroscopy. The cationic surfactant cetyltrimethylammonium bromide (CTAB) and the non-ionic surfactant Triton X-100 (TX-100) were used as exemplars. A wall-jet cell provided controlled hydrodynamics and allowed quantitative modelling of the adsorption kinetics. In this chapter, I show how TIR Raman spectroscopy can be used to follow the interfacial kinetics of two surfactants in a mixed system: CTAB + TX-100. Spectra of binary mixed layers can be decomposed into the two component spectra without the use of selective deuteration and with acquisition times of only 1 second.

Most techniques that can determine the chemical composition of adsorbed surfactant layers are only suitable for equilibrium measurements or the study of slow kinetics on time-scales of minutes to hours. Examples are depletion studies, in which a powder is suspended in a surfactant solution and the surfactant remaining in the solution is measured either in situ or following removal of an aliquot,^{253–255} and neutron reflection,^{256,257} in which selective deuteration allows surfactants to be distinguished. Conversely, techniques that respond rapidly to sub-monolayer changes in the mass of an adsorbed film—such as reflectometry, surface plasmon resonance or quartz crystal microbalance measurements—lack the chemical selectivity to determine the composition of the adsorbed layer. Where one species is irreversibly adsorbed—as is often the case in polymer-surfactant mixtures—the composition of the mixed layer can be inferred from the rinsing behaviour in pure solvent, on the assumption that only the surfactant desorbs. Velegol and Tilton used this approach to study the adsorption of mixtures of CTAB and poly-lysine onto silica²⁵⁸ while Postmus *et al.* studied the adsorption of polyethylene oxide mixed with a variety of different ethylene oxide alkyl

ethers (C_nE_m).⁶² For surfactant-surfactant mixtures it is much more difficult to distinguish the two components. Tiberg and coworkers investigated mixtures of different C_nE_m surfactants by optical reflectometry^{60,64} but needed to assume ideal mixing in order to analyse their data quantitatively.

Spectroscopy is a standard way to distinguish chemical species but most common surfactants do not absorb light in the visible or near-UV, nor do they fluoresce, so UV-vis absorption spectroscopy and fluorescence detection are not generally applicable. All surfactants have molecular vibrations but IR absorption cross-sections are typically three orders of magnitude weaker than those for allowed electronic transitions while Raman scattering cross-sections are around 13 orders of magnitude weaker than fluorescence cross-sections. Consequently, vibrational spectroscopy of nanometre-thick layers of mixed surfactants at interfaces is experimentally challenging. Nevertheless, there are a few cases where attenuated total internal reflection (ATR) IR spectroscopy has been used successfully to monitor adsorption kinetics of mixed surfactant systems. Couzis and Gulari looked at the adsorption of sodium dodecyl sulfate (SDS) and sodium dodecanoate onto Al_2O_3 particles affixed to a germanium ATR crystal.²⁴⁴ The adsorption processes were slow (hours to days, depending on the solution), probably due to a combination of the large surface area of the particles and the slow exchange of solution (the total volume of the cell was exchanged over the course of an hour). Displacement kinetics were determined from the size of the S=O stretching bands of SDS at 1060 and 1200 cm^{-1} although amounts of the two surfactants could not be established quantitatively. An interesting use of ATR-IR was the study by Clark and Ducker of exchange kinetics in adsorbed films of tetradecyl trimethylammonium bromide ($C_{14}TAB$)⁴⁷ and didodecyldimethylammonium bromide⁵¹ on silica. In the first case, they replaced a solution of hydrogenated $C_{14}TAB$ with deuterated $C_{14}TAB$ and measured the rate at which the C-H bonds were replaced by C-D bonds. A multiple reflection Si substrate yielded sufficient signal to take a spectrum every two seconds, but also raised the problem of rapidly exchanging the solution in contact with a large substrate. Clark and Ducker solved this problem ingeniously by inserting an air bubble into the cell between the incoming and outgoing solution, which efficiently displaced the initial solution. Li and Tripp studied the replacement of CTAB on TiO_2 particles by deuterated SDS.²⁵⁹ The exchange process took place over the course of ~ 1 hr and was followed with a time resolution of ca. 5 minutes per spectrum. Shifts in the frequency of the C-H bands were attributed to changes in the packing of surface aggregates and the relative amounts of the two surfactants were determined by integration of peak areas in the C-H and C-D stretching regions. Tabor *et al.* measured fast surfactant adsorption (~ 20 s) with slower surfactant desorption (~ 200 s) for non-ionic surfactant dissolved in a toluene solution.¹⁵⁶ The use of an organic solvent removed the strong IR water peak, but a multiple bounce ATR crystal was still necessary to obtain sufficient signal. Recent work using ATR-IR has looked at the competitive adsorption of chiral adsorbates to chiral surfaces using either deuteration of one enantiomer to distinguish the two²⁶⁰ or modulation excitation spectroscopy, a chemometric technique.^{261,262} Previous work in the group has shown how external reflection FTIR spectroscopy can be used to study the adsorption kinetics of surfactant mixtures at the air-water interface without the


need for selective deuteration of surfactants.²⁶³ The use of a continually expanding liquid surface allowed measurements at surface ages around 0.1 s.

There are no previous reports of the use of Raman scattering to study the kinetics of adsorption of mixed surfactant systems.


The system we have chosen to study here is a binary mixture of a cationic and non-ionic surfactant. No kinetic studies have been carried out on such mixtures and only a limited amount of work has been published on equilibrium adsorption properties, which we summarize briefly here. Huang and Gu looked at the adsorption of a mixture of CTAB and TX-100 onto a silica gel,²⁶⁴ finding that for the individual components ~ 5 times more TX-100 adsorbed than CTAB. TX-100 exhibited a smooth increase to the limiting surface excess while the CTAB showed a plateau in adsorption at both low and high concentration. In the mixed system, adsorption of CTAB was promoted at low total surfactant concentrations; adsorption of both surfactants was inhibited at high total concentrations with the limiting surface excess similar to that of pure CTAB. McDermott *et al.* studied the adsorption of CTAB + C₁₂E₆ on quartz by neutron reflection.²⁶⁵ The CTAB/C₁₂E₆ system is chemically similar to CTAB/TX-100 although the cmc of C₁₂E₆, and hence of the mixed systems, is much lower than for TX-100. They found that small amounts of C₁₂E₆ ($\sim 8\%$ mole fraction) caused a slight decrease in the adsorption of CTAB on quartz whereas for roughly equimolar mixtures the total surface excess was close to that of pure C₁₂E₆. In a subsequent NR study with oxidised silicon in place of quartz, Penfold *et al.* found that for a 0.1 mM equimolar mixture at pH 7 the adsorbed film was strongly enriched in CTAB compared to the bulk.²⁵⁶ Near the isoelectric point of the silica (pH 2.4), mixtures of the two surfactants showed a $\sim 50\%$ increase in adsorption compared to the pure components.²⁵⁷ At low mole fractions of CTAB, the cationic surfactant was enriched in the surface layer while at other concentrations the surface and bulk compositions were approximately equal. Soboleva *et al.* studied the adsorption of a mixture of C₁₄TAB and TX-100 onto quartz sand. They found a small enhancement of the minor component at the surface, but the surface composition was not especially different from the bulk composition.²⁶⁶

This chapter first explores the equilibrium isotherm of CTAB and TX-100 on a planar silica substrate at a total concentration of 2 mM—well above the critical micelle concentration (cmc) of both the pure and mixed systems. This total concentration is used as our benchmark for kinetic measurements. The benchmark compositions were 25%, 50% and 75% CTAB by mole. I also studied the equimolar mixtures at 1 mM and 3 mM concentrations to determine the sensitivity of the adsorption behaviour to total surfactant concentration. I report results from five different scenarios: (i) adsorption of mixed surfactant solutions to bare silica, (ii) desorption of mixed surfactant layers into pure water, (iii) displacement of one pure surfactant by another pure surfactant, (iv) displacement of a mixed surfactant layer by a pure surfactant and (v) the displacement of a pure surfactant layer by a mixture.

Equilibrium measurements on the adsorbed film show strong deviations from ideality and the bulk solution has also been reported to be non-ideal.²⁶⁷ A quantitative kinetic model, similar to that presented for the pure surfactants in chapter 4, would require a full characterization of the mixed surfactant system both in the bulk and



(a)



(b)

Figure 5.1: Illustration of the reproducibility of the data for (a) a desorption process (shown with explanation in figure 5.6(c)) and (b) a replacement process (shown with explanation in figure 5.14(b)).

at the surface at all concentrations and compositions that are encountered during the adsorption or desorption processes. Such a task is a substantial undertaking and we have not attempted it here. The surface measurements alone show a number of unexpected features that can, at a qualitative level, be understood from the non-ideal mixing of the CTAB and TX-100 at the surface.

5.2 Experimental

The reproducibility of individual experiments is illustrated by two typical examples in figure 5.1. Most of the experimental conditions were repeated at least twice (including all systems that showed non-monotonic behaviour), but for clarity only a single example is displayed in the remainder of this chapter for each set of conditions.

A complete kinetic model of adsorption and desorption in CTAB/TX-100 mixtures would be a major undertaking of limited generality and which I have not attempted to

develop. Nevertheless, it is useful to be able to make some quantitative comparison of the rates of adsorption (or desorption) under different experimental conditions. Consequently, we have evaluated the rates of adsorption or desorption between the points when the change is 30% complete and 70% complete, interpolating linearly between data points to establish these two thresholds where necessary. Some of the adsorption curves display an overshoot where the surface excess rises above its equilibrium value; in these cases the change is measured with respect to the highest surface excess. Similarly, some of the desorption processes display a spike before the surface excess begins to fall; when evaluating the rates the change is measured from the top of the spike. Tables showing the maximum rates of flux to and from the surface are given throughout the text. The rate given is a simple average over independent experiments. The variation in the rates observed in repeated experiments on mixtures under nominally identical experimental conditions is typically 20–30% although the qualitative behaviour remains the same. For systems showing the fastest kinetics, the spacing of the data points (every 2 s) limits the accuracy with which the adsorption and desorption rates can be measured. Within the tables, estimated errors are given for those measurements where multiple repeats were recorded based on the standard deviation of the measured rates in the repeats. These values should serve as an indication of the size of the error where only one repeat was recorded (for which no errors are shown).

5.2.1 Modelling

I have attempted to model some of the processes studied in this chapter using the finite-difference mass transport model described in chapter 3. I have used the diffusion coefficients described in chapter 4, and assumed that the diffusion coefficient of mixed micelles is constant with respect to composition. The composition of mixed micelles is described using regular solution theory, with an interaction parameter, β , of -1.0 , from the work of Ruiz and Aguiar.²⁶⁷ Modelling in this chapter is purely concerned with mass transport to and from the surface: I have completely neglected adsorption to the surface. A no-flux boundary condition at the surface ensures that no surfactant can adsorb or desorb. Therefore the effect of the surface acting as a source and sink for monomers will not be seen in this model; I expect that this effect would act to slow or lessen any changes in monomer concentration near the surface (with a typical surface excess of $\sim 4 \mu\text{mol m}^{-2}$, and a bulk concentration of 2 mM, the adsorbed layer accounts for as much surfactant as 2 μm depth of bulk solution). It is still possible to infer trends in the adsorbed amount, since adsorption will increase with the chemical potential of the components (proportional to the log of the monomer concentration). Therefore, the monomer concentration at the subsurface provides a rough guide as to the adsorbed amount. As an example, to model a “desorption” process I would start with a system containing a uniform concentration of (mixed) surfactant and together with a boundary condition of no surfactant at the inlet, and follow the change in concentration with time. The decision not to model adsorption at the surface was taken due to the complexity of the system: even modelling the equilibrium isotherm with respect to composition would be difficult, since the two surfactants have different




Figure 5.2: Surface excess of CTAB (squares) and TX-100 (circles) at a total surfactant concentration of 2 mM as a function of the bulk mole fraction of CTAB, χ_{CTAB} , together with the total surface excess (solid line). The isotherms for the pure components at the same concentrations as that component in the mixture are indicated with dotted lines.

surface areas per molecule and in any case the measured equilibrium isotherm could not be described using a regular solution model. However, as I will show, many of the kinetic effects seen are controlled by mass transport and can be explained by the more restricted model used here.

5.3 Results

5.3.1 Equilibrium

Chapter 4 series described the adsorption isotherms for the individual surfactants, CTAB and TX-100, as a function of concentration. Figure 5.2 presents the surface excess and surface composition for a mixture of CTAB and TX-100 mixture at a constant 2 mM concentration as a function of the mole fraction of CTAB in the mixture.

The most striking aspect of isotherm in figure 5.2 is the increase in TX-100 adsorption when tiny amounts of CTAB are added. Even at bulk mole fraction, $\chi_{\text{CTAB}} = 0.02$, the TX-100 surface excess was double that of the pure component (As a side note, figure C.1 in appendix C shows a concentration isotherm at a constant $\chi_{\text{CTAB}} = 0.05$ composition). The TX-100 surface excess then decreased almost linearly with increasing χ_{CTAB} . In contrast the CTAB surface excess showed a plateau at $\chi_{\text{CTAB}} < 0.6$ before rising smoothly as $\chi_{\text{CTAB}} \rightarrow 1$. Consequently, the total adsorbed amount (illustrated by the solid line in figure 5.2) is a maximum at very low χ_{CTAB} and then decreases monotonically as the CTAB concentration in the mixture increases.

Figure 5.2 also compares the surface excesses in the mixture to those of the pure surfactants at the same bulk concentration. At low χ_{CTAB} there is a synergistic interaction with enhanced adsorption of both TX-100 and CTAB. For $\chi_{\text{CTAB}} > 0.3$, competitive adsorption reduces the surface excess below that of the pure component at the same concentration.




Figure 5.3: A possible interpretation of the dramatic increase in surface excess seen when a small amount of CTAB (red) is added to an adsorbed TX-100 (green) layer. The left-hand picture shows a pure TX-100 layer, while the right-hand side shows the mixed layer. Only a single surface aggregate is shown for clarity.

The interactions between CTAB and TX-100 in mixed micelles are known to be favourable,⁷⁹ but a favourable interaction parameter alone does not explain the complex adsorption isotherm. The dramatic increase in the total surface excess at low χ_{CTAB} suggests that small amounts of adsorbed CTAB in TX-100-rich layers alters the structure of the surface aggregates to allow more efficient coverage of the surface. Conversely, in CTAB-rich layers, the total surface excess is very similar to that in saturated layer of CTAB, suggesting that the two surfactants are competing for sites within aggregates of a similar structure. The doubling of surface excess upon the addition of a small amount of CTAB suggests a transition from a hemimicelle to a full micelle surface structure (or alternatively a monolayer to a bilayer structure). In this tentative explanation, the portion of the micelle in contact with the surface is primarily CTAB, bound electrostatically. The remainder to the micelle is probably mainly TX-100, but the fact that more CTAB is adsorbed in the mixture than in the initial, electrostatically bound, region of the pure CTAB isotherm (see the light blue ‘—’ line of figure 5.2) suggests that some CTAB is incorporated in the remainder of the micelle. Figure 5.3 provides a “cartoon” explanation of the processes described above. This explanation is not entirely consistent with results from small angle neutron scattering that suggest that TX-100²⁶⁸ and other nonionic surfactants²⁶⁹ tend to adsorb to silica as spherical rather than hemispherical aggregates. Further experiments which might shed light on the structure include varying the surface charge on the silica—by adjusting the pH—to promote or hinder the formation of the anchoring CTAB layer; sum-frequency spectroscopy, which is explicitly sensitive to asymmetry in surface structure; and either AFM or surface force measurements to try to determine a layer thickness.

The linear decrease in adsorbed TX-100 with χ_{CTAB} suggests a 1:1 replacement of TX-100 with CTAB as the composition varies. However, the CTAB surface excess does not increase linearly. It is especially surprising that $d\Gamma_{\text{CTAB}}/d\chi_{\text{CTAB}}$ increases at the CTAB-rich end of the isotherm, where one might expect electrostatic repulsion between the charged CTAB molecules to be limiting adsorption. I attribute this increase to an improvement in packing efficiency when TX-100 is removed from the




Figure 5.4: Adsorption of a mixture of CTAB and TX-100 from a 2 mM solution with $\chi_{\text{CTAB}} = 0.5$ onto clean silica. Adsorbed amounts: CTAB (+, red), TX-100 (\times , green), total (*, blue).

surface aggregates: TX-100 has a shorter chain than CTAB as well as an aromatic group and so hydrophobic interactions are likely to be strongest in a homogeneous CTAB aggregate.

5.3.2 Kinetics of adsorption and desorption


In the Experimental Section I have shown how TIR-Raman is able to quantitatively determine the composition of mixed monolayers of CTAB and TX-100 with 1-s acquisition times without the need for deuterated surfactants. Consequently, the rates of adsorption and desorption of the two individual surfactants can be measured separately and not just the rate of change of the total adsorbed amount. Figure 5.4 shows the adsorption kinetics for a 1:1 mixture at 2 mM total concentration. Both species show a smooth increase in surface coverage with time—the slight overshoot in the TX-100 surface excess is probably a consequence of the polydispersity of the surfactant rather than interactions with the CTAB, although it might also arise from mass transport arguments discussed later.

CTAB has a higher cmc than TX-100 (0.9 mM compared with 0.27 mM) so the monomer surfactant is enriched in CTAB compared to the mixture; conversely the micelles are rich in TX-100. Monomers generally diffuse faster than micelles and CTAB diffuses faster than TX-100 ($D_{\text{mon}} = 8.8 \times 10^{-10}$ compared with $2.8 \times 10^{-10} \text{ m}^2 \text{ s}^{-1}$), so one might expect CTAB to show faster adsorption kinetics than TX-100. The fact that we observe the reverse shows that adsorption is not under diffusion control.

Table 5.1 summarizes the adsorption rates (as defined in the Experimental section) for adsorption of various mixtures onto clean silica. The rate of adsorption of TX-100 in mixtures was generally faster than for pure TX-100. One plausible explanation is the favourable hydrophobic interaction of TX-100 with small amounts of CTAB electrostatically bound to the negatively charged silica surface. A second contributing factor may be the increased diffusion coefficient of non-ionic micelles in the presence of a cationic surfactant due to migration fields.

Total surfactant concentration / mM	χ_{CTAB}	CTAB rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$	TX-100 rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$
2.0	0.25	0.08 ± 0.03	0.6 ± 0.2
2.0	0.50	0.10	1.3
2.0	0.75	0.2 ± 0.1	0.3 ± 0.1
2.0	1	0.7	
2.0	0		0.3
3.0	0.5	0.20	0.7

Table 5.1: Rates of adsorption for CTAB/TX-100 mixtures onto a clean silica surface.

Figure 5.5: Modelled subsurface concentrations for the various components of the mixture during the adsorption of a 2 mM solution with $\chi_{\text{CTAB}} = 0.5$.

As discussed earlier, the mass transport of mixed surfactant to the surface during the adsorption process can be modelled (the actual adsorption at the surface is neglected, however). Figure 5.5 shows the modelled time dependence of surfactant concentration at the subsurface as surfactant is injected into the cell. The behaviour appears quite complex, showing both a small overshoot in TX-100 monomer concentration and a dip in CTAB monomer concentration as the total concentration reaches the cmc. We would expect the surface excess to broadly follow the features of the monomer concentrations. Both features are short-lived and could easily be missed with 2 s time resolution (assuming the surface kinetics respond fast enough). The slight overshoot in TX-100 adsorption appears to match the experimental results (figure 5.4), although it is not reproducible enough or large enough to assign definitively; no dip in CTAB adsorption is seen experimentally. The latter is probably unsurprising since CTAB adsorption is less sensitive to bulk conditions than TX-100 adsorption. In brief the model shows ($t \approx 1$ s) CTAB monomers arriving to the surface faster than TX-100 monomers due to the higher diffusion coefficient of CTAB monomers; ($t \approx 2.5$ s) TX-100 monomers arriving at the surface; ($t \approx 6$ s) mixed micelles arrive at the surface enriched in TX-100 (due to the lower cmc of TX-100); almost simultaneously, TX-100 monomers are enriched and CTAB monomers depleted through exchange with the micelles; and ($t \approx 8$ s onwards) the partitioning within the system tends increasingly towards the final equilibrium state as the concentration of both surfactants increases.

The desorption process is more unusual. Figure 5.6 shows experimental data for desorption into pure water of surfactant layers adsorbed from surfactant solutions with 2 mM total concentration and $\chi_{\text{CTAB}} = 0.25, 0.5$ and 0.75 . The CTAB surface excess falls monotonically during the rinsing process but for $\chi_{\text{CTAB}} = 0.5$ and 0.75 (figure 5.6(b) and (c)) the TX-100 surface excess actually increases sharply at the beginning of the desorption process, peaks after about 10 s and then decays monotonically to zero. In the early part of the desorption process the surface excess remains approximately constant, with desorbing CTAB molecules being replaced by TX-100 adsorbing from the sub-surface layer. For the TX-100-rich system shown in figure 5.6(a), the initial surface excess of TX-100 is higher and the peak in the surface excess is less marked; nevertheless, the desorption of TX-100 is delayed compared to CTAB so the surface concentration is enriched in TX-100 during the early part of the desorption process.

Qualitatively similar behaviour was observed for the 3 mM total concentration (figure 5.7), though the increase in TX-100 surface excess was less marked than for the 2 mM case.

This unusual desorption behaviour can be understood from the adsorption isotherm (figure 5.2), where a decrease in the surface excess of CTAB leads to a marked increase in the amount of TX-100 adsorbed. The equilibrium mixing behaviour is augmented by kinetic factors: CTAB diffuses away from the surface faster than TX-100, for the reasons given earlier, so the local surfactant composition becomes enriched in TX-100 as the total surfactant concentration decreases. So long as the sub-surface concentration remains above the cmc (and assuming the micelles and monomers equilibrate rapidly), an increase in the mole fraction of TX-100 is accompanied by an




(a) $\chi_{\text{CTAB}}=0.25$ (b) $\chi_{\text{CTAB}}=0.5$ (c) $\chi_{\text{CTAB}}=0.75$

Figure 5.6: Desorption of mixed layers with an initial solution concentration of 2 mM at various compositions. Pure water was injected into the cell containing the surfactant mixture. Adsorbed amounts: CTAB (+, red), TX-100 (\times , green), total (*, blue).




Figure 5.7: Desorption into pure water of a mixed layers with an initial solution concentration of 3 mM and $\chi_{\text{CTAB}} = 0.5$. Adsorbed amounts: CTAB (+, red), TX-100 (\times , green), total (*, blue).

increase in its chemical potential and hence by an increased driving force to adsorb. Only when the subsurface concentration drops below the cmc, does the chemical potential of the TX-100 begin to fall and the TX-100 desorb from the surface.

Figure 5.8 shows the modelled concentration of the two surfactants at the subsurface over time during the rinsing process (as before, only mass transport is considered and surfactant adsorbed at the surface is neglected). The TX-100 monomer concentration rises as more solution is removed, as described in the previous paragraph. The increase in monomer concentration (and hence chemical potential) explains the spikes in TX-100 surface excess seen experimentally in figure 5.6. Figure 5.9 shows the concentration profiles in the vicinity of the surface at various times, illustrating how a band of increased TX-100 monomer concentration is swept towards the surface around the point where the concentration goes below the cmc.

A related effect was reported by Brinck *et al.*⁶⁴ for the desorption of mixtures of the non-ionic surfactants C_{14}E_6 and C_{10}E_6 . They found an increase in total surface excess upon desorption though their measurement technique did not permit the determination of the surface excesses of the individual surfactants. They noted that C_{10}E_6 is transported away from the surface more rapidly than C_{14}E_6 because it has a much higher cmc and thus a higher monomer concentration. As the C_{10}E_6 near the surface is depleted, the micelles become enriched in C_{14}E_6 and consequently the concentration of monomers of C_{14}E_6 in equilibrium with the micelles increases. Since the equilibrium surface excess of C_{14}E_6 is higher than for the mixed system, the total adsorbed amount initially increases during the desorption process.

5.3.3 Kinetics of displacement of surfactant layers

The previous section described adsorption onto bare silica or desorption into pure water. Here we consider the displacement of one surfactant layer by a surfactant solution of different composition. The observed behaviour is best classified according to the surfactant being injected, since the results seem to be grouped largely on these




Figure 5.8: Modelling the concentration of CTAB and TX-100 at the subsurface for a 2 mM solution with $\chi_{\text{CTAB}} = 0.5$ desorbing. Model parameters are discussed in the text.

lines.

The simplest behaviour is seen in systems where pure CTAB replaces an adsorbed layer (either pure TX-100 or a mixture). The replacement of TX-100 by CTAB, both at 2 mM concentration, is shown in figure 5.10. The equilibrium isotherm shows that a small amount of CTAB nearly doubles the equilibrium adsorbed amount of TX-100. It is not surprising therefore to observe in figure 5.10 that replacement of the TX-100 solution by a CTAB solution initially (though only briefly) increases the adsorbed amount of TX-100. The initial subsurface concentration of TX-100 is 7 times the cmc, so one might expect a significant delay before TX-100 is desorbed from the surface. The CTAB diffusing towards the surface, however, forms mixed micelles with the TX-100 which immediately lowers the chemical potential of the TX-100, driving desorption into solution. The initial rate of adsorption of CTAB to a TX-100 covered surface is initially similar to that on bare silica (figure 5.10 and table 5.2)—a surprising result from the Langmuir perspective that would predict an adsorption rate proportional to the fraction of unfilled sites $(1 - \theta)$. From the Frumkin perspective, however, the $(1 - \theta)$ reduction is offset by an increase in the adsorption rate constant $k_a e^{\omega_a \Gamma / \Gamma_\infty}$ (see chapters 3 and 4) owing to attractive interactions between hydrocarbon chains: it appears that these two effects are roughly in balance. Once the surface concentration of CTAB reaches that of TX-100, the rates of adsorption of CTAB and desorption of TX-100 both decrease markedly. In particular, desorption of TX-100 into a CTAB solution is an order of magnitude slower than into pure water (see chapter 4, figure 4.11). This reduction in rate can be ascribed, at least in part, to the interactions between the two surfactants, which favour a mixed layer relative to a pure layer of either surfactant.

Modelling the mass transport of 2 mM CTAB replacing 2 mM TX-100—adsorption to and desorption from the surface are neglected—shows no unusual behaviour in the monomer concentration (see figure 5.11), which is consistent with the experimental results (figure 5.10). Note that the CTAB in micelles curve goes through a maximum.




Figure 5.9: Concentration profiles respect to distance from the surface at a variety of different times (labelled on each graph).




Figure 5.10: Replacement of 2 mM TX-100 (green \times) with 2 mM CTAB (red $+$). The total surface excess is shown as blue stars. The adsorption of 2 mM CTAB onto a clean silica surface (line) is shown for comparison.

Conc. / mM	System being replaced	χ_{CTAB}	System being injected	CTAB rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$	TX-100 rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$
1	TX-100	—	CTAB	0.14	-0.19
2	TX-100	—	CTAB	0.17	-0.20
3	TX-100	—	CTAB	0.14 ± 0.01	-0.23 ± 0.04
1	mixture	0.5	CTAB	0.09	-0.06
2	mixture	0.25	CTAB	0.14	-0.18
2	mixture	0.5	CTAB	0.18	-0.21
2	mixture	0.75	CTAB	0.10	-0.05
3	mixture	0.5	CTAB	0.15 ± 0.04	-0.10 ± 0.02
2	Water	—	CTAB	0.7	
2	TX-100	—	Water		-1.1

Table 5.2: Rates of replacement of TX-100 or mixtures of TX-100 and CTAB with an equal concentration of CTAB. Comparative rates for the pure surfactant systems are also shown. Negative numbers indicate desorption.




Figure 5.11: Top graph: modelled concentration at the subsurface as 2 mM CTAB replaces 2 mM TX-100. Bottom graph: concentration at the surface of 2 mM TX-100 being rinsed away and 2 mM CTAB being injected, modelled independently (i.e. not in a mixed system, but shown on the same graph for convenience).




Figure 5.12: Modelled concentration profile with respect to distance from the surface during the replacement of 2 mM TX-100 with 2 mM CTAB, 2 s into the process.





Figure 5.13: Injection of a 2 mM solution of CTAB into a 2 mM surfactant mixture with $\chi_{\text{CTAB}} = 0.5$. Adsorbed amounts: CTAB (+, red), TX-100 (\times , green), total (*, blue).


Away from the surface where only small concentrations of the outgoing TX-100 solution remain, the mixed micelles become enriched in TX-100, displacing CTAB from the micelles. However, at the point closer to the surface where the CTAB concentration first begins to rise, the micelles are enriched with CTAB, causing a local maximum in the concentration of CTAB in micelles with respect to the distance from the surface. A concentration profile is shown in figure 5.12, illustrating these effects. As the local maximum reaches the surface it causes a temporary increase in the concentration of CTAB in micelles with respect to time. Figure 5.11 also shows the subsurface concentrations of both CTAB and TX-100 due to mass transport in the absence of the other surfactant (again, neglecting adsorption at the surface). In both cases the processes are fast compared to the mixed system, which agrees with the increased rate of CTAB adsorption seen experimentally and shown in figure 5.10.

For the case when CTAB replaces a mixed layer, there is no transient increase in the TX-100 concentration, but only a smooth replacement of TX-100 by CTAB. The qualitative behaviour is insensitive to either concentration or composition. An illustrative example in which a 1:1 mixture of TX-100 and CTAB is replaced by pure CTAB is shown in figure 5.13. The total adsorbed amount remains approximately constant during the displacement process (the surface remains fully covered with surfactant) and consequently the adsorption and desorption rates are nearly equal (see table 5.2). The rates of exchange of CTAB for TX-100 are typically a factor of 5 slower than for adsorption of CTAB to a bare surface or desorption of TX-100 into water. We note that only a chemically selective technique such as TIR-Raman can reveal the kinetics of exchange in this system.


The reverse situation, in which TX-100 replaces an adsorbed layer of CTAB, shows quite unusual behaviour. Figure 5.14 shows the kinetic traces when the initial CTAB concentration and final TX-100 concentration are 1 mM, 2 mM or 3 mM. At first glance the curves seem to show extreme concentration dependence, but closer inspection shows that in all cases the total concentration (blue stars in figure 5.14) first



(a) 1 mM



(b) 2 mM



(c) 3 mM

Figure 5.14: Replacement of CTAB (red +) by TX-100 (green ×) at three different concentrations. The total surface excess is shown as blue stars. Note that temporal axis in (b) is different from (a) and (c), because of the longer time required for the surface to reach equilibrium.




Figure 5.15: Modelled subsurface concentration during the replacement of 2 mM CTAB with 2 mM TX-100.

drops sharply (though only for 2 s in the 2 mM example) then increases steeply—overshooting the equilibrium adsorbed amount—and finally decays slowly to equilibrium. Variations in the relative timing of the desorption and adsorption (arising from the complex dependence of the mass transport rates and monomer concentrations on the total concentration of surfactant as well as the composition) determine the extent of the dip in total coverage at short times.

Turning to the individual components, the CTAB desorption is similar in all three plots—a steep initial decline followed by a long tail reflecting the plateau in the CTAB adsorption isotherm at low concentrations (figure 5.2). The TX-100 concentration profiles at 2 mM and 3 mM (figure 5.14(b) and (c)) overshoot the final surface excesses before slowly relaxing to equilibrium. This overshoot is most pronounced in the 2 mM case and it takes more than ten minutes for equilibrium to be reached. The overshoot in the adsorption of the non-ionic surfactant can be understood from the equilibrium isotherm in figure 5.2: the incorporation of even a trace of CTAB in the adsorbed layer greatly increases the surface excess of TX-100. Not until the last remnants of CTAB are rinsed from the surface does the TX-100 reach its final coverage. For reasons that are not clear, the 1 mM concentration, while still showing a slow approach to equilibrium, does not show an overshoot in the TX-100 concentration.

The modelled subsurface concentration (based on the modelling of mass transport of surfactant close to the surface) as TX-100 replaces CTAB is shown in figure 5.15. The monomer concentrations—which govern the surface excess—are monotonic, reinforcing my view that the overshoots seen experimentally are primarily a consequence of the equilibrium isotherm. The concentration of CTAB in micelles—but not the total CTAB concentration—goes through a temporary minimum. As the incoming layer of TX-100 moves towards the surface, the small overall mole fraction of TX-100 promotes micelles enriched in TX-100, reducing the amount of CTAB in the micelles. This is the opposite effect to that discussed earlier for CTAB replacing TX-100. The mass transport model is not particularly useful in explaining the long lifetimes (100s of seconds) of the overshoots seen in figure 5.14: according to the model the subsur-




Figure 5.16: Modelled concentration profile with respect to distance from the surface during the replacement of 2 mM CTAB with 2 mM TX-100, 2 s into the process.




Figure 5.17: Replacement of a 2 mM mixed surfactant solution with $\chi_{\text{CTAB}} = 0.5$ by 2 mM TX-100. Adsorbed amounts: CTAB (+, red), TX-100 (\times , green), total ($*$, blue).

face CTAB mole fraction reaches 1% in 23 s, 0.1% in 41 s and 0.01% in 58 s. It is not clear from figure 5.2 what the minimum mole fraction of CTAB needed to cause the increase in TX-100 adsorption, except that it is less than 2%. Comparison with figure 5.11 shows that transport of TX-100 away from the subsurface is actually *slower* than transport of CTAB away from the subsurface, due to the lower cmc (and hence greater incorporation into mixed micelles) of TX-100.

A related set of experiments is the replacement of mixtures of different compositions by pure TX-100. The observed kinetics are similar to the replacement of pure CTAB with TX-100, only starting from a later time where the surface is already of mixed composition. For this reason, the dip in the total adsorbed amount, which occurs in figure 5.14 at low surface excesses of TX-100, is not observed when the starting film already contains a significant amount of TX-100. An example of the displacement of a 2 mM equimolar solution with pure TX-100 is shown in figure 5.17. The

Conc. / mM	System being replaced	χ_{CTAB}	System being injected	CTAB rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$	TX-100 rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$
1	CTAB	—	TX-100	-0.21	0.21
2	CTAB	—	TX-100	-0.27 ± 0.04	0.30 ± 0.05
3	CTAB	—	TX-100	-0.28	0.21
1	mixture	0.5	TX-100	-0.04	0.12
2	mixture	0.25	TX-100	-0.026	0.09
2	mixture	0.5	TX-100	-0.030 ± 0.02	0.19 ± 0.02
2	mixture	0.75	TX-100	-0.042	0.23
3	mixture	0.5	TX-100	-0.04	0.10
2	Water	—	TX-100		0.34
2	CTAB	—	Water	-0.39	

Table 5.3: Rates of adsorption of TX-100 and desorption of CTAB for replacement of either CTAB or a mixed layer by TX-100. Equivalent rates for the individual systems adsorbing and desorbing are shown for comparison.


measured rates are tabulated in table 5.3; they are slower than for the replacement of CTAB by TX-100, especially the rate of desorption. This difference is expected, since the fast initial changes seen in figure 5.14 are bypassed in the mixed initial solutions. CTAB shows a long tail in its desorption, and the processes for mixed systems are starting from a point that is near the start of that tail.

The final permutation on displacement kinetics involves the replacement of a single surfactant by a variety of different mixed surfactant solutions. There are two variations within this category depending on whether the initial layer is composed of CTAB or TX-100.


Where mixtures replace CTAB the results are largely unremarkable: the exchange of components at the surface happens at a rate largely similar to that seen for replacement by pure surfactants (rates are tabulated in table 5.4). The rates of CTAB desorption are generally greater than the rates of TX-100 adsorption. In some cases this disparity in rate leads to a temporary drop in total surface excess as the exchange proceeds. An example is shown in figure 5.18(b) for replacement of a 2-mM CTAB solution by a 2-mM mixture with $\chi_{\text{CTAB}} = 0.5$.

The most interesting interfacial behaviour seems to arise when TX-100 is mixed with just a small amount of CTAB at the surface. We would therefore expect that the replacement of TX-100 by a mixture containing CTAB would result in a sharp increase in the adsorbed amount and this is indeed observed. figure 5.19(b) shows the replacement of a 2 mM TX-100 solution by a 2 mM mixture with $\chi_{\text{CTAB}} = 0.5$. This experiment is the reverse of the experiment shown in figure 5.13. The TX-100 coverage shows a sharp spike as the first CTAB adsorbs to the surface. Since the local concentration of TX-100 is still close to 2 mM, this adsorption occurs very rapidly too rapidly for us to resolve with the current spectrometer. The TX-100 adsorption then decreases quickly towards its equilibrium value as more CTAB competes for space in the adsorbed layer. The total adsorbed amount shows a similar transient spike. Figure 5.19(c) shows the equivalent experiment with 2.5 mM solutions, and shows similar results as for 2 mM.


In the case of TX-100 replacing CTAB (figure 5.14), we noted the absence of an



(a)



(b)



(c)

Figure 5.18: Replacement of a CTAB solution with a surfactant mixture at the same total concentration with $\chi_{\text{CTAB}} = 0.5$. (a) 1 mM, (b) 2 mM and (c) 3 mM. Adsorbed amounts: CTAB (+, red), TX-100 (\times , green), total (*, blue).

System being replaced	Concentration / mM	System injected χ_{CTAB}	CTAB rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$	TX-100 rate / $\mu\text{mol m}^{-2} \text{s}^{-1}$
CTAB	1	0.5	-0.36	0.12
CTAB	2	0.25	-0.38 ± 0.1	0.23 ± 0.3
CTAB	2	0.5	-0.27 ± 0.05	0.17 ± 0.06
CTAB	2	0.75	-0.13	0.06
CTAB	3	0.5	-0.29	0.12
TX-100	1	0.5	0.07	-0.09
TX-100	2	0.25	0.034	-0.11
TX-100	2	0.5	0.10	-0.19
TX-100	2	0.75	0.09 ± 0.02	-0.18 ± 0.1
TX-100	2.5	0.5	0.11	-0.14
TX-100	3	0.5	0.07	-0.13
TX-100	10	0.5	0.15	-0.25


Table 5.4: Rates of adsorption and desorption when mixed solutions replace a single surfactant solution.

overshoot in the TX-100 excess for the lower concentration of 1 mM. This behaviour is repeated here. Figure 5.19(a) shows the replacement of 1 mM TX-100 by an equimolar surfactant mixture: there is a smooth replacement of TX-100 by CTAB with almost no change in the total surface excess. Two higher concentrations were also studied. A 3 mM sample (figure 5.19(d)) showed only a small spike in TX-100 surface excess, compared with 2 mM and 2.5 mM samples, but a pronounced spike re-emerges at 10 mM concentration (figure 5.19(e)). The maximum adsorption and desorption rates are tabulated in table 5.4.


5.4 Conclusions

The first objective of this work was to demonstrate that total internal reflection Raman scattering can be used to study the adsorption kinetics of surfactant mixtures with chemical selectivity. We have shown for the system CTAB + TX-100 that interfacial kinetics can be followed with 2-s time resolution and a typical precision of $< 2 \times 10^{-7} \text{ mol m}^{-2}$ in the adsorbed amount of each component. The spectrometer is not optimized for fast kinetics and technical improvements such as a back-thinned CCD and customized software would improve the time resolution to 0.5 s. This time resolution is still inferior to that achievable by non-chemically selective techniques such as ellipsometry,^{60,64} but is sufficiently fast to follow almost all the adsorption, desorption and displacement experiments presented in this chapter.


Chemometric methods for data analysis are an indispensable aid for processing sets of kinetic spectra and allowed us to distinguish CTAB and TX-100 without selective deuteration of one component. Principal component analysis proved successful even for strongly overlapping spectra, but the surfactant spectra do need to have some distinguishing features. For example, the two surfactants in the study by Brinck and co-workers^{60,64} discussed earlier— C_{10}E_6 and C_{14}E_6 —could not be distinguished by TIR-Raman without deuteration of one component. TIR-Raman is not restricted to silica as a substrate: in work to be presented elsewhere we will show that the silica



(a)




(b)




(c)

Figure 5.19: Replacement of a TX-100 solution with a surfactant mixture at the same total concentration with $\chi_{\text{CTAB}} = 0.5$. (a) 1 mM, (b) 2 mM, (c) 2.5 mM, (d) 3 mM and (e) 10 mM total concentration. Adsorbed amounts: CTAB (+, red), TX-100 (\times , green), total (*, blue). (d) and (e) follow over the page.



(d)



(e)

Figure 5.19: ... continued. Adsorbed amounts: CTAB (+, red), TX-100 (x, green), total (*, blue).

can be coated with a thin organic film and the kinetics of adsorption of surfactant mixtures to the organic surface can still be extracted from the Raman spectra.

CTAB and TX-100 mix non-ideally both in the bulk and at a surface. Although there is no direct interaction between the hydrophobic chains or between the head groups leading to a favourable interaction parameter, the presence of the non-ionic surfactant between cationic surfactants reduces the electrostatic repulsions from the charged head groups while the relatively small trimethyl ammonium head group reduces the steric repulsions between the polymer-like polyethylene oxide chains of the non-ionic surfactant. At a hydrophilic silica surface, the adsorption isotherms of the two surfactants are qualitatively different: while both show a step in the isotherm at a concentration just below the bulk cmc, CTAB shows a plateau in the adsorption isotherm at lower concentrations while TX-100 shows no adsorption. Small amounts of CTAB bound electrostatically to the silica have a dramatic effect on the adsorption isotherm of TX-100. In a 2-mM solution of TX-100, the equilibrium amount of adsorbed TX-100 doubled in the presence of only 2% mole fraction CTAB. These interactions also have a major influence on the interfacial kinetics. When solutions containing CTAB are replaced with pure TX-100, the surface excess of TX-100 overshoots its equilibrium value and only relaxes to that of the pure system after all the CTAB has been washed away from the surface. Similarly, when replacing pure TX-100 with a mixed system the first CTAB to reach the surface induces a sharp rise in TX-100 adsorption. Mass transport effects also produce interesting behaviour. During rinsing of mixed surfactant layers with pure water, the more rapid transport of CTAB away from the surface causes a temporary increase in TX-100 monomer concentration in the subsurface region, which in turn leads to a brief spike in the TX-100 surface excess before the main desorption process commences.

Chapter 6

Cellulose

6.1 Introduction

Cellulose is the core component of cotton and paper. The adsorption of surfactants onto cellulose is a key step in the cleaning of cotton textiles and in the deinking of paper during recycling. Optimisation of formulations is assisted by an understanding of the thermodynamics and kinetics of adsorption of surfactants onto cellulose, both as pure compounds and as mixtures. Here we use total internal reflection (TIR) Raman spectroscopy to study the adsorption of surfactants onto cellulose both as binary solutions in water and as a ternary mixture. TIR Raman spectroscopy was invented in the 1970s^{133,140} but it is only recently that it has been developed into a powerful, surface-sensitive technique for quantitative studies of adsorption at the solid-liquid interface.^{3,40,175} In chapters 4 and 5 we demonstrated the use of TIR Raman spectroscopy in the study of the adsorption of pure and mixed surfactants on a silica substrate. While the technique is at its most straightforward when the substrate is transparent and has vibrational bands that do not overlap with those of the adsorbates, TIR-Raman spectroscopy can also be applied successfully to thin films deposited on transparent substrates even when, as is the case with cellulose, the vibrational spectrum of the thin film strongly overlaps the spectra of the adsorbates of interest.

TIR-Raman exploits the properties of evanescent waves to achieve surface sensitivity and therefore requires the use of surfaces that are flat on the length scale of the wavelength of light. TIR-Raman can be used to study fibrous materials if the material is pressed up against an internal reflection element composed of a high-index material,¹⁹⁹ but quantification is difficult and kinetic studies impractical. Fortunately, there are well-established techniques for preparing thin, flat transparent films of cellulose that are suitable for study by optical and neutron scattering techniques.^{106,270,271} In the work reported here, we first use TIR-Raman spectroscopy to characterise the preparation of a model cellulose surface by hydrolysis of the trimethylsilyl groups in a Langmuir-Blodgett film of hydrophobically modified cellulose. Next we look at the adsorption onto cellulose of the cationic surfactant hexadecyltrimethylammonium bromide (CTAB), which is a commonly used model system. The speed of acquisition of TIR Raman spectra allows us to map the adsorption isotherm much more

thoroughly than has been done before. Third, we look at the adsorption of the non-ionic surfactant Triton X-100 (TX-100) onto cellulose. TX-100 is challenging to study because it removes some of the cellulose layer, however the chemical specificity of Raman allows the two different processes—adsorption of surfactant and degradation of the layer—to be followed independently. Finally, we look briefly at a mixed surfactant system, demonstrating the ability of TIR-Raman spectroscopy to distinguish two surfactants in the presence of a strong cellulose signal.

Previous work on cellulose has taken two different approaches to sample preparation. Fibrous cellulose^{272,273}—for example filter paper—has the advantage of replicating real cellulose substrates closely, but the disadvantages that the cellulose surface contains a wide variety of different environments, that *in situ* characterisation of molecules adsorbed to the cellulose is difficult and that adsorption kinetics are likely to be controlled by transport through the fibres rather than by adsorption onto the surface itself. These drawbacks are overcome with thin, flat cellulose substrates. Such surfaces have allowed the use of a wide range of different experimental methods, including X-ray photoelectron spectroscopy,^{274,275} surface force apparatus,²⁷⁶ atomic force microscopy,²⁷⁷ quartz crystal microbalance,²⁷⁸ ellipsometry¹⁰⁶ and neutron reflectometry^{270,271} as well as the TIR Raman spectroscopy we use here. They are also amenable to quantitative studies of adsorption kinetics under well-defined mass transport conditions.¹⁰⁶

Two different approaches exist for the preparation of thin cellulose films:²⁷⁹ Langmuir-Blodgett (LB) deposition and spin coating. The most common approach to LB deposition, developed by Schaub *et al.*,²⁸⁰ involves the use of the functionalised cellulose derivative trimethylsilyl cellulose (TMSC). The principal advantage of working with TMSC rather than plain cellulose is that—unlike plain cellulose—TMSC can be dissolved in common non-polar solvents such as chloroform, toluene or *n*-hexane. Thin layers of TMSC can then be formed at the air-water interface and transferred onto hydrophobic surfaces such as hydrophobised gold, glass, silicon or mica. Unfunctionalised cellulose can be regenerated by exposure to HCl vapour. The properties of the deposited surface have been characterised extensively by IR spectroscopy, surface plasmon resonance, ellipsometry, surface force measurements and photoelectron spectroscopy,^{276,280,281} providing thicknesses of 10 Å per layer for TMSC and 4 Å per layer for the regenerated cellulose, showing that the charge on the cellulose chains is minimal (based on the absence of a double-layer force) and that the removal of the TMS groups is essentially complete.

Alternatively, TMSC can be spin-coated onto a substrate such as an anchoring polymer attached to silicon,^{106,282} or directly onto silicon or gold,²⁷⁴ followed by hydrolysis to remove the TMS groups. Kontturi and coworkers showed that partial hydrolysis was possible and could be controlled by changing the vapour pressure of HCl and the exposure time.²⁷⁴ Neuman *et al.* showed that it is also possible to spin coat cellulose directly using trifluoroacetic acid as a solvent;²⁸³ the current preferred solvent is N-methylmorpholine-N-oxide (NMMO).²⁸⁴ Cellulose has to be deposited onto an anchoring polymer—rather than directly onto silica—with the choice of polymer affecting the thickness of the surface. For all spin-coating processes, a range of experimental parameters can be used to control the film thickness, with typical values

being 200–1000 Å.

We used Langmuir-Blodgett deposition of TMSC rather than spin coating, principally because the films produced can be made thinner than those from spin-coating, with much finer control over the thickness of the film. A thin film is important in TIR Raman to minimise the Raman signal from cellulose, which overlaps the surfactant spectra.

The adsorption of CTAB onto cellulose has been studied by both neutron reflectometry²⁷⁰ and AFM.²⁷⁷ The two sets of data are not directly comparable since the surfaces were prepared in different ways: the AFM study used unfunctionalised cellulose spin-coated on top of a polymer layer²⁸⁴ whereas the neutron reflectometry used Langmuir-Blodgett deposition of a hydrophobically modified cellulose to produce hydrophobic and hydrophilic cellulose layers. Measured by neutron reflectometry, levels of adsorption onto hydrophilic cellulose and hydrophilic silica were similar ($5.9 \mu\text{mol m}^{-2}$), whereas the level of adsorption onto hydrophobic cellulose was roughly a third lower ($3.9 \mu\text{mol m}^{-2}$). Modelling of the neutron data suggested some intermixing between the hydrophobic cellulose and CTAB, whereas the CTAB self-assembled on top of the hydrophilic cellulose. AFM data revealed the formation of admicelles on a hydrophilic cellulose surface.

The adsorption of nonionic surfactants to cellulose has also been studied. Torn *et al.* followed the adsorption kinetics of a variety of ethylene glycol alkyl ether (C_nE_m) surfactants using optical reflectometry.¹⁰⁶ Although the adsorption processes were initially fast, the surface excess only reached its final value slowly, with C_{12}E_7 showing an inflection point in the adsorption curve, attributed to a surface rearrangement. Similarly, desorption processes also showed a fast initial process followed by a slow decrease towards complete rinsing (complete desorption took more than 20 min). Singh and Notley used AFM to show a mixture of spherical and rodlike micelles for C_{16}E_8 and C_{14}E_6 on cellulose surfaces.⁵⁶ They found that less C_{16}E_8 adsorbed on cellulose than silica, with the isotherm also becoming more step-like on cellulose. AFM imaging of the adsorbed surfactant showed similar spherical (C_{16}E_8) and rod-like (C_{14}E_6) structures on both cellulose and water. Adsorption of C_{12}E_6 and mixtures of CTAB and C_{12}E_6 have been investigated by neutron reflectometry.²⁷¹ C_{12}E_6 appeared to change the structure of both TMS-functionalised and plain cellulose while the original structure was largely recovered on rinsing. In mixtures, the composition of the surface layer on plain cellulose was close to ideal. The surface excess was largely independent of composition on both types of cellulose.

Cohen-Stuart and co-workers used a wall-jet geometry with optical reflectometry as the detection technique.^{282,285} Their studies were confined to the adsorption of polymers onto cellulose.

Several authors have studied adsorption of surfactants onto cellulose fibres. Paria *et al.* looked at the adsorption of TX-100, sodium dodecylbenzenesulfonate and CTAB) onto filter paper:^{254,272} the kinetics were generally slow, taking between 5 and 50 min to complete. Alila *et al.* used oxidation of cellulose fibres to control the surface charge and then investigated the adsorption of different chain lengths alkyl trimethylammonium bromide surfactants.²⁷³ The nature of cellulose fibres makes comparison of these results to those obtained on a thin flat cellulose film almost impossible.




Figure 6.1: Schematic of the deposition of cellulose onto a hemisphere.

6.2 Experimental

This experimental section primarily details the coating method. All other aspects of the experiment remain as described in chapter 2.

The sample preparation was based on the work of Penfold *et al.*,²⁷⁰ with some modifications. The cellulose was coated onto silica hemispheres (10 mm diameter; Global Optics, Bournemouth, UK). The hemisphere was cleaned as described above. The silica surface was hydrophobised by exposure to a 1,1,1,3,3,3-hexamethyldisilazane (purchased from Sigma-Aldrich) atmosphere for >12 h at room temperature, in a dry nitrogen atmosphere and then rinsed with water. TMSC dissolved in dichloromethane at a concentration of 0.7 mg mL^{-1} was spread on the surface of a Langmuir trough (Nima, Coventry, UK) and compressed to a surface pressure of 20 mN m^{-1} . TMSC monolayers were then transferred (figure 6.1) to the hemispheres at dipping and withdrawal rate of 5 mm min^{-1} with a 135-s pause at the end of both the dipping and the withdrawal. During the depositions the hemispheres were mounted so that the flat surface was vertical. Each hemisphere was dipped 5 times to produce a layer approximately 30 \AA thick (based on literature reports that 10 dippings give a $60\text{-}\text{\AA}$ thick film)²⁷⁰. Since parts of the holder and the hemispherical face of the prism also pass through the surface of the trough and may be coated with TMSC, it is not possible to calculate a transfer ratio. The cellulose layer on the curved surface of the hemisphere is too thin to significantly affect incident light entering the hemisphere.

The LB process produces a hydrophobic cellulose surface. Removal of the TMS groups to produce a hydrophilic surface was carried out within the wall-jet cell by exposure of the samples to a 3.5% solution of HCl for 15 mins, at a flow rate of 0.5 mL min^{-1} . We followed the removal of the methyl groups spectroscopically during the acid wash. Following the exposure to acid, the hemispheres were rinsed thoroughly with high purity water. This process deviates from most other work, where the substrate is exposed to the vapour above an HCl solution. The cellulose surfaces prepared are uniform (within the $10\text{--}30 \text{ }\mu\text{m}$ resolution afforded by the laser spot size). The reproducibility in the amount of cellulose deposited is $\pm 20\%$. Similar variability was observed in thickness measurements from neutron reflectivity.^{270,271}

Figure 6.2 illustrates schematically how the coated hemisphere is used in the Raman experiment. The refractive index of wet cellulose is intermediate between that of dry cellulose ($n = 1.53$) and water ($n = 1.33$) and is probably not very different from that of silica ($n = 1.46$).¹⁰⁶ Consequently, total internal reflection of the green




Figure 6.2: Schematic diagram of the cellulose hemisphere and the Raman probe laser (not to scale).




Figure 6.3: TMS cellulose surface damaged by TX-100, viewed under a $50\times$ microscope. The scale-bar represents $10\ \mu\text{m}$. The rainbow coloured horizontal lines show the damage, and are probably due to interference patterns as the cellulose thickness changes. A few scratches to the hemisphere are also visible.

laser beam takes place at the cellulose-water interface rather than the silica-cellulose interface.

All the surfactant isotherms and kinetics presented here were recorded on the unfunctionalised (hydrophilic) cellulose: we found that the TMS-functionalised cellulose was very readily removed by surfactants (especially TX-100, but also CTAB). Figure 6.3 shows a TMS-functionalised cellulose surface damaged by TX-100, illustrating how the removal of cellulose is easily identified. An effective method for identifying a partially removed cellulose surface is to scan the sample in the two axes parallel to the surface while taking short Raman acquisitions to identify if the cellulose signal changes.

In chapter 2 I detailed how the contribution to the Raman spectra from bulk solution could be used to calibrate the surface excess. Unfortunately, the susceptibility of the cellulose films to removal by surfactant limited our ability to use high enough bulk surfactant concentrations to derive an accurate calibration factor. An alternative approach is to compare the intensities of the surfactant peaks (relative to the water background that acts as an internal reference) with those on bare silica.²³⁹ This comparison neglects the water within the cellulose layer and any alignment differences between the two experiments. Consequently, we have presented the surfactant coverages as a component weight and quoted our best estimate for the conversion to surface excess in the figure captions.

6.3 Removal of $-\text{Si}(\text{CH}_3)_3$ from cellulose

Throughout this section we will use TMS to represent the $-\text{Si}(\text{CH}_3)_3$ group, whether it is bonded to cellulose or to the silica substrate. Polarisation-resolved spectra were acquired of a hemisphere hydrophobised with disilazane, after coating with hydrophobic cellulose and following hydrolysis with aqueous HCl. Figure 6.4 shows the Sy and Sx-polarised spectra. For the purposes of this chapter the differences between the




Figure 6.4: TIR Raman spectra in the C–H stretching region of a hydrophobised hemisphere (red solid line), after coating with TMS-cellulose (green dashed line) and after hydrolysis of the TMS groups (blue dotted line): (a) Sy and (b) Sx polarisation. Acquisition time = 300 s. In each case the surface is in contact with water.




Figure 6.5: The three components of the TIR Raman spectra shown in figure 6.4: water and TMS groups on silica (red solid line), plain cellulose (green dashed line) and TMS groups on cellulose (blue dashed line). (a) Sy and (b) Sx polarisation.

polarisations are not important except to note that the relative intensities of peaks differ and so some changes are easier to see in one polarisation than the others.

In order to follow the hydrolysis of the TMS-cellulose in real time by target factor analysis, we decompose the spectra into the three components shown in figure 6.5:

1. water/hydrophobic silica background, including the TMS peak from the hydrophobic coating;
2. unfunctionalised cellulose (the final state of the cellulose with the water/hydrophobic silica background subtracted);
3. TMS covalently attached to cellulose, generated by subtraction of the cellulose spectrum from the TMS-cellulose spectrum.

The difference spectrum between the two forms of cellulose—before and after the acid rinse—is very similar to the difference spectrum between hydrophobised silica and clean silica in water, except that the latter is shifted $\sim 5 \text{ cm}^{-1}$ to higher wavenumber (Figure 6.6). This similarity between the two difference spectra shows that only TMS




Figure 6.6: Comparison of spectra for $-\text{Si}(\text{CH}_3)_3$ groups from functionalised silica (solid red lines) and TMS cellulose (dashed green lines), Sy polarisation, acquisition time = 300 s. Vertical lines indicate approximate peak positions. The peak at 2903 cm^{-1} (TMS-cellulose) or 2908 cm^{-1} (functionalised silica) is the symmetric CH_3 stretch; the peak at 2961 cm^{-1} (TMS-cellulose) or 2966 cm^{-1} (functionalised silica) is the asymmetric CH_3 stretch.²⁷⁵

groups are lost during the acid rinse, and no cellulose. The shift in wavenumber reflects differences in chemical environment and allows the TMS groups on silica and cellulose to be identified independently. The TMS on groups on silica do not appear to be lost during the acid rinse: their contribution to the overall spectrum remains constant. To confirm that the TMS groups attached to silica are not removed by acid, we performed the acid rinse on a hydrophobised, but not cellulose-coated, hemisphere; there was no change between the spectra recorded before and after the acid rinse.

The hydrolysed cellulose spectrum still shows a small peak at 2970 cm^{-1} from the TMS groups (this is seen especially clearly in the Sx polarisation, figure 6.5(b)), indicating that some TMS groups remain after hydrolysis. The peak appears at a higher wavenumber than both the TMS removed from cellulose and the TMS on silica, indicating that the TMS groups on cellulose that survive the acid wash are in a different chemical environment from those removed. The area of the 2970 cm^{-1} peak permits an estimate of the unhydrolysed fraction of TMS groups of 5 to 10%.

Based on the integral of the spectra, the amount of TMSC deposited by the Langmuir-Blodgett process was observed to vary by up to $\pm 30\%$ from its average value but the fraction of TMS groups removed from the TMSC by acid hydrolysis remained constant.

Figure 6.7 shows the kinetics of hydrolysis. Inspection of the abstract principal components showed that only two factors were needed for the target factor analysis: a TMS factor and a factor accounting for the constant background of water and cellulose. The hydrolysis proceeds in two steps: a rapid initial removal of approximately half of the TMS groups, followed by a slower removal of the remaining TMS. These kinetics suggest that the TMS groups exist in two forms, with a difference in accessibility to acid. The two-step kinetics seen here are very reproducible however the fraction of the TMS groups removed in the first step varies between 0.2 and 0.7; the




Figure 6.7: Kinetics of removal of TMS from cellulose. Sy polarisation, 1 s acquisition time per point, flow rate of HCl solution = 0.5 mL min^{-1} . Only the TMS component is plotted; the cellulose/water component varied by less than 20% over the course of the experiment.

kinetic run shown in Fig. 6.7 is in the middle of the range.

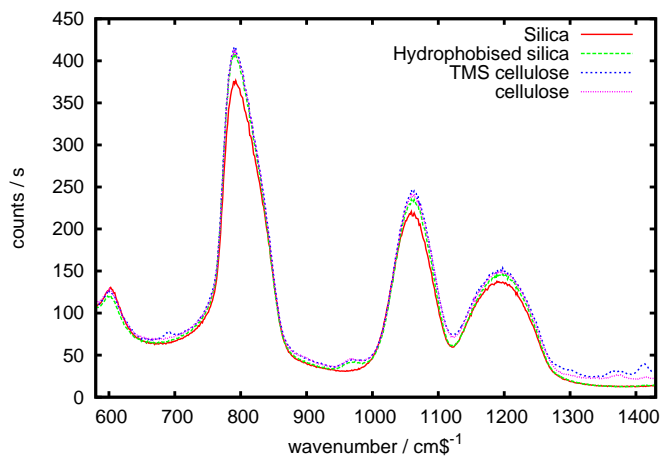
The hydrolysis of TMS cellulose by HCl vapour has previously been followed by XPS, ATR-IR,^{274,275} and static contact angle measurements.²⁷⁵ The XPS data suggested complete removal of the TMS groups, from the part of the film accessible to XPS (approximately the top 5 nm).²⁷⁴ The TMSC film was removed from the vapour after varying exposure times, rinsed with water and then analysed *ex situ*. Only four measurements were taken throughout the complete hydrolysis process, so it was not possible to identify the two-stage kinetics seen here. Very recently the hydrolysis has been followed *in situ* by X-ray reflectivity,²⁸⁶ and fitted to a first-order rate equation with respect to TMSC remaining. The acquisition rate was much slower than possible with TIR-Raman (5 data points over ~ 6 mins), so the 2-step process could easily have been missed, if hydrolysis by vapour proceeds in the same way as hydrolysis by solution.

6.3.1 600–1400 cm^{-1} region


Figure 6.8 shows raw and subtracted spectra of the 600–1400 cm^{-1} spectral region. In ATR-IR spectroscopy the peaks in the region are prominent, with peaks at 1251 and 842 cm^{-1} providing a clear identification of the Si-C bonds to the TMS groups.^{274,275} In TIR-Raman, using a silica hemisphere, the spectrum are dominated by the silica substrate, leading to a poor subtraction since tiny changes in the background spectra show up dramatically. Sx polarisation is shown since this is where the cellulose peaks are most visible.

Key peaks are

- 700 cm^{-1} (Sy, Sx and Px), seen only on TMS cellulose;
- 980 cm^{-1} , (Sy, Sx, Py and Px) seen in all three “functionalised” spectra;
- a possible small peak on the TMS cellulose around 1300 cm^{-1} (Sy, Sx and Px);



(a)



(b)

Figure 6.8: (a) Raw and (b) subtracted spectra of the 600–1400 cm^{-1} region, Sx polarisation. For the subtraction, the unfunctionalised silica spectrum has been used as the base spectrum. Many of the shapes in the subtracted spectra are attributed to slight experimental variation in the background signal from the silica, and should not be overinterpreted.

- two peaks at 1380 and 1415 cm^{-1} (Sy, Sx, Py and Px) on both the TMS- and plain cellulose surfaces, with the 1415 cm^{-1} peak reduced in intensity on the plain cellulose.

Assignment of peaks is difficult since they do not match peaks seen in the much clearer IR-spectra of cellulose on gold surfaces^{274,275} or of the pure compounds when not attached to a surface.²⁷⁸ Use of a crystalline substrate (for example, sapphire) might allow better spectra to be taken of this region. However, on silica, no useful information can be extracted.

6.4 CTAB

An adsorption isotherm of CTAB on hydrophilic cellulose was obtained with the inline mixer to vary continuously the sub-surface concentration of CTAB. Figure 6.9 shows examples of the raw data that generate a CTAB isotherm, together with the two components used in the target factor analysis of the data. When the sets of spectra are processed with TFA, they yield the isotherms shown in figure 6.10. The limiting surface excess from the isotherm was estimated to be $\sim 2 \mu\text{mol m}^{-2}$ from the slope of the CTAB component weight above the cmc (see Experimental Section), however this value is very approximate since measurements at higher concentrations are required for an accurate calibration of the component weights. For this reason we have plotted the isotherm in terms of component weight only. An alternative method for estimating the surface excess is to compare the intensities of the CTAB component adsorbed to cellulose and to silica surface, using the water component as an internal reference. The CTAB component measured on cellulose is approximately half the size of that on silica. Given that the adsorbed amount on the silica has been measured as $5.5 \mu\text{mol m}^{-2}$,⁴⁰ the estimated surface excess on cellulose is about $2.8 \mu\text{mol m}^{-2}$, which is a somewhat higher than the value estimated from bulk surfactant contribution above the cmc but lower than that determined by neutron reflection.²⁷⁰

The isotherms in figure 6.10 show hysteresis between the measurements taken with CTAB concentration increasing (red +), and those with CTAB concentration decreasing (green \times). Changing the flow rate, R , from 0.3 to 0.5 mL min^{-1} had little effect on the hysteresis. Due to the exponential dependence of concentration with time (equations 2.13 and 2.14), the rate of change of concentration in the low concentration region is slowest during the “out” runs and therefore the “out” kinetics are closer to the equilibrium isotherm than the “in”. This interpretation is supported by the difference in the apparent cmcs of the “out” and “in” measurements: for the “out” measurement the onset of desorption appears at 0.9 mM, matching the known cmc of CTAB, whereas for the “in” measurement the limiting surface excess is not reached until the bulk concentration reaches 1.5 mM, which is well above the cmc.

Figure 6.10 shows a plateau in the surface excess in the in measurement at low concentration. A similar plateau has been observed in the adsorption of CTAB on silica and ascribed to an electrostatic interaction between the positively charged surfactant and the negatively charged surface.¹⁹ For CTAB adsorption on cellulose, however, we believe the plateau seen in the “in” measurement is primarily a kinetic feature




Figure 6.9: (a) Example spectra (Sy polarisation) and (b) component spectra used in the analysis of the CTAB isotherm shown in figure 6.10. For part (b) the red solid line shows the water and cellulose background component while the green dashed line is the CTAB component. For part (a) each spectrum shows a 30 s acquisition and only every 10th spectrum is shown. The spectra in part (a) are illustrative so the time-stamps are not individually labelled.




Figure 6.10: Adsorption isotherm of CTAB on hydrophilic cellulose, expressed as CTAB component weight normalised to water. Red '+' represent an "in" measurement; green 'x' represent an "out" measurement. The flow rates were 0.3 mL min⁻¹. The CTAB solution was 5 mM concentration. Each point is from a 30 s spectrum. A component weight of 1.0 is estimated to represent a surface coverage of 2.0–2.8 $\mu\text{mol m}^{-2}$.

since it is not present in the “out” isotherm. The delay in adsorption during the “in” measurement is similar to the slow adsorption region reported for cationic surfactants on silica,^{43,53} which is attributed to the nucleation of surface aggregates. Slow adsorption has not previously been reported on a cellulose surface, but the adsorption kinetics of cationic surfactants on cellulose have not been as extensively studied.

The “out” isotherm is close to linear with concentration. Lattice based isotherms (for example the Langmuir or Frumkin isotherms) cannot provide a good fit to such a linear region. It is unusual for isotherms be so linear and the interpretation is unclear. A linear isotherm could potentially arise from a combination of the wide variety of different adsorption sites present on cellulose, which promote adsorption at low surface excess as the favourable sites are filled first, and the favourable interactions between adjacent CTAB molecules which promote adsorption at higher surface excesses.

The only previous adsorption isotherm of CTAB on cellulose was obtained by Penfold *et al.* using neutron reflection.²⁷⁰ They reported similar levels of adsorption for CTAB on cellulose and silica ($\sim 6 \mu\text{mol m}^{-2}$). Due to the limited availability of neutron beamtime they were only able to measure four concentrations below the cmc, with a minimum concentration of 0.1 mM. When replotted on a linear concentration scale, their data shows $d\Gamma/dc$ decreasing slightly with increasing concentration. They saw no evidence of a plateau at low concentrations, consistent with the plateau in Figure 6.10 being a kinetic rather than thermodynamic effect. They found that adsorption and desorption of CTAB did not change the thickness of the cellulose layer. For thin films (thickness much less than the penetration depth of the evanescent wave), TIR-Raman is sensitive only to the total amount of adsorbed cellulose and not to thickness.

Kinetics of adsorption for CTAB on hydrophilic cellulose are shown in figure 6.11(a). CTAB of the specified concentration was flushed directly into the cell at a flow rate of 0.5 mL min^{-1} . The rates of adsorption increase gradually with increasing concentration before levelling out above the cmc. At the highest concentrations (4 and 10 mM), adsorption is complete within 20 s, which is much faster than reported on filter paper.²⁵⁴ Adsorption onto cellulose is roughly 3 times slower than onto silica (see chapter 4).

The kinetics of desorption are shown in figure 6.11(b). The rates of desorption vary little with concentration. Most of the data are for initial concentrations above or near the cmc. Desorption does not commence until the subsurface concentration drops below the cmc; an initial concentration above the cmc therefore affects the time delay before desorption begins more than the desorption process itself.

6.5 TX-100

An isotherm for the adsorption of TX-100 onto cellulose is shown in figure 6.12(a). The component spectra are shown in figure 6.12(b). The TX-100 caused loss of a small amount of cellulose, which is also shown in figure 6.12(a). The loss of cellulose is a kinetic—rather than an equilibrium—property and should be linked to time rather than the bulk TX-100 concentration (TX-100 concentration increased with time during the “in” measurement and decreases with time on the “out” measurement). The




Figure 6.11: CTAB kinetics for 0.47 mM (red +), 0.75 mM (green \times), 1.0 mM (blue *), 4 mM (pink \square), 10 mM (turquoise \blacksquare). (a) Adsorption kinetics, (b) desorption kinetics. The curves have been offset on the time axis for ease of viewing. The kinetics are presented in terms of component weight (normalised to water) rather than surface excess owing to the uncertainty in the calibration factor.




Figure 6.12: (a) TX-100 adsorption isotherm, expressed as component weight normalised to water. TX-100 is shown in green (“in” measurement: +, “out” measurement: \square); change in cellulose signal during the experiment is shown in blue (“in” measurement: \times , “out” measurement: \blacksquare). Acquisition times were 10 s for the first 100 spectra and 30 s for the remainder; the flow rate was 0.5 mL min^{-1} and the surfactant solution 0.6 mM concentration. For TX-100, a component weight of 1.0 corresponds to a surface excess of approximately $4 \mu\text{mol m}^{-2}$. The change in cellulose signal cannot readily be interpreted as an absolute value. (b) Components used in analysis of the TX-100 isotherm. The background spectrum is shown as solid red lines, the TX-100 spectrum is shown as dashed green lines and the change in cellulose is shown as dotted blue lines.

similarity of the TX-100 isotherm in the ‘in’ and ‘out’ measurements shows that the loss of a small amount of cellulose does not significantly affect the adsorption properties.

We were able to obtain a reasonable straight line increase in TX-100 signal above the cmc (using a higher concentration than in figure 6.12 and not shown here), giving a limiting surface excess for TX-100 ($1.2 \mu\text{mol m}^{-2}$) approximately a quarter of that on pure silica ($4.3 \mu\text{mol m}^{-2}$, see chapter 4). However, comparison of the relative sizes of the TX-100 component for adsorption on silica and cellulose (normalised to water in both cases) yields a much larger value for the maximum surface excess of around $2.6 \mu\text{mol m}^{-2}$ ($\sim 60\%$ of that onto silica). This discrepancy suggests that in the TFA there may be a small amount of mixing of the TX-100 component and the component ascribed to the loss of cellulose. We therefore prefer the higher value as the best estimate of the surface excess. The most important difference between the adsorption of TX-100 onto silica and cellulose is the shape of the isotherm: the isotherm on silica is a step function, while on cellulose the isotherm shows a smooth increase with concentration. Our observations contrast with the results of Singh and Notley who reported that the C_{16}E_8 isotherm was more step-like on cellulose, and saw little change in the isotherm of C_{14}E_7 .⁵⁶ In chapter 4, on adsorption at silica, we attributed the step function to a strong interaction parameter promoting the formation of two phases. The smooth isotherm on cellulose suggests a reduction in the interaction between adjacent TX-100 molecules, possibly because aggregation is disrupted by a rougher surface or due to incorporation of TX-100 into the cellulose layer.

There is a propensity for TX-100 to remove the hydrophilic cellulose from the hydrophobic silica surface. The loss of cellulose is variable from experiment to experiment. During the recording of a single adsorption isotherm (‘in’ or ‘out’) the loss of hydrophilic cellulose is sufficiently slow that the adsorption isotherm is not greatly affected. Greater damage is caused by the three-phase line moving across the sample surface, for example when emptying and refilling the cell, which makes acquiring a large series of adsorption kinetics difficult. TMS-cellulose is particularly labile in TX-100 which prevented us from acquiring any reproducible isotherms from TMSC.

Only a very limited range of adsorption and desorption kinetics have been measured, due to the removal of the cellulose film by TX-100 described above. Examples of both adsorption and desorption kinetics are shown in figure 6.13. Too few adsorption kinetics were acquired to draw useful conclusions. Desorption kinetics are quick and the rate of desorption is essentially independent of concentration. As with the CTAB desorption kinetics this concentration independence is expected and shows that desorption does not start until the bulk concentration has dropped below the cmc.

6.6 Mixed systems

In chapter 5 we have studied the adsorption of mixtures of CTAB and TX-100 mixed systems onto silica, measuring the surface excesses for both surfactants at varying bulk composition with a constant 2 mM concentration. Target factor analysis was used to




Figure 6.13: (a) Adsorption and (b) desorption kinetics of TX-100 onto cellulose. Concentrations are 0.2 mM (red +), 0.4 mM (green ×), 1.0 mM (blue *), 4 mM (pink □) and 10 mM (turquoise ■). The different runs have been offset on the time axis for ease of viewing. The kinetics are presented in terms of component weight (normalised to water) rather than surface excess owing to the uncertainty in the calibration factor.




Figure 6.14: Adsorption of a mixture of TX-100 (green ×) and CTAB (red +) onto cellulose, 2 mM total concentration, with respect to surface excess (given as a mole fraction of CTAB). Total surface excess is also shown (blue *). Recorded while sweeping from pure TX-100 to CTAB with 0.3 mL min^{-1} flow rate, Sy polarisation, 30 s acquisition time. There is considerable uncertainty in the absolute surface excesses, due to the uncertainty in the calibration factors used when analysing the pure surfactant isotherms.

distinguish the two species on the surface. Here, we perform the same experiment on a cellulose surface (Figure 6.14), but using the inline mixer (with 2-mM CTAB replacing 2-mM TX-100) rather than a series of discrete compositions. TFA worked well with two components representing the surfactants and third representing the hydrophobic silica and the cellulose layer, demonstrating that the presence of the overlapping Raman spectrum of cellulose does not prevent quantitative analysis of mixed surfactant adsorption. One of the most striking features of adsorption onto silica is a two-fold increase in adsorbed amount caused by the presence of small (2%) mole fractions of CTAB. The origin of this cooperative effect is likely to be the interaction between the positively charged CTAB and the negatively charged silica surface. On cellulose there is no such promotion of adsorption at low CTAB mole fractions. In contrast there is a decrease in adsorption compared to the pure surfactants throughout most the composition range. Although CTAB and TX-100 do interact favourably, as shown by work on the micellisation²⁶⁷ and our work on silica, it is possible that if the TX-100 adsorbs inside the cellulose surface the opportunity for mixed aggregates to form is limited (this is plausible if TX-100 is able to remove the cellulose layer). The formation of mixed micelles in solution would then decrease the chemical potential of the surfactants and hence decrease the adsorbed amount.

6.7 Conclusions

We have demonstrated the use of TIR-Raman spectroscopy for measuring adsorption onto a model polymer surface (cellulose), and shown that the signal from the cellulose surface does not prevent the determination of the amount of adsorbed surfactants in both pure and mixed systems. TIR-Raman is sensitive enough to follow adsorption kinetics with a 2-s time resolution.

The main experimental difficulty in performing experiments on cellulose deposited on a hydrophobic surface is the propensity of surfactants—especially TX-100—to remove the cellulose during experiments. The TMS-functionalised cellulose film is particularly easily removed. Performing a range of kinetics measurements is especially difficult, since the process of emptying and refilling the cell after each measurement removes the surface especially quickly.

The adsorption isotherms of both CTAB and TX-100 on cellulose are close to linear, in contrast to the isotherms on silica which are strongly influenced by electrostatic interactions with the substrate (in the case of CTAB) and intermolecular interactions between surfactants in surface micelles (for both surfactants). The linearity of the adsorption isotherms on cellulose may arise from a balance of two effects: the heterogeneity of the adsorption sites, which promotes adsorption at low coverages, and the intermolecular interactions that promote adsorption at high coverages.

The surfactant isotherms were acquired continuously under quasi-equilibrium conditions in which slow adsorption/desorption kinetics are manifested by hysteresis between isotherms taken with increasing and decreasing surfactant concentrations. The CTAB isotherm exhibited pronounced hysteresis. A ‘slow adsorption region’ below the cmc has previously been implicated in the adsorption of CTAB on silica and assigned to formation and reorganisation of surface aggregates.^{43,53} One feature of

the inline mixer is that the rate of change of concentration is higher nearer the beginning of the experiment than the end. Consequently, the measurement taken with decreasing concentration is closer to equilibrium in the concentration range below the cmc than is the measurement with increasing concentration.

A mixed adsorption isotherm was acquired at 2-mM total concentration for comparison with adsorption on silica. The mixed layers on cellulose do not show the strong interaction between components that is displayed on silica, with no evidence of cooperative behaviour.

We are also able to follow the removal of TMS functional groups from cellulose during the preparation of the model surface. Although the initial and final states of the cellulose film have previously been studied in considerable detail, only TIR-Raman spectroscopy has the time resolution to follow the hydrolysis as it takes place, revealing a previously unseen two-step hydrolysis.

Chapter 7

Mica

7.1 Introduction

This chapter details the development of TIR Raman spectroscopy on mica. The technique was only partially successful, and therefore the chapter is relatively brief. First I will look at the reasons for using a mica surface and previous work relating to the adsorption of CTAB on mica. Then I will describe the experimental set-up together with initial tests to take spectra from acetonitrile solutions below mica. Finally I discuss an attempt to look at CTAB adsorbing to mica, and also some of the challenges faced in using a mica surface.

Although silica is a convenient surface for TIR Raman spectroscopy, it is not necessarily a good model for commercially relevant surfaces, hence the ability to use TIR Raman with alternative surfaces would be useful. Mica is commonly used in AFM and surface force experiments, and so would allow useful comparison with a wide range of other experiments. One of the main reasons mica has been used extensively is the ease with which clean, flat sheets can be produced. Mica is hydrophilic and negatively charged (allowing cationic surfactants and polymers to bind to it readily).²⁸⁷ The origin of the negative charge is usually attributed to the partial dissolution of K^+ ions on the mica surface and the exchange of K^+ by H^+ or H_3O^+ ions.²⁸⁸ It is commonly used as a model for clays and mineral oxides, although the structures formed by adsorbed surfactant are sometimes quite different.²⁸⁹

The majority of work on CTAB-like cationic surfactants at the water–mica interface has used atomic force microscopy (AFM), and so is more concerned with the shape of aggregates adsorbed in a saturated layer than with the amount adsorbed.

Both $C_{12}TAB$ and $C_{14}TAB$ adsorb to mica as rodlike micelles in the absence of salt.²⁹⁰ When $CsCl$ is added the surface aggregates become shorter with increasing concentration, eventually changing to spherical aggregates. The same applies to $C_{16}TAB$,²⁹¹ although Liu and Ducker reported an absence of structure which only formed cylindrical aggregates with KBr .²⁹² The salt concentration dependence observed is the opposite to that in bulk solution: the micelles grow longer in bulk solution with added salt. Additionally, $C_{16}TAB$ is seen to show long term changes (~ 6 – 20 hrs) when adsorbed to mica, with fast initial adsorption.²⁹³ After adsorption the surface changes from a micellar structure to a bilayer structure with a second—




Figure 7.1: Schematic of TIR-Raman on mica. This is deliberately not to scale (the oil and mica layers are made much thicker than in reality to allow them to be seen).

more weakly bound—bilayer adsorbing on top.²⁹¹ Adding KBr to C_{16} TAB on mica disrupts the structure as K^+ ions compete for negatively charged sites; HBr affects the surface only if present before the surfactant (suggesting H^+ cannot penetrate the surfactant layer), while tetraethylammonium bromide has fairly little effect.²⁹⁴ Liu and Ducker reported a featureless, crystalline phase forming as the temperature was lowered below the “surface Krafft point”.²⁹²

At intermediate concentrations C_{16} TAB,²⁹⁵ C_{18} TAB^{296–298} and C_{18} TAC²⁹⁹ (which has a chloride counterion instead of a bromide) adsorb patchwise onto mica, with the concentration affecting the number of patches, but not their shape or size. The patches were interpreted as bilayers. The patches grow slowly: 10–100 hrs, depending on temperature and the relative size of each patch at the same surface coverage also varies with temperature.²⁹⁸ IR and contact angle data suggests the presence of a liquid-like disordered phase present before the bilayer patches form.^{296,298} In the presence of low concentrations of CTAB (0.01 mM) the mica surface is rendered hydrophobic, even upon rinsing with high purity water.²⁸⁸ The effect is more pronounced for CTAF (fluoride counterion) than CTAB.²⁹⁵

7.2 Experimental

Figure 7.1 shows a rough illustration of the apparatus used to look at mica. An index matching oil is used to couple the light between the hemisphere and the mica. In principle—providing the angle of incidence at the hemisphere–oil interface is not above the critical angle for that interface—it should not matter too much what oil is used: the angle of incidence at the mica should be the same, although the amount of light lost to reflection may vary. However, I was only able to get results with one type of index matching fluid; I am unsure why. Table 7.1 contains refractive indices for a variety of different materials which were tried during my attempts to do TIR Raman on mica.

The mica samples used were 15 mm diameter, 25 μm thick “ruby mica disks”, as

Material	Use	n	Notes	Ref.
Silica	hemisphere	1.46		300
BK7	hemisphere	1.52	potentially fluorescent	300
SF10	hemisphere	1.74		300
hexadecane	ind. match	1.434		301
toluene	ind. match	1.496		301
squalene	ind. match	1.494		301
1-methyl naphthalene	ind. match	1.615		301
1-fluoro naphthalene	ind. match	1.593		301
1-bromo naphthalene	ind. match	1.657		301
		1.563		
mica	—	1.596	muscovite (biaxial)	302
		1.601		

Table 7.1: Refractive indices of materials considered for use with mica. Refractive indexes are given at 532nm where possible. These are all materials tried in the experiments, but most did not work and so are not discussed further in the text.

supplied by Unilever, Port Sunlight.

The first success came from using a BK7 hemisphere with 1-methylnaphthalene as the index matching fluid at an angle of incidence of 65° . Silica and SF10 did not work as hemispheres with 1-methylnaphthalene, although I am not sure why.

Figure 7.2 shows spectra of an acetonitrile solution below a mica surface as well as water below a mica surface. Acetonitrile is a reasonably useful indicator since it has a large Raman cross section (making it easy to identify) and appears even when in solution (useful since it does not adsorb to the surface). In my early attempts to do Raman on mica I had been misled by small amounts of water signal not coming from the cell so acetonitrile provides a useful sign that the spectra acquired do actually come from the sample. As can be seen in figure 7.2 the 1-methylnaphthalene has a very strong Raman spectrum in the CH region, and hence it would be desirable to use a different index matching fluid. Additionally the baseline is reasonably high—probably due to fluorescence from the BK7 hemisphere—although I mitigate the effect of the baseline by subtracting it as a constant value across the whole spectrum before putting the spectra into TFA. Levels of water signal were around 1/10th of what I would expect on silica.

Because of the large Raman signal from 1-methylnaphthalene within the $2700\text{--}3000\text{ cm}^{-1}$ CH region, the similar compounds 1-fluoronaphthalene and 1-bromonaphthalene were tested. By removing the methyl group the amount of CH signal should decrease dramatically as shown in figure 7.3. It is also useful to know what the spectra of mica looks like (shown in figure 7.4), to judge whether any peaks from the mica are likely to overlap. There are very few mica peaks between $1200\text{ and }3500\text{ cm}^{-1}$; at around 3600 cm^{-1} there is a very strong mica peak although this is well outside spectral region I plan to look at.

I found that 1-fluoronaphthalene produced odd interference patterns with the mica (see figure 7.5) and gave a very poorly matched refractive index between the silica and mica meaning that the system was unusable for spectroscopy. The interference patterns were not seen with clean mica and remained even when the 1-fluoronaphthalene was dried off with a tissue, although a thorough washing with methanol removed







Figure 7.2: Spectra—in 4 different polarisations—showing water and a dilute solution of acetonitrile in water below a mica surface. The index matching fluid was 1-methylnaphthalene, For S polarised light (purely perpendicular to the optical axis of mica) the angle of incidence at the hemisphere was 65° , giving a penetration depth of 53 nm (for the squared electric field). P polarised light will have slightly different penetration depth due to the birefringence of the mica. The concentration of acetonitrile was only judged by eye, but was around 5% by volume.



(a) 1-methylnaphthalene



(b) 1-fluoronaphthalene



(c) 1-bromonaphthalene

Figure 7.3: Raman spectra of various naphthalene derivatives. (a) and (c) are taken, with permission, from SDBSWeb: <http://riodb01.ibase.aist.go.jp/sdbs/> (National Institute of Advanced Industrial Science and Technology, 25 May 2011); (b) Reprinted from ref. 303 (*Spectrochimica Acta Part A: Molecular Spectroscopy*, **10**, Shambhu N. Singh and H. S. Bhatti and R. D. Singh, “Vibrational spectra and assignments of 1-halonaphthalenes”, p985–992), Copyright (1978), with permission from Elsevier.




Figure 7.4: Sy polarised spectra of mica, taken by external reflection from the mica-air interface, 50 mW laser power. The peak at 3600 cm^{-1} rises to 170000 counts, but is cropped here so that the rest of the spectrum can be seen more clearly.




Figure 7.5: The mica/1-fluoronaphthalene interface viewed under a microscope. The scale bar is $10\text{ }\mu\text{m}$. The octagonal shape is due to the illuminating light having been passed through a octagonal iris.




Figure 7.6: Raw spectra for CTAB on mica, Sy polarisation, 300s acquisitions, 700 mW laser power. The other polarisations were also recorded but are not shown here.

them. I suspect that the oil was penetrating inside the mica and separating out the layers. An alternative possibility is that the oil might not have been wetting the mica interface. The 1-bromonaphthalene did not seem to match either BK7 or SF10 hemispheres well enough to be take spectra of the water below, although did not show the same problems as the fluorinated variant.

7.3 CTAB

7.3.1 Isotherm

The next step in working with mica was to try to look at the adsorption of a surfactant onto mica. I picked CTAB as being readily available and a surfactant which I would expect to adsorb.

Figure 7.6 shows spectra for CTAB on mica, using 1-methylnaphthalene as the index matching fluid. There is significant variation in the amount of 1-methylnaphthalene signal seen, since it evaporated over the course of the experiment (and eventually needed replacing). Therefore, during the analysis of the data, three components were required (shown in figure 7.7): water, 1-methylnaphthalene and CTAB. The CTAB component came from subtracting water from one of the more oil-free spectra in figure 7.6; the component probably has a small oil contribution, however I do not expect this to significantly affect the results. The CTAB component weight was normalised by the water component weight to account for changes in detection efficiency; The 1-methylnaphthalene component weight was ignored. The base level the spectra appear at also varied significantly; these were all shifted down to so that the average intensity of 2573–2666 cm^{-1} was set to 0 to remove that complication from the analysis. As can be seen in figure 7.7 the strong 1-methylnaphthalene peak at 3050 cm^{-1} has started to appear in other components. I do not think that matters here, but I found it useful to crop out the range 3035–3085 cm^{-1} during some of the more difficult kinetic analysis.

Figure 7.8 shows an attempt to generate an isotherm of the adsorption of CTAB




Figure 7.7: Components used in the analysis of CTAB on mica, S_y polarisation.




Figure 7.8: "Isotherm" of CTAB on mica, showing S_y and S_x polarisations.




Figure 7.9: Components used to analyse the kinetics of CTAB adsorbing onto mica, Sy polarisation. Typical CTAB component weights are very weak ($\sim 1/20$ th of the background component weight).


on mica. Obviously there are significant issues with reproducibility since different points at essentially the same concentration show hugely different adsorbed amounts. Additionally, the two polarisations do not match each other well (and adding the Py and Px polarisations only makes matters worse). Mica is clearly not yet in a usable state.

7.3.2 Kinetics


I also attempted to record a range of kinetic measurements for CTAB adsorbing to mica. In some ways these are likely to be more reliable than the equilibrium measurements since all the data points are recorded within ~ 20 minutes of each other and as part of one process, meaning there is less chance of the alignment or the thickness of index matching fluid changing between measurements. Because the level of 1-methylnaphthalene background signal is likely to stay constant through an individual kinetic run I found it best to use just two components: CTAB and a background component incorporating both water and 1-methylnaphthalene. The appearance of the background component tended to vary from run to run as the level of 1-methylnaphthalene changed, but a typical example is shown in figure 7.9.

Examples of the kinetics of adsorption are shown in figure 7.10. There are obvious issues with where the zero point for a clean surface is; these arise because of the difficulty in finding a representative spectrum of a clean surface with an appropriate amount of 1-methylnaphthalene signal. Instead, I used the most significant factor present as the background, without rotating it by a target spectrum. There is no reason why this should affect the shape of the curve, just the offset.

The kinetics largely seem to show an unusual shape, with an initial rapid increase in adsorbed amount followed by a steady linear increase in adsorbed amount with time. In later experiments I tried running the experiment for longer times (up to




(a) 0.04 mM




(b) 0.3 mM (three separate runs, the third run (in blue and over the extended time) also had a small fraction of acetonitrile added but it was not possible to resolve this separately)

Figure 7.10: Kinetics of CTAB adsorbing to mica.




(c) 3 mM




(d) 10 mM

Figure 7.10: Kinetics of CTAB adsorbing to mica (continued).



(e) 3 mM with acetonitrile



(f) 3 mM with acetonitrile (enlargement on the x-axis)

Figure 7.10: Kinetics of CTAB adsorbing to mica (continued).




Figure 7.11: Components for the 3 mM kinetics with acetonitrile shown in figure 7.10(e), Sy polarisation. The region $3035\text{--}3085\text{ cm}^{-1}$ has been zeroed to ensure that the large peak there is ignored in the fitting.

20 mins) in the hope of seeing when the process stopped, but it appeared to still be ongoing even after this time. The slight downward slope in surface excess before the onset of adsorption is puzzling, and may indicate problems with the analysis of the data. Other kinetics runs—the 0.04 mM and the 3 mM with acetonitrile (Figure 7.10 (e) and (f))—showed much faster kinetics with no rise afterwards. The 0.04 mM was also unusual because of the large amount that adsorbed for the small concentration of surfactant.

Figure 7.10 (e) and (f) shows an attempt to use acetonitrile as a marker to indicate when the solution reaches the surface (the Raman signal from acetonitrile scales linearly with bulk concentration therefore providing a means of telling when the injected solution has reached the surface). The refined component spectra are shown in figure 7.11. The rate at which acetonitrile appears is reasonably similar to the rate at which surfactant appears, which I would expect to make the two components harder to separate. I do not know what causes the shift in component weights between 700 and 800 s.

Figure 7.12 shows the kinetics of desorption for CTAB from mica. Only one example is shown because the process appears to be similar at all the concentrations measured. The process appears to be slow, taking place over the course of 50–100 s.

7.3.3 Conclusions

The slow adsorption of CTAB over long periods of time (figure 7.10) is consistent with the large scatter seen in the isotherm: the spectra for the isotherm were taken before equilibrium was reached and hence the measured surface excess was essentially arbitrary, depending on the time at which it was taken. One interpretation of the slow increases in surface excess could be the slow change in the size of bilayer patches reported from AFM studies.^{295,296} For this explanation to be reasonable the surface




Figure 7.12: Kinetics of desorption of CTAB from mica, 0.3 mM concentration.

excesses of CTAB must be consistent with a partial bilayer coverage. It is difficult to estimate the surface excess from the amount of CH signal: although it is relatively large compared to the water (in comparison with adsorption to silica) we do not understand the optics of the system well enough to know how much water is being probed.

7.4 Other issues

Once the mica surface has been exposed to CTAB, bubbles of air formed quite readily on it, especially directly under the laser spot. It is reasonably well established that CTAB permanently deposits a hydrophobic layer onto mica,²⁸⁸ and so the attachment of air bubbles is unsurprising. I suspect the laser spot may be inducing the formation of these bubbles, since they often appeared there and nowhere else. The attachment of air bubbles could be inhibited in one of two ways: either by degassing the water before the experiment, or by increasing the pressure of the water during the experiment to disfavour the formation of gas bubbles.

I found that once a bubble of air had formed, the laser was liable to burn the index matching fluid, even causing damage to the hemisphere on occasion. This suggests that the laser power was too high, or it was too tightly focused. I suspect the adjacent water was helping to cool the area, and once a bubble had formed that was no longer happening.

7.5 Continuation

Although I have produced some initial results using mica there is still a reasonable amount of work to be done before the results are reproducible, and the system is robust enough to be used routinely. Although the spectrum from the index matching fluid is comparatively strong, I do not believe this presents a huge problem, at least for

single component systems where the analysis is usually straightforward. The primary issue with using mica is that we do not understand where the problems come from; however the knowledge of one working arrangement does provide a starting point for future improvements.

Appendix A

Temperature variation of water


Although none of the experiments in this thesis are directly concerned with temperature variation, many other experiments using TIR-Raman spectroscopy are. Because of the high degree of hydrogen bonding in water, water spectra are quite sensitive to temperature. One solution is to take water backgrounds at all temperatures of interest, however these is typically a 10% variation in spectral signal as a result of slight differences in alignment.

In order to accurately record the variation in water signal I have included a small quantity (0.7% by weight) of acetonitrile to act as an internal marker. Acetonitrile does not form hydrogen bonds so is expected to show much less temperature sensitivity than water. The spectra are then normalised—by eye—to the acetonitrile signal; changes in the shape of the water spectra make factor analysis of little use here. Spectra were recorded at Sy and Sx polarisation with 700 mW power at the laser and 240s total acquisition time.


Figure A.1 shows the spectra of water at a variety of different temperatures. Sy and Sx polarisation are shown; the lower signal levels for Sx polarisation—especially for the acetonitrile signal—make the normalisation spectra more difficult for Sx polarisation. The Sy polarisation shows two overlapping peaks: one at 3200 cm^{-1} and one at 3400 cm^{-1} . The 3200 cm^{-1} peak decreases with temperature, while the 3400 cm^{-1} increases with concentration giving an isosbestic point. The relative heights of the peaks are two peaks shown in figure A.2. For the Sx polarisation, there is one peak at around 3400 cm^{-1} . The trend is less clear, but the peak shows a slight shift to higher wavenumber with increased temperature, shown in figure A.3.

Walrafen *et al.* measured the temperature dependence of the Raman scattering of water.³⁰⁴ The isosbestic point* I measure is at $\sim 3300\text{ cm}^{-1}$, about 100 cm^{-1} lower than that measured by Walrafen *et al.* for the polarisation equivalent to Sy. The literature isosbestic point for the Sx equivalent polarisation is at 3522 cm^{-1} —just on the edge of my recorded spectra—and does not look to correspond to a point in my data. We have neglected to consider the change of the refractive index of water

*according ref. 304 an isosbestic point should really be called an “isoskedastic point” for Raman.



(a) S_y



(b) S_x

Figure A.1: Variation of water spectra with temperature at two different polarisations. Spectra are normalised to the acetonitrile signal (at around 2940 cm^{-1}).




Figure A.2: Peak heights for Sy polarisations, relative to the lowest temperature spectrum.




Figure A.3: Peak position for Sx polarisation. The peak position was selected by finding the wavenumber with the maximum number of counts in the spectrum.

with temperature: for the purpose of being able to perform accurate background subtractions at a variety of temperatures it should not be important.

Appendix B

Anionic surfactants

Anionic surfactants—such as SDS—do not typically adsorb to silica since they are repelled by the negative charge on the substrate. However, when mixed with another surfactant that does adsorb, they may be able to adsorb with it. In this appendix I will briefly present results for SDS mixed with two different surfactants: Triton X-100 (TX-100) and dodecyldimethyl(3-sulfopropyl)ammonium hydroxide (lauryl sulfobetaine, referred to as sulfobetaine from here on). In both cases adsorption was conducted onto a silica substrate.

B.1 SDS–TX-100

Figure B.1(a) shows the surfactant component weight (normalised to water) with respect to concentration for a 1.75:1 SDS:TX-100 mixture in the absence of salt. The data only required two factors to account for it: a water factor and a surfactant factor. The two factors are shown in figure B.2a, together with expected spectra for the two individual surfactants shown in figure B.2b. As can be seen from figure B.1(a), the component weight increases linearly with concentration, suggesting that the signal arises entirely from the bulk solution. Figure B.1(b) illustrates the small size of signal from a 40 mM solution (the highest concentration measured).

A further observation—which I have not investigated thoroughly—is that SDS and mixed surfactant layers containing SDS can adsorb in the presence of CaCl_2 . Adsorption only seems to occur when the cell is filled up from being empty, but not when a surfactant/ CaCl_2 solution replaces a CaCl_2 solution or when a surfactant solution containing CaCl_2 replaces a surfactant solution without salt. This would seem to suggest that the adsorbed layer is deposited by the three phase (water, solid, air) line as it passes across the substrate. The adsorbed layer is not usually stable, and will disappear if left in the cell overnight; I did not succeed in recording the kinetics of this desorption process.

B.2 SDS–lauryl sulfobetaine

On its own the zwitterionic surfactant sulfobetaine adsorbs to silica. Figure B.3a shows the adsorption isotherm for this surfactant, measured in three different polar-




Figure B.1: (a) Normalised component weight with respect to concentration for a 1.75:1 SDS:TX-100 mixture on silica. (b) Sy spectrum of 40 mM solution, indicating the lack of adsorption (compare with the level of C–H signal seen in figure 4.2a, for example).




Figure B.2: (a) Component spectra used in the TFA analysis to produce the data seen in figure B.1(a), Sy polarisation. (b) Spectra of the individual surfactants. The SDS spectrum comes from a high concentration of SDS in the bulk, and is rescaled to be a similar signal level to the TX-100 spectrum.




Figure B.3: (a) Sulfobetaine isotherm expressed as component weight and recorded in 3 different polarisations. (b) Sy spectra used to generate the isotherm in part a.






Figure B.4: Sulfobetaine components for three different polarisations.

isation combinations and expressed as a component weight normalised to water (the slope above the cmc is not good enough for calibration to absolute surface excess). Figure B.3 shows the spectra (Sy polarisation) used to generate the isotherm in (a); the C–H signal level indicates an amount of adsorption similar to that seen for TX-100 in Chapter 4. Figure B.4 shows sulfobetaine spectra in three different polarisations; the spectra show distinct peaks which should make it possible to separate from other surfactants using factor analysis.

The adsorption of a mixture of a 4 mM mixture of SDS and sulfobetaine with respect to composition is shown in figure B.5(a). The data was analysed using the surfactant component shown in figure B.5(b); the spectrum is essentially the same as a sulfobetaine spectrum. The cmc of sulfobetaine is around 3.5 mM (this value agrees reasonably with the data shown in figure B.3) and the cmc of the mixed system decreases slightly (to a minimum of around 2 mM) with increased SDS concentration.³⁰⁵ Therefore, the 4 mM mixture would be expected to contain an appreciable concentration of both monomers and mixed micelles; it is possible that the monomer concentration of sulfobetaine was pushed below the point at which adsorption first occurs, but it is still surprising that no adsorption takes place until the mixture contains



(a)



(b)

Figure B.5: (a) Isotherm of a 4 mM mixture of SDS and sulfobetaine with composition. (b) Surfactant component used in the factor analysis of the isotherm (S polarisation).

less than 3% SDS. Note that ref. 305—the only paper I have found to measure the cmc of a lauryl sulfobetaine–SDS mixture—uses a very unusual definition of “total concentration” (and hence cmc), and has not provided a legend for their graphs so my interpretation of their results might be inaccurate.

Appendix C

Additional data for CTAB/TX-100 mixtures

This appendix presents additional data relating to the CTAB/TX-100 mixed system presented in chapter 5. The data is presented without detailed discussion: that can be found in the main chapter on the subject.




Figure C.1: Isotherm showing the effect of a small mole fraction of CTAB (5%) on the adsorption of TX-100. Signal from the bulk has not been subtracted so may contribute to the increase at high concentrations. An isotherm of pure TX-100 is superimposed for comparison; its concentration has been scaled by 0.95 so that it matches the equivalent concentration of TX-100 in the mixture.




Figure C.2: Composition isotherm (2 mM total concentration) taken with the inline mixer (30 s acquisitions). Red indicates CTAB, green indicates TX-100, and blue indicates total adsorption. “crosses” (+, × and *) indicate data acquired while changing from TX-100 to CTAB and “shapes” (■, □ and ○) indicate data acquired while changing from CTAB to TX-100. Note that the data is slightly flawed in that the “clean” TX-100 surface, at $\chi_{\text{CTAB}} = 0$, was not quite pure and so the surface excess is too high. None-the-less, the figure illustrates the quality of the data that can be obtained using the inline mixer.

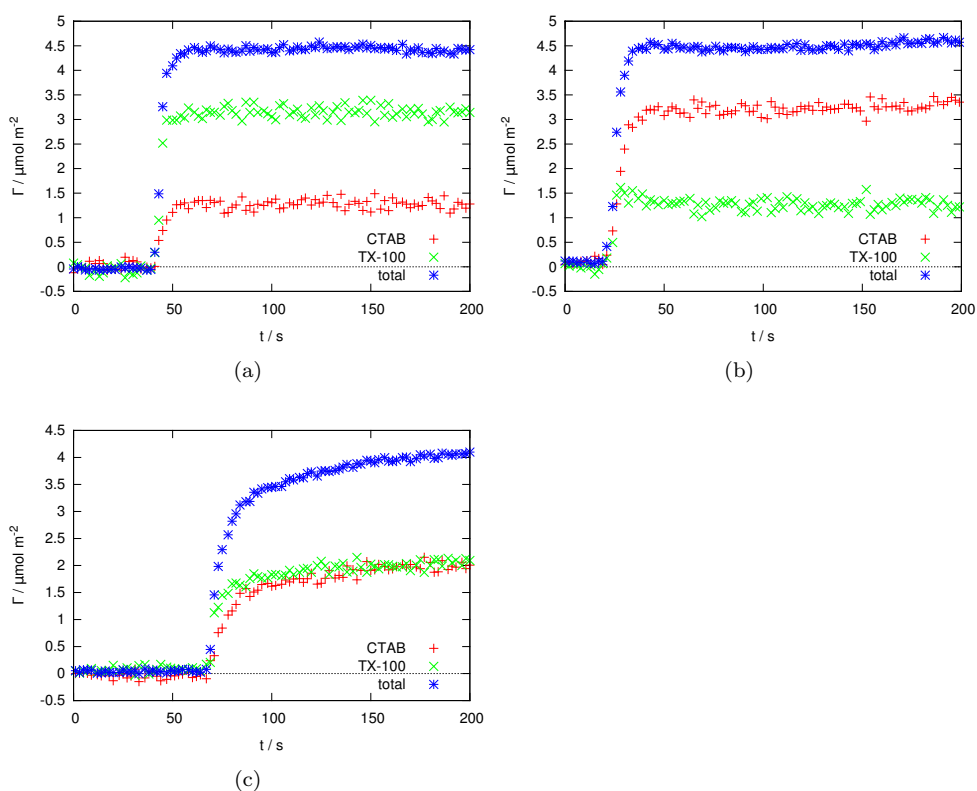


Figure C.3: Adsorption kinetics of CTAB/TX-100 mixtures: 2 mM total concentration $\chi_{\text{CTAB}} =$ (a) 0.25 and (b) 0.75. (c) 3 mM total concentration, $\chi_{\text{CTAB}} = 0.5$.





Figure C.4: CTAB replacing TX-100. Both solutions are (a) 1 mM and (b) 3 mM.


 Figure C.5: CTAB replacing CTAB/TX-100 mixtures: (a) and (b) 2 mM total concentration, $\chi_{\text{CTAB}} =$ (a) 0.25 and (b) 0.75. (c) and (d) $\chi_{\text{CTAB}} = 0.5$, total concentration (c) 1 mM and (d) 3 mM.




Figure C.6: TX-100 replacing CTAB/TX-100 mixtures: (a) and (b) 2 mM total concentration, $\chi_{\text{CTAB}} =$ (a) 0.25 and (b) 0.75. (c) and (d) $\chi_{\text{CTAB}} = 0.5$, total concentration (c) 1 mM and (d) 3 mM.

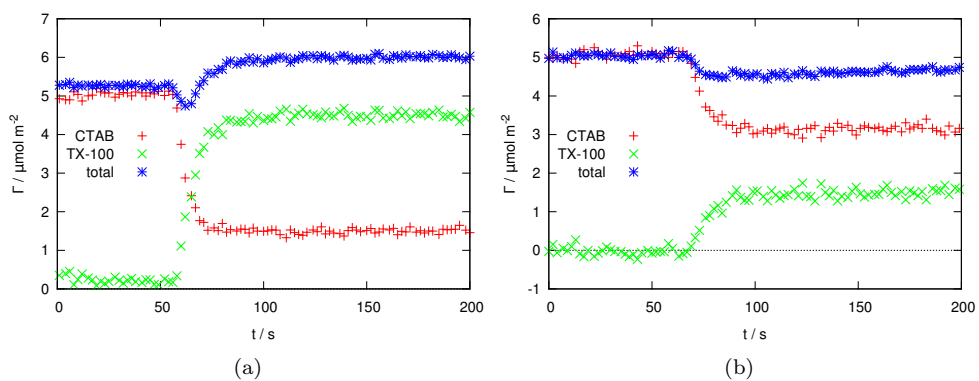


Figure C.7: Replacement of a 2 mM CTAB solution with a 2 mM total mixture of CTAB/TX-100 with $\chi_{\text{CTAB}} =$ (a) 0.25 and (b) 0.75.




Figure C.8: Replacement of a 2 mM TX-100 solution with a 2 mM total mixture of CTAB/TX-100 with $\chi_{\text{CTAB}} =$ (a) 0.25 and (b) 0.75.

Appendix D

Verification of mass transport

It is useful to be able to verify that transport to the surface behaves as described in Chapter 3. To do so I have used a non-surface active marker, and observed the signal arising from the bulk solution in the 100 nm probed by the evanescent wave.

Acetonitrile was chosen as the marker, because it has a strong Raman signal and I have already shown that the Raman response from acetonitrile is linear with concentration (figure 2.10(a)).

D.1 Model

Transport to the subsurface is modelled using the finite-difference model given in Chapter 3, with a no-flux boundary condition at the surface since the acetonitrile does not adsorb. The diffusion coefficient used for dilute ($< 10\%$) acetonitrile³⁰⁶ in water is $1.65 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$, 2 times greater than a typical surfactant monomer. In reality—but not in my model—the diffusion coefficient rises as the mole fraction of acetonitrile rises, but remains approximately constant for mole fractions less than 10%.

Values of $\bar{\alpha}$ at different flow rates were taken from curve 1 in figure 3.4, and converted to α , the dimensional form as described on page 94. The viscosity of a water–acetonitrile mixture does not change significantly with composition at volume fraction of $< 20\%$ acetonitrile³⁰⁷ so variation in viscosity need not be considered here. The actual lookup of the values of α used the MATLAB program “`calculatealpha.m`” (given in the Appendix E); the values are given in table D.1 for the flow-rates used here.

The modelled Raman signal was assumed to be proportional to the concentration at the gridpoint nearest to the surface. In principle the concentration at multiple gridpoints should be weighted according to the intensity of evanescent wave, $e^{-2\beta z}$, however each gridpoint represents a depth of 100 nm—approximately equal to the characteristic penetration depth—and concentration does not vary dramatically with z , so the first gridpoint provides a good approximation.

flow rate / mL min ⁻¹	Re	α / m ⁻¹ s ⁻¹
0.1	0.53	* 210
0.2	1.1	* 840
0.4	2.1	3300
0.5	2.7	4800
0.7	3.7	8800
1.0	5.3	17000
2.0	11	62000
5.0	27	370000

Table D.1: α at different flow rates. Values marked with “*” were outside the range measured in ref. 101, and so was estimated crudely by interpolation beyond the limit calculated.

D.2 Results

D.2.1 Transport to the surface

The transport to the surface of a 4% by mass acetonitrile solution being injected into water was recorded at a number of different flow rates, and the results are plotted in figure D.1.

As a general rule, the initial rate in the modelling and the experimental data match quite well, suggesting that hydrodynamics within the cell is well described by the one-dimensional convection-diffusion equation given in chapter 3. The rate at which the final 1/3 to 1/2 of acetonitrile reaches the surface is usually underestimated by the model. I believe this is probably because acetonitrile is a poor choice of marker—as I will be discuss later in this appendix, mixing between acetonitrile and water is non-ideal—rather than because the hydrodynamics do not match. Ideally the experiment should be repeated with a different marker: either acetonitrile solution being injected into deuterated acetonitrile solution (so the degree of mixing does not change throughout the process), or D₂O into H₂O, or a non-adsorbing organic salt (e.g. sodium tosylate).

D.2.2 Transport from the surface

Modelling of transport from the surface was probed by injecting water into a cell filled with 4% by mass acetonitrile solution. In general, the agreement with the model was *much* poorer than for the reverse experiment; some representative examples are shown in figure D.2. At low flow-rates the removal of acetonitrile was much slower than expected. A reasonable approximation could be obtained by assuming instantaneous mixing throughout the whole cell, using the same equation as governs the inline mixer (2.14). The volume of the cell was estimated as 7 mL, although not measured exactly. At intermediate flow-rates the concentration of acetonitrile shows an unexpected—but reproducible—periodic oscillation accompanied by a gradual decrease in amplitude. The down-slope after each maximum in the oscillation is similar to the slope predicted using the convection-diffusion model. At higher flow-rates the concentration of acetonitrile shows a monotonic decrease, although still not as fast as predicted by the convection-diffusion model.




Figure D.1: Experimental (points) and modelled (lines) transport of acetonitrile to the surface at a variety of different flow rates. The value of α for the two graphs marked with “*” was outside the range measured in ref. 101, and so the modelling for these two graphs has an additional uncertainty. Note that the scale on the x -axis is different for each part. The experimental data has been offset on the x -axis to match the modelled data.




Figure D.2: Experimental (points), convection-diffusion modelled (solid blue lines) and instantaneous mixing modelled (dashed green lines) transport of acetonitrile away from the surface at a variety of different flow rates. Note that the scale on the x -axis is different for each part. The experimental data has been offset on the x -axis to match the modelled data.

I do not believe that the results seen here with acetonitrile are representative of the behaviour for rinsing surfactants from the cell: I have seen no oscillations in the surface excess of any of the surfactants used, and would expect to if the concentration were fluctuating to that extent. Additionally desorption kinetics usual take place rapidly, after a time delay of ~ 70 – 100 s (expected from the length of tubing) suggesting rapid rinsing rather than the slow rinsing seen for acetonitrile.


Where the thermodynamics of mixing is close to ideal Fick's Law of diffusion can be simplified so that it depends on concentration rather than chemical potential:

$$J = -\frac{Dc}{RT} \frac{\partial \mu}{\partial x} \approx -D \frac{\partial c}{\partial x}. \quad (\text{D.1})$$

It is this simplified form that is used in my convection-diffusion model. However, mixing of acetonitrile with water is known to be non-ideal^{308–310}—especially at small mole fractions of acetonitrile—meaning that diffusion does not simply depend on the concentration gradient. The effect of the non-ideal mixing is not obvious, especially in relation to the periodic oscillations. However, ideally the experiment should be conducted using a marker that either mixes ideally, or is governed by the self-diffusion coefficient rather than a mutual diffusion coefficient (when using deuterated markers).

Appendix E

Computer Code

This chapter contains the computer code used in various parts of the thesis. The code is divided by functionality, and presented with explanatory comments. Throughout the code the symbol “” is used to indicate a linebreak introduced while printing the code in the thesis, not present in the original code.

The majority of the code here is MATLAB code. Most of the MATLAB code should work in Octave (a free MATLAB replacement) too although it has not been tested extensively. I have tried to indicate where I know the code does not work with Octave.

E.1 Data analysis

E.1.1 Target factor analysis

Many of these files are adapted from ref. 219. Some of the code has been changed so that it runs faster; the original code is show in the comments throughout the file.

afa.m

“afa.m” produces the abstract principle components, but does not attempt rotation to match the real factors. It is often useful to look at the abstract factors to judge how many factors are needed to describe the system, and to assess what target factors might be involved.

```

% * * * * *
% ** afa.m
% * * * * *
% returns the abstract factors, or optionally the n most important abstract
% factors. Also returns reproduced data based on those n factors
%
% function [r, c, cleanedData] = afa(d, n)
%
% IN:
% d = data matrix
% n = number of significant factors which are important (optional)
% OUT:
% r = rows
% c = columns
% cleanedData = data reproduced with only n factors (r*c)
function [r, c, cleanedData] = afa(d, n)

[rows, columns] = size(d);

smallest = min([rows, columns]);

switch (nargin)
case (1)
    sigNo = smallest;
case (2)
    sigNo = n;
otherwise
    error('afa.m takes one or two arguments.');
```

```

end

if (rows < columns)
    [v, s, u] = svd(d', 0);
else
```

```

    [u, s, v] = svd(d, 0);
end

c = v(:, 1: sigNo)';
r = u(:, 1: sigNo) * s(1: sigNo, 1: sigNo);
cleanedData = r * c;
```

tfa.m

```

% * * * * *
% ** tfa.m
% * * * * *
% tfa.m target factor analysis - a program designed to target test
% suspected factors
%
% function [xp] = tfa(d, n, x)
%
% d = an (r x c) data matrix
% n = number of factors used in target test
% x = an (r x n) matrix composed of n test vectors, each vector with r
% elements
% xp = improved x matrix
%
% no provision is made for handling missing points (blanks) in the targets.
function [xp] = tfa(d, n, x)
format short e

[rx, nx] = size(x);
[r, c] = size(d);
if (rx ~= r)
    error('Target vectors must emulate columns of the data matrix! PROGRAM
    ABORTED');
end

ldngs = r;
s = c;
if (r < c)
    ldngs = c;
    s = r;
    [v, sv, u] = svd(d.', 0); % I'm pretty sure we want .' as transpose, not '
else
    [u, sv, v] = svd(d, 0);
end

% original code - vectorised code is below
% for j = 1: s
%     ev(j) = sv(j, j) * sv(j, j); % ev = lamda^0_j (eqn. 4.17)
%                                     % (with eqn 3.84)
%     rev(j) = ev1(j) / ((r-j+1)*(c-j+1)); % reduced eigenvector (eqn. 4.82)
%     u(:, j) = u(:, j) * sv(j, j); % R = US
% end
ev = (diag(sv).^2).';
rev = ev ./ ((r+1)-(1:s)).*((c+1)-(1:s));
```



```

vbar=v(:,1:n); % v' = C
% vectorised version follows
%for j=1:n
%   t(:,j)=pinv(ubar)*x(:,j); % solving x_load = RT^-1 (eqn 3.136)
%   xp(:,j)=ubar*t(:,j); % test equation the other way round
%   dx=xp(:,j)-x(:,j);
%end
t = pinv(ubar)*x;
xp = ubar*t;
dx = xp - x;

loadings=inv(t)*vbar'; % predict y in eqn 3.132 from transformation
e = d - x*loadings; % x = satisfactor test vector
xx = inv(x'*x); % part of eqn 1.35 (without W)

for j=1:c
    v=xx*(e(:,j)'*e(:,j))/(r-n); % eqns. 4.135 & 4.136
    %for k=1:n
    %   loaderr(k,j)=sqrt(v(k,k));
    %end
    loaderr(:,j) = sqrt(diag(v));
end

```

sfa.m

“sfa.m” attempts to estimate the number of significant factors. I have found it tends to overestimate the number of factors when given noisy data (such as kinetics experiments).

```

%* * * * *
%* sfa.m
%* * * * *
% sfa.m significant factor analysis - a program designed to help determine
% the number of significant factors in a matrix
% all replacements of vectorised code have been checked to ensure they
% match
%
% number of significant factors are returned as n
function [n, ev, re, ind, rev, sl] = sfa(d)
% d = data matrix
format short e

[r, c] = size(d);
if (r < c)
    d = d.'; % I changed ' to .'
[r, c] = size(d);
end

[ u, s, v ] = svd(d,0);

```

```

% original code below. vectorised version follows
%for j= 1:c
%   ev(j) = s(j,j) * s(j,j); %ev = eigenvector
%   rev(j) = ev(j) / ((r-j+1)*(c-j+1)); % reduced eigenvector (eqn. 4.82)
%end
ev = (diag(s).^2)';
rev = ev ./ (((r+1)-(1:c)).*((c+1)-(1:c)));

%sev(c+1) = 0;
%sdf(c+1) = 0;
%for k = c:-1:2
%   sev(k) = sev(k+1) + ev(k); % sev(k) = sum of all evs above level k
%   sdf(k) = sdf(k+1) + (r-k+1) * (c-k+1);
%end
sev = fliplr(cumsum(fliplr(ev)));
sev(1) = 0; sev(c+1) = 0;
sdf(2:c) = fliplr(cumsum(((r+1)-(c:-1:2)).*((c+1)-(c:-1:2))));
sdf(1) = 0; sdf(c+1) = 0;

% vectorised below
%for l = 1:c-1
%   re = real error
%   re(l) = sqrt(sev(l+1) / (r*(c-1))); % eqn 4.44
%   ind(l) = re(l) / (c-l)^2;
%end
re = sqrt(sev(2:c) / (r*(c-1)));
ind = re ./ ((c-(1:c-1)).^2);

semilogy(ind, 'or')
xlabel('FACTOR LEVEL')
ylabel('IND')
pause
clc

[vind, n] = min(ind);
disp(['IND function (eq. 4.63) indicates ', int2str(n), ' significant factors.'])
disp(['The real error (RE) is +/-', num2str(re(n)), ' (see eq. 4.44).'])
pause

re(c) = NaN; ind(c) = NaN;
%for j=1:c
%   t(j,1) = j;
%   t(j,2) = ev(j);
%   t(j,3) = re(j);
%   t(j,4) = ind(j);
%   t(j,5) = rev(j);
%   f = (sdf(j+1) * ev(j)) / ((r-j+1) * (c-j+1) * sev(j+1));
%end
f = (sdf(2:c+1) .* ev) ./ ( ((r+1)-(1:c)) .* ((c+1)-(1:c)) .* sev(2:c+1) );
%for j=1:c
%   f = (sdf(j+1) * ev(j)) / ((r-j+1) * (c-j+1) * sev(j+1));
%   % convert f (see eq. 4.83) into percent significance level
%   if (j<c)
%       tt = sqrt(f);
%       df = c - j;
%       a = tt / sqrt(df);
%       b = df / (df + tt * tt);
%       im = df - 2;
%       jm = df - 2 * fix(df/2);

```

```

%      ss = 1;
%      cc = 1;
%      ks = 2 + jm;
%      fk = ks;
%      if (im - 2) >= 0
%          for k = ks:2:im
%              cc = cc*b*(fk-1)/fk;
%              ss = ss + cc;
%              fk = fk + 2;
%          end
%      end
%      if (df-1) > 0
%          cl = 0.5 + (a*b*ss + atan(a)) * 0.31831;
%      else
%          cl = 0.5 + atan(a) * 0.31831;
%      end
%      if (jm <= 0)
%          cl = 0.5 + 0.5 * a * sqrt(b) * ss;
%      end
%      end
%      sl(j) = 100 * (1-cl);
%      sl(j) = 2*sl(j);
%      t(j,6) = sl;
%      % sl the last point has to be wrong!
%end
tt=sqrt(f(1:c-1));
df = c-(1:c-1);
a = tt ./ sqrt(df);
b = df ./ (df + (tt.*tt));
im = df - 2;
jm = df - (2 * fix(df / 2));
ss = ones(1,c-1);
cc = ones(1,c-1);
ks = 2 + jm;
fk = ks;

imabovzero = ((im-2) >= 0);
reps(imabovzero) = im(imabovzero)-ks(imabovzero);
while any(reps>=0)
    cc(reps>=0) = cc(reps>=0) .* b(reps>=0) .* (fk(reps>=0) - 1) ./ fk(reps
    >=0);
    ss(reps>=0) = ss(reps>=0) + cc(reps>=0);
    fk(reps>=0) = fk(reps>=0) + 2;
    reps(reps>=0) = reps(reps>=0) - 2;
end

dfabovzero = ((df-1) > 0);
cl(dfabovzero) = 0.5 + (a(dfabovzero) .* b(dfabovzero) .* ss(dfabovzero)
+ atan(a(dfabovzero))) * 0.31831;
cl(~dfabovzero) = 0.5 + atan(a(~dfabovzero)) * 0.31831;

jmabovzero = (jm <= 0);
cl(jmabovzero) = 0.5 + 0.5 * a(jmabovzero) .* sqrt(b(jmabovzero)) .* ss(
jmabovzero);

s1 = 2*100*(1-cl);
s1(c) = NaN; % no valid answer for end

t = [ 1:c; ev; re; ind; rev; s1].';

```

```

t(c,6) = 0/0;
disp('SFA RESULTS (note: %SL based on eq. 4.83)')
disp('          n          EV          RE          IND          REV          %SL')
disp(t)
end

```

E.1.2 Other data analysis

orderSpectra.m

The spectrometer software produces binary files, which can be converted into human-readable text files using the “batch file converter” tool provided with the software. For kinetics experiment these human readable files need rearranging further before they can be processed. “orderSpectra.m” performs the rearrangement.

```

function orderedData = orderSpectra(spectra)
% orders a spectrum into rows and columns by times
% assumes all peaks are in the same

[n] = size(spectra,2);

if n ~= 3 % the data is in the wrong format
    error('Data must be in a three column format.');
```

```

end

spectra = spectra.'; % easier to work with data in rows rather than columns (for
our loop)

% get the first time on the spectrum before starting the loop
time = spectra(1,1);
orderedData(2, 1) = time; % write the first column of our data as the first time
processed

peakNumber = 1; % used to store which how far down the spectrum we are
spectraCount = 1; % number of spectra so far processed

for k = spectra
    lastTime = time;
    time = k(1);

    if time ~= lastTime % if this is the case we have moved to a different
    spectrum
        spectraCount = spectraCount + 1;
        orderedData(spectraCount+1, 1) = time;
        peakNumber = 1;
    end
end

```

```

if spectraCount == 1 % on the first iteration
    orderedData(1, peakNumber+1) = k(2); % copy over the wavenumbers to the
    first row
end
% now copy to a matrix where rows contain the spectral response, and
% columns contain each time
orderedData(spectraCount+1, peakNumber+1) = k(3); % (the +1s ensure that the
first row is used for frequencies, % (and the first column for
% times)
peakNumber = peakNumber+1;
end

```

removeCosmicRaysFromTimeSeries.m

Although the spectrometer software can remove cosmic rays (which appear as sharp spikes in the spectra) it requires two extra acquisitions per measurement to do so, which is not feasible for kinetics measurements. This program attempts a similar function without the extra spectra. Because removal of cosmic rays is not always obvious, it prompts for user input before making changes. Sometimes this works better when the spectra are rotated (looking along a time-series at each wavenumber) rather than looking at a spectrum for each time.

```

% function cleanedSpectra = removeCosmicRaysFromTimeSeries(spectra, [bounds, [
droppedPoints]])
%
% takes a time series of spectra, and removes the cosmic rays from them.
% Ignores the first row and column on the assumption that they
% are axis labels. Spectra should be in the rows
% bounds says how many standard deviations to permit (but is optional)
% droppedPoints says how many of the highest points to ignore when
% calculating the standard deviation
function cleanedSpectra = removeCosmicRaysFromTimeSeries(spectra, varargin)
if ~isempty(varargin)
    bounds = varargin{1};
else
    bounds = [];
end
if length(varargin) > 1
    droppedPoints = varargin{2};
else
    droppedPoints = [];
end

```

```

cleanedSpectra = spectra;
repeat = 'y'; % default value
while repeat(1) == 'y'
    cleanedSpectra = removeFromTimeSeriesImpl(cleanedSpectra, bounds,
droppedPoints);
close; figure; plot(cleanedSpectra(2:end,2:end).');
transpose = input('Do you wish to also scan columns for spurious points? (y,
n) [y] ', 's');
if isempty(transpose)
    transpose = 'y';
end
if transpose(1) == 'y'
    cleanedSpectra = removeFromTimeSeriesImpl(cleanedSpectra.', bounds,
droppedPoints); % do transpose too
cleanedSpectra = cleanedSpectra.'; % reverse transpose
close; figure; plot(cleanedSpectra(2:end,2:end).');
end
repeat = input('Repeat? (y,n) [y] ', 's');
if isempty(repeat)
    repeat = 'y';
end
end
function cleanedSpectra = removeFromTimeSeriesImpl(spectra, bounds,
droppedPoints)
cleanedSpectra = spectra(1,:); % copy the first row
rotatedSpectra = spectra.';
for k=rotatedSpectra(:,2:end) % go across columns of rotatedSpectra (.: rows of
spectra)
    replacement = removeCosmicRays([rotatedSpectra(2:end,1), k(2:end)], bounds,
droppedPoints); % do a spectrum at a time
% always passing wavenumbers too
% trim wavenumbers off replacement
replacement = replacement(:,2);
replacement = [k(1), replacement.']; % add back in the first column (and
rotate replacement)
cleanedSpectra = [cleanedSpectra; replacement]; % append the next row to
the result
end

```

removeCosmicRays.m

This function removes cosmic rays from a single spectrum (or alternatively a time-series taken at a single wavenumber). It is not usually used directly, but called from “removeCosmicRaysFromTimeSeries” (discussed above).

```

% function cleanedSpectrum = removeCosmicRays(spectrum, bounds, droppedPoints)
%

```

```

% removes any peaks which are obviously cosmic rays
% should be passed an 'x by 2' or '2 by x' array
% containing the peak heights the peak heights
% bounds says how many standard deviations to permit (default is 10)
% droppedPoints says how many of the highest points to ignore when
% calculating the standard deviation (default is 3)

function cleanedSpectrum = removeCosmicRays(spectrum, varargin)

if ~isempty(varargin)
    bounds = varargin{1};
else
    bounds = [];
end
if length(varargin) > 1
    droppedPoints = varargin{2};
else
    droppedPoints = [];
end

if isempty(bounds)
    bounds = 10;
end
if isempty(droppedPoints)
    droppedPoints = 3;
end

[n,m] = size(spectrum);

if m>n
    spectrum = spectrum.'; % transpose the spectrum
    flipped = 1;
    temp = n; n = m; m = temp; % swap n and m
else
    flipped = 0;
end

if m ~= 2
    error('spectrum should consist of wavenumbers and intensities. Too few columns');
end

cleanedSpectrum = linearRegScan(spectrum, 50, bounds, droppedPoints);

if flipped
    cleanedSpectrum = cleanedSpectrum.';
end % else no rotation needed

function cleanedSpectrum = linearRegScan(spectrum, points, bounds, droppedPoints)
% points defines how big a region to look at once

[rows, columns] = size(spectrum);

if points > rows % process entire spectrum at once
    cleanedSpectrum = linearRegImpl(spectrum, spectrum);
    return;

```

```

end

% for the first part, process a chunk between the start of the spectrum
% and the first 'points' points (for testrange)
spectrum(1:points,:) = linearRegImpl(spectrum(1:points,:), spectrum, bounds,
    droppedPoints);

testRangeStartPoint = points + 1;

validRange = true; % this is true until set false

% now do a loop to cover the rest of the spectrum
while validRange
    testRangeEndPoint = testRangeStartPoint+points;

    if testRangeEndPoint > rows % we have gone out of range
        validRange = false; % stop next go
        testRangeEndPoint = rows; % terminate them both at the same point

        if testRangeStartPoint > rows % check the start is still valid
            break;
        end;
    end

    % if the start is still valid then do the procedure
    % and append the result to cleanedSpectrum
    spectrum(testRangeStartPoint:testRangeEndPoint,:) = linearRegImpl(spectrum(
        testRangeStartPoint:testRangeEndPoint,:), spectrum, bounds,
        droppedPoints);

    % increment the start point
    testRangeStartPoint = testRangeStartPoint+points+1;
end

cleanedSpectrum = spectrum;

function cleanedRange = linearRegImpl(spectrumRange, wholeSpectrum, bounds,
    droppedPoints)
% spectrumRange should be a x*2 array
% wavenumbers and intensities
% in principle a large portion of the data should overlap
% wholeSpectrum is so we can plot it if needs be

% find the largest three point in spectra and remove it since it will be this
% point (if any) that is the cosmic ray, so it's best if it doesn't
% adversely affect the deviation
trimmedSpectrum = spectrumRange;
if numel(trimmedSpectrum(:,2)) > droppedPoints
    for k = 1:droppedPoints
        [value, index] = max(trimmedSpectrum(:,2));
        trimmedSpectrum(index,:) = []; % delete
    end
end

% do a linear regression on the regression range
p = polyfit(trimmedSpectrum(:,1), trimmedSpectrum(:,2), 1);
m = p(1); c = p(2);

```

```

% calculate expected values for each point in the range test range
% based on the results of the linear regression
expected = [ spectrumRange(:,1), m * spectrumRange(:,1) + c ];

trimmedExpected = [ trimmedSpectrum(:,1), m * trimmedSpectrum(:,1) + c ];

% variation from predicted line
% use the trimmed versions, so as to ignore the largest points
differences = trimmedSpectrum(:,2) - trimmedExpected(:,2);

% work out stddeviation for this region
% (based on the deviation from the line)
deviation = std(differences);

% find points that are higher than the expected value plus the deviation
% these are marked as true
comparison = (expected(:,2)+(deviation*bounds)) < spectrumRange(:,2);

% turn comparison into something with 2 columns (both of which are
% identical
comparison = [ comparison, comparison ];

% copy across points where the thing is within range
cleanedRange = spectrumRange;

% check everything else with the user
cleanedRange(comparison) = checkValues(spectrumRange(comparison), expected(
    comparison), wholeSpectrum);

function newValues = checkValues(values, replacements, spectrum)
% values is an array of length x*2, containing either zeros, or values which
% need to
% be changed. Replacements is the same.
% Spectrum is provided so we can plot it

newValues = [];

[n,m] = size(values);

values = reshape(values, n/2, 2);
replacements = reshape(replacements, n/2, 2);

%re-get
[n,m] = size(values);

for k = 1:n
    close; figure; plot(spectrum(:,1),spectrum(:,2), replacements(k,1),
        replacements(k,2),'ro'); % so they can see the point hopefully
    disp('Dubious point found at:');
    disp(values(k,:)); % display the point
    delete=input('Do you want to delete this point? (y,n) [y] ', 's');

    if (isempty(delete))
        delete = 'y';
    end

    if delete(1) == 'y'
        newValues = [newValues; replacements(k,:)];
    else

```

```

        newValues = [newValues; values(k,:)];
    end
end

```

averageOfRangeToZero.m

For each spectrum in a series of spectra “averageOfRangeToZero” subtracts a constant value from every point in the spectrum so that the mean of the range specified is zero. This allows the removal of a constant baseline, which is sometimes useful for factor analysis.

```

% function result = averageOfRangeToZero(input, lowerLimit, upperLimit)
% lower and upper limit are indicies (into zero)
% if upperLimit is undefined it is taken as end

function result = averageOfRangeToZero(input, lowerLimit, upperLimit)

dim1 = size(input,1);

if ~exist('upperLimit','var') || isempty(upperLimit)
    upperLimit = dim1;
end

assert(lowerLimit < upperLimit);

averages = mean(input(lowerLimit:upperLimit,:));

result = bsxfun(@minus,input,averages);

```

E.2 Modelling

This section contains functions relating to Chapter 3.

E.2.1 Basic model

Example script

The modelling program is usually called from a script that sets up the parameters. An example script is included here, containing annotations explaining different possible options that can be passed to the program.

```

clear physconst;
clear initial;
clear frumkinconst;
clear gridconst;
clear langmuirconst;
clear cmc;
clear doNothingFunc;
clear micelleFunc;
clear multicomFunc;
clear boundaryConds;
clear boundaryConst;

% self diffusion coefficients
ctamonomerD = 5.60e-10; % m^2 s-1
ctamicelleD = 5.30e-11; % m^2 s-1
brD = 2e-9; % m^2 s-1
% all taken from Lindman J. Phys. Chem, Vol. 88, No. 21, 1984

gridconst.nn = 2000; % time points
gridconst.jj = 2000; % space points
gridconst.deltaZ = 1e-7; % m
gridconst.deltaT = 1e-2; % s
gridconst.theta = 1;
gridconst.components = 4; % ordered as 'ctab' 'triton' 'ctab micelles' 'triton'
% micelles
physconst.alpha = 5000; % m-1 s-1
physconst.diff = [ 2*ctamonomerD*brD/(ctamonomerD+brD), 2.8e-10, ctamicelleD,
5.3e-11 ]; % m^2 s-1

% Both perfect sink and no-flux boundary conditions, can take
% an optional edge parameter. If it is not passed then it is assumed to be
% 0
sinkconst = [];
sinkconst.edge = 0.23;

langmuirconst.ka = 85/10;
langmuirconst.kd = 40/10;
langmuirconst.gammainfty = 4.45e-6;
frumkinconst = langmuirconst;
frumkinconst.omegaa = +4.3;
frumkinconst.omegad = 1.7;
% where omega is greater than 4, setupFrumkinConsts finds and defines the
% tie lines needed for the Maxwell construction.
frumkinconst = setupFrumkinConsts(frumkinconst);

% initial concentrations (don't worry about monomer micelle equilibrium yet)
initial.conc0 = [ 1,1, 0, 0 ];
% initial adsorbed amounts
initial.gamma0 = [ 0, 0, 0, 0 ];
boundaryConds = { 'frumkin', 'sink', 'flat', 'flat' };
boundaryConst = { frumkinconst, sinkconst, [], [] }; % the langmuir and
% frumkin isotherms required parameters passing to them
% note: the frumkin isotherm does not really work well with the four
% component model; it is shown here for illustrative purposes.

rstConst.beta = -1.0;
rstConst.cmc1 = 0.92; % ctab
rstConst.cmc2 = 0.27;

% for below the cmc, or where multiple components don't interact, doNothingFunc

```

```

% performs no exchange between components in the concs array
doNothingFunc = @(concs, ignore) concs; % return what we put in

% Where both monomers and micelles are present, micelleFunc defines the
% equilibrium
micelleFunc = @(concs) partitionMicelles(concs,cmc); % 0.5 is the cmc

% for a two surfactant mixture with RST, partitionMixtureOfTwo defines the
% equilibrium. However, it is slow, so we prefer to use an interpolation
% based version defined below.
multicomFunc = @(concs) partitionMixtureOfTwo(concs, rstConst);

% the faster interpolation based version of regular solution theory
% requires a calculation at this stage, rather than during the model
disp('Preparing interpolation values for regular solution theory...');
clear concArray; clear alphaArray; clear xls; clear cmcs;
[concArray, alphaArray, xls, cmcs] = prepareRSTInterpolations(rstConst, 10, 60,
50);
disp('Finished preparing interpolation values.');
```

```

multicomFunc2 = @(concs) partitionMixtureOfTwoInterpolation(concs, rstConst,
concArray, alphaArray, xls, cmcs);

% defines the concentration being injected
physconst.concin = multicomFunc2([0,0,0,0]); % mol m^3
% ... and run the initial concentration through the function used to
% partition the components
initial.conc0 = multicomFunc2(initial.conc0);

% now call the actual routine
[ concs, gamma ] = oneDConvDiff(gridconst, physconst, multicomFunc2,
doNothingFunc, [], boundaryConds, boundaryConst, initial,true);

```

oneDConvDiff.m

The routine that performs the main modelling calculation. Note that this will not work with Octave unless “parfor” is replaced with “for”.

```

function [concs, gamma] = oneDConvDiff(grid, physconst, partitionFunc,
gammaPartitionFunc, converg, boundarycond, boundaryconst,
initialConditions, recordConcs)
% *****
% oneDConvDiff
% - solves the one dimensional convection-diffusion equation
%
% -- output --
% concn - concentrations at all timesteps
% gamma - surface excess
% concsmicelles - micelle concentrations at each stage (empty if unused)
% -- input --
% grid = a structure containing:
% - nn, jj = number of timesteps and size of array (both + 1!)
% - deltaZ, deltaT = timesteps in z and t direction
% - components = the number of components the solution is made up of
% (optional, assumed to be 1!)
% - theta = level of implicitness (1 to 0) (optional, assumed 1 if not

```

```

% given)
% physconst is a structure containing
% - diff = diffusion coefficient (in s.i. units) as a vector of length
%   grid.components
% - alpha = constant relating to convection rate as a 1 by 1 matrix
% - concin = incoming concentration as a vector of length grid
%   components (or a 1 by 1 matrix to set them all equal)
% partitionFunc = a function handle to partition between components. It is
%   passed the concs array from a timestep(jj by components), and
%   returns a similar sized array with the components correctly
%   partitioned
% gammaPartitionFunc = partitionFunc, but for the surface components.
%   We also take a vector of gammas for the last time step as the
%   second argument, since I think it's useful to some variants
% converg = maximum value required for convergence to be reached (percent) (↵
%   optional
%   default = 0)
% boundaryCond = text string cell array (components long) giving boundary
%   conditions
% boundaryconsts = structure containing details for implementing boundary
%   conditions (optional - default blank)
% initialConditions = optional structure containing optional fields. Both
%   of these must either be 1 by 1, or a vector of length
%   grid.components
% - conc0 = starting bulk concentration (defaults to zero)
% - gamma0 = starting gamma (defaults to zero)
% recordConcs = optional boolean. If false it saves memory but not
%   recording the concentrations. default = true
% * * * * *
if (~isfield(grid,'nn') || isempty(grid.nn)) % if nn is not given then run to ↵
%   infinity (with convergence conditions)
%   grid.nn = Inf;
%   if isempty(converg) || (0 == converg)
%       error('You're going to be waiting a while since you haven't specified ↵
%         a number of iterations or a convergence condition');
%   end
else
%   grid.nn;
end
if (~isfield(grid,'components')) || isempty(grid.components)
%   grid.components = 1;
end
if (~isfield(grid,'theta')) || isempty(grid.theta) % allow theta not to be ↵
%   specified
%   grid.theta = 1.0;
end
if (grid.theta > 1) || (grid.theta < 0)
%   error(['Theta = ', num2str(grid.theta), '; it must be in the range 0 to 1']) ↵
%   ;
end
if isempty(converg)
%   converg = 0.0;
end
if ~isempty(initialConditions)
%   if isfield(initialConditions,'conc0')

```

```

%   conc0 = zeros(1,grid.jj+1,grid.components);
%   if size(initialConditions.conc0,2) == grid.components && ...
%       size(initialConditions.conc0,1) == 1 && ...
%       size(initialConditions.conc0,3) == 1 % only has size in 2nd axis
%       initialConditions.conc0 = permute(initialConditions.conc0,[1,3,2]);
%       initialConditions.conc0 = repmat(initialConditions.conc0,[1,grid.jj ↵
%         +1,1]);
%   end
%   conc0(1,:,:) = initialConditions.conc0;
%   else
%       conc0 = 0;
%   end
%   if isfield(initialConditions,'gamma0')
%       gamma0 = initialConditions.gamma0;
%   else
%       gamma0 = 0;
%   end
%   conc0 = 0;
%   gamma0 = 0;
end
if isempty(recordConcs)
%   recordConcs = true;
end
% try preallocate for speed
if grid.nn ~= inf
%   preallocateSize = grid.nn;
else
%   preallocateSize = 1000; % guess at reasonable start
end
if recordConcs
%   concs = zeros(preallocateSize, grid.jj+1, grid.components);
%   else
%   concs = zeros(2, grid.jj+1, grid.components); % this step and last step
end
concs(1,:,:) = repmat(conc0, size(concs,2)/size(conc0,2), size(concs,3)/size(conc0 ↵
%   ,3));
gamma(preallocateSize,grid.components) = 0.0; % grows inside loop
gamma(1,:) = gamma0;
concs(1,grid.jj+1,:) = physconst.concin;
for n=2:grid.nn+1
%   if recordConcs
%       nForConcs = n;
%   else
%       nForConcs = 2;
%   end
%   concsForStep = zeros(size(concs,2),grid.components);
%   gammaForStep = zeros(1,grid.components);
%   parfor c=1:grid.components
%       if iscell(boundarycond)
%           bcond = boundarycond{c};
%       else
%           bcond = boundarycond;

```

```

end
if iscell(boundaryconsts)
    bconsts = boundaryconsts{c};
else
    bconsts = boundaryconsts;
end

% write out the physical constants specifically for this value of c
physconstsc = physconstsc;
if (length(physconstsc.diff) > 1)
    physconstsc.diff = physconstsc.diff(c);
end
if (length(physconstsc.concin) > 1)
    physconstsc.concin = physconstsc.concin(c);
end

[ concsForStep(:,c), gammaForStep(c) ] = oneDConvDiffStep(concs(↵
    nForConcs-1,:,c), gamma(n-1,c), grid.deltaZ, grid.deltaT, grid.↵
    theta, physconstsc, bcond, bconsts);
end

concsForStep = permute(partitionFunc(concsForStep),[3,1,2]);
gammaForStep = gammaPartitionFunc(gammaForStep,gamma(n-1,:));

concs(nForConcs,,:) = concsForStep;
gamma(n,:) = gammaForStep;

differencesConcs = 100*sqrt(sum((concs(nForConcs,:) - concs(nForConcs-1,:))↵
    .^2)/sum(concs(nForConcs-1,:))); % as a percentage
if all(gamma(n-1,:) == 0)
    differencesGammas = 100*abs(gamma(n,:) - gamma(n-1,:))/gamma(n-1,:);
else if all(gamma(n,:) == 0)
    differencesGammas = 100*abs(gamma(n,:) - gamma(n-1,:))/gamma(n,:);
    else
        differencesGammas = 0;
    end
end
if all(differencesConcs < converg) && all(differencesGammas < converg)
    concs = concs(1:n,,:);
    gamma = gamma(1:n,,:);
    break;
end

if ~recordConcs
    concs(1,,:) = concs(2,,:); % copy across
end
end
end

```

oneDConvDiffStep.m

Performs a single time step in the modelling calculation; not designed to be used directly.

```

function [concn, gamma] = oneDConvDiffStep(concnminus1, gammanminus1, deltaZ, ↵
    deltaT, theta, physconstsc, boundaryCond, boundaryconsts)

```

```

% * * * * *
% oneDConvDiff
% - performs one step in solving the one dimensional convection-diffusion ↵
% equation
%
% --- output ---
% concn - concentration at timestep n
% gamma - a surface excess value (for boundary conditions that need it)
% --- input ---
% concnminus1 = concentrations at timestep n-1
% gammanminus1 = a surface excess value (for some boundary conditions)
% deltaZ, deltaT = timesteps in z and t direction
% theta = level of implicitness (1 to 0) (optional - defaults to 1)
% physconstsc = a structure containing fields
% - diff = diffusion coefficient (in s.i. units)
% - alpha = constant relating to convection rate
% - concin = incoming concentration
% boundaryCond = boundary condition to apply - case sensitive text string
% allowed values:
% * 'sink' = perfect sink boundary condition (possibly above an edge)
% * 'source' = perfect source boundary condition - fixes the
% subsurface concentration to a value, unless it's naturally above
% that value (when we act as flat)
% * 'matchisotherm' = compare gamma (at last time step) with an an
% isothermFunc (takes subsurface conc). If above, then surface is a
% source, if below, surface is a sink.
% * 'flat' = sets dconc/dz to zero on the boundary - not physically
% realistic but a useful test of what happens (realistic for micelles)
% * 'langmuir' = langmuir isotherm
% * 'frumkin' = frumkin isotherm
% boundaryConsts = structure containing details relating to boundary
% conditions (not needed for sink)
% * * * * *
%
% subfunction, to account for my silly style of array indexing
% anonymous functions are much quicker than defining them as full functions
% (it would appear)
jIndex = @(j) (j+1);
revjIndex = @(index) (index-1);

boundaryextras = [];

if isempty(theta) % allow theta not to be specified
    theta = 1.0;
end
if (theta > 1.0)
    error('Theta must be less than 1');
else if (theta < 0)
    error('Theta must be more than 0');
end
end

if isempty(boundaryCond)
    boundaryCond = 'sink';
end
switch boundaryCond
case {'sink','source'}
    if ~isfield(boundaryconsts,'edge')
        boundaryconsts.edge = 0; % defaults to zero so acts as a true ↵
        perfect sink
    end
end

```

```

end
case {'matchisotherm'}
% valid
case {'flat'}
% valid
case {'langmuir', 'frumkin'}
if strcmpi(boundaryCond, 'frumkin') % a couple more constants for this
case
frconverginfty = 0.000001;
frconverggamma = 0.0000001;
end
% valid
otherwise
mesg = [boundaryCond, ' is not a valid boundary condition.'];
error(mesg);
end

%alpha = physconst.alpha;
diff = physconst.diff;
concin = physconst.concin;

jj = revjIndex(size(concminus1));
jj = jj(2);

% cache some commonly called values
zerotojj = jIndex(0:jj);
%onetojj = jIndex(1:jj);
j0 = jIndex(0);

% preallocate for (a bit of) speed
concn(jIndex(jj)) = 0;
switch boundaryCond
case 'sink'
boundaryextras.edgeconc = concnminus1(j0); % use mirror type boundary
condition
boundaryextras.asFlat = boundaryextras.edgeconc < boundaryconst.edge;
if ~boundaryextras.asFlat
boundaryextras.edgeconc = boundaryconst.edge; % don't allow it to
exceed edge amount
end
case 'source'
boundaryextras.edgeconc = concnminus1(j0);
boundaryextras.asFlat = (boundaryextras.edgeconc > boundaryconst.edge);
if ~boundaryextras.asFlat
boundaryextras.edgeconc = boundaryconst.edge; % don't allow it to
drop below amount
end
case 'matchisotherm'
boundaryextras.edgeconc = concnminus1(j0);
idealgamma = boundaryconst.isothermFunc(boundaryextras.edgeconc);
if gammanminus1 > idealgamma
elseif gammanminus1 < idealgamma
else
boundaryextras.asFlat = true; % do nothing if conc is right
end
case 'flat'
case {'langmuir', 'frumkin'}
conc0est = concnminus1(j0); % just use last value
if strcmpi(boundaryCond, 'frumkin')
[ gammaest, dadtn, dddtn ] = frumkinGamma(gammanminus1, conc0est,

```

```

boundaryconst, deltaT, frconverginfty, frconverggamma);
boundaryextras.dadtn = dadtn;
boundaryextras.dddtn = dddtn;
% get the fluxes to and from the surface from this too (this
% is based on previous step)
if ((boundaryconst.omegaa - boundaryconst.omegad) > 4) && ... %
require critical behavior
(gammanminus1 > boundaryconst.tielower) && (gammanminus1 <
boundaryconst.tieupper) % and to be within the 2
phase region
fractionUpperPhase = (gammanminus1 - boundaryconst.tielower) /
(boundaryconst.tieupper - boundaryconst.tielower);
kaPrefactor = boundaryconst.ka * (fractionUpperPhase * ...
exp(boundaryconst.omegaa*boundaryconst.tieupper /
boundaryconst.gammainfty) * ...
(boundaryconst.gammainfty - boundaryconst.tieupper) + ...
(1 - fractionUpperPhase) * ...
exp(boundaryconst.omegaa*boundaryconst.tielower /
boundaryconst.gammainfty) * ...
(boundaryconst.gammainfty - boundaryconst.tielower));
kdPrefactor = boundaryconst.kd * (fractionUpperPhase * ...
exp(boundaryconst.omegad*boundaryconst.tieupper /
boundaryconst.gammainfty) * boundaryconst.tieupper +
...
(1 - fractionUpperPhase) * ...
exp(boundaryconst.omegad*boundaryconst.tielower /
boundaryconst.gammainfty) * boundaryconst.tielower);
else
kaPrefactor = boundaryconst.ka * exp(boundaryconst.omegaa*
gammanminus1/boundaryconst.gammainfty) * ...
(boundaryconst.gammainfty - gammanminus1);
kdPrefactor = boundaryconst.kd * exp(boundaryconst.omegad*
gammanminus1/boundaryconst.gammainfty) * ...
gammanminus1;
end
boundaryextras.surfaceFluxFromLastStep = kaPrefactor * concnminus1(
j0) ...
- kdPrefactor;
else
gammaest = (boundaryconst.ka * conc0est * boundaryconst.gammainfty
+ (gammanminus1/deltaT)) / ...
((1/deltaT) + boundaryconst.ka*conc0est + boundaryconst.kd);
boundaryextras.dadtn = boundaryconst.ka*(boundaryconst.gammainfty
- gammaest);
boundaryextras.dddtn = boundaryconst.kd*gammaest;
boundaryextras.surfaceFluxFromLastStep = boundaryconst.ka *
concnminus1(j0) * (boundaryconst.gammainfty - gammanminus1)
...
- boundaryconst.kd * gammanminus1;
end
otherwise
error(['You shouldn't be here (' , num2str(boundaryCond), ')!']);
end
[a,b,c,d,boundaryextras]=generateOneDConvDiffParams(concnminus1,gammanminus1,
deltaZ,deltaT,theta,physconst,boundaryCond,boundaryextras);

```

```

% I've implemented two methods of matrix inversion - The "just do it in
% matlab" method and the Thomas algorithm. Timing on a 5000x5000 matrix
% with matlab 2009b gives pretty much no difference in speed between
% them (~2s for 200 time points).

% matrix inversion method
%{
% this has to do with how spdiag works, and getting the zero elements
% to the right square
a = [ a(2:end), a(1) ];
c = [ c(end), c(1:end-1) ];

M = sparse(length(d),length(d));
M = spdiags([a.',b.',c.],[-1,0,1],M);
% d = M*concn
% inv(M)*d = inv(M)*M*concn = concn
% (this is equivalent to the below)
concn = M\d.';
concn = concn.'; % rotate back to format used elsewhere
%}

% thomas algorithm for inversion
[e, f, elow, flow] = thomasalgorithm(a, b, c, d, 0, concn, jIndex, revjIndex);

switch boundaryCond
case 'flat'
    concn(j0) = flatConc0(elow, flow);
case {'sink', 'source', 'matchisotherm'}
    if boundaryextras.asFlat
        concn(j0) = flatConc0(elow, flow);
    else
        concn(j0) = flow;
    end
otherwise
    concn(j0) = flow; % set by lower boundary condition
end

% concn is already preallocated earlier
for jInd=1:(length(concn)-1)
    concn(jInd+1) = e(jInd)*concn(jInd) + f(jInd);
end

switch boundaryCond
case {'sink', 'source', 'matchisotherm'}
    if boundaryextras.asFlat
        flux = 0;
    else
        flux = diff*(concn(jIndex(1))-concn(j0))/deltaZ;
    end
    gamma = gammanminus1 + flux*deltaT;
    return;
case 'flat'
    gamma = gammanminus1;
    return; % from the while loop
case 'langmuir'
    %gamma = gammaest;
    gamma = (boundaryconsts.ka * concn(j0) * boundaryconsts.gammainfty +
            gammanminus1/deltaT) ...
            / (1/deltaT + boundaryconsts.ka*concn(j0) + boundaryconsts.kd);
    %boundaryextras.dadtn = boundaryconsts.ka*(boundaryconsts.gammainfty -

```

```

            gamma); % recalculate in terms of the new values of gamma
    %boundaryextras.dadtn = boundaryconsts.kd*gamma;
case 'frumkin'
    %gamma = gammaest;
    [ gamma, dadtn, dddtn ] = frumkinGamma(gammanminus1, concn(j0),
        boundaryconsts, deltaT, frconverginfty, frconverggamma);
    %boundaryextras.dadtn = dadtn;
    %boundaryextras.dadtn = dddtn;
end

n = 0;
% the vast majority of these terms stay exactly the same
% therefore move them out of the loop
Ja(zeroTojj+1) = a;
Jb(zeroTojj+1) = b;
Jc(zeroTojj+1) = c;

while true
    n = n+1;
    switch boundaryCond
    case 'langmuir'
        dadtn = boundaryconsts.ka*(boundaryconsts.gammainfty - gamma); %
            recalculate in terms of the new values of gamma
        dddtn = boundaryconsts.kd*gamma;

        JkaTerm = -boundaryconsts.ka;
        JkdTerm = boundaryconsts.kd;
        boundaryextras.dadtn = dadtn;
        boundaryextras.dadtn = dddtn;
    case 'frumkin'
        JkaTerm = boundaryconsts.ka*(boundaryconsts.omegaa ...
            - 1 ...
            - gamma*boundaryconsts.omegaa/boundaryconsts.gammainfty) ...
            *exp(boundaryconsts.omegaa*gamma/boundaryconsts.gammainfty);
        JkdTerm = boundaryconsts.kd*(1 + boundaryconsts.omegad/
            boundaryconsts.gammainfty*gamma) ...
            *exp(boundaryconsts.omegad*gamma/boundaryconsts.gammainfty);

        % dependency is much simpler when critical, cos it's just
        % linear
        if ((boundaryconsts.omegaa - boundaryconsts.omegad) > 4) && ... %
            require critical behavior
            (gamma > boundaryconsts.tielower) && (gamma <
                boundaryconsts.tieupper) % and to be within the 2
                phase region

            fractionUpperPhase = (gamma - boundaryconsts.tielower) /
                (boundaryconsts.tieupper - boundaryconsts.tielower);

            JkaTerm = boundaryconsts.ka / (boundaryconsts.tieupper -
                boundaryconsts.tielower) * ...
                (exp(boundaryconsts.tieupper*boundaryconsts.omegaa/
                    boundaryconsts.gammainfty)*(boundaryconsts.gammainfty
                    - boundaryconsts.tieupper) ...
                - exp(boundaryconsts.tielower*boundaryconsts.omegaa/
                    boundaryconsts.gammainfty)*(boundaryconsts.gammainfty
                    - boundaryconsts.tielower));
            JkdTerm = boundaryconsts.kd / (boundaryconsts.tieupper -
                boundaryconsts.tielower) * ...
                (exp(boundaryconsts.tieupper*boundaryconsts.omegad/

```

```

        boundaryconsts.gammainfty)*boundaryconsts.tieupper ...
        -exp(boundaryconsts.tielower*boundaryconsts.omegad/
        boundaryconsts.gammainfty)*boundaryconsts.tielower);

    dadtn = boundaryconsts.ka*(fractionUpperPhase*exp(boundaryconsts
    .omegaa*boundaryconsts.tieupper/boundaryconsts.gammainfty)
    *(boundaryconsts.gammainfty-boundaryconsts.tieupper) + ...
    (1-fractionUpperPhase)*exp(boundaryconsts.omegaa*
    boundaryconsts.tielower/boundaryconsts.gammainfty)*(
    boundaryconsts.gammainfty-boundaryconsts.tielower));
    dddtn = boundaryconsts.kd*(fractionUpperPhase*exp(boundaryconsts
    .omegad*boundaryconsts.tieupper/boundaryconsts.gammainfty)*
    boundaryconsts.tieupper + ...
    (1-fractionUpperPhase)*exp(boundaryconsts.omegad*
    boundaryconsts.tielower/boundaryconsts.gammainfty)*
    boundaryconsts.tielower);
else
    dadtn = boundaryconsts.ka*exp(boundaryconsts.omegaa*gamma/
    boundaryconsts.gammainfty)*(boundaryconsts.gammainfty -
    gamma);
    dddtn = boundaryconsts.kd*exp(boundaryconsts.omegad*gamma/
    boundaryconsts.gammainfty)*gamma;
end
boundaryextras.dadtn = dadtn;
boundaryextras.dddtn = dddtn;
otherwise
    error(['In Newton-Rhapson loop for a boundary condition (',
    boundaryCond, ') that requires no iteration']);
end

% need to regenerate a,b,c,d at j=0
physconst2 = physconst; physconst2.concin = concn(jIndex(1));
% don't use concin at end
%[a,b,c,d,boundaryextras]=generateOneDConvDiffParams(concminus1,
    gammanminus1,deltaZ,deltaT,theta,physconst,boundaryCond,boundaryextras
    );
[a(j0),b(j0),c(j0),d(j0),boundaryextras]=generateOneDConvDiffParams(
    concnminus1(j0),gammanminus1,deltaZ,deltaT,theta,physconst2,
    boundaryCond,boundaryextras,false);

% do Newton-Rhapson iteration here
% see http://compton.chem.ox.ac.uk/john/Thesis/2/2.html#H2.7.1
concnjminus1 = [ 0, concn(1:end-1) ];
concnjplus1 = [ concn(2:end), concin ];
F(zeroTojj+1) = ... % indexing is one higher to give room for gamma
    a.*concnjminus1 + b.*concn + c.*concnjplus1 - d;

F(1) = ... % gamma term
    gammanminus1 - gamma + ...
    deltaT*dadtn*concn(j0) - deltaT*dddtn;
Faccuracy = 0.00001;
FmeanAccuracy = 1e-10;

if max(abs(F)) < Faccuracy && mean(abs(F)) < FmeanAccuracy % this constants
    are a bit arbitrary
    break;
end
end

```

```

Ja(j0+1) = ... % df0/dgamma
    theta*(JkaTerm*concn(j0) - JkdTerm)/deltaZ;
Jb(1) = ... % df-1/dgamma
    -1 + (JkaTerm * concn(j0) - JkdTerm) * deltaT;
Jc(1) = ... % df-1/dc0
    dadtn*deltaT;

% Ju(z+1) = Ju(z) - F
uj = [gamma, concn];
ujminus1 = [0, uj(1:end-1)];
ujplus1 = [uj(2:end), concin];

% i suspect this can be reduced to calculate indices 2 and 1 only
Jd = Ja.*ujminus1 + Jb.*uj + Jc.*ujplus1 - F;

[enew, fnew, elow, flow] = thomasalgorithm(Ja, Jb, Jc, Jd, 0, concin, jIndex
    , revjIndex);

uj(1) = flow; % set by lower boundary condition
for jInd=1:(length(uj)-1) % first row is j, second is j+1
    uj(jInd+1) = enew(jInd)*uj(jInd) + fnew(jInd);
end
concn = uj(2:end);
gamma = uj(1);

if(n>100)
    msg = [ 'Exceeded 100 iterations. max(abs(F)) = ', num2str(max(abs(F))),
    '; mean(abs(F)) = ', num2str(mean(abs(F))), ' ' ];
    error('DW:iters',msg);
end
end

end
% done

function conc0 = flatConc0(elow, flow)
% elow and flow are the values of e and f on the lower boundary

% conc0 = conc-1 + 0
% conc0 = E.conc-1 + F
% conc-1 = E.conc-1 + F
% conc-1 = F/(1-E)
% conc0 = E.F/(1-E) + F
conc0 = (elow * flow)/(1-elow) + flow;
end

```

thomasalgorithm.m

Generates the e and f parameters of the Thomas Algorithm (i.e. the first sweep). The majority of this function is implemented in the C function

“thomasAlgorithmMainLoop”, however, the equivalent MATLAB code is shown here but commented out.

```
function [e, f, elow, flow] = thomasalgorithm(a, b, c, d, ebound, fbound, jIndex
    , revjIndex)
% * * * * *
% implements thomasalgorithm
%
% --- output ---
% e, f = matrices of output
% elow, flow = single values for the lowest j value of e and f, which
%           can't be stored in the standard matrices
% --- input ---
% a, b, c, d = matrices of input
% ebound, fbound = values of high e and f
% jIndex, revjIndex = function handle to map a j value to an array index
% * * * * *

jja = revjIndex(size(a));
jja = jja(2);

if (length(a) == length(b) && (length(b)==length(c)) && (length(c)==length(d))
    jj = jja;
else
    msg = [ 'size of input arrays a,b,c and d are not identical: a = ', size(a)
        , ' b = ', size(b), ' c = ', size(c), ' d = ', size(d) ];
    error(msg);
end

% lower boundary condition
lowerj = (revjIndex(1));

indexes = [ jIndex(jj:-1:lowerj+1); jIndex(jj-1:-1:lowerj) ]; % first row is j,
    second is j-1

% use marginly faster c version which just does the loop below
[e, f] = thomasalgorithmmainloop(a,b,c,d,indexes,ebound,fbound);
%e(jIndex(jj)) = ebound;
%f(jIndex(jj)) = fbound;

%for jInd=[ jIndex(jj:-1:lowerj+1); jIndex(jj-1:-1:lowerj) ] % first row is j,
    second is j-1
%     denom = c(jInd(1))*e(jInd(1)) + b(jInd(1));
%     e(jInd(2)) = -a(jInd(1))/denom;
%     f(jInd(2)) = (d(jInd(1)) - c(jInd(1)) * f(jInd(1)))/denom;
%end

jlowerj = jIndex(lowerj);
denom = c(jlowerj)*e(jlowerj) + b(jlowerj);
elow = -a(jlowerj)/denom;
flow = (d(jlowerj) - c(jlowerj) * f(jlowerj))/denom;
```

thomasalgorithmmainloop.c

```
#include "mex.h"
#include <math.h>

unsigned int rnd(double x)
{
    return floor(x+0.5);
}

void mexFunction(int nlhs, mxArray* plhs[],
    int nrhs, const mxArray* prhs[])
{
    const int A_IN = 0;
    const int B_IN = 1;
    const int C_IN = 2;
    const int D_IN = 3;
    const int INDEXES_IN = 4;
    const int EBOUND_IN = 5;
    const int FBOUND_IN = 6;
    const int E_OUT = 0;
    const int F_OUT = 1;
    mwSize sizeAm, sizeBm, sizeCm, sizeDm, sizeAn, sizeBn, sizeCn, sizeDn,
        sizeIndexesm, sizeIndexesn;
    double *e, *f;
    const double *a, *b, *c, *d, *indexes;
    unsigned int i;

    if (nrhs != (FBOUND_IN+1))
    {
        mexErrMsgTxt("7 input arguments required");
    }
    if (nlhs > 2)
    {
        mexErrMsgTxt("Too many output arguments (>2)");
    }

    sizeAm = mxGetM(prhs[A_IN]);
    sizeBm = mxGetM(prhs[B_IN]);
    sizeCm = mxGetM(prhs[C_IN]);
    sizeDm = mxGetM(prhs[D_IN]);
    sizeAn = mxGetN(prhs[A_IN]);
    sizeBn = mxGetN(prhs[B_IN]);
    sizeCn = mxGetN(prhs[C_IN]);
    sizeDn = mxGetN(prhs[D_IN]);
    if ( (sizeAm != sizeBm) || (sizeBm != sizeCm) || (sizeCm != sizeDm) || (sizeAn
        != sizeBn) || (sizeBn != sizeCn) || (sizeCn != sizeDn) )
    {
        mexErrMsgTxt("Sizes of A, B, C and D arrays are not equal");
    }
    if ( !mxIsDouble(prhs[A_IN]) || !mxIsDouble(prhs[B_IN]) || !mxIsDouble(prhs[
        C_IN]) || !mxIsDouble(prhs[D_IN]) ||
        mxIsComplex(prhs[A_IN]) || mxIsComplex(prhs[B_IN]) || mxIsComplex(prhs[
        C_IN]) || mxIsComplex(prhs[D_IN]) )
    {
        mexErrMsgTxt("A, B, C and D must be double and noncomplex");
    }
    sizeIndexesm = mxGetM(prhs[INDEXES_IN]);
    sizeIndexesn = mxGetN(prhs[INDEXES_IN]);
    if ( (sizeIndexesn != sizeAn-1) || (sizeIndexesm != 2) )
    {
        mexErrMsgTxt("Sizes of index array are wrong");
    }
}
```

```

}
if ( !mxIsDouble(prhs[INDEXES_IN]) || mxIsComplex(prhs[INDEXES_IN]) )
{
    mexErrMsgTxt("indexes must be uint32 and noncomplex");
}

/* make output variables now */
plhs[E_OUT] = mxCreateDoubleMatrix(sizeAm, sizeAn, mxREAL);
plhs[F_OUT] = mxCreateDoubleMatrix(sizeAm, sizeAn, mxREAL);

e = mxGetPr(plhs[E_OUT]);
f = mxGetPr(plhs[F_OUT]);
indexes = mxGetPr(prhs[INDEXES_IN]);
a = mxGetPr(prhs[A_IN]);
b = mxGetPr(prhs[B_IN]);
c = mxGetPr(prhs[C_IN]);
d = mxGetPr(prhs[D_IN]);

e[(rnd(indexes[0])-1)] = mxGetScalar(prhs[EBOUND_IN]);
f[(rnd(indexes[0])-1)] = mxGetScalar(prhs[FBOUND_IN]);

for (i = 0; i < sizeIndexesn; ++i)
{
    unsigned int idx = round(indexes[i*2])-1;
    unsigned int idxm1 = round(indexes[i*2+1])-1;
    double denom = c[idx]*e[idx] + b[idx];

    e[idxm1] = -a[idx]/denom;
    f[idxm1] = (d[idx] - c[idx] * f[idx])/denom;
}
}

```

partitionMicelles.m

```

function [concsOut] = partitionMicelles(concsIn, cmc)
% *****
% PARTITIONMICELLES
% - given a cmc and a set of monomer and monomer-in-micelle
% concentrations, produces an adjusted set of concentrations
%
% -- output --
% - concsOut = an n by 2 array, the first column of which contains
% monomer concentrations, and the second column of which contains
% micelle concentrations
%
% -- input --
% - concsIn = as concsOut
% - cmc = the value of the cmc (si units)
% *****

totalConcs = sum(concsIn, 2);

micelleconc = totalConcs - cmc; % concs is used as concsmonomers
abovecmc = micelleconc >= 0;
micelleconc(~abovecmc) = 0;
concsbelowcmc = totalConcs;

```

```

concsbelowcmc(abovecmc) = cmc;
concsOut = [ concsbelowcmc, micelleconc ];

```

E.2.2 Frumkin isotherm

This section contains functions relating to the Frumkin isotherm, both for equilibrium fitting and the kinetic model.

fitFrumkin.m

Attempts to fit a Frumkin isotherm, by varying the three parameters Γ_∞ , K_L and ω .

fitFrumkin.m

```

function [ gammainfty, Kl, omega ] = fitFrumkin(actualData, gammainfty)
% function [ gammainfty, Kl, omega ] = fitFrumkin(actualData)
%
% Outputs are the three parameters for the frumkin isotherm.
% There are no limits on the range omega can be (i.e. it can readily be >4)
% (the function also plots a graph showing the fit)
% "actualData" should be provided in two columns: concentration (mM) and surface
% coverage
%
% if gammainfty is provided then fitting is only done in Kl and omega

concs = actualData(:,1).';
gammass = actualData(:,2).';

if ~exist('gammainfty', 'var') || isempty(gammainfty)
    minimumgamma = 0.5*max(gammass); %1e-6;
    maximumgamma = 2*max(gammass); %1e-5;
    gammastep = (maximumgamma - minimumgamma)/50;
    testGammass = minimumgamma:gammastep:maximumgamma;
else
    testGammass = gammainfty;
end
rms = zeros(size(testGammass));
Kls = zeros(size(testGammass));
omegas = zeros(size(testGammass));
for n = 1:length(testGammass)
    g = testGammass(n);
    [ rms(n), Kls(n), omegas(n) ] = fitWithSetGamma(concs, gammass, g);
end

```

```

figure(1); plot(testGammas,rms);
xlabel('\Gamma'); ylabel('rms error');
figure(2); plot(testGammas,Kls)
ylabel('K_L');
figure(3); plot(testGammas,omegas)
ylabel('\omega')

[ minGammaRms, index ] = min(rms);
gammainfty = testGammas(index);
Kl = Kls(index);
omega = omegas(index);

logconcs = floor(log(min(concs))-1):0.2:ceil(log(max(concs))+1);
fconcs = exp(logconcs);
gs = equilibriumFrumkinGamma(fconcs,gammainfty,Kl,omega);
figure(4); plot(log(concs),gammas,'+');
xlabel('log(conc)'); ylabel('\Gamma');
hold on
plot(log(fconcs),gs);
hold off

function [rms,Kl,omega] = fitWithSetGamma(concs,gammas,gammaInfty)

KlFirstGuess = 10;
omegaFirstGuess = -1;

func = @(vec) difference(concs,gammas,vec(1),vec(2),gammaInfty);

[vec,fval,exitflags] = fminsearch(func,[KlFirstGuess,omegaFirstGuess]);

Kl = vec(1); omega = vec(2);
rms = difference(concs,gammas,Kl,omega,gammaInfty);

function [rms] = difference(concs,gammas,Kl,omega,gammaInfty)

predGamma = equilibriumFrumkinGamma(concs,gammaInfty,Kl,omega);

rms = sum((predGamma - gammas).^2);

```

equilibriumFrumkinGamma.m

Calculates the surface excess at equilibrium using the Frumkin isotherm.

```

function [ gamma, n ] = equilibriumFrumkinGamma(conc, gammainfty, KL, omega0)
% function [ gamma, n ] = equilibriumFrumkinGamma(conc, gammainfty, KL, omega0)
% conc is in mM

plusminus = 0.0001 * gammainfty;

gamma = [];
n = [];

if (omega0 > 4)

```

```

[s,e] = findfrumkintieline(omega0,gammainfty);
discontinuity = true;
tieConc = e/(gammainfty * (1 - e/gammainfty) * KL * exp(omega0*e/gammainfty));
else
discontinuity = false;
end
if(KL < 0)
warning('DW:badKL','KL for Frumkin isotherm is negative');
gamma = inf*ones(size(conc));
return
end
for c = conc % iterate over concentrations
startPoint = 0;
endPoint = gammainfty; % region to search over
if (discontinuity)
if (c>tieConc)
startPoint = e;
else
endPoint = s;
end
end
diffFun = @(gammaest) getRhs(gammaest, c, gammainfty, KL, omega0) - gammaest;

% if there's no sign change in the region between the start and end
% point it's usually because we're dealing with tiny fractions. Return
% the average and give a warning
if (diffFun(startPoint) >= 0) == (diffFun(endPoint) >= 0)
msg = ['No sign change in equilibriumFrumkinGamma between gamma = ',
num2str(startPoint), ' and ', num2str(endPoint), '.
The difference values at the points are: ', num2str(diffFun(
startPoint)), ' and ', num2str(diffFun(endPoint)), '
respectively.'];
warning('DW:nosignchange',msg);
gamma = [ gamma, (startPoint+endPoint)/2 ];
n = [ n, 0 ];
continue;
end

try
[g,fval,exitflag,output] = fzero(diffFun,[startPoint,endPoint],optimset(
'TolX',plusminus));
catch e
disp(['Error in equilibriumFrumkinGamma while searching between gamma = ',
num2str(startPoint), ' and ', num2str(endPoint), '.']);
disp(['The difference values at the start and end points are: ', num2str(
diffFun(startPoint)), ' and ', num2str(diffFun(endPoint)), '
respectively.']);
disp(['conc=', num2str(c)]);
disp(['and input arguments: gammainfty=', num2str(gammainfty), ', KL=',
num2str(KL), ', omega0=', num2str(omega0)]);
rethrow(e);
end

if (exitflag ~= 1)
warning('DW:nonconv',output.message);

```

```

end

gamma = [ gamma, g ];
n = [ n, output.iterations ];
end

function rhs = getRhs(gammaest, conc, gammainfty, KL, omega)
KfC = (KL * exp(omega * gammaest/gammainfty) * conc);
rhs = KfC .* gammainfty ./ (1 + KfC);

```

findfrumkintieline.m

Where $\omega > 4$, the Frumkin isotherm has a Maxwell construction with a tie line linking the upper and lower phases. The function finds the tie line. The internals of the function are slightly more complicated than they need to be, since they also calculates gradients, which are not used in the minimisation process.

```

function [startOfTie,endOfTie] = findfrumkintieline(omega,GammaInfty)

func = @(input) integratedAmount(input,omega,GammaInfty);

[minimum,maximum]=findSurfaceEnergyMinMax(omega);
startGuess = minimum/2;
endGuess = (maximum+1)/2;

if exist('OCTAVE_VERSION') ~= 0
    [output, fval] = fminsearch(func,[startGuess,endGuess]); % octave doesn't
    return 3rd argument
else
    [output, fval, exitflag] = fminsearch(func,[startGuess,endGuess]);
    if (exitflag ~= 1)
        disp(['omega: ', num2str(omega), ' exitflag: ', num2str(exitflag)])
    end
end

startOfTie = output(1)*GammaInfty;
endOfTie = output(2)*GammaInfty;

if isnan(startOfTie) || isnan(endOfTie)
    msg = ['Neither start of tie nor end of tie should be NaN. Values are ',
        num2str(startOfTie), ' and ', num2str(endOfTie)];
    msg = [msg, ' with omega=', num2str(omega), ' and gammaInfty=', num2str(
        GammaInfty)];
    error(msg);
end

function [g] = calculateSurfaceEnergy(x,K)

```

% note that we are missing gamma0, and the RT Gammainfty term

```

% since these cancel in our later calculations
g = log(1 - x) + (K/2)*x.^2;

function [val] = calculateIntegratedSurfaceEnergy(x,K,GammaInfty)

% again, ditching gamma0 and RT Gammainfty
a = 1./(x*GammaInfty); % area of molecule
val = a .* log(1-x) - log(1./x - 1)/GammaInfty - K*x/(2*GammaInfty);
% disp(val)

function [difference,gradient] = calculateIntegratedAmount(xl,xh,K,GammaInfty)
% using two values of x
%
differenceInA = 1./(xh*GammaInfty) - 1./(xl*GammaInfty);
ddifferenceInAdxl = 1./(xl^2*GammaInfty);
ddifferenceInAdxh = -1./(xh^2*GammaInfty);

lhs = calculateIntegratedSurfaceEnergy(xh,K,GammaInfty) -
    calculateIntegratedSurfaceEnergy(xl,K,GammaInfty);
dlhsdxl = (1/GammaInfty)*(1/xl^2 * log(1-xl) + K/2);
dlhsdxh = (-1/GammaInfty)*(1/xh^2 * log(1-xh) + K/2);

gx1 = calculateSurfaceEnergy(xl,K);
rhs = gx1.*differenceInA;
drhsdxl = dSurfaceEnergydx(xl,K)*differenceInA + gx1.*ddifferenceInAdxl;
drhsdxh = gx1.*ddifferenceInAdxh;
gxh = calculateSurfaceEnergy(xh,K);
otherRhs = gxh.*differenceInA;
dotherrhsdxh = dSurfaceEnergydx(xh,K)*differenceInA + gxh.*ddifferenceInAdxh;
dotherrhsdxl = gxh.*ddifferenceInAdxl;
errorInRhs = (rhs-otherRhs)^2;
errorInRhsdxl = 2*(rhs-otherRhs)*(drhsdxl-dotherrhsdxl);
errorInRhsdxh = 2*(rhs-otherRhs)*(drhsdxh-dotherrhsdxh);

lhsmrhs = (lhs - rhs)^2;
dlhsmrhsdxl = 2*(lhs-rhs)*(dlhsdxl-drhsdxl);
dlhsmrhsdxh = 2*(lhs-rhs)*(dlhsdxh-drhsdxh); % drhsdxh = 0;
difference = lhsmrhs + errorInRhs; % error in rhs is just to make it necessary
    for the surface tension to match at both ends
gradient = [dlhsmrhsdxl + derrorInRhsdxl, dlhsmrhsdxh + derrorInRhsdxh];
%dlhsmrhsdxl + derrorInRhsdxl

function [difference,gradient] = integratedAmount(input,K,GammaInfty)
% input(1) is xl, input(2) is difference to make xh
gradient = [nan,nan];
if (input(1) > 1) || (input(2) > 1) || (input(2) < input(1)) || (input(1) <= 0)
    difference = inf; % attempt to bound it
else
    [minimum,maximum] = findSurfaceEnergyMinMax(K);
    if (input(2) < maximum) || (input(1) > minimum)
        difference = inf; % we know we have to start at lower x than the minimum
        and go to higher x than the maximum
    else
        [difference,gradient] = calculateIntegratedAmount(input(1),input(2),K,
            GammaInfty);
    end
end

function [minimum,maximum] = findSurfaceEnergyMinMax(K)
% dg/dx = -1/(1-x)+K*x = 0

```

```
% -1 + Kx - Kx^2 = 0
% x = ((+/- sqrt(K^2 - 4K)) + K)/2K
maxMin1 = (sqrt(K^2 - 4*K) + K)/(2*K);
maxMin2 = (-sqrt(K^2 - 4*K) + K)/(2*K);
% the maximum is the one at highest x
maximum = max(maxMin1, maxMin2);
minimum = min(maxMin1, maxMin2);

function [gr] = dSurfaceEnergydx(x,K)
gr = -1./(1-x)+K*x;
```

setupFrumkinConsts.m

Optionally adds the tie line constants where necessary, for preparing the kinetic Frumkin constants for the convection-diffusion calculation.

```
function newConsts = setupFrumkinConsts(oldConsts)
% function newConsts = setupFrumkinConsts(oldConsts)
% calculates tieline starts and ends where necessary

omega = oldConsts.omegaa - oldConsts.omegad;
if (omega) > 4 % critical
    [s,e] = findfrumkintieline(omega, oldConsts.gammainfty);
    oldConsts.tielower = s;
    oldConsts.tieupper = e;
end

newConsts = oldConsts;
```

frumkinGamma.m

Used in the kinetics calculation to get an estimate of what the surface excess, Γ , should be based on the last timestep and the subsurface concentration.

```
function [ gamman, dadtn, dddtn, n ] = frumkinGamma(gammanminus1, conc0, %
    constants, deltaT, converginfty, converggamma)

% * * * * *
% frumkingamma
% - estimates the value of gamman in the frumkin isotherm by an
% iterative process
%
% --output--
% - gamman - answer
% - dadtn - flux to the surface dependent on timestep n
% - dddtn - flux from the surface dependent on timestep n
% - n - number of iterations
%
% --input--
```

```
% - gammanminus1 - gamman at last time step
% - conc0 - concentration at subsurface
% - constants - a structure containing the fields
% * gammainfty - \Gamma constant
% * ka, kd - rate constants
% * omegaa, omegad - omega parameters
% * setupFrumkinConsts should also be run on constants
% - deltaT - timestep
% - converginfty - convergence criteria (as a fraction of gammainfty)
% - converggamma - convergence criteria (as a fraction of the current gamma)
% * * * * *

ka = constants.ka; kd = constants.kd;
gammainfty = constants.gammainfty;
wa = constants.omegaa; wd = constants.omegad;

if conc0 < 0
    conc0 = 0;
    warning('DW: zeroconc', 'Concentration for Frumkin isotherm is below 0. %
        Setting to 0');
end

func = @(gamman) getRhsLhs(gamman, gammanminus1, conc0, constants, deltaT);

plusminusinfty = gammainfty * converginfty;
plusminusgamma = gammanminus1 * converggamma;

plusminus = max(plusminusinfty, plusminusgamma);

if ((constants.omegaa - constants.omegad) > 4) && ... % require critical %
    behavior
    (gammanminus1 > constants.tielower) && (gammanminus1 < constants.tieupper) %
        and to be within the 2 phase region

    fractionUpperPhase = (gammanminus1 - constants.tielower) / (constants. %
        tieupper - constants.tielower);
    explicitgamman = fractionUpperPhase * explicitGamma(constants.tieupper, conc0 %
        , constants, deltaT) + ...
        (1 - fractionUpperPhase) * explicitGamma(constants.tielower, conc0, %
        constants, deltaT);
else
    explicitgamman = explicitGamma(gammanminus1, conc0, constants, deltaT);
end

if abs(func(explicitgamman)) < plusminus % answer is good enough
    n = 0;
    gamman = explicitgamman;
else % search approach
    try
        [gamman, ignore, exitflags, outputinfo] = fzero(func, [0, gammainfty], %
            optimset('TolX', plusminus));
    catch e
        disp('Error while handling fzero for Frumkin. Dumping assorted stuff:');
        disp(['gammanminus1 = ', num2str(gammanminus1)]);
        disp(['conc0 = ', num2str(conc0)]);
        if strcmp(e.identifier, 'MATLAB: fzero: ValuesAtEndPtsSameSign')
            disp('End values error caught. Values are:');
            disp(['f(0): ', num2str(func(0))]);
            disp(['f(gammainfty): ', num2str(func(gammainfty))]);
        end
    end
```

```

    rethrow(e)
end
if (exitflags == 0)
    warning('DW:nonconv','Too many iterations used... returning gamman
    anyway');
end
n = outputinfo.iterations;
end
if ((constants.omegaa - constants.omegad) > 4) && ... % require critical
    behavior
    (gamman > constants.tielower) && (gamman < constants.tieupper) % and to be
    within the 2 phase region

    fractionUpperPhase = (gamman - constants.tielower) / (constants.tieupper -
    constants.tielower);
    dadtn = ka*(fractionUpperPhase*exp(wa*constants.tieupper/gammainfy)*
    gammainfy-constants.tieupper) + ...
    (1-fractionUpperPhase)*exp(wa*constants.tielower/gammainfy)*(gammainfy-
    constants.tielower);
    dddtn = kd*(fractionUpperPhase*exp(wd*constants.tieupper/gammainfy)*
    constants.tieupper + ...
    (1-fractionUpperPhase)*exp(wd*constants.tielower/gammainfy)*constants.
    tielower);
else
    dadtn = ka*exp(wa*gamman/gammainfy)*(gammainfy - gamman);
    dddtn = kd*exp(wd*gamman/gammainfy)*gamman;
end

function [a, dadg] = getRhsLhs(gamman, gammanminus1, conc0, constants, deltaT)
[a, dadg] = getRhs(gamman, gammanminus1, conc0, constants, deltaT);
a = a-gamman;
dadg = dadg-1;

function [rhs, drhsdg] = getRhs(gamman, gammanminus1, conc0, consts, deltaT,
    noCritical)
if (nargin < 6) || isempty(noCritical)
    noCritical = false;
end
if length(gamman) > 1;
    disp('oops')
end
if ~noCritical && ... % haven't disabled critical behaviour
    ((consts.omegaa - consts.omegad) > 4) && ... % require critical behavior
    (gamman > consts.tielower) && (gamman < consts.tieupper) % and to be within
    the 2 phase region

    fractionUpperPhase = (gamman - consts.tielower) / (consts.tieupper - consts.
    tielower);
    upperRhs = getRhs(consts.tieupper, gammanminus1, conc0, consts, deltaT, true);
    lowerRhs = getRhs(consts.tielower, gammanminus1, conc0, consts, deltaT, true);
    rhs = fractionUpperPhase * upperRhs + (1-fractionUpperPhase) * lowerRhs;

    drhsdg = 1/(consts.tieupper-consts.tielower)*(upperRhs-lowerRhs);
else
    rhs = deltaT * consts.ka * conc0 * exp(consts.omegaa*gamman/consts.
    gammainfy) .* (consts.gammainfy - gamman) ...
    - deltaT * consts.kd * exp(consts.omegad*gamman/consts.gammainfy) .*
    gamman + gammanminus1;
end

```

```

drhsdg = deltaT * consts.ka * conc0 * consts.omegaa * exp(consts.omegaa*
    gamman/consts.gammainfy) ...
- deltaT * consts.ka * conc0 * exp(consts.omegaa*gamman/consts.
    gammainfy) ...
- deltaT * consts.ka * conc0 * consts.omegaa * exp(consts.omegaa*gamman/
    consts.gammainfy)*gamman/consts.gammainfy ...
- deltaT * consts.kd * exp(consts.omegad*gamman/consts.gammainfy) ...
- deltaT * consts.kd * consts.omegad * exp(consts.omegad*gamman/consts.
    gammainfy)*gamman/consts.gammainfy;
end

function g = explicitGamma(gammanminus1, conc0, consts, deltaT)
g = (consts.ka * conc0 * exp(consts.omegaa * gammanminus1 / consts.gammainfy) *
    (consts.gammainfy - gammanminus1) ...
    - consts.kd * exp(consts.omegad * gammanminus1 / consts.gammainfy) *
    gammanminus1) * deltaT ...
    + gammanminus1;
end

```

calculatealpha.m

Can be used to look up the value of α for a given flow rate, based on the values calculated by Dąbroś and van de Ven;¹⁰¹ provides a warning when the requested flow rate is outside the range given but still returns an estimate.

```

function [ alpha, reynolds ] = calculatealpha(flowrate, volumeunit, pertimeunit, R)
% function [ alpha, reynolds ] = calculatealpha(flowrate, volumeunit, pertimeunit,
    R)
% calculates alpha from a specified flow rate
% volumeunit and pertimeunit are text strings
% R = radius in metres

% convert flow rate to m^3 s^-1
flowrate = flowrate * volumeconvertfactor(volumeunit) / timeconvertfactor(
    pertimeunit); % m^3 s^-1

U = flowrate / (pi*R*R); % m s^-1
kinematicvisc = 1e-6; % m^2 s^-1
reynolds = U * R / kinematicvisc;

% disp(['Re = ', num2str(reynolds)]);

knownRes = [ 2, 10, 20, 30, 35 ];
% knownalphabars = [ 0.0, 5, 9.5, 15, 19, 22 ]; % read crudely from Dabros and
    van der Venn (for h/R = 1.0)
knownalphabars = [ 0.0, 1.5, 5.5, 11.5, 15.5, 17 ]; % (for h/R = 1.7)

if reynolds > max(knownRes)
    warning('DW:alphacalc', ['Re of ', num2str(reynolds), ' is greater than the
    maximum known value of ', num2str(max(knownRes)) ]);
elseif reynolds < min(knownRes)
end

```

```

warning('DW:alphacalc', ['Re of ', num2str(reynolds), ' is less than the
maximum known value of ', num2str(min(knownRes)) ]);
end
% now add 0.0 to known Res (after the test)
knownRes = [0, knownRes ];
alphabar = interp1(knownRes, knownalphabars, reynolds, 'linear', 'extrap');
alpha = alphabar .* reynolds * kinematicvisc / (R*R*R); % m-1 s-1

function factor = volumeconvertfactor(volumeunit)
% returns the factor needed to get m3
volumeunit = lower(volumeunit);
switch volumeunit
case { 'ml' }
factor = 1e-6;
case { 'l', 'dm3' }
factor = 1e-3;
case { 'm3', 'm3' }
factor = 1;
otherwise
error(['volumeunit, ' is not recognised as a valid unit of volume.']);
end

function factor = timeconvertfactor(timeunit)
% returns factor needed to be in seconds
timeunit = lower(timeunit);
switch timeunit
case { 's' }
factor = 1;
case { 'min' }
factor = 60;
case { 'h', 'hour' }
factor = 60*60;
otherwise
error(['timeunit, ' is not recognised as a valid unit of time.']);
end

```

E.2.3 Regular solution theory

partitionMixtureOfTwo.m

Partitions a two surfactant mixture into monomers and mixed micelles. This function is too slow to be used in the main body of the convection-diffusion calculation, so an interpolated approximation is used instead.

```

function [concsOut] = partitionMixtureOfTwo(concsIn, constants)
% *****
% PARTITIONMIXTUREOFTWO

```

```

% - given concentrations for a two component mixture, partitions them
% appropriately between monomers and micelles
%
% -- output --
% - concsOut = an n by 4 array, the first two columns contain monomer
% concentrations for components 1 and 2 respectively, and the next two
% columns contain micelle concentrations for the two components
%
% -- input --
% - concsIn = as concsOut
% - constants = an object containing the fields
% + beta - the rst constant (note that mixedcmc or alpha cmc can be
% specified instead)
% + cmc1 and cmc2 - the mixed cmcs
% *****

% divide into the two components
totalConc1 = concsIn(:,1) + concsIn(:,3);
totalConc2 = concsIn(:,2) + concsIn(:,4);

totalConcBoth = totalConc1 + totalConc2;

alpha1 = totalConc1 ./ (totalConcBoth);
alpha1(isnan(alpha1)) = 0;

if any(alpha1 > 1 & alpha1 < 0)
error(['alpha1 outside specified range: ', num2str(alpha1)]);
end

x1 = rstX1AboveCMC(constants, totalConcBoth, alpha1);
if any(x1 > 1 & x1 < 0)
error(['x1 outside specified range: ', num2str(x1)]);
end

cmc = rstCMC(constants, alpha1);

lessThanCmc = totalConcBoth < cmc;

c1mon = x1.*exp(constants.beta*(1-x1).^2)*constants.cmc1;
belowCmc1 = alpha1.*totalConcBoth;
c1mon(lessThanCmc) = belowCmc1(lessThanCmc);

c2mon = (1-x1).*exp(constants.beta*x1.^2)*constants.cmc2;
belowCmc2 = (1-alpha1).*totalConcBoth;
c2mon(lessThanCmc) = belowCmc2(lessThanCmc);

c1mic = x1.*(totalConcBoth-(c1mon+c2mon));
c2mic = (1-x1).(totalConcBoth-(c1mon+c2mon));

concsOut = zeros(size(concsIn));

concsOut(:,1) = c1mon;
concsOut(:,2) = c2mon;
concsOut(:,3) = c1mic;
concsOut(:,4) = c2mic;

```



```

% --- input ---
% - concsIn = as concsOut
% - constants = an object containing the fields
%   + beta - the rst constant (note that mixedcmc or alpha cmc can be
%   specified instead)
%   + cmc1 and cmc2 - the mixed cmcs
% - takes the list of concs and alphas used to produce the x1s and cmcs
% given as big x by y arrays (must be linearly spaced)
% - the x1s as an x by y array
% - cmcs as a vector y long
% * * * * *

% divide into the two components
totalConc1 = concsIn(:,1) + concsIn(:,3);
totalConc2 = concsIn(:,2) + concsIn(:,4);

totalConcBoth = totalConc1 + totalConc2;

alpha1 = totalConc1 ./ (totalConcBoth);
alpha1(isnan(alpha1)) = 0;

if any(alpha1 > 1 | alpha1 < 0)
    msg = 'alpha1 outside specified range.';
    maximum = max(alpha1);
    if maximum > 1
        msg = [ msg, ' Maximum is ', num2str(maximum), '.' ];
    end
    minimum = min(alpha1);
    if minimum < 0
        msg = [ msg, ' Minimum is ', num2str(minimum), '.' ];
    end
    warning('DW:InvalidAlpha1',msg);
    alpha1(alpha1 > 1) = 1;
    alpha1(alpha1 < 0) = 0;
end
if max(max(totalConcBoth)) > max(max(concArray))
    error(['Concentration ', max(max(totalConcBoth)), ' is greater than '
precomputed array concentration ', max(max(concArray)) ])
end

if exist('OCTAVEVERSION') ~= 0
    method = 'linear'; % octave seems to have issues with spline
else
    method = '*spline';
end
x1 = interp2(concArray, alphaArray, x1s, totalConcBoth, alpha1, method);
if any(x1 > 1 | x1 < 0)
    warning('DW:x1range', 'x1 outside 0 to 1 range (but conc is probably small and
thus it may not matter).')
end

% select the direction with the increasing values
if alphaArray(1,1) ~= alphaArray(end,1)
    alphaVect = alphaArray(:,1);
else
    alphaVect = alphaArray(1,:);
end
cmc = interp1(alphaVect, cmcs, alpha1, '*spline');

```

```

if any(any(isnan(cmc)))
    error('cmc isnan')
end

lessThanCmc = totalConcBoth < cmc;

% same as partitionMixtureOfTwo
c1mon = x1.*exp(constants.beta*(1-x1).^2)*constants.cmc1;
belowCmc1 = alpha1.*totalConcBoth;
c1mon(lessThanCmc) = belowCmc1(lessThanCmc);

c2mon = (1-x1).*exp(constants.beta*x1.^2)*constants.cmc2;
belowCmc2 = (1-alpha1).*totalConcBoth;
c2mon(lessThanCmc) = belowCmc2(lessThanCmc);

c1mic = x1.*(totalConcBoth-(c1mon+c2mon));
c2mic = (1-x1).(totalConcBoth-(c1mon+c2mon));

concsOut = zeros(size(concsIn));

concsOut(:,1) = c1mon;
concsOut(:,2) = c2mon;
concsOut(:,3) = c1mic;
concsOut(:,4) = c2mic;

if any(any(any(isnan(concsOut)))) || any(any(any(isinf(concsOut))))
    errorMsg = ['Infinite or NaN concentration in concOut', num2str(reshape(
concsOut,1,[]))];
    if any(any(isnan(concsIn))) || any(any(isinf(concsIn)))
        errorMsg = [ errorMsg, ' (present in concsIn too)'];
    end
    errorMsg = [errorMsg, '. totalConc = ', num2str(totalConc1.)];
    error(errorMsg);
end

```

prepareRSTInterpolations.m

Prepares the two arrays used by “partitionMixtureOfTwoInterpolation.m”. This is run once, at the start of the calculation.

```

function [concArray, alphaArray, x1s, cmcs] = prepareRSTInterpolations(constants,
maxConc, concSteps, alphaSteps)
% * * * * *
% * prepareRSTInterpolations:
% * prepares a precalculated grid of x1 values using RST to be
% * interpolated between quickly at runtime
% *
% * --- OUTPUT ---
% * - list of concs and alphas used to produce the x1s and cmcs
% * given as big x by y arrays (linearly spaced)
% * - the x1s as an x by y array
% * - cmcs as a vector y long
% *
% * --- INPUT ---

```

```

% * - constants = an object containing the fields
% * + beta - the rst constant (note that mixedcmc or alpha cmc can be
% * specified instead)
% * + cmc1 and cmc2 - the mixed cmcs
% * - maxConc - maximum concentration to calculate
% * - concSteps, alphaSteps - number of steps to use for conc and alpha
% *
% * * * * *
%
[ concArray, alphaArray ] = meshgrid(linspace(0,maxConc,concSteps), linspace(
(0,1,alphaSteps));

x1s = rstX1AboveCMC(constants,concArray,alphaArray);
cmcs = rstCMC(constants, linspace(0,1,alphaSteps));

```

rstX1atCMC.m

```

function [x1] = rstX1atCMC(constants,alpha, stepsize)
% out:
% * x1 - the mixed micelle composition
% in:
% * constants, an object containing the fields
% * + beta - the rst constant (note that mixedcmc or alpha cmc can be
% * specified instead)
% * + cmc1 and cmc2 - the mixed cmcs
% * alpha - composition
% * stepsize - (optional), the accuracy to which x1 should be calculated

% assume correct constants specified
if ~exist('stepsize','var') || isempty(stepsize)
    stepsize = 1e-4;
end

x1 = zeros(size(alpha));

for n = 1:length(alpha) % cycle through alphas (and associated x1s)
    al = alpha(n);
    if al == 0;
        x1best = 0;
    elseif al == 1;
        x1best = 1;
    else
        fun = @(x1) newX1calc(constants.beta, constants.cmc1, constants.cmc2, x1,
            , al);

        x1best = fminbnd(fun, stepsize, 1-stepsize, optimset('TolX', stepsize));
    end
    x1(n) = x1best;
end

function [diff] = newX1calc(beta, cmc1, cmc2, x1test, alpha)
% Holland and Rubingh, J. Phys Chem. 87, p1984, eqn 13,16,17
newX1 = alpha*cmc2*exp(beta*x1test^2)*(1-x1test)/(cmc1*(1-alpha)*exp(beta*(1-x1test)^2));
diff = abs(x1test - newX1);

```

E.2.4 Visualisation routines

Although not directly related to the performing the convection-diffusion calculation, it is still useful to be able to visualise the results.

plotTimeSeries.m

Displays a two graph layout showing the change in surface excess and concentration profile with time. Optionally returns a frame-by-frame recording that can be saved as a video file. This function is unlikely to work with Octave.

```

function [mov] = plotTimeSeries(concs, gamma, delaytime, frameskip, deltaZ,
    deltaT, legendstr)

% plots a time series
% -- input --
% - concs = concentrations
% - gamma = surface excess
% - delaytime = pause between each concentration (s)
% - number of frames to step forward each time (optional)
% - deltaZ (optional, for correct labeling of axes)
% - deltaT (optional, for correct labeling of axes)
% - legendstr (array of strings for the legend)

if nargin < 4 || isempty(frameskip)
    frameskip = 1;
end
if nargin < 5 || isempty(deltaZ)
    deltaZ = 1;
end
if nargin < 6 || isempty(deltaT)
    deltaT = 1;
end
if nargin < 7 || isempty(legendstr)
    legendstr = {'show'};
end

[n, m, o] = size(concs);

idx = 1; % set this at the start of the sequence
dying = false;
paused = false;

xaxis = (1:m)*deltaZ;
timeaxis = (1:n)*deltaT;

bottom = 0.01;
height = 0.05;

```

```

gap = 0.03;
position = 0.02;
smallButtonSize = 0.07;
largeButtonSize = 0.2;

fig = figure;

% restartButton
uicontrol(fig, 'Style', 'pushbutton', 'String', '<', ...
    'Units', 'normalized', 'Position', [position, bottom, smallButtonSize, height], ...
    'Callback', @restart_Callback);
position=position+smallButtonSize+gap;

% skipLeftButton
uicontrol(fig, 'Style', 'pushbutton', 'String', '<<', ...
    'Units', 'normalized', 'Position', [position, bottom, smallButtonSize, height], ...
    'Callback', @skipLeft_Callback);
position=position+smallButtonSize+gap;

% leftButton
uicontrol(fig, 'Style', 'pushbutton', 'String', '<', ...
    'Units', 'normalized', 'Position', [position, bottom, smallButtonSize, height], ...
    'Callback', @left_Callback);
position=position+smallButtonSize+gap;

% stopButton
uicontrol(fig, 'Style', 'pushbutton', 'String', 'stop', ...
    'Units', 'normalized', 'Position', [position, bottom, largeButtonSize, height], ...
    'Callback', @stop_Callback);
position=position+largeButtonSize+gap;

% pauseButton
pauseButton = uicontrol(fig, 'Style', 'pushbutton', 'String', 'pause', ...
    'Units', 'normalized', 'Position', [position, bottom, largeButtonSize, height], ...
    'Callback', @pause_Callback);
position=position+largeButtonSize+gap;

% rightButton
uicontrol(fig, 'Style', 'pushbutton', 'String', '>', ...
    'Units', 'normalized', 'Position', [position, bottom, smallButtonSize, height], ...
    'Callback', @right_Callback);
position=position+smallButtonSize+gap;

% skipRightButton
uicontrol(fig, 'Style', 'pushbutton', 'String', '>>', ...
    'Units', 'normalized', 'Position', [position, bottom, smallButtonSize, height], ...
    'Callback', @skipRight_Callback);

maxconc = max(max(max(concs)));
minconc = min(min(min(concs)));
if maxconc==0
    maxconc=1;

```

```

end

if nargin >= 1 % record frames
    frameno = 1;
    mov(1:(n/frameskip)) = getframe(gcf); % preallocate
end

function [conccaxis, gammaaxis] = drawInitially()
    subplot(2,3,3); plot(0,zeros(1,0)); legend(legendstr{:});
    axis off
    conccaxis = subplot(2,3,4:6);
    gammaaxis = subplot(2,3,1:2);
end

[conccaxis, gammaaxis] = drawInitially();

function redraw()
    subplot(conccaxis);
    plottedconcs = squeeze(concs(idx*frameskip,:,:));
    plot(xaxis, plottedconcs); axis([0 max(xaxis) minconc maxconc]);
    subplot(gammaaxis);
    plot(timeaxis, gamma); %legend('show')
    hold on;
    plot(idx*frameskip*deltaT, gamma(idx*frameskip,:), '+r');
    hold off;
end

while true
    try
        idx = min(n/frameskip, idx + 1);
        redraw();
        pause(delaytime);
        if nargin >= 1 % record frames
            mov(idx) = getframe(gcf);
            frameno = frameno+1;
        end
    catch e
        % this handles the case where the figure is closed by swallowing
        % the error and quitting
        if strcmp(e.identifier, 'MATLAB:subplot:InvalidAxesHandle')
            break;
        else
            rethrow(e)
        end
    end

    while paused % wait
        pause(0.1);
        if dying
            break;
        end
        if ~ishandle(fig) % it's been closed
            break;
        end
    end

    if dying
        clf(fig); % get rid of buttons
        [conccaxis, gammaaxis] = drawInitially();
        redraw();
        break;
    end
end
end

```

```

function restart_Callback(hObject,eventdata)
    resetPosition(1);
end

function stop_Callback(hObject,eventdata)
    dying = true;
end

function pause_Callback(hObject,eventdata,specifyPause)
    if exist('specifyPause','var') && ~isempty(specifyPause) && islogical(~
        specifyPause)
        paused = specifyPause;
    else
        paused = ~paused; % toggle
    end
    if paused
        set(hObject,'String','resume');
    else
        set(hObject,'String','pause');
    end
end

function skipLeft_Callback(hObject,eventdata)
    resetPosition(idx - ceil(n/(10*frameskip)));
end

function left_Callback(hObject,eventdata)
    resetPosition(idx - 1);
end

function right_Callback(hObject,eventdata)
    resetPosition(idx + 1);
end

function skipRight_Callback(hObject,eventdata)
    resetPosition(idx + ceil(n/(10*frameskip)));
end

function resetPosition(newIdx)
    % tests if idx is valid
    idx = min(n/frameskip,max(1,newIdx));
    pause_Callback(pauseButton,[],true);
    redraw();
end
end

```

Bibliography

- [1] E. Kissa, *Textile Research Journal*, 1996, **66**, 660–668.
- [2] S. Basu, J. Luthra and K. D. P. Nigam, *Journal of Environmental Science and Health, Part B: Pesticides, Food Contaminants, and Agricultural Wastes*, 2002, **37**, 331–344.
- [3] D. Beattie, M. Lidström Larsson and A. R. Holmgren, *Vibrational Spectroscopy*, 2006, **41**, 198–204.
- [4] N. Chen and B. Bhushan, *Journal of Microscopy*, 2006, **221**, 203–215.
- [5] D. Vollhardt and G. Czichocki, *Langmuir*, 1990, **6**, 317–322.
- [6] J. R. Lu, I. P. Purcell, E. M. Lee, E. A. Simister, R. K. Thomas, A. R. Rennie and J. Penfold, *Journal of Colloid and Interface Science*, 1995, **174**, 441–455.
- [7] M. J. Rosen, *Surfactants and Interfacial Phenomina*, John Wiley & Sons Inc., Hoboken, New Jersey, 3rd edn., 2004.
- [8] L. Quintero, *Journal of Dispersion Science and Technology*, 2002, **23**, 393–404.
- [9] M. A. Bos and T. van Vliet, *Advances in Colloid and Interface Science*, 2001, **91**, 437–471.
- [10] W. H. Briscoe, S. Titmuss, F. Tiberg, R. K. Thomas, D. J. McGillivray, and J. Klein, *Nature*, 2006, **444**, 191–194.
- [11] J. N. Israelachvii, *Intermolecular and Surface Forces*, Academic Press, 3rd edn., 2011.
- [12] P. Atkins and J. de Paula, *Atkins' Physical Chemistry*, Oxford University Press, 7th edn., 2002.
- [13] M. J. Rosen, in *Surfactants and Interfacial Phenomina*, John Wiley & Sons Inc., Hoboken, New Jersey, 3rd edn., 2004, ch. 1. Characteristic Features of Surfactants.
- [14] D. Colegate, *Ph.D. thesis*, University of Durham, 2009.
- [15] R. Nagarajan and E. Ruckenstein, *Langmuir*, 1991, **7**, 2934–2969.
- [16] E. A. G. Aniansson and S. N. Wall, *The Journal of Physical Chemistry*, 1974, **78**, 1024–1030.

- [17] E. A. G. Aniansson and S. N. Wall, *The Journal of Physical Chemistry*, 1975, **79**, 857–858.
- [18] E. A. G. Aniansson, S. N. Wall, M. Almgren, H. Hoffmann, I. Kielmann, W. Ulbricht, R. Zana, J. Lang and C. Tondre, *The Journal of Physical Chemistry*, 1976, **80**, 905–922.
- [19] R. Atkin, V. S. J. Craig, E. J. Wanless and S. Biggs, *Advances in Colloid and Interface Science*, 2003, **103**, 219–304.
- [20] S. Paria and K. C. Khilar, *Advances in Colloid and Interface Science*, 2004, **110**, 75–95.
- [21] S. K. Parida, S. Dash, S. Patel and B. Mishra, *Advances in Colloid and Interface Science*, 2006, **121**, 77–110.
- [22] R. Atkin, J. Eastoe, E. J. Wanless and C. D. Bain, in *Dynamics of Surfactant Self-Assemblies*, ed. R. Zana, Taylor & Francis, 2005, vol. 125, ch. Dynamics of Adsorption of Cationic Surfactants at Air–Water and Solid–Liquid Interfaces.
- [23] R. Zhang and P. Somasundaran, *Advances in Colloid and Interface Science*, 2006, **123–126**, 213–229.
- [24] G. H. Bolt, *The Journal of Physical Chemistry*, 1957, **61**, 1166–1169.
- [25] B. H. Bijsterbosch, *Journal of Colloid and Interface Science*, 1974, **47**, 186–198.
- [26] Y. Gao, J. Du and T. Gu, *J. Chem. Soc., Faraday Trans. 1*, 1987, **83**, 2671–2679.
- [27] P. Somasundaran and D. W. Fuerstenau, *The Journal of Physical Chemistry*, 1966, **70**, 90–96.
- [28] A. Fan, P. Somasundaran and N. J. Turro, *Langmuir*, 1997, **13**, 506–510.
- [29] P. Chandar, P. Somasundaran and N. J. Turro, *Journal of Colloid and Interface Science*, 1987, **117**, 31–46.
- [30] T. P. Goloub, L. K. Koopal, B. H. Bijsterbosch and M. P. Sidorova, *Langmuir*, 1996, **12**, 3188–3194.
- [31] A. de Keizer, M. Böhmer, T. Mehrian and L. Koopal, *Colloids and Surfaces*, 1990, **51**, 339–357.
- [32] T. P. Goloub and L. K. Koopal, *Langmuir*, 1997, **13**, 673–681.
- [33] S. Manne and H. E. Gaub, *Science*, 1995, **270**, 1480–1482.
- [34] S. B. Velegol, B. D. Fleming, S. Biggs, E. J. Wanless and R. D. Tilton, *Langmuir*, 2000, **16**, 2548–2556.
- [35] V. Subramanian and W. A. Ducker, *Langmuir*, 2000, **16**, 4447–4454.
- [36] C. Ström, P. Hansson, B. Jönsson and O. Söderman, *Langmuir*, 2000, **16**, 2469–2474.

- [37] J. C. Schulz, G. G. Warr, P. D. Butler and W. A. Hamilton, *Phys. Rev. E*, 2001, **63**, 041604.
- [38] P. Somasundaran, J. T. Kunjappu, C. V. Kumar, N. J. Turro and J. K. Barton, *Langmuir*, 1989, **5**, 215–218.
- [39] P. Chandar, P. Somasundaran, K. C. Waterman and N. J. Turro, *The Journal of Physical Chemistry*, 1987, **91**, 148–150.
- [40] E. Tyrode, M. W. Rutland and C. D. Bain, *Journal of the American Chemical Society*, 2008, **130**, 17434–17445.
- [41] P. L. Hayes, E. H. Chen, J. L. Achtyl and F. M. Geiger, *The Journal of Physical Chemistry A*, 2009, **113**, 4269–4280.
- [42] P. L. Hayes, A. R. Keeley and F. M. Geiger, *The Journal of Physical Chemistry B*, 2010, **114**, 4495–4502.
- [43] R. Atkin, V. S. J. Craig and S. Biggs, *Langmuir*, 2000, **16**, 9374–9380.
- [44] R. Atkin, V. S. J. Craig and S. Biggs, *Langmuir*, 2001, **17**, 6155–6163.
- [45] R. Atkin, V. S. J. Craig, E. J. Wanless and S. Biggs, *Journal of Colloid and Interface Science*, 2003, **266**, 236–244.
- [46] R. Atkin, V. S. J. Craig, E. J. Wanless and S. Biggs, *Journal of Physical Chemistry B*, 2003, **107**, 2978–2985.
- [47] S. C. Clark and W. A. Ducker, *The Journal of Physical Chemistry B*, 2003, **107**, 9011–9021.
- [48] O. Theodoly, L. Cascão-Pereira, V. Bergeron and C. J. Radke, *Langmuir*, 2005, **21**, 10127–10139.
- [49] T. D. Curwen, J. A. Warner, C. D. Bain, R. G. Compton and J. K. Eve, *Journal of Physical Chemistry C*, 2007, **111**, 12289–12304.
- [50] C. Hodges, S. Biggs and L. Walker, *Langmuir*, 2009, **25**, 11503–11508.
- [51] A. Khan, W. A. Ducker and M. Mao, *The Journal of Physical Chemistry B*, 2006, **110**, 23365–23372.
- [52] R. F. Tabor, J. Eastoe and P. J. Dowding, *Journal of Colloid and Interface Science*, 2010, **346**, 424–428.
- [53] S. C. Howard and V. S. J. Craig, *Soft Matter*, 2009, **5**, 3061–3069.
- [54] J.-F. Liu, G. Min and W. A. Ducker, *Langmuir*, 2001, **17**, 4895–4903.
- [55] L. M. Grant, F. Tiberg and W. A. Ducker, *The Journal of Physical Chemistry B*, 1998, **102**, 4288–4294.
- [56] S. K. Singh and S. M. Notley, *The Journal of Physical Chemistry B*, 2010, **114**, 14977–14982.

- [57] P. K. Misra, B. K. Mishra and P. Somasundaran, *Journal of Colloid and Interface Science*, 2003, **265**, 1–8.
- [58] F. Tiberg, B. Jönsson and B. Lindman, *Langmuir*, 1994, **10**, 3714–3722.
- [59] F. Tiberg, *J. Chem. Soc., Faraday Trans.*, 1996, **92**, 531–538.
- [60] J. Brinck and F. Tiberg, *Langmuir*, 1996, **12**, 5042–5047.
- [61] J. Brinck, B. Jönsson and F. Tiberg, *Langmuir*, 1999, **15**, 7719–7724.
- [62] B. R. Postmus, F. A. M. Leermakers, L. K. Koopal and M. A. Cohen Stuart, *Langmuir*, 2007, **23**, 5532–5540.
- [63] J. Brinck, B. Jönsson and F. Tiberg, *Langmuir*, 1998, **14**, 1058–1071.
- [64] J. Brinck, B. Jönsson and F. Tiberg, *Langmuir*, 1998, **14**, 5863–5876.
- [65] C. Stubenrauch, V. B. Fainerman, E. V. Aksenenko and R. Miller, *The Journal of Physical Chemistry B*, 2005, **109**, 1505–1509.
- [66] J. H. Clint, *J. Chem. Soc., Faraday Trans.*, 1975, **71**, 1327–1334.
- [67] P. M. Holland and D. N. Rubingh, *Journal of Physical Chemistry*, 1983, **87**, 1984–1990.
- [68] H. Lange and K. H. Beck, *Kolloid Z. Z. Polym.*, 1973, **251**, 424.
- [69] P. Letellier, A. Mayaffre and M. Turmine, *Journal of Colloid and Interface Science*, 2011, **354**, 248–255.
- [70] J. D. Hines, *Current Opinion in Colloid & Interface Science*, 2001, **6**, 350–356.
- [71] D. N. Rubingh, in *Solution chemistry of surfactants*, ed. K. L. Mittal, Plenum, New York, 1979, p. 337.
- [72] J. R. Lu, R. K. Thomas and J. Penfold, *Advances in Colloid and Interface Science*, 2000, **84**, 143 – 304.
- [73] M. J. Hey, J. W. MacTaggart and C. H. Rochester, *J. Chem. Soc., Faraday Trans. 1*, 1985, **81**, 207–213.
- [74] H. P. M., in *Effects of Surfactant Structure on the Thermodynamics of Mixed Micellization*, American Chemical Society, 1984, ch. 11, pp. 141–151.
- [75] L. Huang and P. Somasundaran, *Langmuir*, 1997, **13**, 6683–6688.
- [76] M. J. Rosen, in *Surfactants and Interfacial Phenomina*, John Wiley & Sons Inc., Hoboken, New Jersey, 3rd edn., 2004, pp. 380–382.
- [77] M. J. Rosen and X. Y. Hua, *Journal of Colloid and Interface Science*, 1982, **86**, 164–172.
- [78] C. D. Eads and L. C. Robosky, *Langmuir*, 1999, **15**, 2661–2668.
- [79] C. C. Ruiz and J. Aguiar, *Molecular Physics*, 1999, **97**, 1095–1103.

- [80] J. D. Hines, R. K. Thomas, P. R. Garrett, G. K. Rennie and J. Penfold, *Journal of Physical Chemistry B*, 1998, **102**, 8834–8846.
- [81] T. Tominaga, in *Mixed Surfactant Systems*, ed. M. Abe and J. F. Scamehorn, Marcel Dekker, New York, 2nd edn., 2005, vol. 124, pp. 135–164.
- [82] D. G. Leaist, *J. Colloid Interface Sci.*, 1986, **111**, 230–239.
- [83] K. MacEwan and D. Leaist, *Journal of Physical Chemistry B*, 2002, **106**, 10296–10300.
- [84] R. Heger, C. Ikier, L. Belkoura, D. Woermann and H. Klein, *J. Chem. Soc. Faraday Trans.*, 1995, **91**, 3385–3388.
- [85] D. G. Leaist and L. Hao, *Journal of Physical Chemistry*, 1993, **97**, 7763–7768.
- [86] D. G. Leaist and L. Hao, *J. Chem. Soc. Faraday Trans.*, 1993, **89**, 2775–2782.
- [87] D. G. Leaist and S. M. Abdu, *Journal of Chemical & Engineering Data*, 2001, **46**, 922–926.
- [88] Y. Nogami, M. Iwata and T. Tominaga, *Journal of Molecular Liquids*, 2005, **119**, 83–87.
- [89] D. M. Colegate and C. D. Bain, *Phys. Rev. Lett.*, 2005, **95**, 198302.
- [90] R. Miller, P. Joos and V. B. Fainerman, *Advances in Colloid and Interface Science*, 1994, **49**, 249–302.
- [91] C. D. Bain, *Advances in Colloid and Interface Science*, 2008, **144**, 4–12.
- [92] P. A. Cuypers, J. W. Corsel, M. P. Janssen, J. M. Kop, W. T. Hermens and H. C. Hemker, *Journal of Biological Chemistry*, 1983, **258**, 2426–2431.
- [93] J. W. Corsel, G. M. Willems, J. M. M. Kop, P. A. Cuypers and W. T. Hermens, *Journal of Colloid and Interface Science*, 1986, **111**, 544–554.
- [94] T. D. Curwen, *Ph.D. thesis*, University of Oxford, 2006.
- [95] A. A. Levchenko, B. P. Argo, R. Vidu, R. V. Talroze and P. Stroeve, *Langmuir*, 2002, **18**, 8464–8471.
- [96] J. A. Cooper and R. G. Compton, *Electroanalysis*, 1998, **10**, 141–155.
- [97] E. M. Furst, E. S. Pagac and R. D. Tilton, *Industrial & Engineering Chemistry Research*, 1996, **35**, 1566–1574.
- [98] E. S. Pagac, D. C. Prieve and R. D. Tilton, *Langmuir*, 1998, **14**, 2333–2342.
- [99] B. K. Lok, Y.-L. Cheng and C. R. Robertson, *Journal of Colloid and Interface Science*, 1983, **91**, 104–116.
- [100] J. C. Dijt, M. A. Cohen Stuart, J. E. Hofman and G. J. Fleer, *Colloids and Surfaces*, 1990, **51**, 141–158.

- [101] T. Dąbrosz and T. G. M. van de Ven, *Colloid & Interface Science*, 1983, **261**, 694–707.
- [102] N. V. Rees, O. V. Klymenko, B. A. Coles and R. G. Compton, *The Journal of Physical Chemistry B*, 2003, **107**, 13649–13660.
- [103] J. L. Melville, N. Simjee, P. R. Unwin, B. A. Coles and R. G. Compton, *J. Phys. Chem. B*, 2002, **106**, 10424–10431.
- [104] D. Xu, C. Hodges, Y. Ding, S. Biggs, A. Brooker and D. York, *Langmuir*, 2010, **26**, 18105–18112.
- [105] H. D. Bijsterbosch, M. A. Cohen Stuart and G. J. Fleer, *Macromolecules*, 1998, **31**, 9281–9294.
- [106] L. H. Torn, L. K. Koopal, A. de Keizer and J. Lyklema, *Langmuir*, 2005, **21**, 7768–7775.
- [107] C. Geffroy, M. A. Cohen Stuart, K. Wong, B. Cabane and V. Bergeron, *Langmuir*, 2000, **16**, 6422–6430.
- [108] K. Sakai, E. Smith, G. Webber, C. Schatz, E. Wanless, V. Butun, S. Armes and S. Biggs, *Journal of Physical Chemistry B*, 2006, **110**, 14744–14753.
- [109] K. Sakai, M. Vamvakaki, E. G. Smith, E. J. Wanless, S. P. Armes and S. Biggs, *Journal of Colloid and Interface Science*, 2008, **317**, 383–394.
- [110] E. Bitziou, N. C. Rudd, M. A. Edwards and P. R. Unwin, *Analytical Chemistry*, 2006, **78**, 1435–1443.
- [111] T. D. Curwen, C. D. Bain and J. K. Eve, *Journal of Physical Chemistry C*, 2007, **111**, 12305–12314.
- [112] J. Gooding, B. Coles and R. Compton, *Journal of Physical Chemistry B*, 1997, **101**, 175–181.
- [113] J. Gooding, C. Brennan, J. Atherton, B. Coles and R. Compton, *Journal of Physical Chemistry B*, 1997, **101**, 182–188.
- [114] E. Hecht, *Optics*, Addison-Wesley, 2nd edn., 1987.
- [115] S. G. Lipson, H. Lipson and D. S. Tannhauser, *Optical Physics*, Cambridge University Press, 3rd edn., 1995.
- [116] M. Born and E. Wolf, *Principles of Optics*, Pergamon Press, 6th edn., 1980.
- [117] D. A. Beattie, S. Haydock and C. D. Bain, *Vibrational Spectroscopy*, 2000, **24**, 109–123.
- [118] P. R. Greene, *Ph.D. thesis*, University of Oxford, 2003.
- [119] P. W. Atkins and R. S. Friedman, in *Molecular Quantum Mechanics*, Oxford University Press, 3rd edn., 1997, ch. 10. Molecular rotations and vibrations.

- [120] P. Atkins and J. de Paula, in *Atkins' Physical Chemistry*, Oxford University Press, 7th edn., 2002, ch. 16. Spectroscopy 1: rotational and vibrational spectra.
- [121] C. V. Raman and K. S. Krishnan, *Nature*, 1928, **121**, 501–502.
- [122] C. V. Raman and K. S. Krishnan, *Proc. Roy. Soc. (London) A*, 1928, **122**, 23–35.
- [123] A. Smekal, *Naturwissenschaften*, 1923, **11**, 873–875.
- [124] B. Crawford, Jr. and D. Swanson, in *Infrared and Raman Spectroscopy, Part A*, ed. E. G. Brame, Jr. and J. G. Grasselli, Marcel Dekker, Inc., 1976, ch. 1. An Introduction to Molecular Vibrations.
- [125] D. A. Long, *The Raman Effect: A Unified Treatment of the Theory of Raman Scattering by Molecules*, John Wiley & Sons, Ltd., 2002.
- [126] G. Davidson, in *Group Theory For Chemists*, Macmillan Education Ltd., 1991, ch. 8. Group theory and vibrational spectroscopy.
- [127] S. G. Kazarian and K. L. A. Chan, *Applied Spectroscopy*, 2010, **64**, 135A–151A.
- [128] R. A. Dluhy, S. M. Stephens, S. Widayati and A. D. Williams, *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, 1995, **51**, 1413–1447.
- [129] B. L. Mojet, S. D. Ebbesen and L. Lefferts, *Chem. Soc. Rev.*, 2010, **39**, 4643–4655.
- [130] D. Axelrod, in *Fluorescence Microscopy of Living Cells in Culture Part B. Quantitative Fluorescence Microscopy–Imaging and Spectroscopy*, ed. D. L. Taylor and Y.-L. Wang, Academic Press, 1989, vol. 30, pp. 245–270.
- [131] P. L. Stiles, J. A. Dieringer, N. C. Shah and R. P. Van Duyne, *Annual Review of Analytical Chemistry*, 2008, **1**, 601–626.
- [132] A. Kudelski, *Surface Science*, 2009, **603**, 1328–1334.
- [133] T. Ikeshoji, Y. Ono and T. Mizuno, *Appl. Opt.*, 1973, **12**, 2236–2237.
- [134] M. Fujihira and T. Osa, *Journal of the American Chemical Society*, 1976, **98**, 7850–7851.
- [135] T. Takenaka and K. Yamasaki, *Journal of Colloid and Interface Science*, 1980, **78**, 37–43.
- [136] T. Takenaka and H. Fukuzaki, *Journal of Raman Spectroscopy*, 1979, **8**, 151–154.
- [137] T. Takenaka and T. Nakanaga, *The Journal of Physical Chemistry*, 1976, **80**, 475–480.
- [138] T. Nakanaga and T. Takenaka, *The Journal of Physical Chemistry*, 1977, **81**, 645–649.

- [139] R. Iwamoto, M. Miya, K. Ohta and S. Mima, *Journal of the American Chemical Society*, 1980, **102**, 1212–1213.
- [140] R. Iwamoto, M. Miya, K. Ohta and S. Mima, *Journal of Chemical Physics*, 1981, **74**, 4780–4790.
- [141] R. Iwamoto, K. Ohta, M. Miya and S. Mima, *Applied Spectroscopy*, 1981, **35**, 584–587.
- [142] W. Hölzer, O. Schröter and A. Richter, *Journal of Molecular Structure*, 1990, **217**, 253–264.
- [143] W. Carius and O. Schröter, *physica status solidi (a)*, 1980, **59**, K115–K118.
- [144] W. Carius, W. Hölzer and O. Schröter, *Die Makromolekulare Chemie, Rapid Communications*, 1983, **4**, 469–470.
- [145] G. Mattei, B. Fornari and M. Pagannone, *Solid State Communications*, 1980, **36**, 309–312.
- [146] M. Ohsawa, K. Hashima and W. Suëtaka, *Applications of Surface Science*, 1984, **20**, 109–120.
- [147] L. D’Hooge and J. Vigoureux, *Chemical Physics Letters*, 1979, **65**, 500–506.
- [148] Y. Levy, C. Imbert, J. Cipriani, S. Racine and R. Dupeyrat, *Optics Communications*, 1974, **11**, 66–69.
- [149] J. F. Rabolt, N. E. Schlotter and J. D. Swalen, *The Journal of Physical Chemistry*, 1981, **85**, 4141–4144.
- [150] J. S. Kanger, C. Otto, M. Slotboom and J. Greve, *The Journal of Physical Chemistry*, 1996, **100**, 3288–3292.
- [151] T. Plowman, S. Saavedra and W. Reichert, *Biomaterials*, 1998, **19**, 341–355.
- [152] E. V. Efremov, J. B. Buijs, C. Gooijer and F. Ariese, *Applied Spectroscopy*, 2007, **61**, 571–578.
- [153] L. G. Tisinger and A. J. Sommer, *Microscopy and Microanalysis*, 2004, **10**, 1318–1319.
- [154] C. A. Michaels, *Journal of Raman Spectroscopy*, 2010, **41**, 1670–1677.
- [155] G. M. Lerman, A. Israel and A. Lewis, *Applied Physics Letters*, 2006, **89**, 223122.
- [156] R. F. Tabor, J. Eastoe and P. Dowding, *Langmuir*, 2009, **25**, 9785–9791.
- [157] G. W. Faris and R. A. Copeland, *Appl. Opt.*, 1997, **36**, 2686–2688.
- [158] M. O. Trulson and R. A. Mathies, *The Journal of Chemical Physics*, 1986, **84**, 2068–2074.
- [159] W. K. Thompson, *Trans. Faraday Soc.*, 1965, **61**, 2635–2640.

- [160] A. E. Klingbeil, J. B. Jeffries and R. K. Hanson, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 2007, **107**, 407 – 420.
- [161] M. L. Larsson, A. Holmgren and W. Forsling, *Journal of Colloid and Interface Science*, 2001, **242**, 25–30.
- [162] Y. R. Shen, *Nature*, 1989, **337**, 519–525.
- [163] C. D. Bain, *J. Chem. Soc., Faraday Trans.*, 1995, **91**, 1281–1296.
- [164] G. L. Richmond, *Chemical Reviews*, 2002, **102**, 2693–2724.
- [165] Y. R. Shen and V. Ostroverkhov, *Chemical Reviews*, 2006, **106**, 1140–1154.
- [166] H. C. Allen, N. N. Casillas-Ituarte, M. R. Sierra-Hernandez, X. Chen and C. Y. Tang, *Phys. Chem. Chem. Phys.*, 2009, **11**, 5538–5549.
- [167] F. M. Geiger, *Annual Review of Physical Chemistry*, 2009, **60**, 61–83.
- [168] Z. S. Nikolov, J. C. Earnshaw and J. J. McGarvey, *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 1993, **76**, 41–49.
- [169] Z. S. Nikolov, J. C. Earnshaw and J. J. McGarvey, *Journal of Raman Spectroscopy*, 1993, **24**, 411–416.
- [170] H. Yui, H. Fujiwara and T. Sawada, *Chemical Physics Letters*, 2002, **360**, 53–58.
- [171] J. S. Kanger, C. Otto and J. Greve, *The Journal of Physical Chemistry*, 1996, **100**, 16293–16297.
- [172] G. I. Stegeman, R. Fortenberry, C. Karaguleff, R. Moshrefzadeh, W. M. Hetherington, III, N. E. Van Wyck and J. E. Sipe, *Opt. Lett.*, 1983, **8**, 295–297.
- [173] W. M. K. P. Wijekoon, Z. Z. Ho and W. M. Hetherington, *J. Chem. Soc., Faraday Trans.*, 1993, **89**, 1067–1069.
- [174] E. Okamura, J. Umemura and T. Takenaka, *Journal of Raman Spectroscopy*, 1991, **22**, 759–762.
- [175] C. Lee and C. D. Bain, *Biochimica et Biophysica Acta (BBA) - Biomembranes*, 2005, **1711**, 59–71.
- [176] C. Lee, H. Wacklin and C. D. Bain, *Soft Matter*, 2009, **5**, 568–575.
- [177] D. J. Neivandt, M. L. Gee, M. L. Hair and C. P. Tripp, *The Journal of Physical Chemistry B*, 1998, **102**, 5107–5114.
- [178] T. Morikawa, E. Shirai, J. Tanno, H. Takanashi, A. Yasuda and K. Itoh, *Molecular Crystals and Liquid Crystals Science and Technology. Section A. Molecular Crystals and Liquid Crystals*, 1998, **312**, 69–94.
- [179] I. Kaplan-Ashiri, E. J. Titus and K. A. Willets, *ACS Nano*, 2011, **5**, 1033–1041.
- [180] A. J. McQuillan, *Advanced Materials*, 2001, **13**, 1034–1038.

- [181] K. Fujiwara and H. Watarai, *Langmuir*, 2003, **19**, 2658–2664.
- [182] S. Yamamoto and H. Watarai, *The Journal of Physical Chemistry C*, 2008, **112**, 12417–12424.
- [183] T. Takenaka, *Chemical Physics Letters*, 1978, **55**, 515–518.
- [184] J. Perera and G. Stevens, *Analytical and Bioanalytical Chemistry*, 2009, **395**, 1019–1032.
- [185] J. P. R. Day and C. D. Bain, *Phys. Rev. E*, 2007, **76**, 041601.
- [186] R. A. Dluhy, N. A. Wright and P. R. Griffiths, *Applied Spectroscopy*, 1988, **42**, 138–141.
- [187] R. A. Campbell, S. R. W. Parker, J. P. R. Day and C. D. Bain, *Langmuir*, 2004, **20**, 8740–8753.
- [188] G. R. Bell, C. D. Bain and R. N. Ward, *J. Chem. Soc., Faraday Trans.*, 1996, **92**, 515–523.
- [189] T. Kawai, J. Umemura and T. Takenaka, *Chemical Physics Letters*, 1989, **162**, 243–247.
- [190] D. A. Beattie, R. Fraenkel, S. A. Winget, A. Petersen and C. D. Bain, *The Journal of Physical Chemistry B*, 2006, **110**, 2278–2292.
- [191] P. Cann and H. Spikes, *Tribology Letters*, 2005, **19**, 289–297.
- [192] R. Fraenkel, G. E. Butterworth and C. D. Bain, *Journal of the American Chemical Society*, 1998, **120**, 203–204.
- [193] D. A. Beattie, S. Winget and C. D. Bain, *Tribology Letters*, 2007, **27**, 159–167.
- [194] S. Webster, D. N. Batchelder and D. A. Smith, *Applied Physics Letters*, 1998, **72**, 1478–1480.
- [195] J. Grausem, B. Humbert, M. Spajer, D. Courjon, A. Burneau and J. Oswald, *Journal of Raman Spectroscopy*, 1999, **30**, 833–840.
- [196] G. Mattei, M. Pagannone, B. Fornari and L. Mattioli, *Solid State Communications*, 1982, **44**, 1495–1498.
- [197] Y. Q. Sheng, E. Recknagel, A. Weidinger, Z. G. Gu, Z. Y. Lai and Z. C. Zhuang, *physica status solidi (b)*, 1988, **145**, 151–156.
- [198] F. Ishizaki and M. Kim, *Japanese Journal of Applied Physics*, 2008, **47**, 1621–1627.
- [199] P. R. Greene and C. D. Bain, *Spectroscopy Europe*, 2004, Aug/Sep, 8–15.
- [200] P. R. Greene and C. D. Bain, *Colloids and Surfaces B: Biointerfaces*, 2005, **45**, 174–180.
- [201] N. H. Fontaine and T. E. Furtak, *J. Opt. Soc. Am. B*, 1997, **14**, 3342–3348.

- [202] J. F. Power, *Review of Scientific Instruments*, 2002, **73**, 4057–4141.
- [203] P. C. Hansen, *Inverse Problems*, 1992, **8**, 849.
- [204] P. C. Hansen, *SIAM Review*, 1992, **34**, 561–580.
- [205] N. H. Fontaine and T. E. Furtak, *Phys. Rev. B*, 1998, **57**, 3807–3810.
- [206] K. J. McKee and E. A. Smith, *Review of Scientific Instruments*, 2010, **81**, 043106.
- [207] K. L. A. Chan and S. G. Kazarian, *Applied Spectroscopy*, January 2007, **61**, 48–54.
- [208] T. Frosch, K. L. A. Chan, H. C. Wong, J. T. Cabral and S. G. Kazarian, *Langmuir*, 2010, **26**, 19027–19032.
- [209] L. J. Fina, *Applied Spectroscopy Reviews*, 1994, **29**, 309–365.
- [210] L. J. Fina and G. Chen, *Vibrational Spectroscopy*, 1991, **1**, 353–361.
- [211] R. A. Shick, J. L. Koenig and H. Ishida, *Applied Spectroscopy*, 1 August 1993, **47**, 1237–1244.
- [212] R. A. Shick, J. L. Koenig and H. Ishida, *Applied Spectroscopy*, 1 August 1996, **50**, 1082–1088.
- [213] N. J. Overall, *Applied Spectroscopy*, 2009, **63**, 245A–262A.
- [214] N. J. Overall, *Analyst*, 2010, **135**, 2512–2522.
- [215] *Renishaw Raman Imaging Microscope WiRE V1.2 users' notes, Issue 1.0*.
- [216] S. Maussang, Renishaw plc, *private communication with Eric Tyrode*.
- [217] R. G. Snyder, H. L. Strauss and C. A. Elliger, *The Journal of Physical Chemistry*, 1982, **86**, 5145–5150.
- [218] T. J. O'Leary and I. W. Levin, *The Journal of Physical Chemistry*, 1984, **88**, 1790–1796.
- [219] E. R. Malinowski, *Factor Analysis in Chemistry*, John Wiley & Sons, 2nd edn., 1991.
- [220] D. R. Mason and E. L. Piret, *Industrial & Engineering Chemistry*, 1950, **42**, 817–825.
- [221] N. G. Anderson, *Organic Process Research & Development*, 2001, **5**, 613–621.
- [222] E. Santacesaria and S. Carr, *Applied Catalysis*, 1983, **5**, 345–358.
- [223] I. Aizpurua and M. J. Barandiaran, *Polymer*, 1999, **40**, 4105–4115.
- [224] T. F. Svitova, M. J. Wetherbee and C. J. Radke, *Journal of Colloid and Interface Science*, 2003, **261**, 170 – 179.

- [225] T. F. Svitova and C. J. Radke, *Industrial & Engineering Chemistry Research*, 2005, **44**, 1129–1138.
- [226] V. B. Fainerman, M. E. Leser, M. Michel, E. H. Lucassen-Reynders and R. Miller, *The Journal of Physical Chemistry B*, 2005, **109**, 9672–9677.
- [227] H. Schlichting, *Boundary-Layer Theory*, McGraw-Hill Inc., 1979.
- [228] R. L. Panton, in *Incompressible Flow*, John Wiley & Sons, Inc., 2nd edn., 1996, pp. 282–289.
- [229] N. Frössling, *Lunds Univ. Årsskr. N. F. Afd. 2*, 1940, **35**, No. 4.
- [230] J. Melville, N. Simjee, P. R. Unwin, B. A. Coles and R. G. Compton, *Journal of Physical Chemistry B*, 2002, **106**, 2690–2698.
- [231] M. D. Deshpande and R. N. Vaishnav, *J. Fluid Mech.*, 1982, **114**, 213–236.
- [232] J. V. Macpherson, M. A. Beeston and P. R. Unwin, *J. Chem. Soc., Faraday Trans.*, 1995, **91**, 899–904.
- [233] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*, Dover Publications, Inc., 1972.
- [234] J. Alden, *Ph.D. thesis*, Oxford University, 1998.
- [235] M. B. Abbott and B. D. R., *Computational Fluid Dynamics: An Introduction for Engineers*, Longman Scientific & Technical, 1989.
- [236] I. Langmuir, *Journal of the American Chemical Society*, 1916, **38**, 2221–2295.
- [237] A. Frumkin, *Zeitschrift für Physik A Hadrons and Nuclei*, 1926, **35**, 792–802.
- [238] J. K. Ferri and K. J. Stebe, *Journal of Colloid and Interface Science*, 1999, **209**, 1–9.
- [239] D. A. Woods, J. Petkov and C. D. Bain, *The Journal of Physical Chemistry B*, 2011, **115**, 7341–7352.
- [240] A. Sadkowski, *Journal of Electroanalytical Chemistry and Interfacial Electrochemistry*, 1986, **208**, 69–76.
- [241] M. J. Avena and L. K. Koopal, *Environmental Science & Technology*, 1999, **33**, 2739–2744.
- [242] L. K. Koopal and M. J. Avena, *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 2001, **192**, 93–107.
- [243] C. F. Wertz and M. M. Santore, *Langmuir*, 2002, **18**, 1190–1199.
- [244] A. Couzis and E. Gulari, in *Adsorption from Aqueous Binary Surfactant Mixtures onto the Solid-Liquid Interface*, American Chemical Society, 1992, ch. 25, pp. 354–365.
- [245] R. Thomas, *Annual Review of Physical Chemistry*, 2004, **55**, 391–426.

- [246] G. Fragneto, J. R. Lu, D. C. McDermott, R. K. Thomas, A. R. Rennie, P. D. Gallagher and S. K. Satija, *Langmuir*, 1996, **12**, 477–486.
- [247] J. Penfold, I. Tucker and R. K. Thomas, *Langmuir*, 2005, **21**, 11757–11764.
- [248] H. Gharibi, S. Javadian, B. Sohrabi and R. Behjatmanesh, *Journal of Colloid and Interface Science*, 2005, **285**, 351–359.
- [249] X. W. Fang, S. Zhao, S. Z. Mao, J. Y. Yu and Y. Du, *Colloid Polym. Sci.*, 2003, **281**, 455–560.
- [250] B. Lindman, M. C. Puyal, N. Kamenka, R. Rymden and P. Stilbs, *The Journal of Physical Chemistry*, 1984, **88**, 5048–5057.
- [251] F. J. Trogus, R. S. Schechter and W. H. Wade, *Journal of Colloid and Interface Science*, 1979, **70**, 293–305.
- [252] Q-Sense Instruments, http://www.q-sense.com/qcm_d_faq--33.asp, accessed November 2010.
- [253] S. Partyka, S. Zaini, M. Lindheimer and B. Brun, *Colloids and Surfaces*, 1984, **12**, 255–270.
- [254] S. Paria, C. Manohar and K. C. Khilar, *Industrial & Engineering Chemistry Research*, 2005, **44**, 3091–3098.
- [255] S. C. Biswas and D. K. Chattoraj, *Journal of Colloid and Interface Science*, 1998, **205**, 12–20.
- [256] J. Penfold, E. J. Staples, I. Tucker and L. J. Thompson, *Langmuir*, 1997, **13**, 6638–6643.
- [257] J. Penfold, E. J. Staples, I. Tucker and R. K. Thomas, *Langmuir*, 2000, **16**, 8879–8883.
- [258] S. B. Velegol and R. D. Tilton, *Langmuir*, 2001, **17**, 219–227.
- [259] H. Li and C. P. Tripp, *The Journal of Physical Chemistry B*, 2004, **108**, 18318–18326.
- [260] A. Häbich, G. G. Qiao and W. Ducker, *Langmuir*, 2010, **26**, 13944–13953.
- [261] R. Wirz, T. Bürgi and A. Baiker, *Langmuir*, 2003, **19**, 785–792.
- [262] M. Bieri and T. Bürgi, *The Journal of Physical Chemistry B*, 2005, **109**, 10243–10250.
- [263] J. P. R. Day, R. A. Campbell, O. P. Russell and C. D. Bain, *The Journal of Physical Chemistry C*, 2007, **111**, 8757–8774.
- [264] Z. Huang and T. Gu, *Colloids and Surfaces*, 28, **1987**, 159–168.
- [265] D. C. McDermott, D. Kanelleas, R. K. Thomas, A. R. Rennie, S. K. Satija and C. F. Majkrzak, *Langmuir*, 1993, **9**, 2404–2407.

- [266] O. A. Soboleva, A. A. Yaroslavtsev, G. A. Badun and B. D. Summ, *Colloid Journal*, 2004, **66**, 470–476.
- [267] C. C. Ruiz and J. Aguiar, *Langmuir*, 2000, **16**, 7946–7953.
- [268] G. Despert and J. Oberdisse, *Langmuir*, 2003, **19**, 7604–7610.
- [269] D. Lugo, J. Oberdisse, M. Karg, R. Schweins and G. H. Findenegg, *Soft Matter*, 2009, **5**, 2928–2936.
- [270] J. Penfold, I. Tucker, J. Petkov and R. K. Thomas, *Langmuir*, 2007, **23**, 8357–8364.
- [271] I. Tucker, J. Petkov, J. Penfold and R. K. Thomas, *Langmuir*, 2010, **36**, 8036–8048.
- [272] S. Paria, C. Manohar and K. C. Khilar, *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 2005, **252**, 221–229.
- [273] S. Alila, S. Boufi, M. N. Belgacem and D. Beneventi, *Langmuir*, 2005, **21**, 8106–8113.
- [274] E. Kontturi, P. C. Thüne and J. W. H. Niemantsverdriet, *Langmuir*, 2003, **19**, 5735–5741.
- [275] T. Mohan, R. Kargl, A. Doliska, A. Vesel, S. Kstler, V. Ribitsch and K. Stana-Kleinschek, *Journal of Colloid and Interface Science*, 2011, **358**, 604–610.
- [276] M. Holmberg, J. Berg, S. Stemme, L. Ödberg, J. Rasmusson and P. Claesson, *Journal of Colloid and Interface Science*, 1997, **186**, 369–381.
- [277] S. M. Notley, *J. Phys. Chem. B*, 2009, **113**, 13895–13897.
- [278] T. Tammelin, T. Saarinen, M. Österberg and J. Laine, *Cellulose*, 2006, **13**, 519–535.
- [279] E. Kontturi, T. Tammelin and M. Österberg, *Chem. Soc. Rev.*, 2006, **35**, 1287–1304.
- [280] M. Schaub, G. Wenz, G. Wegner, A. Stein and D. Klemm, *Advanced Materials*, 1993, **5**, 919–922.
- [281] C. Wegner, V. Buchholz, L. Ödberg and S. Stemme, *Advanced Materials*, 1996, **8**, 399–402.
- [282] C. Geffroy, M. P. Labeau, K. Wong, B. Cabane and M. A. Cohen Stuart, *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 2000, **172**, 47–56.
- [283] R. D. Neuman, J. M. Berg and P. M. Claesson, *Nordic Pulp & Paper Research Journal*, 1993, **8**, 96–104.
- [284] S. Gunnars, L. Wågberg and M. A. Cohen Stuart, *Cellulose*, 2002, **9**, 239–249.

- [285] L.-E. Enarsson and L. Wågberg, *Biomacromolecules*, 2009, **10**, 134–141.
- [286] E. Kontturi and A. Lankinen, *Journal of the American Chemical Society*, 2010, **132**, 3678–3679.
- [287] P. Linse and P. M. Claesson, *Macromolecules*, 2009, **42**, 6310–6318.
- [288] O. Teschke and E. F. de Souza, *Langmuir*, 2003, **19**, 5357–5365.
- [289] J. C. Schulz and G. G. Warr, *Langmuir*, 2000, **16**, 2995–2996.
- [290] F. P. Duval, R. Zana and G. G. Warr, *Langmuir*, 2006, **22**, 1143–1149.
- [291] R. E. Lamont and W. A. Ducker, *Journal of the American Chemical Society*, 1998, **120**, 7602–7607.
- [292] J.-F. Liu and W. A. Ducker, *The Journal of Physical Chemistry B*, 1999, **103**, 8558–8567.
- [293] Y. L. Chen, S. Chen, C. Frank and J. Israelachvili, *Journal of Colloid and Interface Science*, 1992, **153**, 244–265.
- [294] W. A. Ducker and E. J. Wanless, *Langmuir*, 1999, **15**, 160–168.
- [295] S. Perkin, N. Kampf and J. Klein, *The Journal of Physical Chemistry B*, 2005, **109**, 3832–3837.
- [296] W. A. Hayes and D. K. Schwartz, *Langmuir*, 1998, **14**, 5913–5917.
- [297] M. Fujii, B. Li, K. Fukada, T. Kato and T. Seimiya, *Langmuir*, 1999, **15**, 3689–3692.
- [298] J. M. Mellott, W. A. Hayes and D. K. Schwartz, *Langmuir*, 2004, **20**, 2341–2348.
- [299] C. E. McNamee, H.-J. Butt, K. Higashitani, I. U. Vakarelski and M. Kappl, *Langmuir*, 2009, **25**, 11509–11515.
- [300] *Luxpop refractive index database* (<http://www.luxpop.com/>).
- [301] *Sigma-Aldrich product specifications*.
- [302] *The Physics Hypertextbook* (<http://physics.info/refraction/>).
- [303] S. N. Singh, H. S. Bhatti and R. D. Singh, *Spectrochimica Acta Part A: Molecular Spectroscopy*, 1978, **34**, 985–992.
- [304] G. E. Walrafen, M. S. Hokmabadi and W.-H. Yang, *Journal of Chemical Physics*, 1986, **85**, 6964–6969.
- [305] P. Wydro and M. Paluch, *Journal of Colloid and Interface Science*, 2005, **286**, 387–391.
- [306] A. J. Easteal, *Aust. J. Chem.*, 1980, **33**, 1667–75.
- [307] J. W. Thompson, T. J. Kaiser and J. W. Jorgenson, *Journal of Chromatography A*, 2006, **1134**, 201–209.

- [308] R. H. Stokes, *The Journal of Chemical Thermodynamics*, 1987, **19**, 977–983.
- [309] H. T. French, *The Journal of Chemical Thermodynamics*, 1987, **19**, 1155–1161.
- [310] J. Catalan, C. Diaz and F. Garcia-Blanco, *Org. Biomol. Chem.*, 2003, **1**, 575–580.