Poking Vesicles: What Molecular Dynamics Can Reveal About Cell Mechanics

Benjamin Barlow

Thesis submitted to the
Faculty of Graduate and Postdoctoral Studies
In partial fulfillment of the requirements
For the M.Sc. degree in Physics

Department of Physics
Faculty of Science
University of Ottawa

Acknowledgements

Béla Joós: without whom none of this would have been possible. I am truly grateful for the guidance and opportunities he has provided me.
Martin Bertrand: after all, this was his idea! Aside from being an excellent coach, Martin provided the essential computational infrastructure —the foundation on which this project was built. His imagination & technical aptitude in molecular dynamics and computational physics were indispensable.
Kristina Haase: for providing lovely images of cells deformed under AFM, and for her guidance in reviewing relevant experimental work on cells.
Louis Jacques: for mathematical rigour, humour, and debate.
Alison Harman: for helpful discussion on common issues we both confronted simulating vesicles.

Contents

List of Figures

Preface

Because cells are machines, their *structure* determines their *function* (health). But their structure *also* determines cells' *mechanical* properties. So if we can understand how cells' mechanical properties are influenced by *specific* structures, then we can observe what's happening inside of cells via $\underbrace{\text{mechanical measurements}}_{\text{less invasive}^{0}}$.

The Atomic Force Microscope (AFM) has become a standard tool for investigating the mechanical properties of cells. In many experiments, an AFM is used to 'poke' adherent cells with nanonewton forces, and the resulting deformation observed via, e.g. Laser Scanning Confocal Microscopy[17, 18]. Results of such experiments are often interpreted in terms of continuum mechanical models which characterize the cell as a linear viscoelastic solid[26].

This "top-down" approach of poking an intact cell —complete with cytoskeleton, organelles etc.— can be problematic when trying to measure the mechanical properties and response of a *single* cell component. Moreover, how are we to know the *sensitivity* of the cell's mechanical properties to partial modification of a single component (e.g. reducing the degree of cross-linking in the actin cortex)?

In contrast, the approach taken here[1] —studying the deformation and relaxation of lipid bilayer vesicles— might be called a "bottom-up" approach to cell mechanics. Using Coarse-Grained Molecular Dynamics simulations, we study the deformation and relaxation of bilayer vesicles, when poked with constant force. The relaxation time, equilibrium area expansion, and surface tension of the vesicle membrane are studied over a range of applied forces. Interestingly, the relaxation time exhibits a strong force-dependence. Force-compression curves for our simulated vesicle show a strong similarity to recent experiments[34] where giant unilamellar vesicles were compressed in a manner nearly identical to that of our simulations.

[0] compared to e.g. electron microscopy, where one must kill the cell to be able to image it.

[1] And by [34, 11, 33, 9, 22, 6, 28, 16] to name a few.

Introduction

This work is about lipid bilayer vesicles, their mechanical properties, and how said properties change when the vesicle contents are altered. To investigate the mechanical properties of vesicles we ran computer simulations wherein a vesicle is squeezed between two plates. We do this because vesicles serve as prototypes of simple cells and knowledge of their mechanical properties —and how said properties are determined by their contents/composition— is fundamental to understanding living cells.

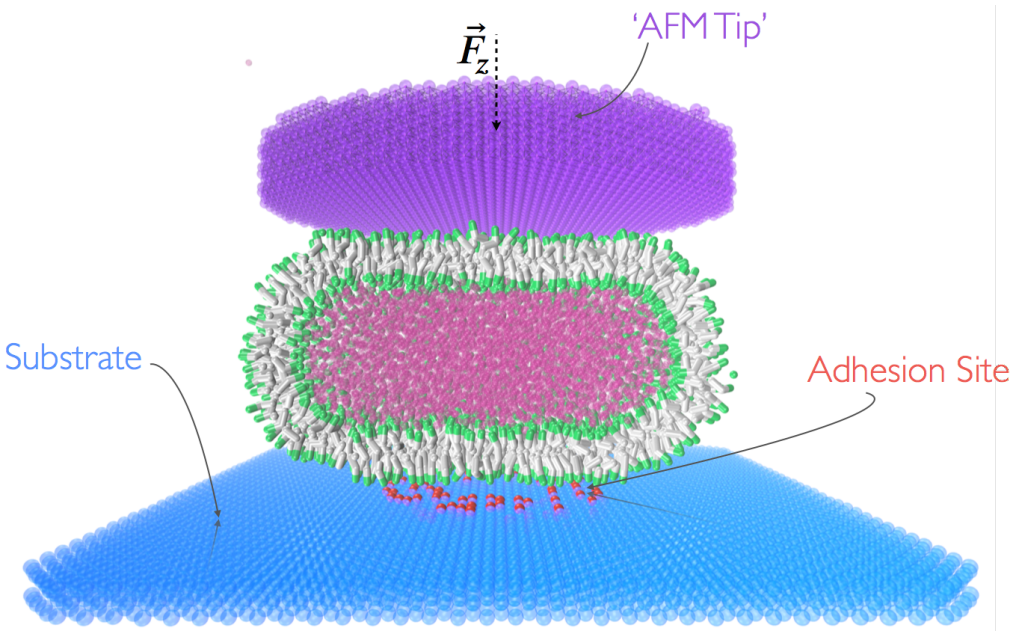

Figure 1.1: We squeeze a simulated vesicle between two plates. This is analogous to real world experiments where Atomic Force Microscopes (AFMs) have been used to poke/squeeze/deform real cells and vesicles.

This procedure is relevant to real[1] experiments[18, 17, 1, 24, 35, 34] which use an Atomic Force Microscope (AFM) to poke and squeeze and stretch living cells and vesicles (see Figure 1.2).

[1]Not to say that the findings of simulations aren't *real*! We use the words 'real' and 'simulated' to distinguish experiments that use molecules in aqueous solution from experiments in which molecules are *represented* via electrons in silicon.

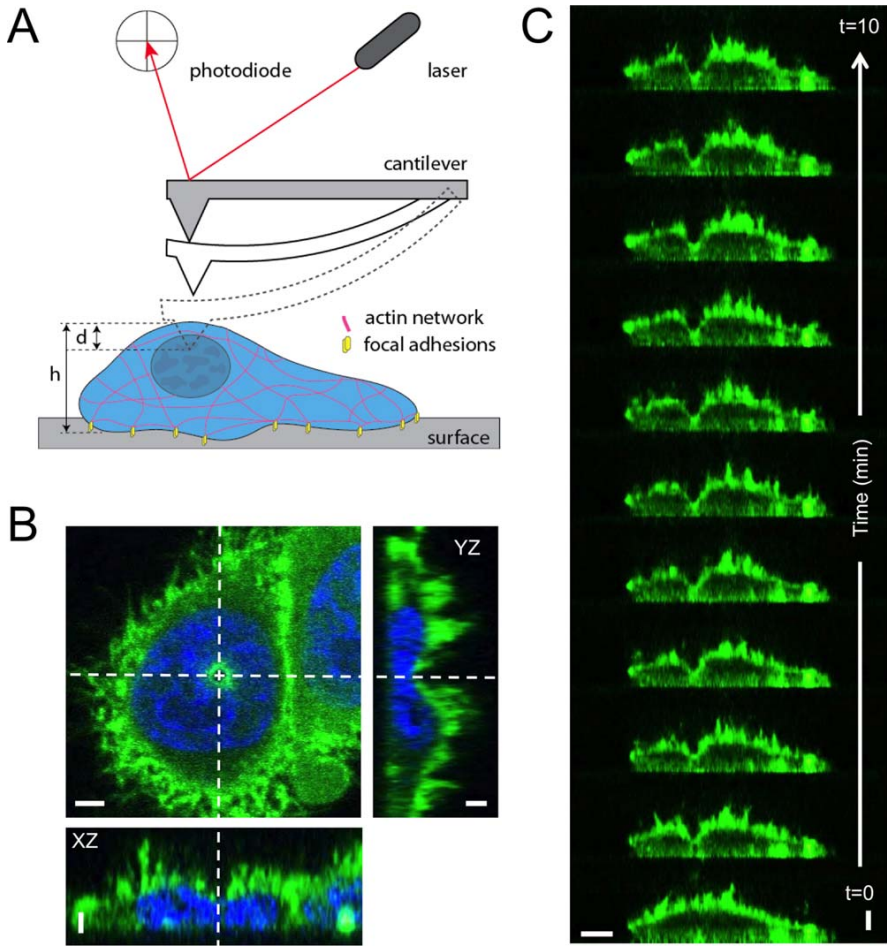

Fig. 1. AFM setup and LSCM volume acquisition. A: Schematic diagram of the AFM set-up used as a nanoindentor. The AFM is employed to impose an indentation of a given depth (d) into a cell of a given height (h). **B**: A living HeLa cell deformed by a 10-nN force applied by an AFM tip above the cell's nucleus. Orthogonal views show that both the membrane and underlying nucleus were deformed. **C:** Deformation over a 10 min duration – the depth of deformation increases as demonstrated from $t = 0$ to $t = 10$. Green: PH-PLCδ-EGFP, Blue: Hoechst 33342, Scale bars = 5 μm. [Color figure can be viewed in the online issue which is available at wileyonlinelibrary.com.]

Figure 1.2: Adapted with permission[2] from "Resiliency of the Plasma Membrane and Actin Cortex to Large-Scale Deformation" by Kristina Haase and Andrew Pelling[18]. In 'C' we observe a cell (fuzzy green blob) being poked by an AFM (notice the growing triangular indentation, left of the cell's centre in C).

In fact, practically equivalent experiments were recently performed on giant liposomes by Edith Schäfer et. al [34] (published in 2013). But cells and vesicles are not deformed only in the lab. Inside our own bodies, every time the heart beats, every time we breathe, every time we flex a muscle of any kind —at every moment in cells all over the body, mechanical deformation of the membrane, cytoskeleton and cell contents is occurring.

Why Simulate Vesicles?

This question can be split in two: "Why simulate?" and "Why vesicles?". Daniel C. Dennett provides a philosophical answer to the former:

> "You *know* that if you succeed in getting a computer program to model some phenomenon, there are ***no causes*** at work in the model ***other than*** the causes that are composed of all the ***arithmetical operations***."

In computer models we know *exactly* what effects are at play —in experiment we often don't, and the challenge is to find which of the many factors dominate. That is, unlike other types of experiment, *computer experiments* such as molecular dynamics simulations offer *complete knowledge* of the model system. We know the positions, momenta, and forces associated with *all* particles in the system at any given moment. This is because computers do *exactly* as they are told. So long as we are careful to make sure we understand what it is that we told the computer to do, we know exactly what it is doing! The circular character of this knowledge means that a simulation can't tell us anything about the world unless we have studied nature herself, and poured some of this physical truth into the code.

If we have captured some aspect of nature in the equations governing the particles' interactions, then simulation can be a powerful supplement to experiment, analytical calculation and human imagination. It is the responsibility of the researcher to know what real physical system his simulation corresponds to, and to ensure that it actually does.

Biology answers the latter. Lipid bilayer vesicles (vesicles) are essential to life on earth. These closed quasi-spherical membranes self-assemble in nature, driven to do so by the hydrophobic effect.

[2] No additional uses are granted (such as derivative works or other editions).

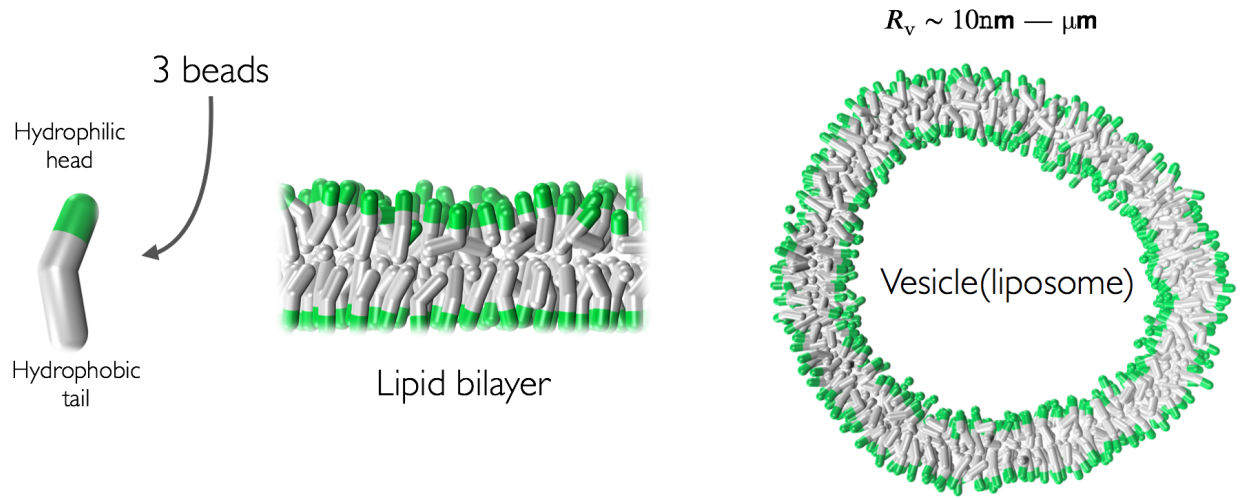

Figure 1.3: Self-assembly of amphiphiles into closed quasi-spherical vesicles. In *nature*, the closed configuration is entropically favourable because it minimizes the hydrophobic tails' exposure to the surrounding fluid. In our *simulations*, the hydrophobic interaction is modelled by making exposure of the tails energetically costly. Thus in contrast to real vesicles, the simulated closed vesicle configuration is stable because it is *energetically* favourable. A major benefit of the molecular dynamics approach to simulating bilayers (such as those shown here) is that this method reproduces key properties of real bilayers —e.g. in-plane fluidity, ability to open and close pores, etc.

The cell membrane is at its foundation a lipid bilayer versicle —which makes the fact that vesicles self-assemble without any help from living organisms very interesting![3](While the phenomenon of self-assembly is ubiquitous in cells, what makes the self-assembly of the lipid bilayer particularly significant to the origin of life, is that *amphiphilic molecules* that form such bilayers *naturally occur in absence of living organisms*. Evolution by natural selection explains the step-by-step development and honing of biological machines which, if they instead had to 'appear all at once', would be prohibitively improbable. But before evolution can operate, it requires a population of organisms [with talents varying among individuals] on which to act. So the question is "How did life appear?". To have life [as it appears on earth], you need cells. Cells require closed membranes. Since amphiphilic molecules —occurring in absence of life— self-assemble into such membranes, they seem to answer part of *the question*.)

From a biophysical perspective the most basic prototype of a living cell is the lipid bilayer vesicle —the *spherical cow* of cell mechanics. This makes vesicles a good place to begin down the road toward simulating more complicated systems. Once an 'empty' vesicle (containing

[3]If you wanted to make a cell from scratch, you would start with a closed membrane, and then add components to the interior and membrane. Otherwise you would have to deal with the problem of cell components diffusing away, in absence of a membrane to confine them. Lipid bilayer vesicles are just such membranes —before you have life, the foundation of the cell is provided for free.

only fluid) has been characterized, new components can be added step by step. In this way, the mechanical properties of empty vesicles serve as a baseline against which to compare the results of more complicated systems. When observing the changes which result from adding new components, computer simulation allows us to control the character of said components very delicately. (For example, to compare against results of experiments which alter the cytoskeleton (perhaps by decrosslinking the actin cortex) one might simulate a vesicle with its own cortex. Then by carefully controlling the number of cross-links which remain in the simulated cortex, one could study the sensitivity of the vesicle's bulk mechanical properties to the degree of crosslinking.)

But vesicles live rich lives of their own. They are used all over the body as transport containers for important chemicals such as neurotransmitters. e.g. Inside every neuron, there are vesicles filled with neurotransmitter, each carried down the axon by a motor protein. Once a vesicle[4] arrives at the terminal button (where the axon meets another neuron, forming a synapse) it will wait nearby the cell membrane, until an action potential shoots down the axon and causes it to fuse with the cell membrane, and spill its neurotransmitter cargo into the synapse. This process happens at *every* synapse, so that a sizeable fraction of the $\sim 10^{11}$ neurons in your brain are *each* using many vesicles every second to move your eyes, think these thoughts, and view these letters.

Vesicles are also useful for drug delivery since a vesicle filled with medicine will diffuse much more slowly than the individual drug molecules —as a result of its larger size. This means that medicine (e.g. chemotherapy drugs) encapsulated in vesicles will spend more time in the location where they are injected (e.g. near a tumour) before diffusing into other tissues. Vesicles can also be prepared with activated membranes containing e.g. antibodies in order to deliver drugs to specific tissues in a targeted manner[25]. Liposomes rupture under sufficient shear stress (resulting from the flow-velocity profile in blood vessels), with larger liposomes being more delicate. Constriction of blood flow will tend to produce greater shear forces and liposome encapsulated drugs have also been used to treat ischemia as in the cases of stroke and heart attack. By choosing the appropriate vesicle size to rupture selectively in flow conditions resulting from the unusual vasculature of tumours, chemotherapy drugs could be delivered to where they're needed without harming healthy tissues throughout the body.

[4] now called a "synaptic vesicle"

Summary

If we want to understand cell mechanics, and the meaning of AFM measurements, we must understand the mechanical properties of the cell membrane. The most basic prototype of a cell membrane, and indeed of a cell, is the lipid bilayer vesicle. The purpose of this project is to begin the journey of building up a virtual cell, with components added step by step. This way, we can learn how the mechanical properties of cells are determined by their composition.

Molecular Dynamics

The type of simulation performed in this research is called Molecular Dynamics (MD). In molecular dynamics, one has an initial distribution of particles in space, and a list of interaction potentials between the particles. The program then integrates Newton's equations of motion for all particles, and the time-evolution of the system is recorded for later analysis. The physical phenomena exhibited by the system emerge purely from the initial conditions (initial particle positions and velocities), the interaction potentials between the particles, and the thermostat[5] —plus any external forcing included in the simulation routine.

There may be many different types of particles used, so that different interaction potentials can be defined for different pairs of particles types. That is, the force exerted by one particle on another may depend not only on the distance between them, but also which types of particles make up the pair. For example, one might set up a simulation box filled with fluid particles as well as some randomly distributed amphiphiles[6] mixed in. An interaction potential which mimics the hydrophobic effect would then cause the amphiphiles to self-assemble into closed vesicles, just as in nature[28]. This is an example of a *real* physical phenomenon *emerging* in simulation as a direct result of the interaction potentials used in the code —absent any external forcing. One could also begin a simulation with such a vesicle already set up, and then apply an external force to the vesicle and observe its reaction to this stimulus (the topic of this thesis).

The previous two examples are meant to highlight why it is reasonable to think of MD simulations as *computer experiments*. The built-in assumptions are the laws of mechanics and thermodynamics, the particle-particle interaction potentials, and the initial positions and momenta of the particles. The rich behaviour of such systems *emerges* from these basic physical assumptions —and nothing more.

[5] random forces imposed on particles to simulate thermal motion.
[6] e.g. phospholipids found in eukaryotic cell membranes
[7] from Wikipedia

Coarse-Grained Molecular Dynamics (CGMD)

To get an idea of what coarse grained molecular dynamics is, let's consider the following diagram of a lipid molecule[7]. This molecule (dipalmitoylphosphatidylcholine, or 'DPPC') is a very common type of lipid which is found in cell membranes all over the human body, and in eukaryotic cells in general.

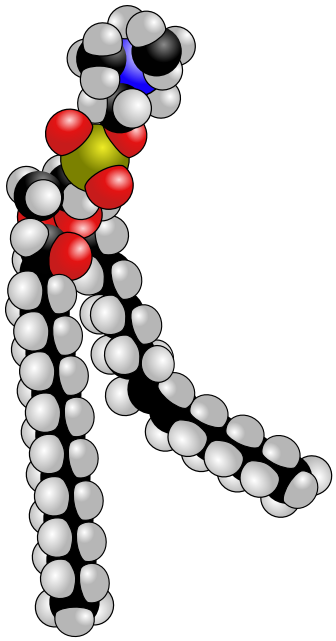

Figure 1.4: This Atomistic representation of a phospholipid (phosphatidylcholine graphic: Wikipedia) is itself coarse-grained, because atoms are themselves bound many-particle systems. The lipid's polar head group (hydrophilic) is shown in the upper portion of the figure (large blue bead and surrounding molecules), and its two hydrophobic tails are clearly visible.

In Figure 1.4 the molecule is already coarse-grained because the atoms are treated as simple spheres. But we know that atoms really are composed of electrons orbiting nuclei. And the nuclei in turn are neutrons and protons which themselves are made up of quarks and so on...

Likewise in CGMD we coarse grain the *molecule* itself. In the simulations discussed here, the molecule pictured above in Figure 1.4 is coarse-grained down to just three beads.

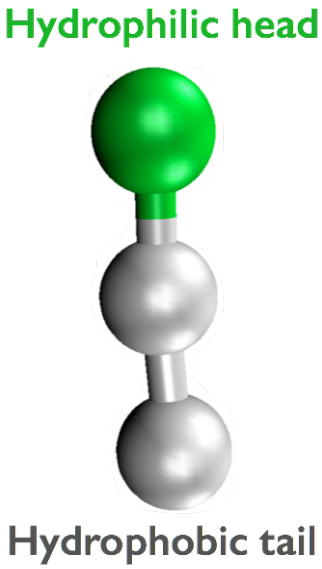

Figure 1.5: Coarse-grained lipid molecules used in this work.

The top green 'head' bead takes the place of the lipid's *hydrophilic* polar head group, and the remaining two 'tail' beads mimic the lipid's *hydrophobic* tail groups. This way the *amphiphilic* character of the more complicated molecule is captured using many fewer particles.

The motivation for coarse-graining is that the *duration* required to compute a given simulation is proportional to the number of particles it contains. Coarse-graining allows us to greatly reduce the number of particles in a given simulation (compared to the atomistic approach), and thus allows us to simulate larger systems —systems representative of many more molecules— with the same limited number of particles.

Background

Everything cells do, they must do through physics. This may seem like an unnecessary statement, but it is easy to get used to thinking of plants, animals and their cells as having some kind of 'will'. One might casually say "the paramecium swims about, looking for food" as though that were what suited it. It swims because it is a machine with a finite set of programmed capabilities, responding to the physical environment surrounding it. This same physical environment is what, over much longer timescales, moulded it through natural selection. The source of the novel information which produced its talents was not foresight, it was the randomness of gene mutation. This same source of novelty is always operating in all living things —usually to their detriment, occasionally to their benefit. Life is matter in motion, converting lower entropy energy (incoming photons of highly parallel sunlight) into higher entropy energy (blackbody radiation from waste heat, leaving the earth in all directions).[30]

Every brilliant 'decision' a cell makes, each sophisticated action it undertakes, *has*[1] to happen by the laws of physics —just as the sun *has* to burn. Understanding the phenomenon of life as it truly is, an orchestra of atoms moving blindly yet brilliantly under the natural laws of motion, is a fundamental challenge of science. This challenge is often referred to as biophysics.

Cells: Membranes and Filaments

The human body as a dynamic physical structure

Imagine a large multicellular organism (e.g. you).[2]
If, with a magic wand, you removed all of the cells[3], what would remain? A skeleton might come to mind, perhaps with some tendons, ligaments et cetera. But multicellular organisms like ourselves are made of much more than cells and bones. We are *full* of *scaffolding*, without which our tissues would fall apart (among other things).

[1] It is true that there are many probabilistic phenomena in physics, but this probabilistic character is also outside the control of living things.

Having vanished *all* of the cells, the object before you would actually look very much like the human you started out with: it would be the complete ***extracellular matrix*** of an entire person.

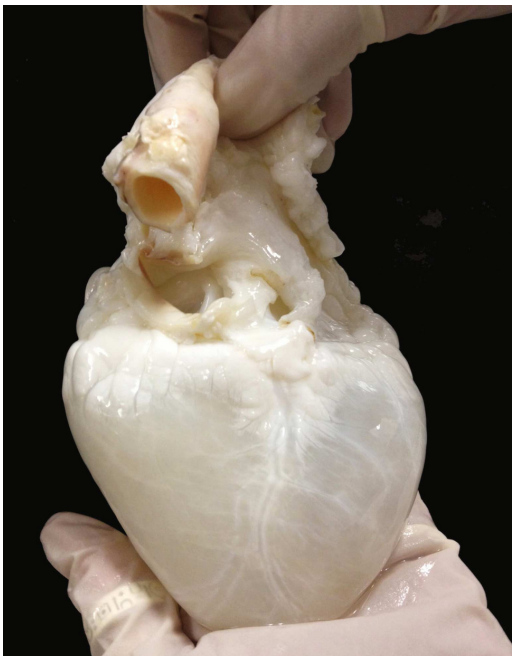

Figure 2.1: A "ghost heart": The cells have been washed away, leaving the heart's extracellular matrix intact —including its valves, vascular tree etc. Every organ, every tissue in the body has such a structure associated with it. (Image copyright Texas Heart Institute, http://www.texasheart.org)[4]

The extracellular matrix (ECM) is made up of a variety of filaments and matrix proteins which are excreted and structured by the cells that inhabit it.[2] This system of interconnected filaments is the scaffolding which holds us together. It stabilizes the physical structure of our tissues and influences their bulk mechanical properties. The ECM does not only provide structural support, it also influences the "survival, development, migration, proliferation, shape, and function"[2] of the cells in contact with it.

Our cells *also* have their own *internal* system of filaments —an '*intra*cellular matrix' as it were,[5] called the cytoskeleton. So the body is filled with an intricate mesh of nanofilaments,

[3] If what follows disturbs the reader, imagine something like a cloned arm.

[3] in practise they can be washed away with detergent[31]

[4] No additional uses are granted (such as derivative works or other editions).

[5] I've referred to the cytoskeleton as an '*intra*cellular matrix' here for the sake of conceptual continuity with the *extra*cellular matrix, to which the cytoskeleton is not only conceptually but *physically* connected in real life. While I'm treating 'intracellular matrix' as a made-up term here, a quick google search suggests that it is in use —though somewhat loosely depending on the source.

extending among different cells and even connecting with the cytoskeleton at the cell membrane. Taken together, the ECM and cytoskeletons constitute a physical connection running from the organelles of different cells all the way up to the macroscopic scale of tissues and organs that make up the human body.

Eukaryotic cells as dynamic physical structures[6]

The parts of the cell which will be of most interest to us are the membrane and cytoskeleton(particularly the actin cortex).

Cell membrane

A cell is distinguished from the outside world by its membrane. The membrane is the most basic feature of any living cell —without it, the very notion of a cell is meaningless. Its crucial function is to keep the cell contents isolated from the outside world, enabling the cell to maintain and control its delicate internal chemistry. At its foundation the cell membrane is a phospholipid bilayer[7] $\sim 5\text{nm}$ thick[30, 2] which self-assembles in absence of anything we would call life. Said self-assembly is driven by the hydrophobic effect, which results from the decrease in entropy of the lipid-water system when the lipids' non-polar tails are exposed to the surrounding water molecules. But cell membranes are much more than lipid bilayers —they are highly equipped, selectively permeable, ion-pumping, information transducing, self-organizing, automatically repairing machines! Animal cell membranes are composed $\sim 50\%$ of lipids by mass[2], the remainder being the various fantastic proteins which serve as pumps, voltage-gated channels, anchors to the cytoskeleton, receptor sites etc. The in-plane fluidity of lipid bilayer membranes allows cells to rearrange membrane components freely.

[6] Throughout this work 'cell' should be taken to mean 'eukaryotic cell' and not cells in general, unless otherwise indicated by context.

[7] See Figures 1.3 and 1.4

Cytoskeleton

As mentioned above eukaryotic cells possess their own intricate, dynamic system of fibres referred to collectively as the ***cytoskeleton***. The cytoskeleton strongly influences the mechanical properties of cells, and produces the amazing variety of cell shapes —from the simplicity of red blood cells to the intricacy of neurons. There is a 'cortex' of actin filaments running along the membrane which serves to reinforce it, control its shape, and mediate the transduction of mechanical information from the membrane to the inner cytoskeleton. The actin cortex is a gel —a web-like structure of actin filaments cross-linked to one another at large angles. There are also highly rigid microtubules[14] running outward from the centrosome[8] to the membrane. Among other functions, these tubules serve as bridges for the transport of important materials (like the neurotransmitters already mentioned) which are encapsulated in vesicles and pulled along the microtubules by motor proteins such as kinesin.[2]

The cytoskeleton is a highly active machine. It controls the shape, internal structure, and mechanical properties of the cell in a dynamic manner, either in response to external stimuli or due to self-initiating processes within. For example, it is through remodelling of the cytoskeleton that cells are able to extend filaments that adhere to their surroundings (focal adhesions) and then reel them in to fasten the cell into position. By continuously casting out new adhesions[9] at one end of the cell, and reeling them in while letting go of old adhesions, the cytoskeleton enables cells to crawl (cell motility).

Muscle cell contraction is effected by arrays of long cytoskeletal filaments which slide against each other lengthwise[2] —when you flex your muscles, you're really "flexing your cytoskeletons"[10]. [To see the motion more clearly, clasp your hands together with fingers interlocking (leave the thumbs free). Then slowly pull your hands apart and push them back together, with the interlocked fingers sliding along each other lengthwise. The fingers are the cytoskeletal filaments. Pushing the hands together is contractile motion, pulling them apart is relaxation.]

The specific cytoskeletal filaments responsible for muscle contraction are called actin filaments and myosin filaments. The sliding action is powered by ATP, the "fuel molecule", which is consumed by the individual myosin molecules[11] which pull themselves along their neighbouring actin filament in discrete steps where they (attach to)-(tug on)-(detach from) the actin filament [30]. With enough myosins working together, there will be enough attached guys pulling at any moment to sustain the motion while the detached guys reposition —so that a net contractile motion is established. Because of the discrete nature of the molecular

[8] a structure near the nucleus in animal cells from which microtubules are nucleated

[9] (using much of the same machinery as for focal adhesions)

[10] Hopefully this abuse of language remains a worthy mnemonic on account of its clarity.

[11] myosin molecules bundle together to form long myosin filaments

mechanism underlying muscle contraction, it is interesting to consider that (1) the apparent smooth motion of muscle contraction is actually the result of a massive number of sharp tugs, just as the steady rush of a rain storm is the result of a great number of individual splashes and (2) like the rain storm the pulsed nature of muscle contraction should become evident when observed with sufficient time/force resolution.

Given that it is easy to halve the length of a muscle in $\sim 0.1\mathbf{s}$ (e.g. flexing your bicep) it is reasonable to assume that the individual muscle cells are halving their lengths (more or less) in the same short period. Which means that the myosin molecules (composing sarcomeres $\sim 4\mu\mathbf{m}$) which are on the order of tens of $\mathtt{nm}$ may be sliding a distance of a few $\mu\mathbf{m}$ in that time. So the myosin molecules may be travelling ~ 1000 times their own length per second, which suggests that the timescale of grab, pull, release may be on the order of $\sim 1\mathtt{ms}$.

The cellular microenvironment

On its own a cell is already an impressive creature, but in animals like ourselves trillions of cells must cooperate to build, and function as, a higher-level organism —working together to fulfil their collective goal of passing on the genes they all share. Aside from the knowledge contained in their DNA, the only information that is available to cells is *local* information. As such they are faced with the task of coordinating their actions to produce intricate and complex structures such the nervous system, heart, liver et cetera —guided only by what they can sense in their microenvironment. The more salient information a cell is able to gather from its local environment, the better able it will be to answer said challenge.
The information available to the cell can be divided into 2 rich categories: mechanical and chemical. The ability of cells to sense, react to, and communicate through, their chemical environment is a worthy topic on its own. However we will focus on the *mechanical* aspect of cells' interaction with their environment.

Mechanotransduction —sensing the local mechanical environment

Cells sense the character of their environments not only chemically, but also physically[1, 2]. They *feel* out their location to sense the mechanical properties (e.g. stiffness) of their microenvironment and notice when these properties change in space and time. Usually cells are closely neighboured by other cells which together with the extracellular matrix determine the character of the local mechanical environment. Thus cells' ability to sense the mechanical properties of their microenvironment gives them access to important information about what kind of tissue they inhabit and what is happening —allowing them to know what is the right thing to do. For example, when the mechanical properties of the microenvironment vary with time, the reactions triggered at the cellular level can lead to essential processes such as muscle building (myogenesis).[1] Understanding the influence of the mechanical environment

on stem cell differentiation and gene expression is essential to success in using stem cells for regenerative medicine[24].

In addition, cells respond actively to mechanical stimulation. Epithelial cells are able to detect shear stresses from fluid flow in blood vessels. Cells can also detect *localized* stresses (e.g. at focal adhesion sites[1, 17] or when being poked with a sharp (high-aspect ratio) AFM tip[18, 35]). Localized stresses can lead to cytoskeletal remodelling and translocation of organelles *far away* from the 'point of contact' (i.e. localized source of the stress)[17]. e.g. In cells which were poked near their centre with a force of **10nN** using a sharp AFM tip, mitochondria way out at the cell edge (originally tens of **μm** away from the point of contact) were observed to displace significantly in response to the relatively distant source of mechanical stimulation[35]. Since mitochondria are tied to the cytoskeleton, this is evidence that localized forces can result in large scale cytoskeletal remodelling.

Starting Simple: Vesicles & Lipid Bilayers

Clearly cells are very complicated mechanical objects —eukaryotic cells especially so. As model physical systems lipid bilayer membranes/vesicles have been an attractive starting point for theoretical work, simulations and experiments.

There are multiple ways to organize the relevant theoretical, experimental, and simulation work —how should we go about it? The invention of the optical microscope at the beginning of the 1600s predates the conception of Charles Babbage's Analytical Engine[12] by over 200 years —and the construction of the first electronic computers by over 300 years. Cells and vesicles were therefore accessible to *in-vitro* study long before computer simulation would show potency in this realm. Similarly, most of the relevant theory describing membranes and bilayer vesicles was derived prior to (though sometimes within a few years of) MD simulation of self-assembling amphiphiles. So the temporal arrangement would be: cells, vesicles, theory of bilayer membranes, simulation.

But conceptual coherence overrules historical context. For this reason, we begin with theory, then discuss experiments and simulations.

[12] The first ever design for a Turing-complete computer —described in 1837, though construction was abandoned.

Theory

Area expansion

In 1984 Helfrich & Servuss [23] derived an expression

$$\alpha_{proj} \overset{\text{def}}{\equiv} \left(\frac{\Delta A_{proj}}{A_{\circ proj}} \right)_{\gamma>0} = \underbrace{\frac{k_B T}{8\pi\kappa} ln\left(\frac{\frac{\zeta}{A} + \frac{\gamma}{\kappa}}{\frac{\zeta}{a} + \frac{\gamma}{\kappa}} \right)}_{entropic} + \underbrace{\frac{\gamma}{K_A}}_{direct} , \tag{2.1}$$

relating the increase in a membrane's projected area, to its surface tension γ.

In the above equation, κ is the bending rigidity, K_A is the area compressibility modulus, a is the area per surfactant molecule ('area per lipid') and A_{proj} is the 'projected' area of the membrane (defined below). ζ is a parameter which depends on the membrane shape —for flat membranes $\zeta = \pi^2$, whereas for closed quasispherical membranes $\zeta = 24\pi$.

Equation 2.1 arises from the equipartition of energy among thermally excited normal modes of the bilayer —out-of-plane waves, pumped by the Brownian kick of the surrounding water.

Equation 2.1 has two contributions —one entropic (logarithmic term) due to flattening of undulations, and the other due to direct expansion of the membrane. In the low tension regime, an increase in tension will flatten out undulations thereby increasing the projected area of the membrane. Once most undulations have been flattened, the continued increase in the apparent surface area of the membrane is dominated by a true, direct expansion of the membrane.

Projected Area versus Triangulated Area

The expression derived in [23] uses the symbol A to refer to the 'projected area' (A_{proj}), *not to be confused* with the actual total surface area of the membrane. We denote the *actual* surface area of the membrane A_{tri} (triangulated area), since it is measured using a triangulation algorithm (written by Martin Bertrand[5, 6], see Figure 4.1).

To see the distinction between these two 'areas', it helps to imagine a bed sheet held horizontal to the ground, flapping gently in the wind as it would when you lay it down before a picnic.[13] As the flag flaps in the wind, forming its own undulating hills and valleys, its surface area need not change at all. If at any instant in time, you carefully added up piece by piece the surface area of the whole undulating sheet, it would be approximately what you would have measured were the sheet rolled out flat —with a bit of variation depending on the

[13] (For simplicity, the picnic is taking place on a particular hill chosen so that the ground is normal to the incoming sunlight.)

tension in the sheet.

On the other hand, as it undulates in the wind the sheet casts a shadow on the ground.

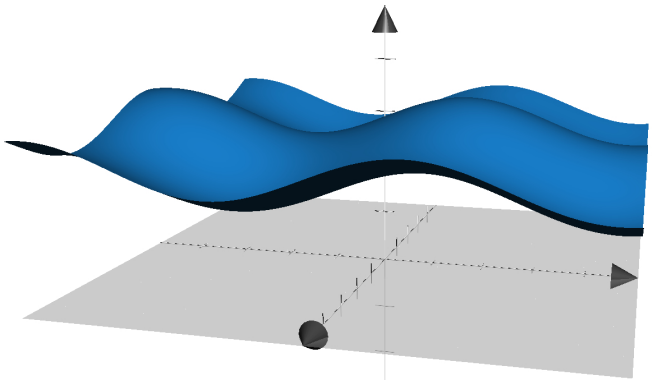

Figure 2.2: The projected area of an undulating sheet is the area of its shadow.

Its shadow may be an irregular shape, and the surface area of this shadow may indeed be *much* less than the total area of the sheet. Relating our sheet to the more general case of a mathematical membrane, A_{tri} measures the actual surface area of the sheet while A_{proj} is like the surface area of its shadow. To remove the complication of the sheet's edges, picture it as being a delineated portion of a much larger membrane which extends in all directions —as a finite opaque square painted onto an infinite clear membrane. As the infinite membrane undulates, the projected area of the painted square will be less than the square's actual surface area.

While the comparison to a shadow is helpful in describing the projected area of a an undulating planar membrane, it breaks down for closed membranes —like vesicles. A_{proj} is defined more robustly as the area of a 'mean' or 'effective' surface associated with the time-averaged shape of the membrane. Consider Figures 2.3 and 2.4:

At any given moment almost *none* of the membrane lies on the mean surface,

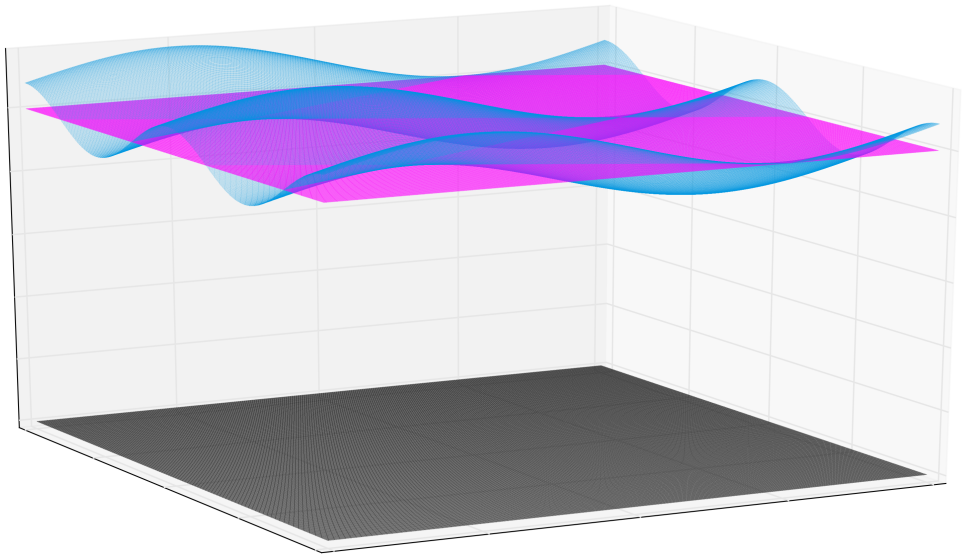

Figure 2.3: Local segment of an undulating membrane, and its corresponding time-averaged mean surface —whose area is the projected area of the membrane. A curved surface can be treated as locally flat (the mean surface is curved on a much larger length scale than the undulations). For a planar membrane surface area is to projected area, what the length of a curved line is to that of a straight line connecting the same two points.

but averaging over timescales much longer than that of the individual fluctuations, the amount of membrane lying to one side of the mean surface is ballanced by a commensurate portion on the opposite side.

In Figure 2.4, we look at the lower-dimensional case of 'projected circumference'.

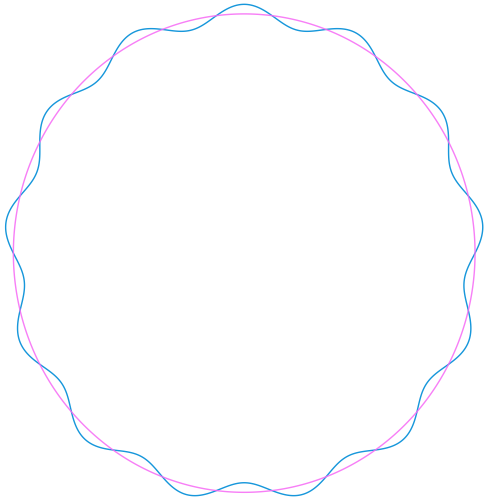

Figure 2.4: The 'mean contour' of an undulating contour. This example is intended to help extend the notion of projected area to closed membranes. Circumference is to a contour what area is to a surface.

While the projected circumference is less than the actual circumference of the contour at any given instant, both contours *enclose the same surface area*. Likewise for a vesicle, the mean surface should be calculated in such a way that it *encloses the same volume* as the undulating membrane.[14]

For any finite tension $\alpha_{proj} < \alpha_{tri}$, since some short wavelength small amplitude undulations must persist. But as tension increases the projected area approaches the actual surface area of the membrane. i.e. If the membrane were unbreakable (able to withstand arbitrary surface tension) then

$$\lim_{\gamma \to \infty} \frac{A_{proj}}{A_{tri}} = 1.$$

[14] While the volume of the vesicle can be compressed somewhat, at any given compression the mean surface and undulating surfaces still must both enclose said volume.

Shape equation

Because we compress our vesicle right up to the point of lysis (inducing very large curvature at the equator) we require a means of calculating the importance of bending when we measure the surface tension. The 'shape equation' (2.4), derived by Hefrich and Zhong-can, provides such a tool.

The shape of a vesicle at mechanical equilibrium will be that which *minimizes* the 'shape energy' $\mathcal{F}_{shape}$ —the free energy of the vesicle as a function of its shape. By setting the first variation of the shape energy equal to 0,

$$0 = \delta^{(1)}\mathcal{F}_{shape} = \delta^{(1)}\left\{\oint dA\left(\frac{1}{2}\kappa\left(c_1 + c_2 - c_\circ\right)^2 + \gamma\right) + \Delta P\int dV\right\}, \tag{2.2}$$

Helfrich and Zhong-can[38] derived the equilibrium shape equation for closed, fluid membranes (e.g. vesicles):

$$\Delta P - 2\gamma H + \underbrace{\kappa(2H + c_\circ)(2H^2 - 2G - c_\circ H) + 2\kappa\Delta_S H}_{bending} = 0.^{15} \tag{2.3}$$

Equation 2.3 augments the Young-Laplace law by adding a correction to the surface tension to include the effect of bending elasticity.

sign conventions

In (2.3) Helfrich defines $\Delta P = (P_{out} - P_{in})$, and $H = -\frac{1}{2}(c_1 + c_2)$. These definitions (as well as the sign convention used for curvatures) differ by a sign from those used in this thesis.[16] However, the change in sign conventions does not alter the equation: Multiplying Equation 2.3 by (-1) gives

$$(-\Delta P_{Helf.}) + 2\gamma(-H_{Helf.}) - \kappa\left(\left(2(-H_{Helf.}) + (-c_{\circ_{Helf.}})\right)\left(2H^2_{Helf.} - 2G - (-c_{\circ_{Helf.}})(-H_{Helf.})\right)\right) \because$$

$$\because -2\kappa\Delta_S(-H_{Helf.}) = 0,$$

which under our standard sign convention —$\Delta P = (P_{in} - P_{out})$, $H = \frac{1}{2}(c_1 + c_2)$ and $c_\circ = -c_{\circ_{Helf.}}$— becomes

$$\boxed{\Delta P - 2\gamma H + \kappa(2H + c_\circ)(2H^2 - 2G - c_\circ H) + 2\kappa\Delta_S H = 0}. \tag{2.4}$$

Therefore, the form of the shape equation is identical under both conventions.

[15] Δ_S is defined in Equation 4.19 —it is called the covariant surface Laplacian (a.k.a. the Laplace-Beltrami operator) on the surface of the vesicle. H and G are the local mean and Gaussian curvatures (respectively).

[16] (Since the Gaussian curvature $G = c_1 c_2$, it is unaffected by the change in sign convention.)

Experiment

Area expansion

In 1990, Evans and Rawicz [11] applied Helfrich and Servuss' formula to the study of real giant unilamellar vesicles. Membrane tension was controlled by partially aspirating the vesicles into micropipettes.
The form of Equation 2.1 used by Evans and Rawicz,

$$\alpha_{E.R.} \cong \frac{k_B T}{8\pi\kappa} \ln\left(1 + \frac{\gamma A_{proj}}{C_o}\right) + \frac{\gamma}{K_A}, \tag{2.5}$$

arises from choosing the tensionless ($\gamma = 0$) state of the bilayer as the 'reference' ($\alpha_{E.R.} = 0$) state.[17]

Evans and Rawicz studied the area expansion of vesicles subject to tensions $10^{-7} \leqslant \gamma \leqslant 10^{-3} \frac{N}{m}$, and observed a logarithmic dependence at low tension followed by a linear dependence at larger tensions —consistent with the model proposed by Helfrich and Servuss.

Further compelling experimental support for the Helfrich model was provided by Dimova et al.[9] in 2009. In this case, giant unilamellar vesicles (GUVs) were deformed using electric fields, and their area expansion plotted against the resulting membrane tension.

Vertical compression

Edith Schäfer et al. [34] published experiments on *real* vesicles in 2013 which were practically *identical* to the simulations performed here. In their experiments, Schäfer et al. used an AFM to compress GUVs between parallel plates and studied the GUVs' deformation at different forces (see Figure 2.5).

Our simulations show the same vesicle contour as observed by Schäfer for GUVs (Figure 4.18). We also obtain a similar force-compression curve, in spite of the massive difference in scale between the giant vesicles used by Schäfer et al. and our tiny vesicle.

A working assumption made by Schäfer is that adhesion to the substrate causes the vesicles' undulations to be largely flattened. Based on our calculations of the surface tension, it appears that this may also be the case for our simulated vesicle. Unfortunately, Schäfer et al. did not report on any dynamic properties (such as relaxation time)[18] of the vesicles.

[17] see Deriving Evans-Rawicz in the Appendix for details.
[18] which will be described in the Discussion section

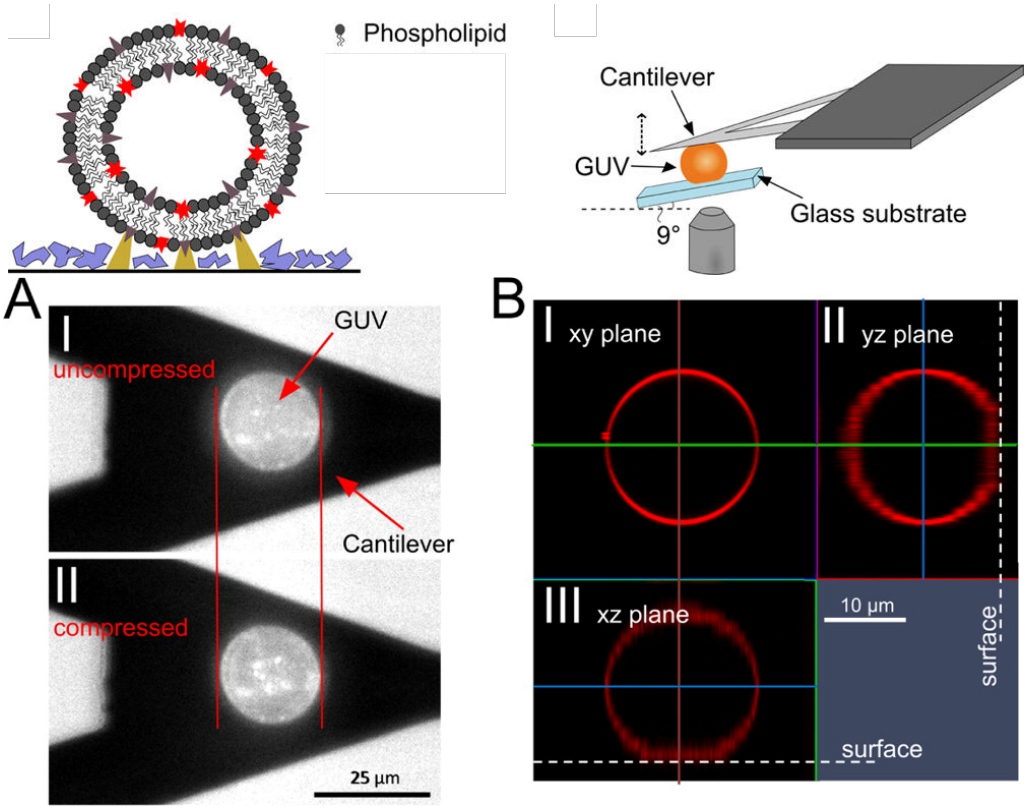

Figure 3. (A) Two optical micrographs (fluorescence combined with bright-field microscopy) illustrating the experimental compression of giant liposomes between a tipless cantilever and a glassy substrate. (I) Vesicle prior to compression. (II) Vesicle at 10% compression. The image was taken through the glass substrate with an inverted microscope beneath the AFM setup. The GUV is labeled with TR-DHPE (0.5 mol %) and adsorbed on a protein-coated glass slide. (B) Confocal laser scanning micrograph showing the contour of a sessile DOPC liposome adhered to an avidin-coated coverslip. The equatorial *xy* plane (I) and the two cross sections (*yz* (II) and *xz* (III)) are shown.

Figure 2.5: Real-world experiments on GUVs analogous to the simulations performed here were carried out by Edith Schäfer et al. in 2013. (Adapted with permission[19] from [34].)

Relaxation time

By subjecting GUVs to pulsed electric fields, Dimova et al.[33, 10] observed relaxation times $\tau \sim 100\mu\mathbf{s}$ for tensions on the order of $\sim 5\mathbf{mN/m}$ (approaching lysis). As with our simulated vesicles, they observed exponential relaxation in response to a step-function type stress.[20]

Relying on dimensional analysis[33], Dimova et al. express the relaxation time as

$$\tau \sim \frac{\eta_m}{\gamma}, \tag{2.6}$$

with η_m being the membrane viscosity and γ the total tension. They find that the dilatational surface viscosity η_d provides a relaxation time which is in agreement with their τ-measurements.[21] However, Dimova et al. did not report any correlation between relaxation times (*below* lysis tension) and the magnitude of the tension-inducing electric field. Instead they present Equation 2.6 as an intuitive relation using dimensional analysis —which gives a good estimate of the 'typical' observed relaxation time. They explain the absence of said correlation in terms of the confounding effect of variable initial (pre-pulse) tension among the vesicles.

Contents

Just as our simulated vesicle can be used as a model system to which more complicated components are added, contents can be added to real giant vesicles in vitro. This is what was done by Helfer[21, 22], when they added a system similar to an actin cortex to the inside of giant vesicles, and compared their rheological properties to those of empty vesicles.

Simulation

From the perspective of simulation work most relevant to our own, there is a spectrum of granularity[22].

Coarse-graining

The scarcity of time means that the best speed of scientific discovery is *as fast as we can go*. Meanwhile, there is no sense in doing something useless, no matter how *quickly*. The tradeoff in simulation, is between detail and speed.

So the cost of granularity is system size —though the price is always going down. Some

[19] No additional uses are granted (such as derivative works or other editions).

[20] We observe the relaxation of the vesicle following sudden *application* of a constant stress, whereas Dimova et al. observe relaxation following the sudden *release* of a previously applied stress.

[21] whereas the in-plane shear viscosity of the membrane was two orders of magnitude too small.

[22] i.e. completeness of detail in terms of the number of particles per molecule and the complexity of the potentials governing their motion

membrane simulations coarse-grain the lipids down to *single* particles. Others have about the same granularity as the present work, treating lipids as simple polymers, surrounded by Lennard-Jones fluid. Introducing more complicated molecular structure and combinations of pair potentials, the Lennard-Jones fluid can be replaced by polarizable fluid —mimicking the polar character of water molecules— and the simple lipid membrane can be replaced by specific, detailed molecules. Eventually we arrive at atomistic simulations, involving many fewer molecules and hence much smaller systems. As computer power continues to grow, the space between simulated vesicles and the real thing seems destined to dwindle. Conceptually, a real vesicle may be thought of as the limiting case of simulations with increasing granularity and system size. But not all detail is worth having. In fact, one strength of the simulation approach is the ability to study simplified model systems where the dominating effects are known.[23]

In *our* case we must optimize the ratio of vesicle diameter to membrane thickness, while preserving the physical character of the bilayer and keeping the number of particles low enough that we can run enough simulations (each of sufficient duration) to measure the relaxation time.

Goetz and Lipowski

In 1998, using a combination of Monte Carlo (MC) and Molecular Dynamics (MD) simulations, Goetz and Lipowski[16, 15] simulated and characterized self-assembling bilayers composed of various types of coarse grained lipids. This study of the self-assembly of coarse-grained amphiphiles was the basis for the bilayer used in these simulations —both in terms of the geometry of the lipids and the choice of effective pair potentials.

Goetz and Lipowski focused on three separate coarse-grained lipid geometries —flexible single-tailed lipids, single-tailed lipids made semiflexible by imposing an angle potential, and semiflexible two-tailed lipids. For all three geometries they observed self assembly into spherical micelles, cylindrical micelles, and bilayers.

[23] As mentioned in Why Simulate Vesicles?

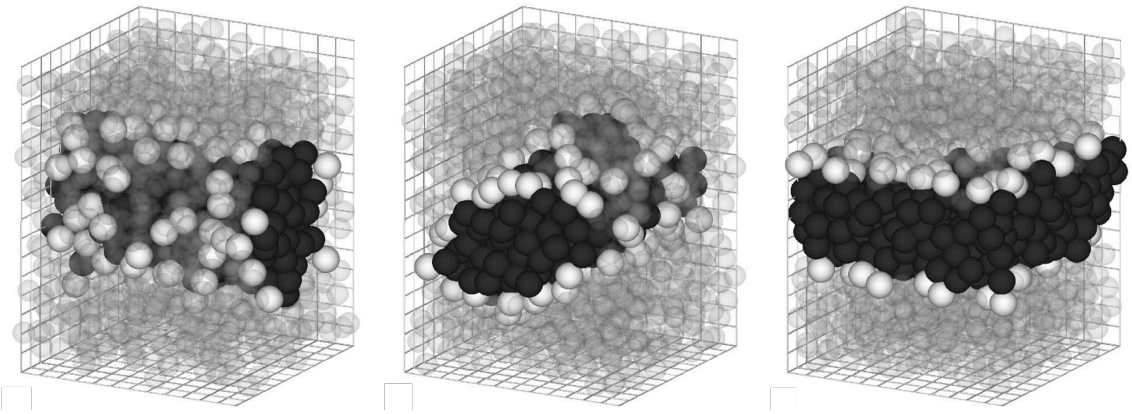

Figure 2.6: Self assembly of coarse-grained lipids into cylindrical micelles(left, centre) and bilayers(right). (Adapted with permission[24] from [16].)

By observing the in-plane diffusion of lipids in the bilayer, Goetz and Lipowski demonstrated that the coarse-grained lipids were indeed forming genuine fluid bilayers. In showing that much of the physical character of lipid bilayers emerges from the self-assembly of such simple coarse-grained molecules, Goetz and Lipowski helped open up the possibility of performing larger scale MD simulations (like the ones discussed here) on a practical time/budgetary scale.

Identical vesicles

As mentioned earlier, one of the major current and potential applications of vesicles is drug encapsulation/delivery. Vesicle sizes range from tens of $\mu\mathbf{m}$ to tens of $\mathrm{n}\mathbf{m}$. Since a vesicle's size is a fundamental parameter affecting its performance in the human circulatory system, control over the size distribution of manufactured vesicles (monodispersity) is a basic requirement for medical application.[25]

One technique for producing monodisperse vesicles is by extrusion through large arrays of nanofluidic channels.In order to improve the understanding of this process and thus the ability to predict the size of vesicles subjected to it, Martin Bertrand[6] published molecular dynamics simulations of the extrusion of vesicles through a variety of pore sizes.

Alison Harman also has run simulations of vesicles in confined Poiseuille flow, investigating the deformation of vesicles which results from flow conditions similar to those of the human blood stream.[20]

[24] No additional uses are granted (such as derivative works or other editions).
[25] Imagine cars were manufactured with an rms length deviation of 50 feet!

Marrink

In 2003 Marrink *&* Mark[28] demonstrated the *self-assembly* of coarse-grained lipids into small unilamellar *vesicles* via MD simulations. The self-assembled vesicles had diameters corresponding to $\cong$ 20nm —approximately the same size as the vesicles simulated in this thesis.

Their goal was to simulate specific biologically relevant lipids with sufficient detail that the mechanical properties (bending rigidity, area compressibility, etc.) and structural features (packing fractions of the two leaflets, stress profile, etc.) of the simulated bilayers would be as quantitatively similar as possible to those of real lipids. As such the simulations included much more molecular detail than our highly coarse-grained 3-bead lipids. For example, their bilayers were composed mostly of dipalmitoylphosphatidylcholine. Each coarse-grained DPPC lipid consisted of 12 beads —a 4-bead head groups and two 4-bead tails. While our simulations use a soft-sphere potential to mimic the hydrophobic effect, Marrink et. al use a set of five Lennard-Jones potentials of varying energy well depths (ε). The lipids also included angle potentials as well as a screened coulomb interaction to approximate the charge distributions of the head groups.

As a result, their solvent density, bilayer thickness, surface area per-lipid and area compressibility agree with experimental values to within $\lesssim 10\%$ —the bending modulus to within an order of magnitude.These simulations also provided detailed information about intermediate stages of the vesicle self-assembly process. They observed vesicle formation occurring on a ns timescale, and studied stress, fluctuations and asymmetry in the bilayer.

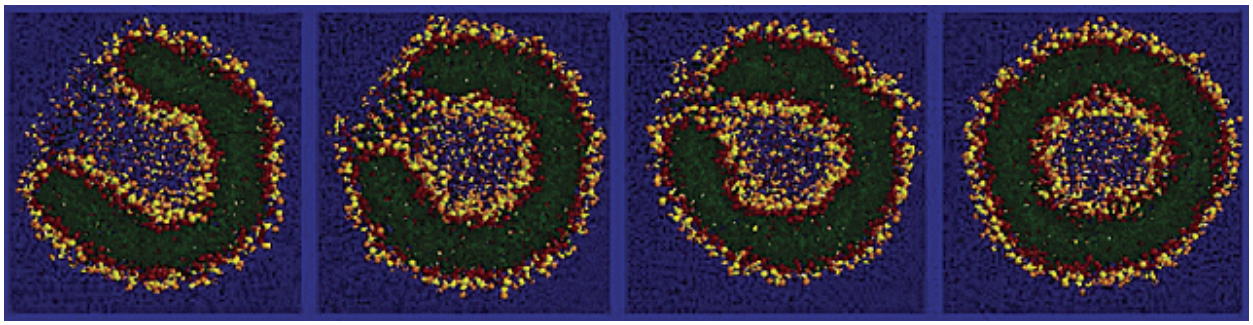

Figure 2.7: Self-assembly of small unilamellar vesicles in purely MD simulations performed by Marrink *&* Mark. (Adapted with permission[26] from [28].)

[26]No additional uses are granted (such as derivative works or other editions).

Hydrophobic Effect[27]

In real water, the hydrophobic effect is actually entropic. Assuming constant temperature as usual, the change in free energy $\Delta\mathcal{F}$ that results from immersing non-polar atoms (e.g. the hydrocarbon tails of phospholipids) in H_2O

$$\Delta\mathcal{F} = \Delta E - T\Delta S \tag{2.7}$$

is dominated by the entropic term.

In liquid water there exists a fluctuating network of hydrogen bonds. When a phospholipid molecule's non-polar tails are immersed in water, the surrounding H_2O molecules will reorient in order to keep the hydrogen bond network intact. They do this by forming a 'cage' of hydrogen bonds enclosing the surface of the lipid tail.[30]. Preserving the network in this way decreases the entropy of the lipid-water system, therefore exposure of the tails to surrounding water molecules increases the free energy $\mathcal{F}$. For entropy to dominate $\Delta\mathcal{F}$ we require

$$\frac{\Delta E}{T} < \Delta S, \tag{2.8}$$

so the hydrophobic effect should grow stronger with increasing temperature (at least in the regime near room temperature where water is far from boiling) if this is the case. As pointed out by [16][28] the simulated 'hydrophobic effect' used here is *energetic*, *not entropic*, because it is simulated using an additional interaction potential $\mathbf{V}_{ss}$ —see Equation 3.3. In order for our simulations to scale with temperature like the true hydrophobic effect we would have to build a temperature dependence into $\mathbf{V}_{ss}$. However in all of these simulations $T = 1$, so temperature scaling was not investigated.

[27] adapted largely from [30]

[28] the work on which the set of potentials used in this paper is based

Surface Tension γ

What is surface tension?

Surface tension is defined in differing ways, sometimes confusingly[4] —because there is much overlap, clarity of communication is essential. Much of the confusion seems to arise from using the terms "surface tension" and "cohesive forces" interchangeably, as well as failing to emphasize the differences between interfaces and membranes. The two main notions to distinguish are those of *interfacial* tension (often called surface tension) at the boundary between two different fluids, and the tension in a *membrane* (also often called surface tension, sometimes called membrane tension) —e.g. skin of a balloon. Interfacial tension is constant (i.e. *independent* of surface area of the interface) —it is determined fully by the temperature, pressure and composition of the interfacing substances. When the area expands new molecules appear at the surface to fill in the space. Membrane tension need not be constant and *no new molecules* incorporate into the membrane as its area expands[29]. With respect to membrane tension verses interfacial tension:

Membranes *stretch*, but interfaces *grow*.

Consider a glob of water floating inside the international space station (ISS)[30]. The cohesive force between two water molecules is much greater than the attraction between water and the gasses of which air is composed. A water molecule inside the glob will be surrounded by cohesive neighbouring water molecules. But a $\boldsymbol{H_2O}$ molecule at the surface(air-water interface) of the glob has as its neighbours both water and air species. As a result water molecules at the surface of the blob have a greater potential energy than those deeper inside —the cohesive forces among $\boldsymbol{H_2O}$ molecules are constantly trying to pull molecules at the surface inward. The potential energy of the blob is minimized by minimizing its surface area, and in the absence of gravity a blob of *any* size[31] will be driven by surface tension to take on

[29] Of course there are more complicated situations wherein e.g. a vesicle in a solution of the lipids from which it is made will accept new molecules into the membrane as it is stretched.

[30] This choice of environment is appealing because we have a familiar air-water interface without the complication of gravity (ignore the world outside ISS).

[31] Within reason...make the blob of water too massive and it will collapse under its own gravity into a black hole!

a spherical shape.

Now replace our floating glob of water with a water-balloon (an idealized water balloon, with no knot). The air-water interface has been replaced by a rubber membrane, which has its own tension —stress tangential to the balloon's surface. The water-balloon will also take on a spherical shape in order to minimize its surface tension.

A vesicle is much more like the water-balloon than the bare glob of water.

Due to the difficulty in measuring the γ we have focused on the more modest goal of determining how γ *varies* with $\boldsymbol{F}_z$ —rather than knowing the absolute value of the tension with high accuracy.

Calculating surface tension: Work approach

In what follows, we derive the formula

$$\boxed{\gamma(F_z) = F_z \cdot \left(\frac{dz}{dA}\right) + \sum_j P_j \frac{dV_j}{dA}}, \tag{2.9}$$

for the surface tension as a function of the squeezing force[32], and discuss its numerical implementation. The sum over j reads

$$\sum_j (\,)_j = (\,)_{\text{inner-fluid}} + (\,)_{\text{membrane}} + (\,)_{\text{outer-fluid}}. \tag{2.10}$$

We obtain this measure of the tension in terms of the work done on the vesicle.

Merits of 'work approach' to γ

An important feature of the 'work approach' ('free energy approach') to calculating the surface tension, is that the *entire membrane* (as well as inner/outer fluid) is included in the calculation.

In contrast, when estimating the tension via geometric means we are forced to consider only a small section of the membrane (ring at the equator). We then infer the tension of the entire membrane from the local curvature and pressure gradient near the equator —using formulae such as the Young-Laplace equation (4.8). However formulae like Equation 4.8 assume a membrane of *negligible* thickness, when the actual curvature and pressure measurements occur on either side of a relatively *thick membrane* —thereby ignoring the *tension profile* across it. The existence of a tension profile across the membrane makes it difficult to properly weight the available curvature and pressure measurements[33] in a way that reliably reproduces the correct tension from formulae such as (4.8).

Because the 'work approach' takes the entire membrane into account, we believe that it is a superior estimate of the true tension in the vesicle membrane, when compared with the local-geometry based estimates.

[32] z, A, P_j and V_j are all functions of F_z

[33] We do not know the local stress profile in the membrane (stress profiles have been measured for similar membranes [16]). We make do with pressure data for the membrane as a whole, and curvature measurements for the bounding surfaces of the membrane.

Work and Free Energy

When we compress the vesicle, we are doing work on the system. Some of this work goes into compressing the inner fluid of the vesicle, some of it goes into stretching the membrane, et cetera. As will be shown momentarily, the total mechanical work done on the system by the AFM Tip in squeezing the vesicle is equal to the change in the system's free energy $\mathcal{F}$,

$$dW = d\mathcal{F}. \tag{2.11}$$

This is useful, since the surface tension γ can be defined in terms of the differential free energy,

$$\boxed{d\mathcal{F} = \gamma\, dA - \sum_j P_j\, dV_j},^{34} \tag{2.12}$$

of the system (i.e. vesicle and solvent).

We want to isolate for γ. Since $\mathcal{F}(F_z)$, $A(F_z)$, and the $V_j(F_z)$ all vary smoothly as functions of the force, we may write

$$d(\langle\rangle) = \frac{d(\langle\rangle)}{dF_z}\, dF_z, \tag{2.13}$$

where $\langle\rangle = V_j$, A, or $\mathcal{F}$. There is therefore no problem writing

$$d(\langle\rangle) = \frac{\left(\frac{d(\langle\rangle)}{dF_z}\cancel{dF_z}\right)}{\left(\frac{dA}{dF_z}\cancel{dF_z}\right)}\left(\frac{dA}{dF_z}\cancel{dF_z}\right) = \left(\frac{d(\langle\rangle)}{dA}\right) dA.^{35} \tag{2.14}$$

By the logic of (2.14), we rewrite (2.12) as

$$\gamma\cancel{dA} = \frac{d\mathcal{F}}{dA}\cancel{dA} + \sum_j P_j \frac{dV_j}{dA}\cancel{dA},$$

[34] In Equation 2.12, $dA \equiv d(\textit{surface area})$ —but *what* surface area exactly?. Because each side of the bilayer is an interface, each monolayer contributes to $d\mathcal{F}$. We therefore choose $dA = d\left(A^{\textit{inner leaflet}} + A^{\textit{outer leaflet}}\right)$ in calculating the surface tension.

[35] Equation 2.14 may seem like a deliberate waste of ink. It isn't. The point it's meant to convey is that we are *able to compute* each $d(\langle\rangle)$ as $d(\langle\rangle) = \frac{d(\langle\rangle)}{dA}\, dA$ *only because* we know every $\langle\rangle$ as a function of the same independent variable —specifically F_z.

and finally get to the point of all this:

$$\boxed{\gamma = \frac{d\mathcal{F}}{dA} + \sum_j P_j \frac{dV_j}{dA}}. \tag{2.15}$$

Equation 2.15 is saying that if we can measure the relevant changes in the pressure, volume, and free energy of the system, then we'll know the surface tension. The important thing to notice about (2.15) is that it lacks partial derivatives —it contains only total derivatives and can therefore be evaluated numerically.

We will now obtain Equation 2.11, which becomes very useful when inserted into Equation 2.15. All simulations were performed at fixed temperature,

$$T = \textit{const.}$$

So the differential free energy is

$$d\mathcal{F} = d(U - TS) \tag{2.16}$$

$$= dU - \left\{ (\cancel{dT})^{0} S + T(dS) \right\},$$

that is

$$\boxed{d\mathcal{F} = dU - T\,dS}. \tag{2.17}$$

Now

$$dU = \underbrace{T\,dS}_{\text{“}dQ\text{”}} + dW,$$

where W denotes work done *on* the system, and Q is "heat".

Plugging this new dU into Equation 2.17 gives

$$d\mathcal{F} = \underbrace{(\cancel{T\,dS} + dW)}_{dU} - \cancel{T\,dS},$$

$$\longrightarrow \boxed{d\mathcal{F} = dW}. \tag{2.18}$$

Hence, *because* $\boldsymbol{T = const}$ *the free energy change* $\boldsymbol{d\mathcal{F}}$ *is the mechanical work* $\boldsymbol{dW}$ *done on the system.* We are then free to rewrite (2.15) as

$$\boxed{\gamma = \underbrace{\frac{dW}{dA} + \sum_j P_j \frac{dV_j}{dA}}_{\text{know how to measure this}}}, \tag{2.19}$$

and thus, the question becomes "what is dW?".

But the mechanical work is something *we can readily calculate.* We know the equilibrium position of the AFM Tip at each given applied force. Because we allow the vesicle to equilibrate at each force, a fit to the force-compression data gives us a curve which is equivalent to what we would observe for quasistatic deformation of the vesicle. Therefore, we can *directly* calculate the work W done on the system (using the data shown in Figures 4.9, 4.10, 4.11 and 4.12) since

$$\textit{Work} = \textit{Force} \times \textit{distance}, \tag{2.20}$$

$$\longrightarrow W = \int F(z)\,dz; \tag{2.21}$$

and the differential work may then be computed as

$$dW = \left(\frac{dW}{dz}\right) dz,$$

$$= \left(\frac{d}{dz}\int F(z)\,dz\right) dz,$$

$$\therefore \boxed{dW = F(z)\,dz}. \tag{2.22}$$

Therefore the total derivative $\frac{dW}{dA}$ is readily calculated from the height and surface area measurements at each force:

$$\frac{dW}{dA} = \left(\frac{dz}{dA}\right)\left(\frac{dW}{dz}\right)$$

$$\longrightarrow \boxed{\frac{dW}{dA} = F(z)\left(\frac{dz}{dA}\right)}. \tag{2.23}$$

By plugging (2.23) into (2.19), we finally arrive at Equation 2.9:

$$\gamma = F(z)\frac{dz}{dA} + \sum_j P_j \frac{dV_j}{dA}.$$

Viscosity[36]

What is viscosity?

Viscosity or 'thickness' is a measure of a fluid's resistance to flow. Put another way, the viscosity represents the rate at which momentum in a fluid diffuses radially outward perpendicular to the relative motion of the fluid particles. The following example is meant to illuminate the previous sentence: Consider water sitting stationary in a long cylindrical pipe (forget about the ends of the pipe). Divide the water into thin imaginary concentric cylindrical layers. Now accelerate the pipe along its axis (z-direction). What happens to the fluid? Initially, almost nothing. Only the fluid right next to the pipe wall will notice as this outermost 'layer' is pulled along with the wall —closer to the middle of the pipe, the fluid is undisturbed. But now that the outermost layer of fluid is moving along with the pipe wall, the next layer in toward the centre will begin to move too. Slowly, the velocity "works its way in" toward the centre of the pipe, until momentum has been imparted to every particle in the system. Momentum has diffused through the fluid, perpendicular to the relative velocity of the fluid layers. The more viscous the fluid, the more rapidly this momentum reaches fluid at the centre of the pipe.

If we repeat this experiment with a stationary rod placed at the centre of the moving pipe, and simply pull the pipe in the z-direction with a velocity v_z relative to the rod, the steady-state solution will be one in which the fluid velocity matches the motion of each boundary, and varies smoothly between the boundaries.

To think more clearly about viscosity, consider the simpler case of two parallel plates. Viscosity can be measured as follows: Take two plates of surface area A sufficiently large that edge effects can be ignored. Bring the two plates close together and lubricate the space between them with a fluid whose viscosity you would like to know. Holding one of the places fixed, pull on the other with a force F so that the two places slide across one another face to

[36] The stress and viscosity calculational schemes detailed in this section were borrowed from [19] and [29].

face with velocity Δv —their separation Δy kept constant. The viscosity of the fluid relates the shear-force per unit area to the velocity gradient ($\frac{\Delta v}{\Delta y}$):

$$\frac{F}{A} = \eta \frac{v}{y}, \tag{2.24}$$

or

$$\boxed{\frac{F}{A} = \eta \frac{\partial v_x}{\partial y}}. \tag{2.25}$$

Note that Equation 2.25 has the dimensions of force per unit area like pressure, but the force vector $\vec{F}$ is in the plane of the plate —rather than perpendicular to the plate as would be the case for pressure. The stress related to viscosity is the shear-force analogue of pressure.

Stress tensor $J_{\alpha\beta}$

Much of the following discussion is adapted from 'Molecular Dynamics Simulation' by J.M. Haile.[19]

We wish to know the viscosities of the various Lennard-Jones fluids and polymer melts which were used as our solvent and vesicle contents. We have available two mathematically equivalent methods for calculating the viscosity of a MD fluid —the Green-Kubo method and Einstein Mean Squared Displacement (MSD) method. Both methods require the microscopic stress tensor —$J_{\alpha\beta}$[37]. Both approaches are outlined here, but the method which was used was indirect Einstein MSD.

"The components $J_{\alpha\beta}$ measure the rate at which β-directed momentum is transported in the α-direction."[38]

Viscosities of different fluids were calculated in separate simulations, each consisting of a periodic box containing only the fluid of interest. This was done for the sake of simplicity and computational efficiency. Computing the stress tensor of a periodic simulation box of homogeneous fluid is much easier than doing so for a single component (e.g. fluid inside vesicle) surrounded by other components (e.g. vesicle membrane) in a more complicated system. In the former we simply compute a single stress tensor for the entire periodic simulation box at each point in time. In the latter, rather than one simulation box we have multiple closed sub-boxes —with stress tensors $J_{\alpha\beta}(\vec{x}_j)$[39]. This *greatly* increases computational complexity

[37] Another commonly used term is 'pressure tensor' which is given as (pressure tensor) = −(stress tensor). However one must be careful, since different sign conventions are used when defining the stress tensor. i.e. (one person's stress tensor) = ±(another person's).

[38] paraphrasing Haile

[39] where $\vec{x}_j$ denotes the centre of the j^{th} sub-box

of the task, as it requires that we accurately portion the total stress among all the sub-boxes.

A procedure for allocating stress among the sub-boxes would be something like the following. Picture a 3-dimensional rectangular grid, so that the entire simulation box is portioned into cubes ('sub-boxes'). For every pair of particles in the system:

- Take the line ($\vec{\ell}$) joining the two particles —This line will cross a certain set of sub-boxes,
- Determine *which* sub-boxes, and calculate the length (ℓ_j) of the *portion* of $\vec{\ell}$ which falls within each sub-box,
- The ratio $\frac{\ell_j}{\ell}$ gives the *portion* of the force between the two particles which is to be allocated when calculating the stress tensor in sub-box j.

Aside from being somewhat complicated the above 'local stress calculation' would also be prohibitively slow, unless incorporated into the simulation engine(HOOMD-blue). While knowledge of the local stress tensor is very desirable —for it would have permitted direct calculation of e.g. the surface tension in the membrane— this calculation was beyond the scope of this work.

Therefore, to calculate each of the various LJ-fluids' viscosities, a separate simulation was run which consisted solely of box of said material. In such a setup, the stress tensor for the entire simulation box is:

$$J_{\alpha\beta}(t) = \sum_{i=1}^{N} \frac{1}{m_i} p_{i\alpha}(t) p_{i\beta}(t) + \frac{1}{2} \sum_{i \neq j}^{N} r_{ij\beta}(t) F_{ij\alpha}(t), \tag{2.26}$$

which, since $p_i = m_i v_i$ and $m_i \equiv 1$ (in these simulations, all particles have mass=1), simplifies to

$$\boxed{J_{\alpha\beta}(t) = \sum_{i=1}^{N} v_{i\alpha}(t) v_{i\beta}(t) + \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} r_{ij\beta}(t) F_{ij\alpha}(t)} \tag{2.27}$$

where $v_{i\alpha}$ is the α-component of particle i's velocity, and $r_{ij\beta}$ is the β-component of the vector pointing *to* particle i *from* particle j .

Green-Kubo

The Green-Kubo method for calculating transport coefficients (viscosity, diffusion coefficient, etc.):

$$(\text{transport coefficient}) = \int \left(\begin{array}{c} \text{time-correlation function of} \\ \text{appropriate microscopic current} \end{array} \right) dt\,. \tag{2.28}$$

In the case of $\overbrace{\text{shear viscosity } \eta}^{(\text{transport coefficient})}$, the time-correlation function is the 'stress autocorrelation function' denoted $\eta(t)$

$$\eta = \int_0^\infty dt\, \eta(t). \tag{2.29}$$

with

$$\eta(t) = \frac{\rho}{N k_B T} \frac{1}{3} \sum_{\alpha<\beta} \langle J_{\alpha\beta}(t_\circ) J_{\alpha\beta}(t_\circ + t) \rangle_{t_\circ}. \tag{2.30}$$

N is the number of particles in the simulation box, and ρ is fluid density.

Equation 2.30 needs some unpacking. The notation "$\langle \textit{function} \rangle_{t_\circ}$" means we're averaging over $t_\circ$ (initial time), so that

$$\langle J_{\alpha\beta}(t_\circ) J_{\alpha\beta}(t_\circ + t) \rangle_{t_\circ} = \frac{1}{M} \sum_{n=1}^{M} J_{\alpha\beta}(t_n) J_{\alpha\beta}(t_n + t) \tag{2.31}$$

averages over M initial times t_n. The $\frac{1}{3}\sum_{\alpha<\beta}$[40] comes into (2.30) because we are averaging the off-diagonal terms of a matrix. Said matrix has components $\langle J_{\alpha\beta}(t_\circ) J_{\alpha\beta}(t_\circ + t) \rangle_{t_\circ}$ and *it is symmetric* because $J_{\alpha\beta}$ is[41] —so we need *only* average the xy, yz, xz components.

Putting this all together gives

$$\boxed{\eta(t) = \frac{\rho}{3 k_B T} \frac{1}{N} \sum_{\alpha<\beta} \left\{ \frac{1}{M} \sum_{n=1}^{M} J_{\alpha\beta}(t_n) J_{\alpha\beta}(t_n + t) \right\}}. \tag{2.32}$$

[40] to be clear: $\sum_{\alpha<\beta} (\,)_{\alpha\beta} \equiv (\,)_{xy} + (\,)_{yz} + (\,)_{xz}$

[41] see Equation 2.27

And the explicit expression for the viscosity (2.29) is

$$\boxed{\eta = \frac{\rho}{3k_B T}\frac{1}{NM}\int_0^\infty dt \left\{\sum_{\alpha<\beta}\sum_{n=1}^{M} J_{\alpha\beta}(t_n)J_{\alpha\beta}(t_n+t)\right\}}. \tag{2.33}$$

Equations 2.33 and 2.29 can be remembered compactly as

$$\boxed{\eta = \langle \eta_{\alpha\beta} \rangle_{\alpha\beta}}, \tag{2.34}$$

with

$$\eta_{\alpha\beta} = \frac{1}{Vk_B T}\int_0^\infty dt \, \langle J_{\alpha\beta}(t_\circ)J_{\alpha\beta}(t_\circ+t)\rangle_{t_\circ}. \tag{2.35}$$

Indirect Einstein approach to obtain η

Mathematically equivalent to Green-Kubo is the 'indirect Einstein' approach to calculating viscosity.[42] We have opted to use the latter due to ease of implementation, and ease of comparison with other work (e.g. [29]) wherein this technique is used.

Einstein relation

The 'indirect Einstein' mean squared displacement approach states that

$$\eta_{\alpha\beta} = \frac{1}{2Vk_B T}\lim_{t\to\infty}\frac{d}{dt}\left\langle \left(\int_{t_\circ}^{t} dt \, J_{\alpha\beta}(t)\right)^2 \right\rangle_{t_\circ}, \tag{2.36}$$

so that the viscosity may be extracted from a linear fit to the slope of

$$\frac{1}{2Vk_B T}\left\langle \left(\int_{t_\circ}^{t} dt \, J_{\alpha\beta}(t)\right)^2 \right\rangle_{t_\circ}, \tag{2.37}$$

at relatively long simulation times (see Figure 4.27). As in Equation 2.34, the viscosity calculation is performed for each distinct $\alpha\beta$-pair, and averaged.

[42] See [29], as well as 'Molecular Dynamics Simulation' by J.M. Haile.

Simulation Methodology

Simulation Box

The model consists of approximately 140,000 particles in a simulation box with periodic boundary conditions. Therein we find a small unilamellar vesicle filled and surrounded by Lennard-Jones fluid (the explicit solvent) at an initial density of $0.8\frac{particles}{\sigma^3}$.

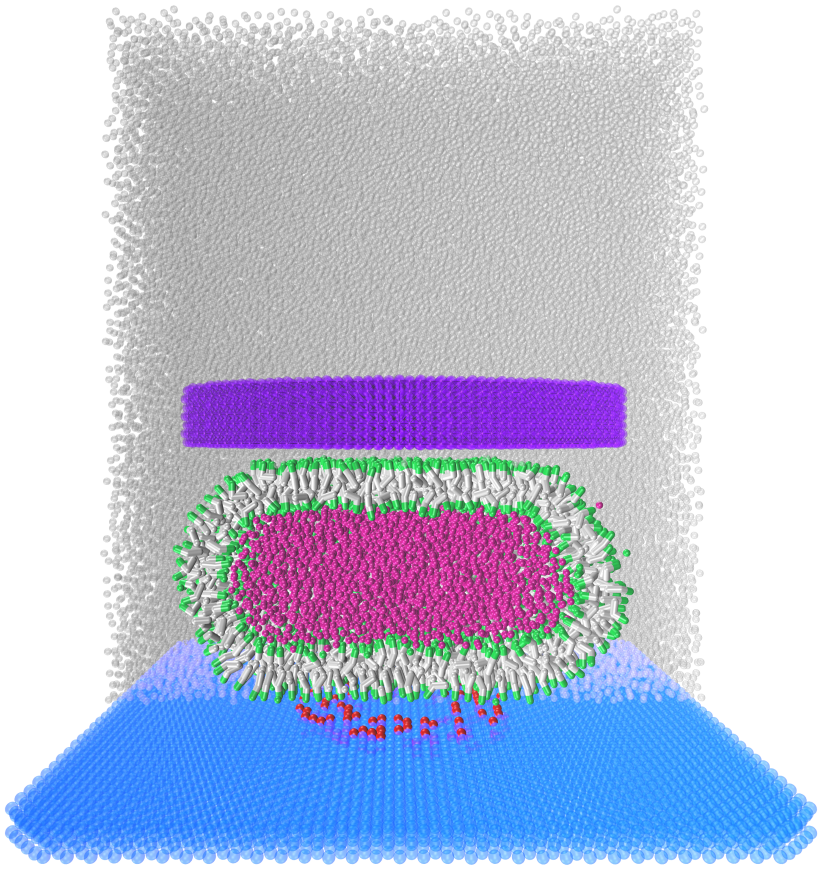

Figure 3.1: (Colour online) The simulation box contains ≅140,000 particles. The 'AFM Tip' is shown in violet. Substrate particles are coloured light blue. The bilayer is shown with lipid heads/tails coloured green/white, respectively. Inner/outer fluid particles are shown in pink/grey respectively. On the upper right-hand-side of the vesicle, a single fluid particle which has leaked out from inside of the vesicle is visible.

The vesicle is sandwiched between a substrate and an AFM Tip —both consisting of

fluid-like particles, arranged in an fcc lattice. The vesicle membrane is composed of 3000 coarse-grained *lipids* —each lipid having one hydrophilic 'head' particle and two hydrophobic 'tail' particles, connected by harmonic bonds. Adding its 4859 inner fluid particles, the vesicle consists of 13859 particles in total. By comparison, the AFM tip consists of 7740 particles.

Inertia

The typical mass of an AFM cantilever is $m_{Tip} \sim 10\text{ng}$[36], and the vesicles squeezed by Schäfer and coworkers[34] had radii $R_{ves} \sim 10\mu\text{m}$, giving a mass $m_{ves} \sim 4\text{ng}$. So in [34] the relative inertia of the tip was $\left(\frac{m_{Tip}}{m_{ves}}\right)_{real} \approx 2.5$. Using the particle numbers listed above, the corresponding ratio for the simulations is $\left(\frac{m_{Tip}}{m_{ves}}\right)_{sim} = \frac{7740}{13859} \cong 0.56$.[1]

The M.D. Model — in detail

AFM Apparatus

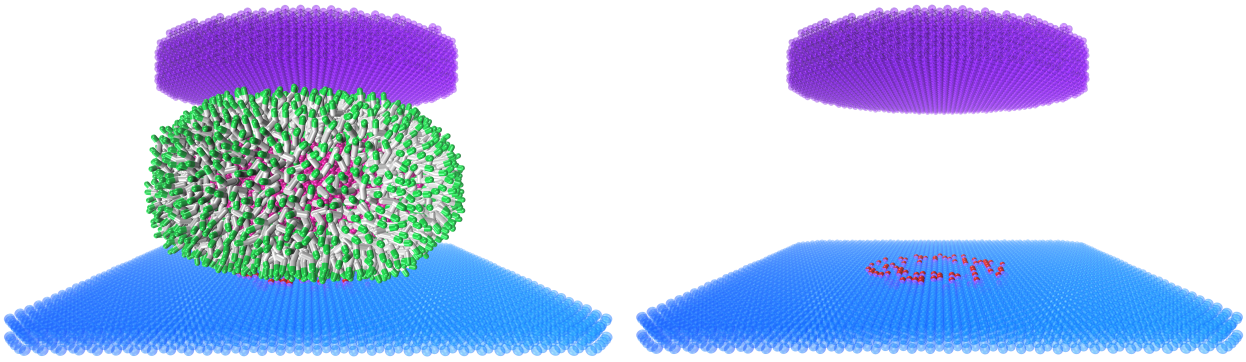

Figure 3.2: AFM Apparatus. In addition to the ordinary substrate particles, a bullseye of randomly distributed 'sticky' particles was placed at the centre of the substrate to ensure adhesion. Without this *adhesion site*, the vesicle would slip out from underneath the AFM Tip.

[1] All particles in the simulation have the same mass $m_j \equiv 1\mu$.

Thermostat —Dissipative Particle Dynamics

In biophysics, soft-matter and physical chemistry we are almost always studying systems that are *warm*. In fact, it takes constant work to thermally isolate a system of particles from the ambient heat of its surroundings. In nature, absent any externally powered thermal regulation mechanism, a system of particles will constantly exchange heat with its surroundings until thermal equilibrium is reached —the temperature of the surroundings matches the temperature of the system of particles of interest.

Even *at* equilibrium, the exchange of heat —kinetic energy associated with the particles' randomly directed velocities— between the system and its surroundings continues. However since on average the same momentum flows out of the system as flows into it from the environment, no *net* flow of thermal kinetic energy (heat) occurs —hence "thermal equilibrium".

Initially equilibrated systems can be pushed out of thermal equilibrium by external forcing. However the surroundings will absorb excess heat from (or provide heat to) the system such that the average kinetic energy (temperature) of all particles in the system returns to equilibrium with the environment[2]. The immensity of the surroundings relative to the system means that the temperature of the surroundings is practically unaffected by exchange of heat with the system. For this reason, in statistical mechanics and thermodynamics the surrounding environment is often referred to as a 'Heat reservoir'.

Thus in the local systems we wish to simulate, it is not energy that is conserved, but *temperature* $\boldsymbol{T}$. But in simulations there *are* no surroundings. So we need to inject(extract) kinetic energy into(out of) the system, in the form of randomly directed momentum distributed among the particles *as if* there were a surrounding thermal environment —this is the role of the thermostat.

Benefits of DPD

[2] e.g. A syringe partially filled with air. We plug the end of the syringe and wait for it to equilibrate. Next we force the plunger down, holding it in the new compressed position. The pressure increases inside the (sealed) syringe. But so does the temperature since the air molecules receive an extra bit of momentum when they ricochet off of the descending plunger. The heat which has been injected into the air immediately begins to heat the inner walls of the syringe, as the rapidly moving inner air molecules collide with the bits of plastic that make up the wall, causing them to jiggle. This jiggling is of course transmitted through the plastic, and so the heat works its way through to the outer surface of the syringe and heats the surrounding air —until the syringe and the air within it are back in thermal equilibrium with the environment. We now pull back on the plunger. The internal pressure falls, and along with it the inner air temperature (now the air molecules that ricochet off of the plunger *lose* momentum). Now it is the surroundings that heat the walls of the syringe and eventually the low temperature air contained within, restoring thermal equilibrium. Because of the vast difference in size between the surroundings (earth's atmosphere) and the system (air contained in syringe), the temperature of the surroundings is practically unchanged throughout the entire process —it acts as a heat *reservoir*, absorbing and providing thermal energy to the system without changing its own temperature provided that we don't pump heat into the system too rapidly.

The thermostat used in our simulations is HOOMD-blue's implementation[32] of the Dissipative Particle Dynamics (DPD) thermostat. We have used the DPD thermostat because it conserves momentum[32] both locally and globally, ensuring proper hydrodynamics. Momentum transport throughout the simulation box is enabled via local velocity rescaling. While this DPD-thermostat can accommodate a time-varying temperature, we keep $\boldsymbol{T}$ fixed so that the simulations are $\boldsymbol{NVT}$ —the ***N***umber of particles, ***V***olume and ***T***emperature of the system are *all held constant*.

System setup — initial state

Each simulation began from the same initial state, albeit with a different random seed. The initial state was prepared using both the ESPResSo and HOOMD-blue simulation engines. In ESPResSo the positions of all particles in the system were initialized, and bonds between them declared. Once setup, the system was warmed up and run briefly, then saved in the form of an XML file. Next, the XML file detailing the system was read into HOOMD-blue, where the simulation ran for $\sim 10^7$ time steps —to allow the vesicle to equilibrate. Finally, this equilibrated system was itself output as an XML file, which was used as the starting configuration for future simulations.

While ESPResSo was essential to setup the original vesicle and AFM apparatus, future modifications such as filling the vesicle with polymers and pumping fluid into the vesicle were done using python scripts written by Ben Barlow. Said scripts depend heavily on the MDAnalysis package.

Potentials

Interactions were governed by the same potentials used in [16, 6], plus two additional potentials —one constrains the motion of the AFM Tip, the other ensures localized adhesion of the vesicle to the substrate. The former of these additional potentials is a cylindrical constraint. Each 'Tip particle' is subjected to a cylindrical harmonic potential, preventing the particle from deviating significantly from its equilibrium position in the xy-plane. This keeps the AFM Tip centred and level —as each of its constituent particles is constrained to move along an axis in the z-direction— while allowing motion perpendicular to the substrate. The entire AFM Tip crystal is effectively "riding on rails". The latter potential is a modified version of the same shifted Lennard-Jones potential as used in [16, 6], where the strength of the attractive term has been tuned via an α-parameter

$$V_{LJ} = 4\varepsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \alpha\left(\frac{\sigma}{r}\right)^{6}\right], \tag{3.1}$$

provided by HOOMD-blue. In these simulations, the above pair potential was implemented for the interaction of a small number of substrate particles (in the centre of the substrate, coloured red in Figure 3.2) with the lipid heads. The enhanced attraction caused the lipid heads to adhere to this 'bullseye', keeping the vesicle centred under the AFM Tip.

The downward force exerted by the AFM Tip on the vesicle was implemented by applying to each individual particle in the Tip crystal an external force equal to the total desired force $F_z \in [10\text{F}, 260\text{F}]$ divided by the number of particles in the Tip.

Pair-wise interactions

The two key interaction potentials in these simulations are the Lennard-Jones potential V_{LJ}

$$V_{LJ} = 4\varepsilon\left(\left(\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right)\right), \tag{3.2}$$

and the soft-sphere potential

$$V_{SS} = ar^{-9}. \tag{3.3}$$

The default interaction potential in these simulations is V_{LJ} —it governs the interaction

between any two particles of the same type, as well as most interactions involving two particles of differing type. Exceptions are as follows: Tail particles interact with other particle types via V_{SS}, but still interact with each other via V_{LJ}. In this way, V_{SS} is used to give the coarse-grained lipids their amphiphilic character —i.e. it mimics the hydrophobic effect which holds the vesicle together. Adhesive substrate particles interact with the lipid heads via Equation 3.1, as discussed above.

Bonded interactions

Bonds between monomers in lipid molecules (as well as in the polymers placed inside the vesicle in certain simulations) are governed by a harmonic potential

$$\text{Harmonic:} \quad \mathbf{V}_{bond} = \frac{k}{2}(r - r_\circ)^2, \tag{3.4}$$

with $k = 5000$. Bonds among particles making up the AFM Tip crystal, as well as the bonds between the substrate particles and their anchor points are implemented using $\mathbf{V}_{bond}$ with $k = 3000$.

Origin of the Lennard-Jones potential

The Lennard-Jones potential

$$V_{LJ} = 4\boldsymbol{\varepsilon}\left(\left(\frac{\boldsymbol{\sigma}}{r}\right)^{12} - \left(\frac{\boldsymbol{\sigma}}{r}\right)^{6}\right) \tag{3.5}$$

contains two competing terms —an attractive term which goes as $-r^{-6}$ and a repulsive term which goes as r^{-12}. Where do these terms come from? The attractive (r^{-6}) term has a strong physical basis. The repulsive term is empirical —and a convenient approximation for computers since $r^{-12} = (r^{-6})^2$. As such we're only going to derive the attractive term.

Derivation of attractive term —dispersion forces, induced dipoles

Much of what follows is adapted from Fritz London's 1936 paper "The General Theory of Molecular Forces"[27]. The attractive term in the LJ potential originates from the interaction of instantaneous dipoles (resulting from quantum fluctuations) with dipoles induced by them in nearby atoms. Due to the uncertainty principle, a neutral atom which has on average zero dipole moment will at any instant of time have some random net dipole moment.

We can approach this problem from two directions —classical electrodynamics and quantum mechanics. Consider two neutral atoms, separated by a distance several times their radii. Think of their dense positively charged nuclei and sparse negatively charged electron clouds. Classically, we might do the following: Imagine that atom$^{\#}1$ undergoes a fluctuation wherein

the electrons' centre of mass moves a distance x_1 away from that of the nucleus —momentarily creating a dipole $\mathbb{p}_1$ with electric field $\mathbb{E}_1$. The dipole field $\mathbb{E}_1$ will now be felt by atom$^{\#}2$, displacing *its* negative charges' centre of mass away from that of its positives by a distance x_2. Assuming a weak electric field, the restoring force on atom$^{\#}2$'s electrons is

$$\mathbb{F}_2 = -kx_2 = -q_2\mathbb{E}_1,$$

and the induced dipole moment is

$$\mathbb{p}_2 = qx_2 = \frac{q_2^2\mathbb{E}_1}{k}.$$

From the multipole expansion we take the approximation for the potential energy

$$V_{12} = \left(q_{\text{tot}}\Phi - \mathbb{p}\cdot\mathbb{E} - (\text{quadrupole}) + \cdots\right) \propto -\mathbb{p}_2\cdot\mathbb{E}_1 \quad \Big|_{\text{total charge: } q_{\text{tot}}=0}$$

$$\implies V_{12} \propto -\mathbb{E}_1^2(\vec{r}_{21}),^{3}$$

and since the dipole field $\mathbb{E}_1(\vec{r}_{21})$ goes as r_{21}^{-3},

$$V_{12} \propto -r_{21}^{-6} \quad \textit{Q.E.D.}$$

But truly, the attractive term in V_{LJ} is quantum mechanical in nature. Quantum mechanics tells us that atoms are actually fluctuating dipoles, since due to the uncertainty principle their charges can never be at rest. 'Fluctuations' of the type mentioned above are a consequence of this zero-point motion. Returning to our two atoms, the instantaneous dipole moment of one atom produces an electric field, which is felt by the other, creating an induced dipole. As a result, there is a kind of synchronization of the two atoms' oscillating dipole moments. The two atoms' charge distributions are correlated in their motion.

"...the dispersion energy arises from intermolecular correlation of charge fluctuations..."[4]

The r^{-6} dependence can also be obtained via perturbation theory.

[3] $\vec{r}_{21} \overset{\text{def}}{\equiv} \vec{r}_2 - \vec{r}_1$; so that the electric field felt by atom$^{\#}2$ due to atom$^{\#}1$ is $\mathbb{E}_1(\vec{r}_{21})$.

[4] A. David Buckingham

Units

Computer simulations are intrinsically dimensionless. In order to draw an analogy between events occuring *in silico* and the dynamics of a real-life physical system, we must relate the particle mass $\boldsymbol{\mu}$, Lennard-Jones length $\boldsymbol{\sigma}$, and the unit of simulation-time $\boldsymbol{\tau}$ to their real-world counterparts.[5] Once expressed in terms of physical units such as metres $\mathbf{m}$, kilograms $\mathrm{k}\mathbf{g}$ and seconds $\mathbf{s}$, the results (predictions) of a computer simulation become amenable to experimental testing in the laboratory. Here we follow the method laid out by Goetz and Lipowski [16] to convert our units from virtual to physical. This procedure (given in *complete* detail below) yields

$$\begin{cases} \boldsymbol{\mu} \approx 6 \times 10^{-26} \mathrm{k}\mathbf{g} \quad \cong (\textit{mass of 2 water molecules, or 3 carbon atoms}) \\ \boldsymbol{\sigma} \approx 0.5 \mathrm{n}\mathbf{m} \\ \boldsymbol{\tau} \approx 2 \mathrm{p}\mathbf{s} \\ \boldsymbol{\varepsilon} \approx 4 \times 10^{-21} \mathbf{J} \\ \mathbf{F} \approx 10 \mathrm{p}\mathbf{N}, \end{cases} \tag{3.6}$$

where $\mathbf{F}$ is the dimensionless unit of force.

Mass —$\boldsymbol{\mu}$

In these simulations, every particle is assigned the same mass —$m_i \equiv \boldsymbol{\mu}$. Since the fluid beads are meant to represent water particles, a lower-bound on $\boldsymbol{\mu}$ can be found by calculating the weight of a single H_2O molecule. Wolfram|Alpha gives the molecular weight of water as $m_{H_2O} \cong 18.0153 \frac{\mathbf{g}}{\mathbf{mol}}$, we have:

$$\boldsymbol{\mu} \geqslant m_{H_2O} \cong \left(18.0153 \frac{\mathbf{g}}{\mathbf{mol}}\right) \times \frac{1}{\mathbf{N_A}} = \left(\frac{18.0153}{6.022141 \times 10^{23}}\right) \mathbf{g} \cong 2.992 \times 10^{-26} \mathrm{k}\mathbf{g}.$$

Beads making up the tails of the simulated phospholipids provide an upper bound, as they may represent up to four CH_2 molecules. The molecular weight of a single CH_2 group is $\cong 14.0266 \frac{\mathbf{g}}{\mathbf{mol}}$, so that

[5] that is, to the dimensions of the material system which the experimenter is attempting to model.

$$\mu \leqslant \left(\frac{4 \times 14.0266}{6.022141 \times 10^{23}}\right) g \cong 9.317 \times 10^{-26} kg.$$

Length —σ

Here, the Lennard-Jones length[6] is taken to be the nearest-neighbour distance between pairs of particles in the simulation.[7] Thus, a lower bound for σ is just the separation between two fluid beads —that is, the effective radius of a H_2O molecule.

The mass-density of water *defines* the gram —$\rho_{H_2O} = \frac{N}{V} = 1 \frac{g}{[cm]^3}$. Using the above calculated value for the mass of a single water molecule, the number density of water molecules is, roughly

$$\rho_\# = \frac{\rho_{H_2O}}{m_{H_2O}} \cong \frac{\left(\frac{1g}{[10^{-2}m]^3}\right)}{2.992 \times 10^{-26} \frac{kg}{\text{molecule}}} = \frac{\left(\frac{10^6 g}{m^3}\right)}{2.992 \times 10^{-23} \frac{g}{\text{molecule}}},$$

$$= \frac{\left(\frac{10^{29}}{m^3}\right)}{2.992 [\text{molecule}]^{-1}} = \frac{10^{29}}{2.992} \frac{\text{molecules}}{m^3}.$$

Thus

$$\sigma \geqslant \left(\frac{1}{\rho_\#}\right)^{\frac{1}{3}} \cong \left(\frac{2.992}{10^{29}}\right)^{\frac{1}{3}} m \cong 0.31 nm.$$

If a lipid tail bead represents at most 4 CH_2 groups, then by inspection of the following figure, the maximum distance between these beads along the lipid chain must be on the order of 4 carbon-carbon bond lengths:

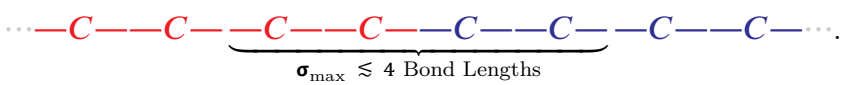

According to Wolfram|Alpha, the $C——C$ bond length in butane is approximately $0.1536 nm$. We thus have

$$\sigma \leqslant 4 \times 0.1536 nm = 0.6144 nm.$$

[6] We let the Lennard-Jones length σ equal the bond length

[7] excluding the AFM Tip, of course

Energy —$\boldsymbol{\varepsilon}$

The energy units in these simulations are defined as $\boldsymbol{\varepsilon} = k_B T = 1$ —where k_B denotes Boltzmann's constant, and $\boldsymbol{\varepsilon}$ is the Lennard-Jones energy. Since the Lennard-Jones fluid which fills the simulation box is supposed to represent liquid water at SATP,[8] we have

$$\boldsymbol{\varepsilon} \sim k_B \times (25^\circ\mathbf{C}) = \left(1.38065 \times 10^{-23} \frac{\mathbf{J}}{\mathbf{K}}\right) \times (298.15\mathbf{K}) \cong \begin{cases} 4.116 \times 10^{-21}\mathbf{J} \\ 2.479 \dfrac{\mathbf{kJ}}{\mathbf{mol}} \end{cases}.$$

Time —$\boldsymbol{\tau}$

The unit of simulation time is

$$\boldsymbol{\tau} = \sqrt{\frac{\boldsymbol{\mu}\boldsymbol{\sigma}^2}{\boldsymbol{\varepsilon}}}.$$

Plugging in the above conversions yields

$$\begin{cases} \boldsymbol{\tau} \geqslant \sqrt{\dfrac{(2.992 \times 10^{-26}\mathbf{kg})(0.31\mathbf{nm})^2}{4.116 \times 10^{-21}\mathbf{J}}} \cong 0.84\mathbf{ps} \\ \vdots \\ \boldsymbol{\tau} \leqslant \sqrt{\dfrac{(9.317 \times 10^{-26}\mathbf{kg})(0.6144\mathbf{nm})^2}{4.116 \times 10^{-21}\mathbf{J}}} \cong 2.92\mathbf{ps} \end{cases}.$$

Force —$\mathbf{F}$

Now that the units for mass, length and time have been obtained, we can determine a conversion for the dimensionless force unit $\mathbf{F}$:

$$\mathbf{F} = \boldsymbol{\mu}\frac{\boldsymbol{\sigma}}{\boldsymbol{\tau}^2} \Longrightarrow 6.714\mathbf{pN} \leqslant \mathbf{F} \leqslant 13.15\mathbf{pN}.$$

[8] Standard Ambient Temperature and Pressure

Dimensions Used

In summary, the conversion from dimensionless to physical units yields values in the range

$$\boxed{\begin{array}{c} 2.992\times10^{-26}\mathrm{kg} \leqslant \mu \leqslant 9.317\times10^{-26}\mathrm{kg} \\ 0.31\mathrm{nm} \leqslant \sigma \leqslant 0.61\mathrm{nm} \\ 0.84\mathrm{ps} \leqslant \tau \leqslant 2.92\mathrm{ps} \\ 6.714\mathrm{pN} \leqslant F \leqslant 13.15\mathrm{pN} \end{array}} \leftarrow \begin{cases} \mu \approx 6\times10^{-26}\mathrm{kg} \\ \sigma \approx 0.5\mathrm{nm} \\ \tau \approx 2\mathrm{ps} \\ F \approx 10\mathrm{pN}, \end{cases}$$

and $\varepsilon \approx 4\times10^{-21}\mathrm{J}$. Throughout this work, measurements are given in dimensionless Lennard-Jones units, unless indicated otherwise.

Results

Measurements — Observing The Vesicle

Part of the beauty of working with computer simulations is that, since we have complete knowledge of the positions & momenta and the forces acting on all particles in the system at each instant of simulation time, one can measure practically anything —given sufficient effort and ingenuity.

There are however some quantities which we know with absolute certainty (e.g. applied force), others which are known to very high accuracy (e.g. vertical compression of vesicle) and still others that are much harder to determine accurately (e.g. surface tension). There are lots of different measurements to analyze, and in this analysis the 'independent variable' on which a measurement depends will affect the physical information gleaned. Because we are studying a membrane the surface tension γ is a much more interesting and appropriate choice as the independent variable than the applied force $\boldsymbol{F}$. However when setting up a MD simulation of the kind discussed here, one can*not* directly *choose* the surface tension as a parameter —rather a squeezing *force* $\boldsymbol{F}$ *is chosen*, and a surface tension results from this force. Until we actually run the experiment, we don't know what the resulting tension will be. In all of the data to be presented $\boldsymbol{F}$ is the *true* independent variable (i.e. input) on which all other variables depend —e.g. $\gamma(\boldsymbol{F})$, $\alpha(\boldsymbol{F})$, $\tau(\boldsymbol{F})$, $\varepsilon(\boldsymbol{F})$. For this reason we first present the various measurements in terms of how they vary with the applied force, and only then go on to analyze the more physically interesting relations *among* composite quantities (e.g. $\alpha(\boldsymbol{F})$ vs. $\gamma(\boldsymbol{F})$).

Stresses and Strains

We will always want to measure stress-strain pairs, as these variables allow us to describe the mechanical properties of materials in a scale independent way. For instance, wood is compressible —like all materials. If we place a wooden beam end-to-end in a giant vice[1] and squeeze it with a force w, the length ℓ of the beam will decrease by an amount $\Delta\ell$.[2]

[1] Imagine that this vice is as big as we want it to be, so that it can squeeze arbitrarily long beams.

[2] Choose the z-axis as that running parallel to the beam, so that ℓ is the z-length of the beam.

Naturally, $\Delta\ell$ increases with f. But if we repeat this experiment for many different beam lengths ℓ_j subject to the same lengthwise force f, and measure their resulting $\Delta\ell_j$ we find that $\Delta\ell_j$ *depends* on ℓ_j. So $\Delta\ell$ is not a satisfactory description of the deformation (strain) occurring in the wood which makes up the beam —we want a description of deformation which is *independent* of the size of the object being deformed. In this case that measure is $\varepsilon = \frac{\Delta\ell}{\ell}$, generally referred to as a 'strain'. While $\Delta\ell$ will vary from beam to beam, ε which describes the *local* deformation at every point inside the wood which makes up the beam does not. The applied force f is unsatisfactory in the same way $\Delta\ell$ is —the same *force* applied to a tree trunk will yield a very different result ($\frac{\Delta\ell}{\ell}$) when used to squeeze a toothpick! Rather than the force squeezing the beam, we should measure the *stress* $\Sigma_{beam} = \frac{f}{A}$ —the ratio of the squeezing force divided by the beam's cross sectional area[3]. Whether it be tree-trunk or toothpick, a beam subjected to the same Σ will undergo the same ε.[4] In this way, stress/strain variables are the natural language to use when describing mechanical properties of *materials* themselves, rather than some particular object made out of them.

The commensurate stress-strain variables we will focus on are

$$(\text{stress}, \text{strain}) = \begin{cases} (\Sigma, \varepsilon) \equiv (\text{applied force} \div \text{contact area}, \text{vertical compression}) \\ (\gamma, \alpha) \equiv (\text{surface tension}, \text{area expansion}), \end{cases} \tag{4.1}$$

where

$$\alpha \overset{\text{def}}{=} \frac{A - A_\circ}{A_\circ}. \tag{4.2}$$

Surface tension γ and area expansion α are the natural variables to describe the state of a membrane, whereas Σ and ε are the natural choice for a beam. However, the latter will not be ignored entirely as they do still provide some useful information. Consider a vesicle being squeezed along the z-axis, in the same manner as the beam discussed above. What is Σ in this case? We still know the squeezing force f, but to get the correct expression for $\Sigma_{vesicle} = \frac{f}{?}$, we need to know what quantity to use in place of the beam cross sectional area. To determine the stress commensurate with vertical compression of the vesicle, we need to divide the downward force applied to the vesicle by the *surface area* A_c *of the region of the membrane in contact with the AFM Tip*. The stress $\frac{F}{A_c}$ is actually just the pressure difference ΔP between the inner and outer fluid which is induced by squeezing the vesicle.

For the sake of making comparisons with experiments on real vesicles we will also examine how the vertical compression of the vesicle (as a fraction of the maximum vertical compression

[3] 'stress' and 'pressure' have the same units, and are often used to refer to the same thing, with pressure = −(stress)

[4] That is, if made of similar materials with the same elastic constants.

the vesicle can withstand without rupturing) varies with the applied force. NOTE: this latter pair are *not* a stress-strain pair, though they give some of the same information.

Area Expansion

Surface tension drives the vesicle to minimize its surface area, constrained by its volume. The shape which minimizes the ratio of surface area to volume is a sphere. Because the vesicle contents resist compression, any deviation from a spherical shape will increase its surface area. As a result, *the vesicle's surface area expands when we squeeze it.*

Time evolution

For each simulation, the surface area of the vesicle is calculated as a time-series.

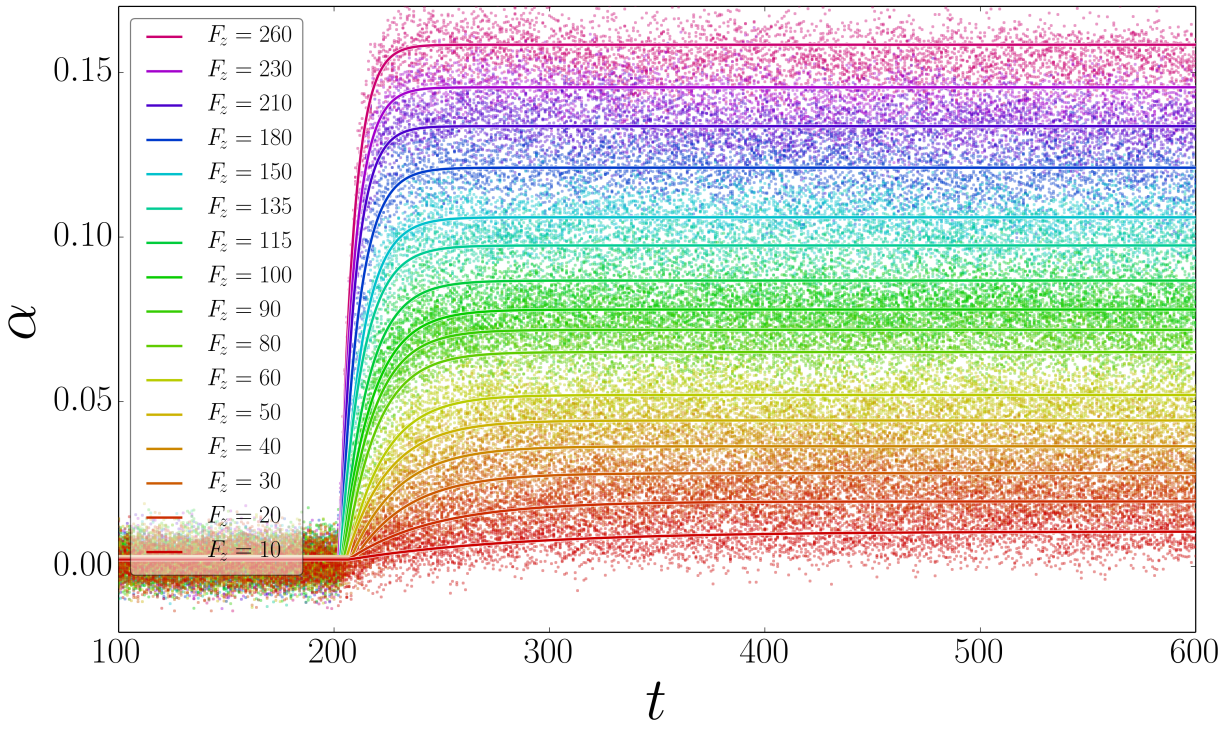

Figure 4.1: Triangulated area expansion timeseries of vesicle at different forces. The calculation is performed using a script written by Martin Bertrand[6], which implements Nina Amenta's 'crust' algorithm [3] to triangulate the inner and outer leaflets of the vesicle. Squeezing force is indicated by colour, with red= 10_F and violet= 260_F. Solid lines indicate fits to $\alpha(t)$ from which the relaxation time is obtained.

Each of the vesicle's inner and outer leaflets have their surface area calculated separately.

We can curve-fit both data-sets independently or combine them into a single weighted[5] average. Figure 4.1 tells us that our simulated vesicle can withstand much greater area expansion than e.g. red blood cells, which will undergo lysis for $\alpha \gtrsim 4\%$.[6]

Equilibrium

Using the data selection procedure depicted in Figure 4.9, we obtain the equilibrium (i.e. post-relaxation) surface area at a variety of forces —see below.

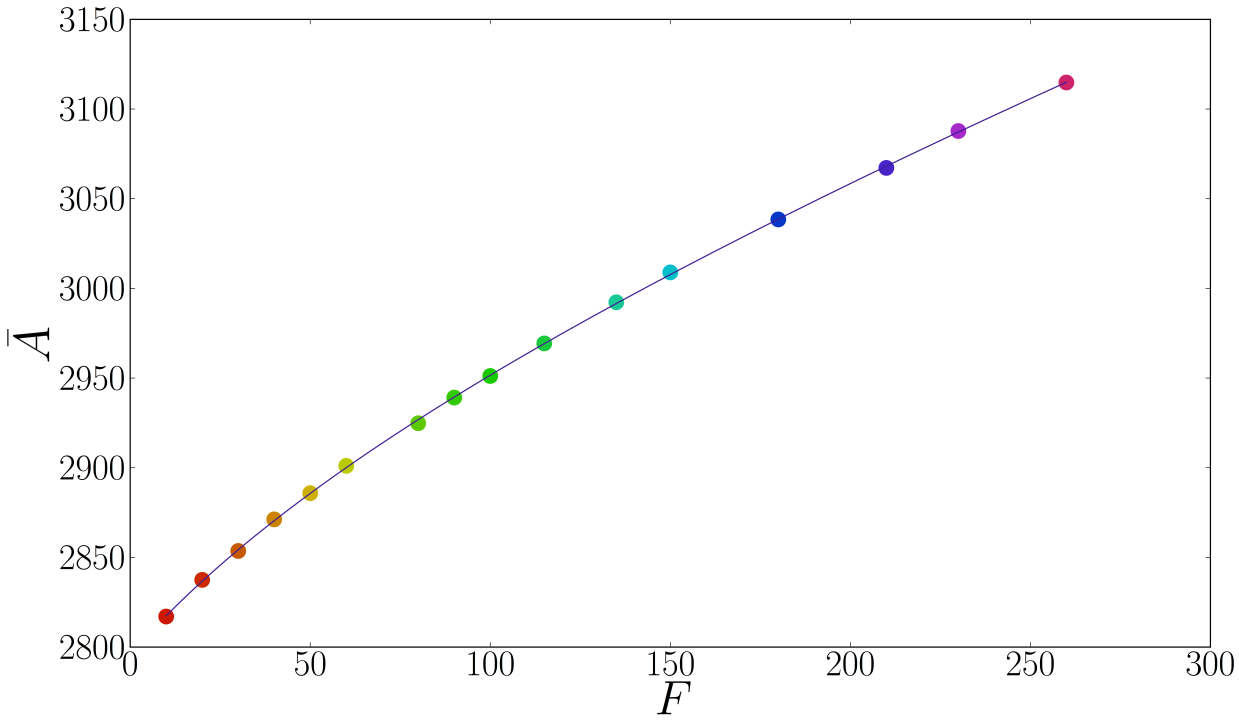

Figure 4.2: *Post-relaxation* triangulated surface area of vesicle (weighted average of inner/outer leaflet surface area) at different forces. Each point represents an ensemble average of ~ 10 simulations. The marker size is roughly one standard-deviation. (Colouring is redundant in this figure, as colour corresponds to squeezing force, which is indicated on the horizontal axis. Colour was merely included for ease of comparison with other figures which use the same colour-scheme, e.g. Figure 4.1.)

[5] weighted by the number of lipids in each leaflet as a fraction of the total number of lipids in the membrane

[6] This may be an artifact of the hydrophobic effect (which is entropic in nature) having been simulated using a soft-sphere potential.

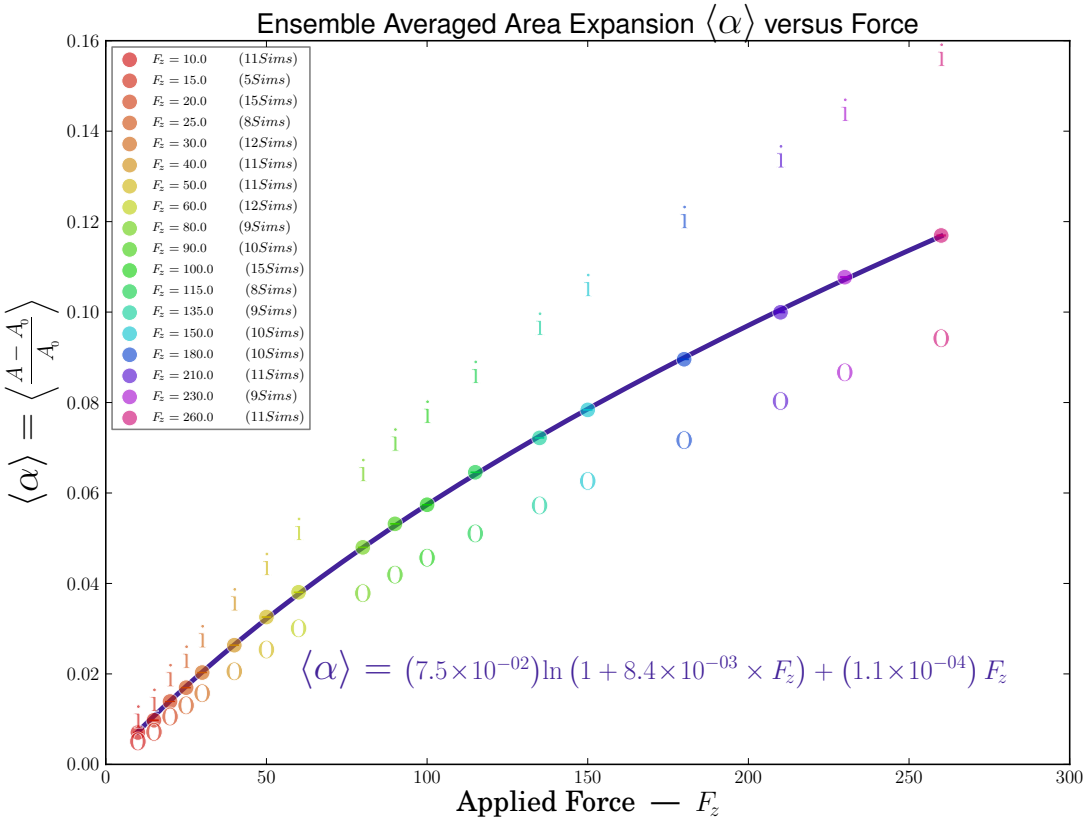

Figure 4.3: Area expansion of vesicle versus force. This plot shows that the inner leaflet expands more than the outer leaflet when the vesicle is squeezed. The expansion of the inner(outer) leaflet is plotted using i's(o's) and their weighted average is plotted with filled circles. For $F \gtrsim 270\text{pN}$ the vesicle is unstable and will undergo lysis. (The inset formula comes from the Helfrich model —see "Background" and "Discussion" sections.)

Because the relative area expansion of the inner surface is greater than that of the outside[7] (see Figure 4.3), the signal-to-noise ratio is greatest for the inner leaflet area expansion timeseries. As a result, the relaxation time is more easily obtained from the inner leaflet data.

[7] This is probably specific to squeezing *small* vesicles.

Vertical Compression

Time evolution

The motion along the z-axis of the AFM Tip's centre of mass is displayed throughout relaxation. Simulations at different forces are indicated by different colours in the figure below.

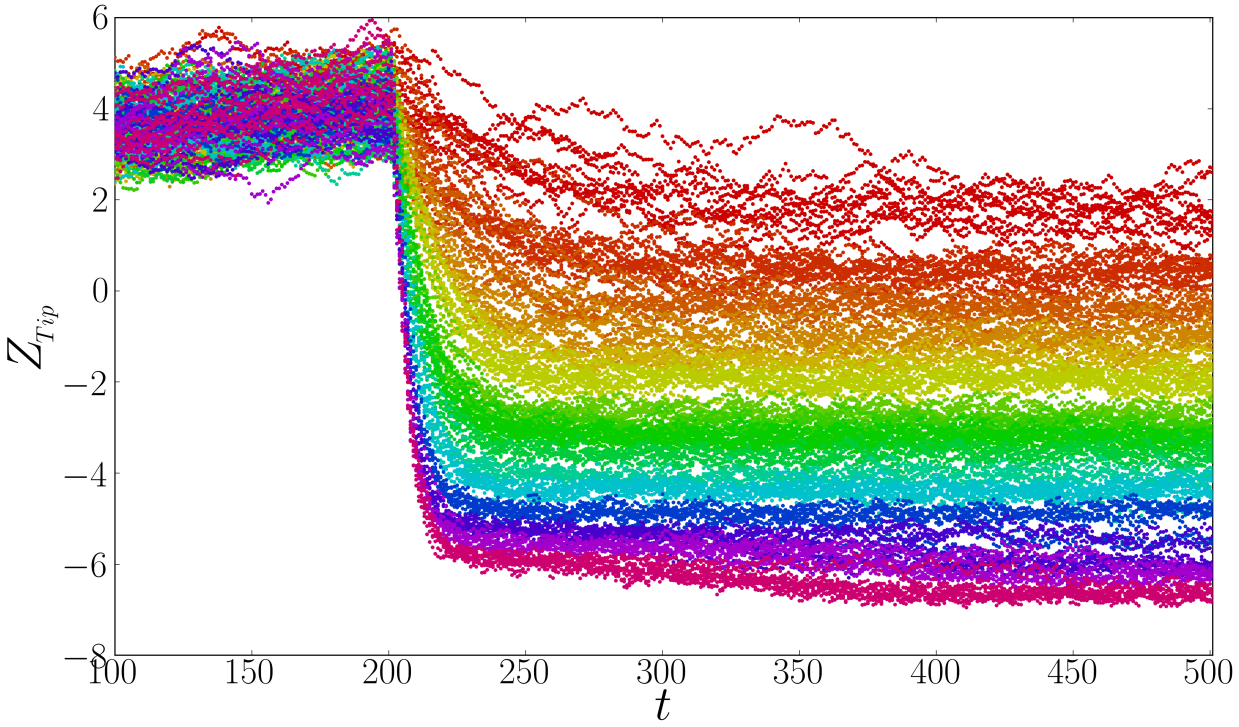

Figure 4.4: Timeseries of AFM Tip z-coordinate. This is *not* the distance between the AFM Tip and the substrate, but merely shows how the Tip moves during relaxation. (Colour indicates squeezing force —legend is identical to that of Figure 4.1.)

Relative to the vesicle, the inertia of the simulated AFM tip is $1/5^{th}$ the relative inertia of the real tip used in [34] (see Inertia under Simulation Methodology). Since the inertia of the AFM tip dampens undulations, the smaller relative inertia of the simulated AFM tip ought to alter the power spectrum of fluctuations in both α and z_{Tip}.

Equilibrium

Using the data selection procedure depicted in Figures 4.9 and 4.10, we measure the equilibrium (post-relaxation) height of the AFM tip. The change in said height, as a function of the applied force corresponds to 'vertical compression' of the vesicle, shown below.

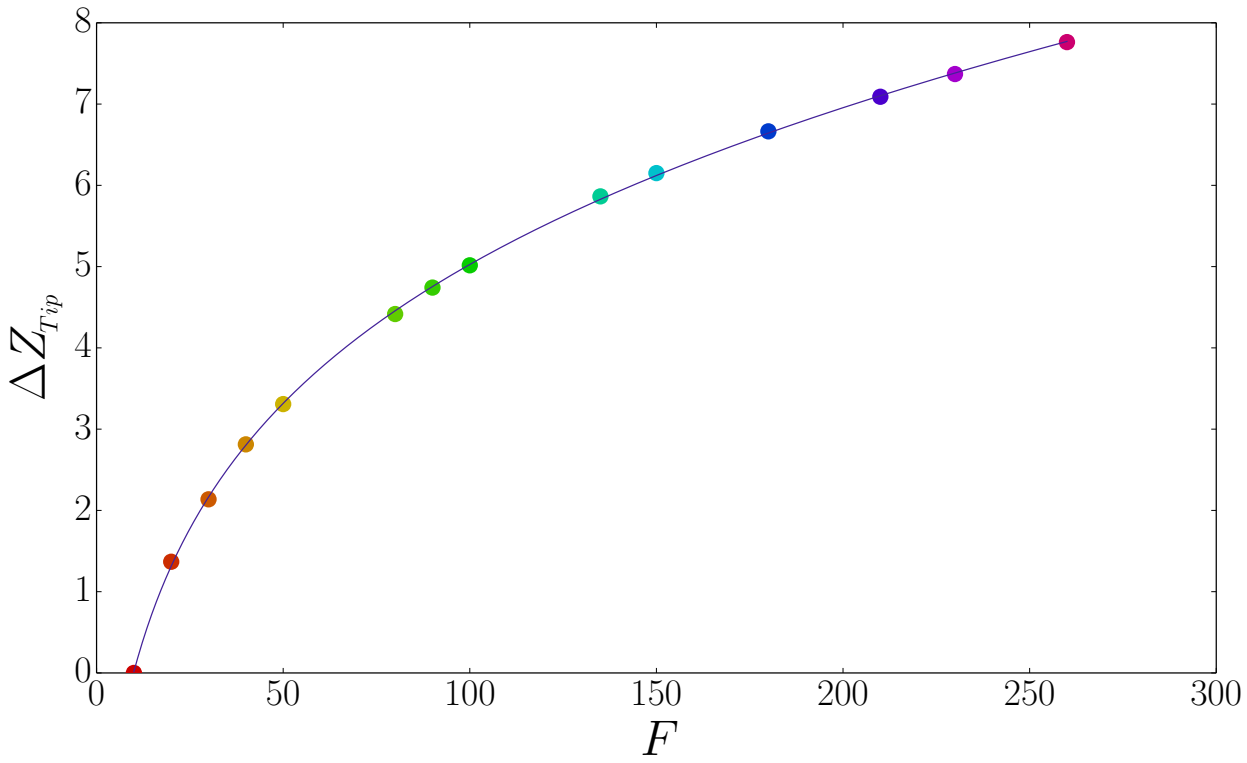

Figure 4.5: Vertical compression of vesicle, relative to the height at $F_z = 10$. (Obtained from the raw data shown in Figure 4.4, using the averaging procedure shown in Figure 4.10.) Said compression is measured via the change in the (equilibrium) z-coordinate of the AFM Tip's centre of mass, as a function of the force. The marker size is $\gtrsim$ one standard-deviation. (Colouring is redundant in this figure, as colour corresponds to squeezing force, which is indicated on the horizontal axis. Colour was merely included for ease of comparison with other figures which use the same colour-scheme, e.g. Figure 4.1.)

The data displayed in Figure 4.5 is analogous —but not identical[8] to— the measurements performed by Schäfer et al.[34] on giant vesicles.

[8] Schäfer and coworkers control the *compression* of the vesicle (by adjusting the height of the AFM tip), and measure the resulting force. In contrast, we control the *force* and observe the resulting compression. So while the data presented in [34] includes measurements at zero force, we require a nonzero minimum applied force before we can meaningfully measure the AFM tip height.

Relaxation Time

Relaxation[9] at a given force is well described as an exponential saturation

$$\alpha(t) \propto \left(1 - e^{-\frac{t}{\tau}}\right), \tag{4.3}$$

see Figures 4.1 and 4.7. This type of viscoelastic creep response corresponds to the 'Kelvin-Voigt' model, or the more general 'Standard Linear Solid' (SLS) model.

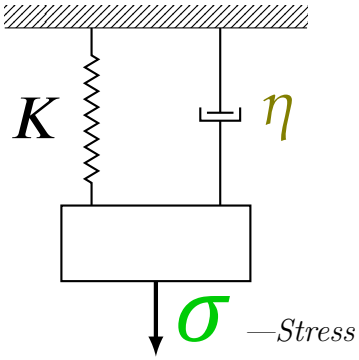

$$\varepsilon(t) = \frac{\sigma_\circ}{K}\left(1 - e^{-\frac{t}{\tau}}\right) \quad \text{—Strain}$$

Figure 4.6: Kelvin-Voigt model of a linear viscoelastic solid.

In this model $\tau \sim \frac{\eta}{K}$, i.e. the relaxation time is presumed to be independent of the stress.

Even at relatively high forces (>100pN), fluctuations in the vesicle's surface area are fairly large —on the same order of magnitude as the mean area expansion. For this reason, when fitting for the relaxation time at a given force, data from multiple equivalent simulations are combined and then fit as one large data set (See Figure 4.7). This way, the influence of random fluctuations from any particular simulation on the outcome of the curve-fit is greatly reduced. That is, the relaxation time is fit to an ensemble of simulations.

[9] That is, the time-evolution of the system after we activate the squeezing force.

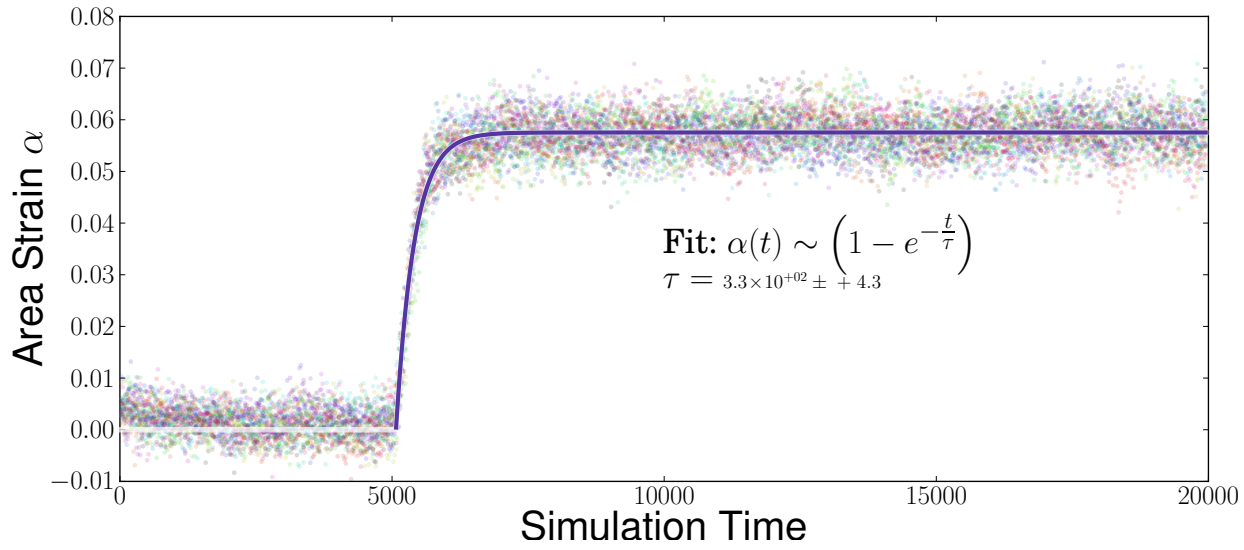

Figure 4.7: Ensemble fit to creep response of bilayer at $F_z = -100\text{pN}$. For each value of the applied force, data from multiple simulations are combined and then fit as one large data set. This helps to reduce the uncertainty on the relaxation time, by reducing influence of noise from any particular simulation on the fit.

As it turns out, the relaxation time shows a *strong* force dependence (see Figure 4.8), which is discussed in the Discussion section.

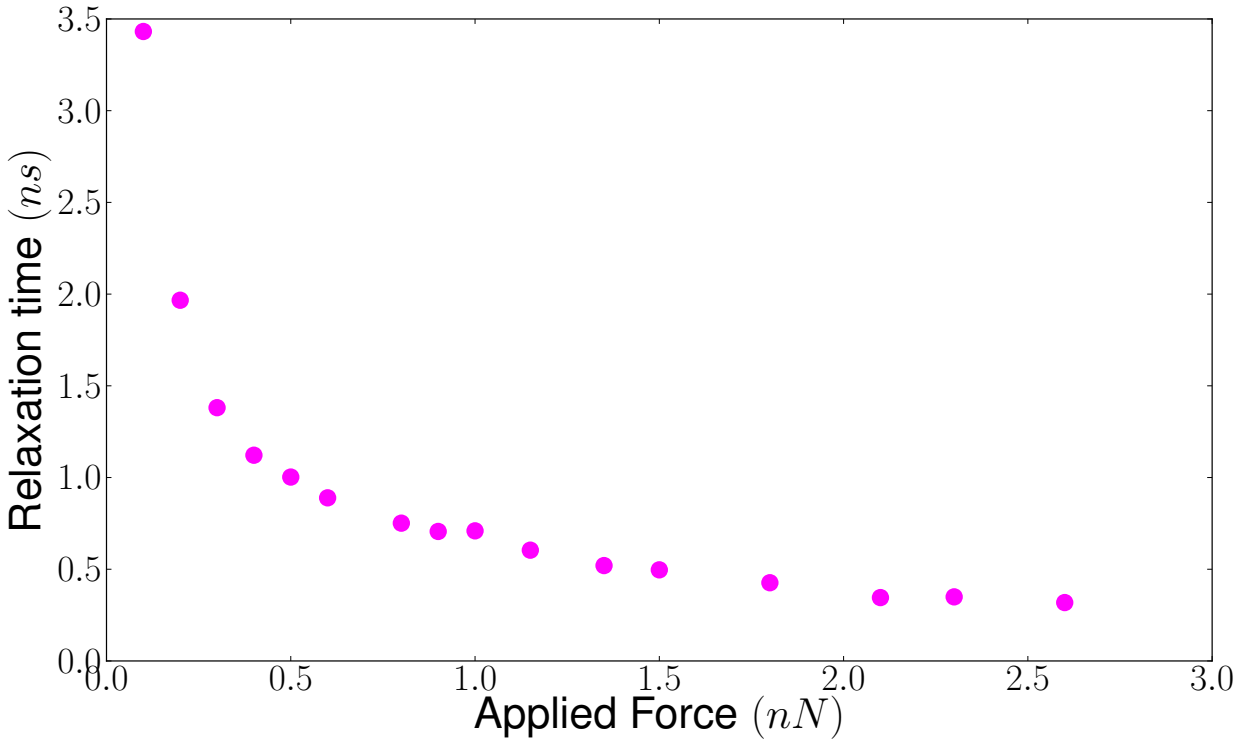

Figure 4.8: Relaxation time plotted as a function of the squeezing force.

Data selection

After the vesicle has been compressed, on longer timescales it will tend to slide (slightly) off-centre in the xy-plane —out from underneath the AFM Tip.[10] The effect becomes significant only at higher forces, and is mitigated by the adhesion site[11] on the substrate beneath the vesicle. In order that we measure the correct equilibrium area expansion, vertical compression, et cetera, we must select data *before* the vesicle has had time to drift significantly, but *after* the initial relaxation. This is done in a consistent manner, by using the relaxation time to define the boundaries of the data selection window. That is, when computing the post-relaxation averages, we select data beginning several relaxation times after the onset of deformation. The selection window then spans a set number of relaxation times (see Figure 4.9).

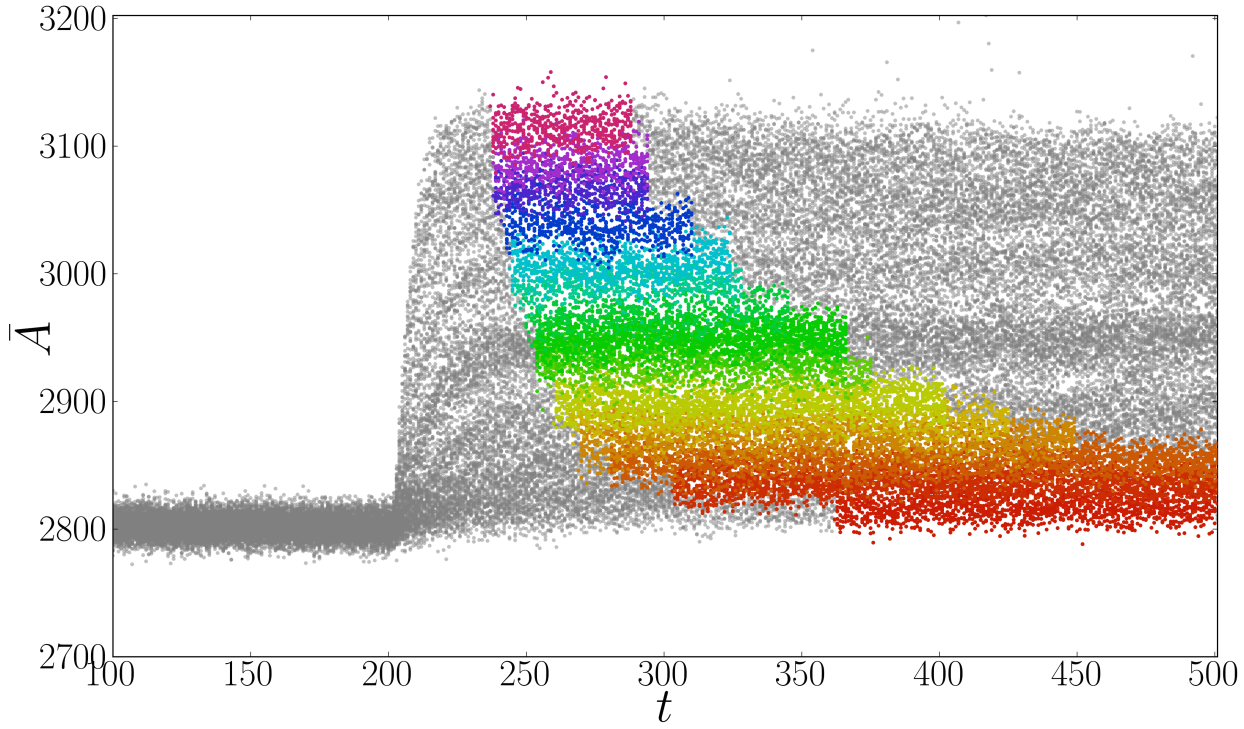

Figure 4.9: Data-selection windows (coloured regions) for weighted average of inner/outer leaflet surface area. Post-relaxation data is selected beginning $\gtrsim 3\tau$ after deformation begins, and ending $\lesssim 10\tau$ later. (Colour indicates squeezing force —legend is identical to that of Figure 4.1.)

[10] The adhesion site (underneath the vesicle, see Figure 3.2) does a good job of keeping the vesicle centred when the squeezing force is $\lesssim$ 200pN. However at forces exceeding $\sim$ 200pN, the vesicle will 'slip' or 'drift' off-centre (i.e. away from the z-axis) to a sufficient extent that the vertical compression (and area expansion) is altered. This effect can be seen at the bottom of Figure 4.4, where the vesicle drift shows up as an apparent secondary, longer timescale relaxation.

[11] see Figure 3.2

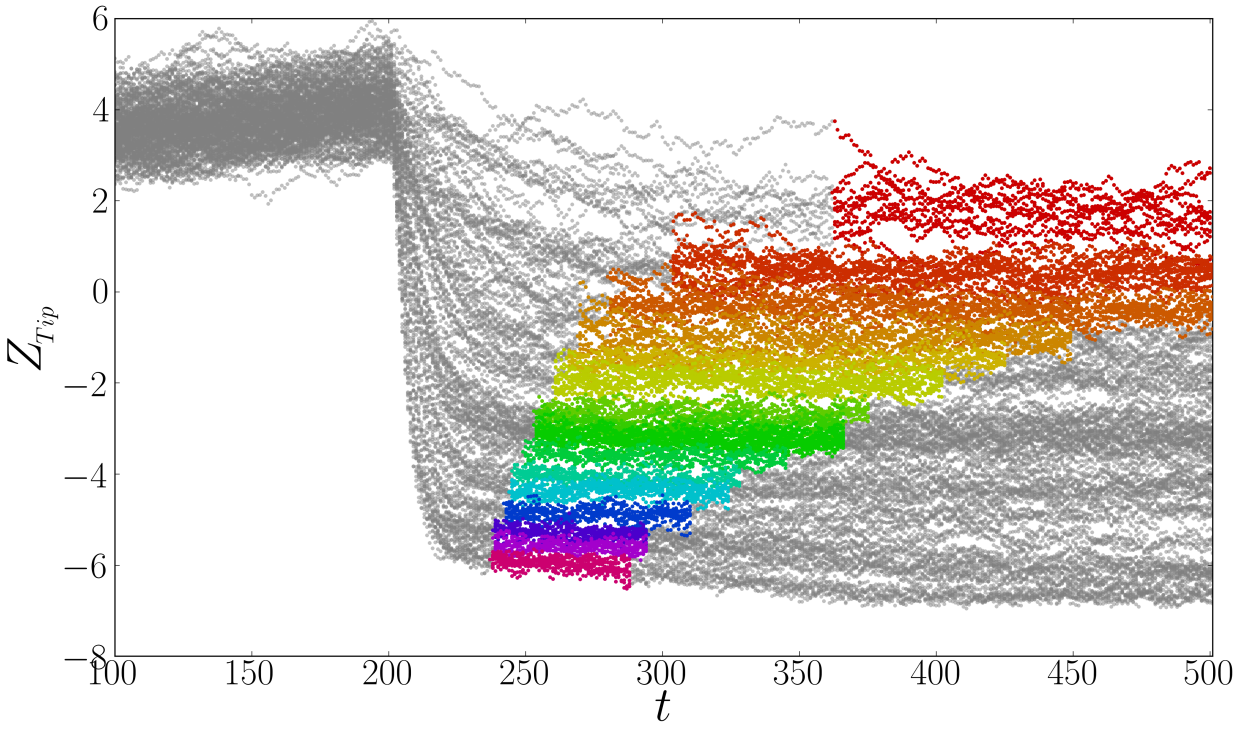

Figure 4.10: Timeseries of AFM Tip z-coordinate. Post-relaxation data is selected beginning $\gtrsim 3\tau$ after deformation begins, and ending $\lesssim 10\tau$ later. The necessity of this selection procedure is most obvious for this data, where the effect of the vesicle drifting out from underneath the Tip (at higher forces) shows up as a longer timescale 'compression' of the vesicle that isn't actually happening! That is, the vesicle is still being deformed, but it may be *less* compressed as a result of having drifted off-centre —even though the AFM tip position would measure a greater compression. (Colour indicates squeezing force —legend is identical to that of Figure 4.1.)

We chose a starting time $\gtrsim 3\tau$ because after 3τ, the vertical compression, area expansion, etc., has relaxed to $\gtrsim 95\%$ of its asymptotic value (cf. Equation 4.3). That is, once $t \gtrsim 3\tau$ most of the relaxation has already happened.

Surface Tension γ

Computing $\frac{dW}{dA}$

We compute $\frac{dW}{dA}$ by interpolating and then numerically differentiating the available simulation data (show below in Figures 4.11, 4.12 and 4.13).

The actual measurements of vertical compression and surface area expansion are far too sparse for the purpose of computing a numerical derivative; so we interpolate using a curve-fitting procedure. It turns out that if we add some extra free paramaters to the Helfrich model (including shifting and scaling of the independent variable), we can fit many of the various measurements quite well.

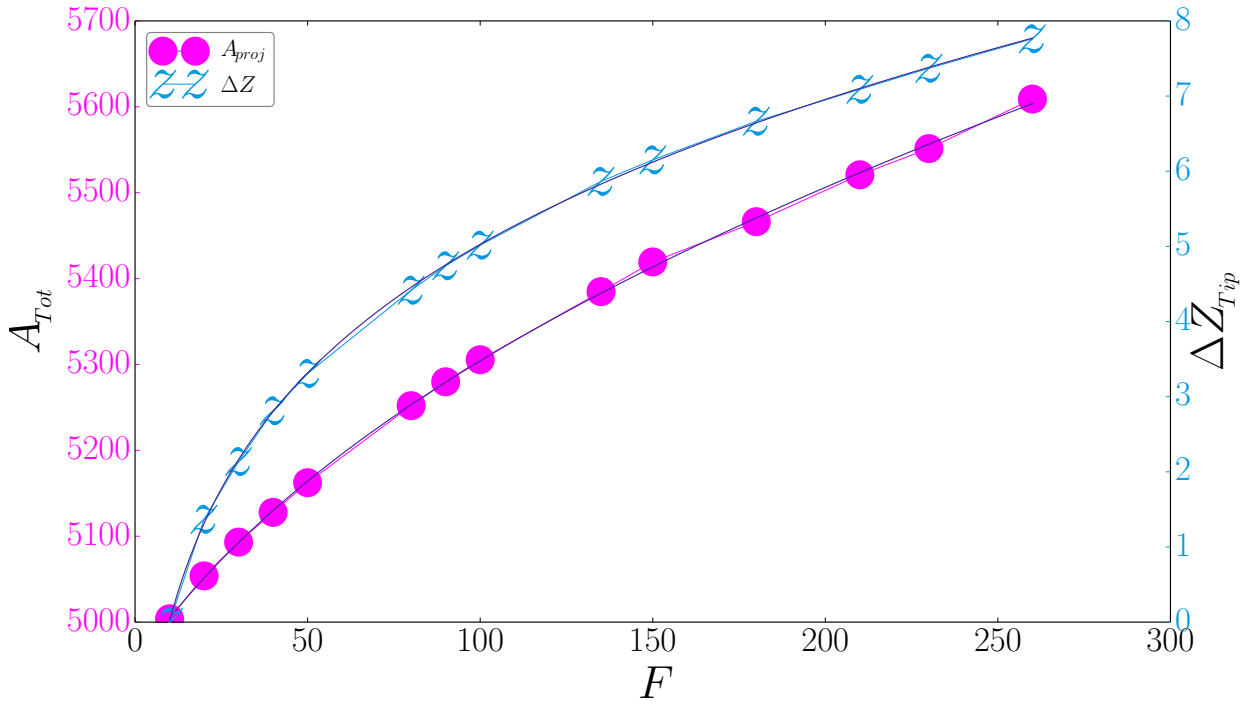

Figure 4.11: Data interpolation: vertical compression and surface area measurements are each fit independently as functions of the applied force. (Note: $A_{tot} = A^{inner\,leaflet} + A^{outer\,leaflet}$)

Once fit independently, we are able to combine separate measurements since everything has been measured initially as a function of the same variable —the applied force F_z. In this case since

$$\frac{dW}{dA} = F(z)\frac{dz}{dA}$$

we need to re-express the compression ΔZ_{tip} as a function of the projected area A_{proj}.

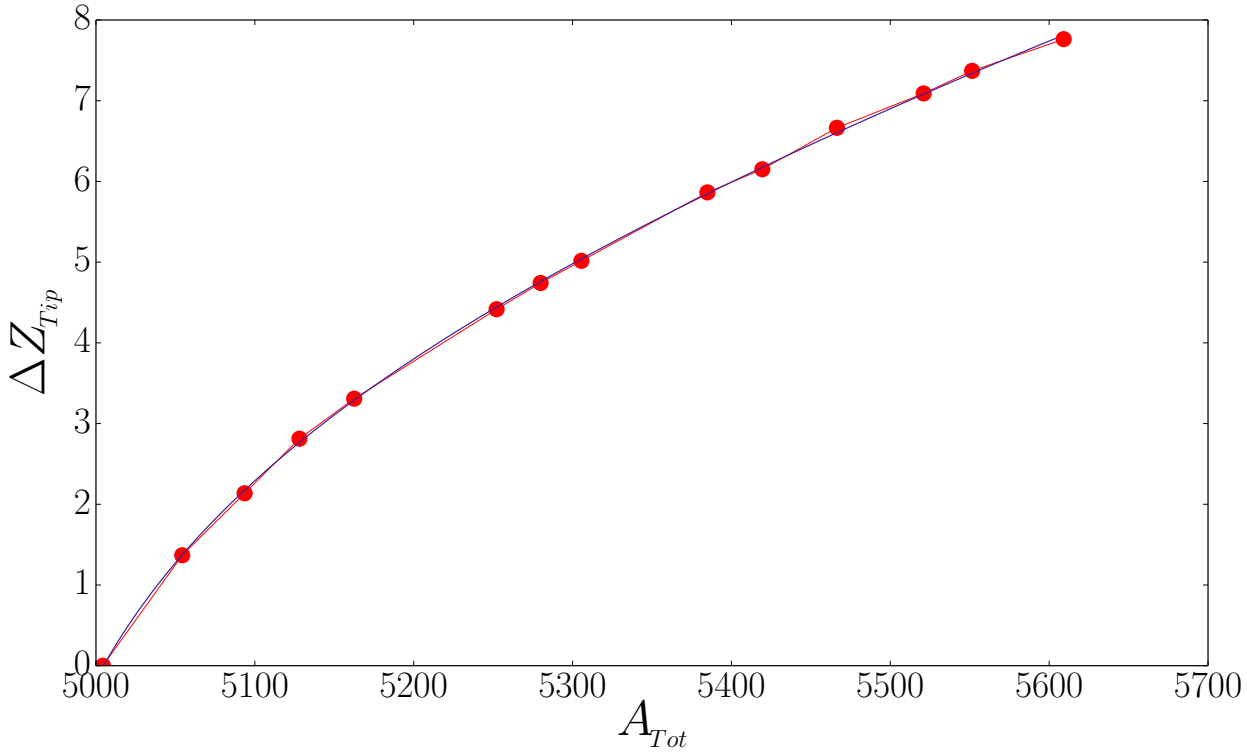

Figure 4.12: Vertical compression versus surface area: the two sets of interpolated data obtained in Figure 4.11 are plotted against each other to express ΔZ_{tip} as a function of A_{proj}. Numerical differentiation of this curve yields Figure 4.13.

Now we can go ahead and numerically differentiate (4.12) to reap the curve (4.13).

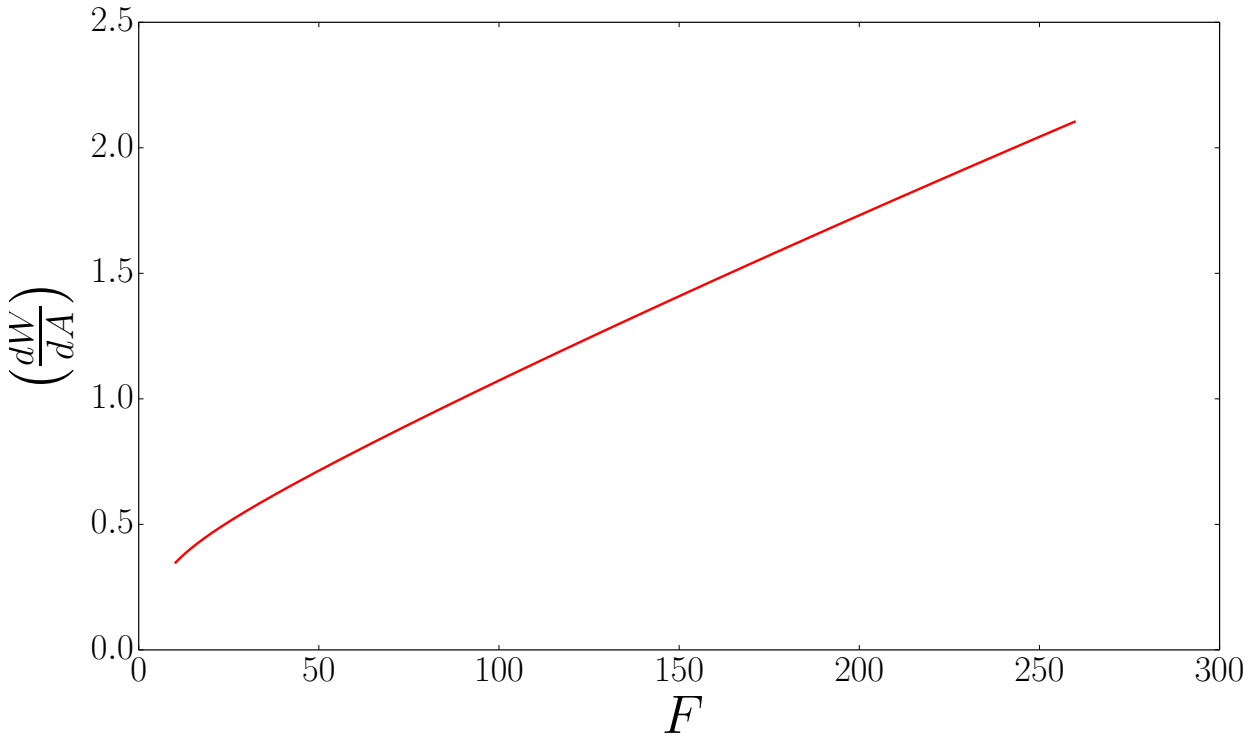

Figure 4.13: We numerically differentiate the work done on the system (by the AFM Tip) with respect to the vesicle's (projected) surface area.

Note that in and of itself $\frac{dW}{dA}$ is an interesting quantity. As we push the vesicle further and further away from equilibrium, the membrane is stretched —expanding in both surface area and volume— while other fluid is compressed. $\frac{dW}{dA}$ contains information about all of the various forms of deformation which occur simultaneously as the vesicle is squeezed.

Finding $\sum_j P_j \frac{dV_j}{dA}$

So far we have the first piece of Equation 2.19, but we can't calculate the tension from $\frac{dW}{dA}$ alone. In order to extract γ, we still require the pressure and volume data. Pressure data for specific particle groups is output by HOOMD-blue, except that it calculates the pressure using a volume term which is set equal to the volume of the entire simulation box. Therefore, to get the correct pressure data one must volume-adjust the raw pressure P_j data output by HOOMD-blue —which is a trivial step once the volume V_j of the region occupied by particle group j is known.

Since the outer fluid density ρ is readily calculated and the number of outer fluid particles

is known, we calculate

$$V^{outer\cdot fluid} = \frac{N^{outer\cdot fluid}}{\rho^{outer\cdot fluid}}.$$

The volume of the inner fluid is calculated in the same manner.[12] Having obtained the volumes occupied by the inner & outer fluid, we are able to volume-correct the pressure data output by HOOMD-blue as shown in Figure 4.14.

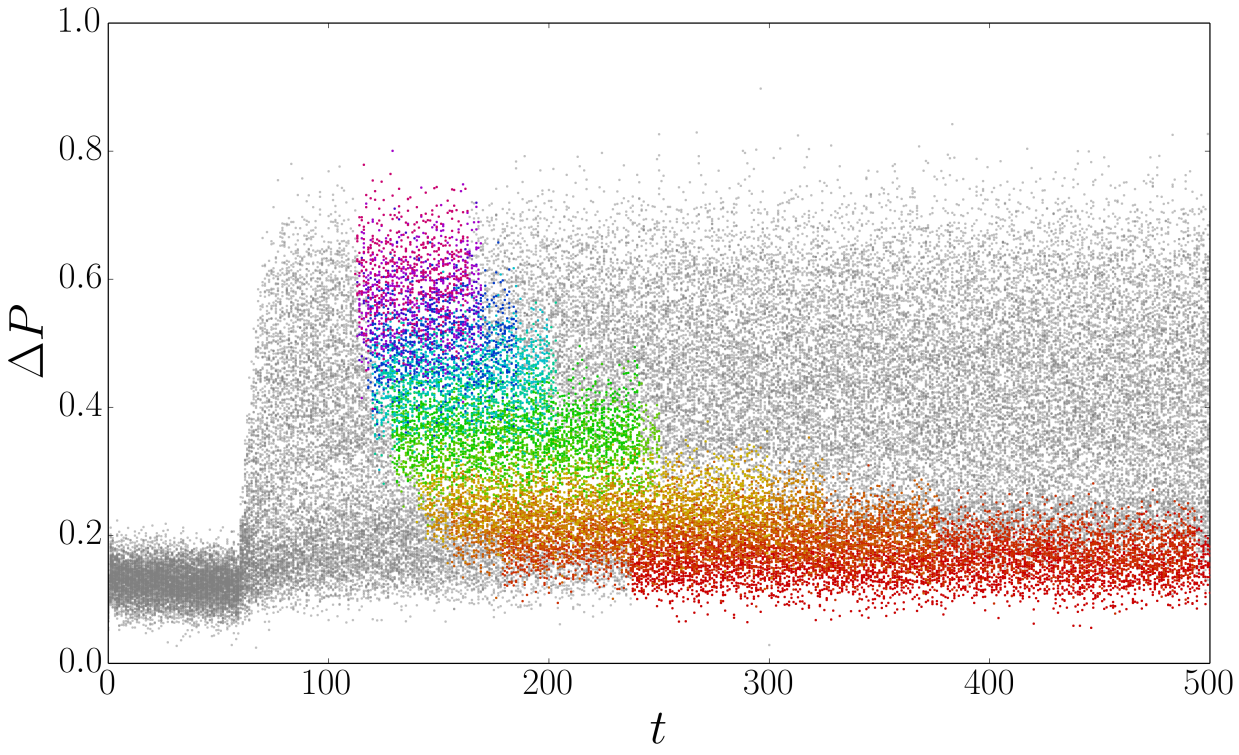

Figure 4.14: Pressure difference across the membrane —timeseries at various forces. This time-series is an example of the type of data used to obtain the pressure as a function of the force —cf. Figure 4.15. Colour highlights the data-selection window at each force.

We also need to know the volume & pressure of the fluid in the region occupied by the membrane. The volumes enclosed by the inner/outer leaflets of the vesicle can be measured using fits to the vesicle profile (Figures 4.18 and 4.23). We refer to said volumes as 'projected volumes', since a given V_{proj} is enclosed by that same shell whose surface area is taken to be A_{proj}. The volume enclosed by the inner leaflet (volume of vesicle contents) is denoted $V_{proj}^{inner\cdot leaflet}$, and $V_{proj}^{outer\cdot leaflet}$ that enclosed by the outer leaflet. The membrane volume is then

$$V_{proj}^{membrane} = V_{proj}^{outer\cdot leaflet} - V_{proj}^{inner\cdot leaflet}. \tag{4.4}$$

[12] Figure 4.26 demonstrates the accuracy of this procedure, as well as its agreement with the vesicle profile measurements.

Having obtained the volume and pressure data for the membrane as well as inner and outer fluid, we plot the ensemble-averaged measurements at each force in Figures 4.15 and 4.16.

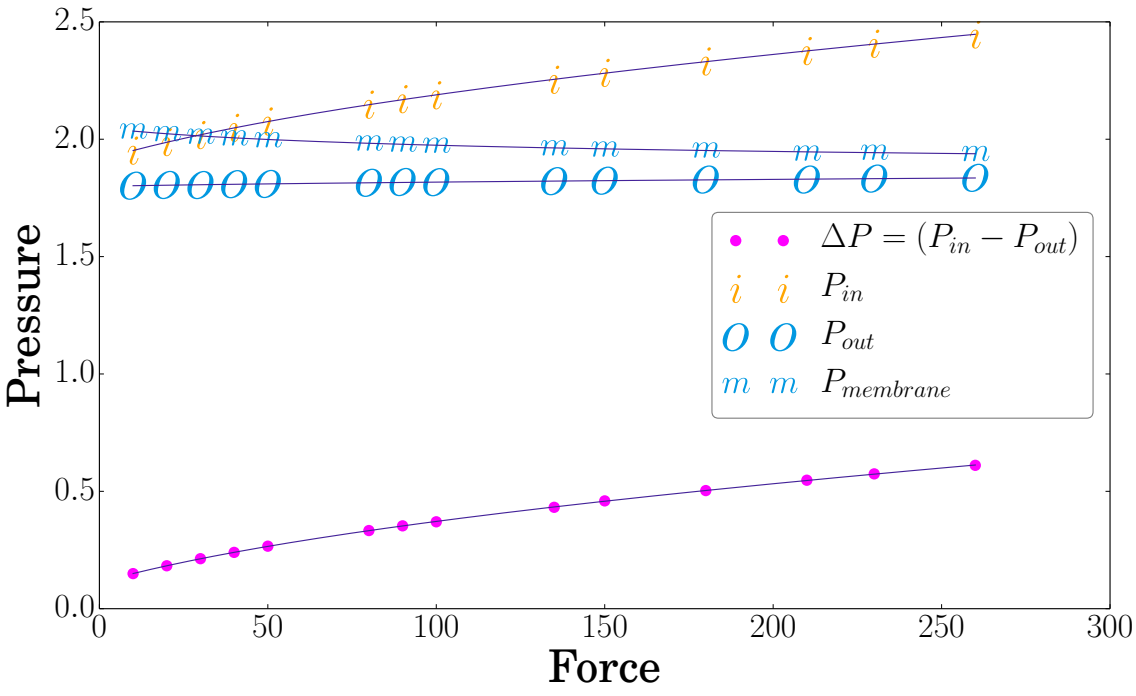

Figure 4.15: Fitted pressure data inside/outside of vesicle, and within vesicle membrane. Points are data, lines are fits. We also show the pressure difference $\boldsymbol{\Delta P}$ across the membrane —while both inner and outer pressure increase, so does $\boldsymbol{\Delta P}$.

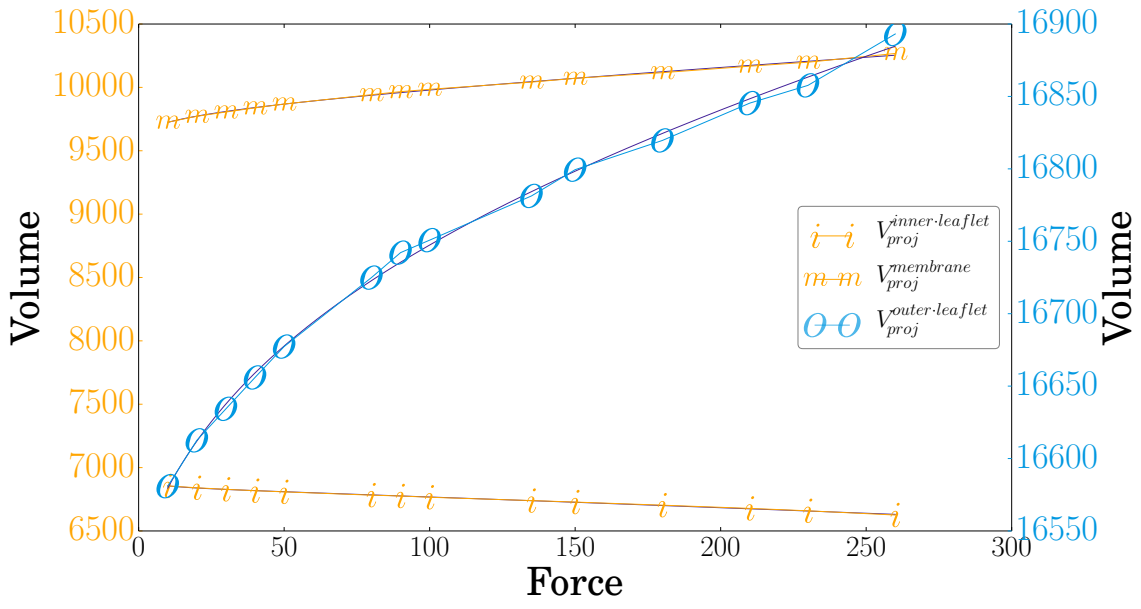

Figure 4.16: Volume of various vesicle regions as a function of the applied force. Points are data, lines are fits.

Recall that we are setting ourselves up to calculate each $P_j \frac{dV_j}{dA}$, and therefore need to know the $\frac{dV_j}{dA}$. Figure 4.16 directly provides the necessary data for the vesicle contents and membrane. And since we know that the volume of the system is conserved, we are able to indirectly obtain the derivative of the outer volume as well —using the outer-leaflet volume data in Figure 4.16. The derivative of the volume of the fluid surrounding the vesicle is then

$$\frac{d}{dA} V^{outer\cdot fluid} = -\left(\frac{d}{dA} V_{proj}^{outer\cdot leaflet} \right). \tag{4.5}$$

As was done for the compression & area expansion data, we have fit the sparse pressure and volume measurements to enable the numerical differentiation.

Tension versus $\frac{dW}{dA}$

Now that the PdV contribution has been measured, we are finally able to plot the surface tension as estimated via the work-approach (Figure 4.17).

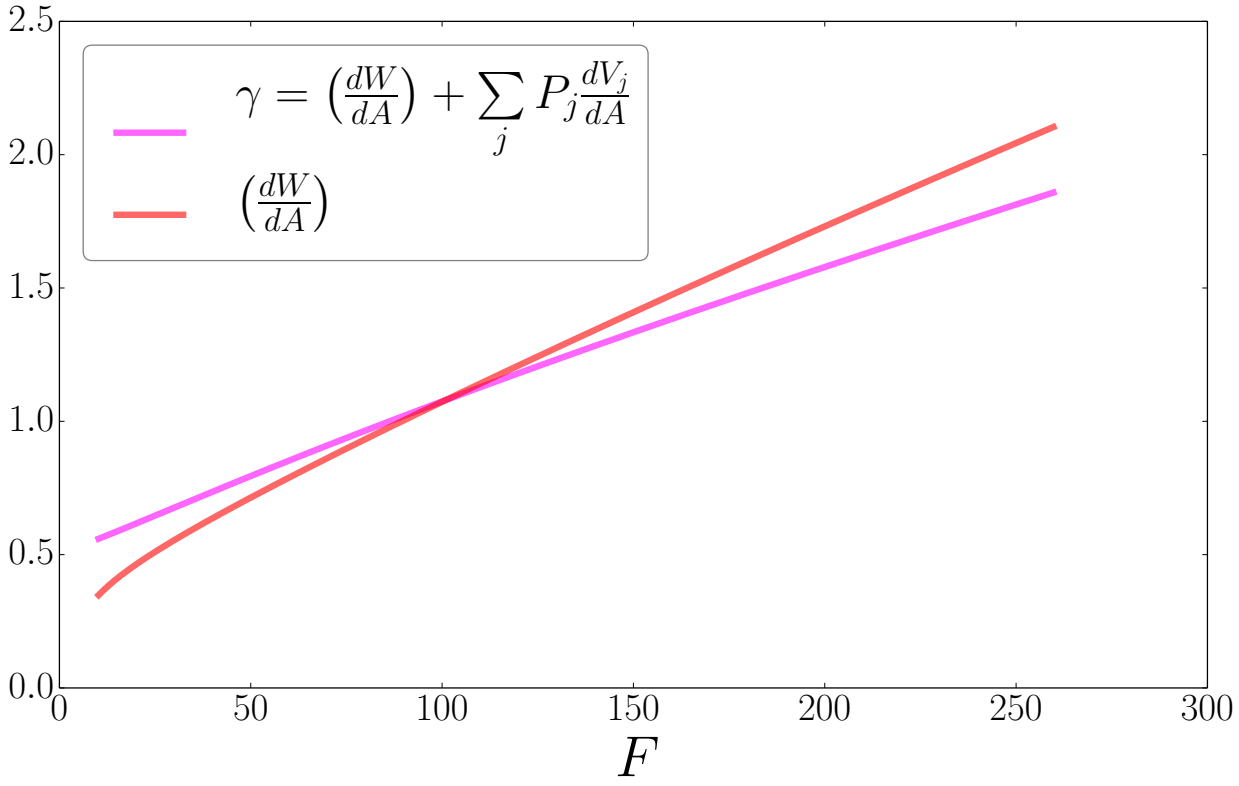

Figure 4.17: Surface tension (pink) estimated via the 'work approach' (i.e. via the free energy of the system). The red line is the total derivative $\frac{dW}{dA}$ of the work done on the system with respect to the projected area.

Before we can get comfortable with this estimate of the tension, there is a final detail to consider: bending energy. The free energy as written in Equation 2.12 lacked mention of any contribution due to bending. We must estimate the importance of bending in our simulations —to find out whether bending can be ignored without undermining the surface tension calculation shown in Figure 4.17. More important than the magnitude of the bending correction is the degree to which it varies over the force range studied here. If it varies so little (compared to the change in tension) that it can be treated as approximately constant, then no important change in the curve $\alpha(\gamma)$ will be produced by including bending in the estimation of γ. As we will see in Figure 4.24 & Equation 4.27 (in the subsection titled 'Shape equation'), this is precisely the case —the change in the bending correction is less than 0.2% that of the tension, and is therefore negligible in our estimate of the curve $\gamma(F)$.

Calculating surface tension: Geometric approach

Primarily, we invoke the geometric approach as a means of estimating the importance of bending effects relative to the surface tension. Bending and surface tension are related to each other via the shape (geometry) of the vesicle. It stands to reason that if this alternative approach yields the same magnitude for the tension as that obtained via the 'work approach', then the geometric approach likely provides a decent estimate of the relative magnitude of bending effects.

The following methods rely almost entirely on the local membrane *curvature* at the equator in order to estimate the surface tension.

Young-Laplace equation

The first approach uses the Young-Laplace relation

$$\Delta P = \gamma(-\nabla \cdot \hat{\mathbf{n}}) = \gamma \cdot (2H) \tag{4.6}$$

which relates the surface tension γ at some point in a membrane to its mean curvature[13] $H = \frac{1}{2}\nabla \cdot \hat{\mathbf{n}}$ at that point and the pressure difference ΔP[14] across said membrane. Since the mean curvature is

$$H = \frac{1}{2}(c_1 + c_2) \tag{4.7}$$

—where $c_j \equiv \frac{1}{R_j}$ and the R_j are the two principle radii of curvature— we can rewrite Equation 4.6 as

$$\boxed{\gamma_{Young\text{-}Laplace} = \frac{\Delta P}{c_1 + c_2}}, \tag{4.8}$$

thus giving the surface tension γ as the pressure difference ΔP divided by the sum of the two principle curvatures c_j. Assuming that the tension is uniform throughout the membrane, we need only calculate the curvature of the free surface. Figures 4.18 and 4.19 (below) describe how such curvature measurements are performed.

[13] $\nabla \cdot \hat{\mathbf{n}}$ = divergence of the normal vector to the surface at that point

[14] To be clear: ΔP = (the difference in pressure of the two regions which the membrane divides)

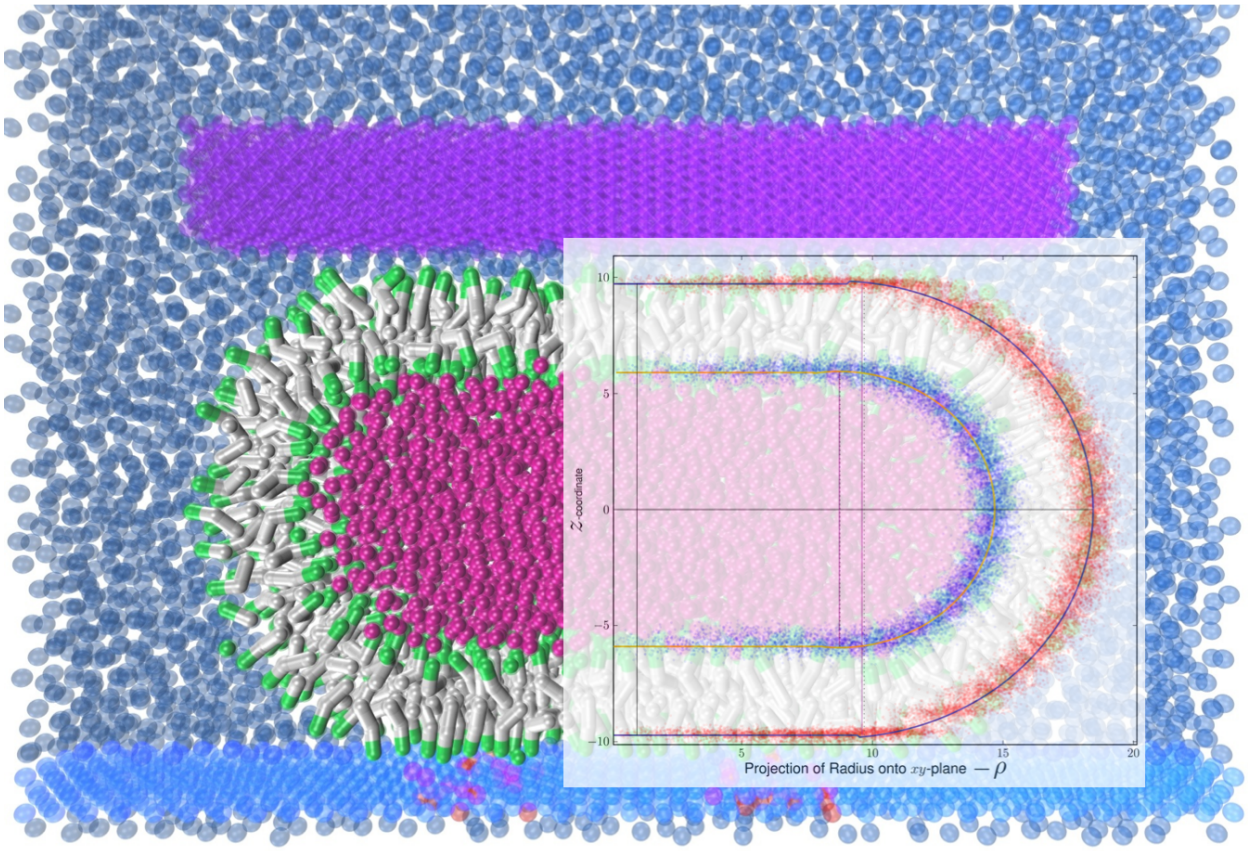

Figure 4.18: Vesicle profile and corresponding fit used to measure the curvature of the vesicle's free surface —the surface in contact with neither AFM Tip nor substrate. (cf. Figures 4.19 and 4.23)

Figure 4.18 is not entirely honest, since the blue and red arcs lying on top of the inner and outer leaflets (respectively) actually represent *all* head particles, collapsed onto (ρ, z) coordinates $(\rho = \sqrt{x^2 + y^2})$. That is, the θ information about the position of each head particle is discarded so that the arc of the free surface can be represented in two dimensions. This way, when we curve-fit the red and blue arcs we are fitting to a θ-average of the free surface.

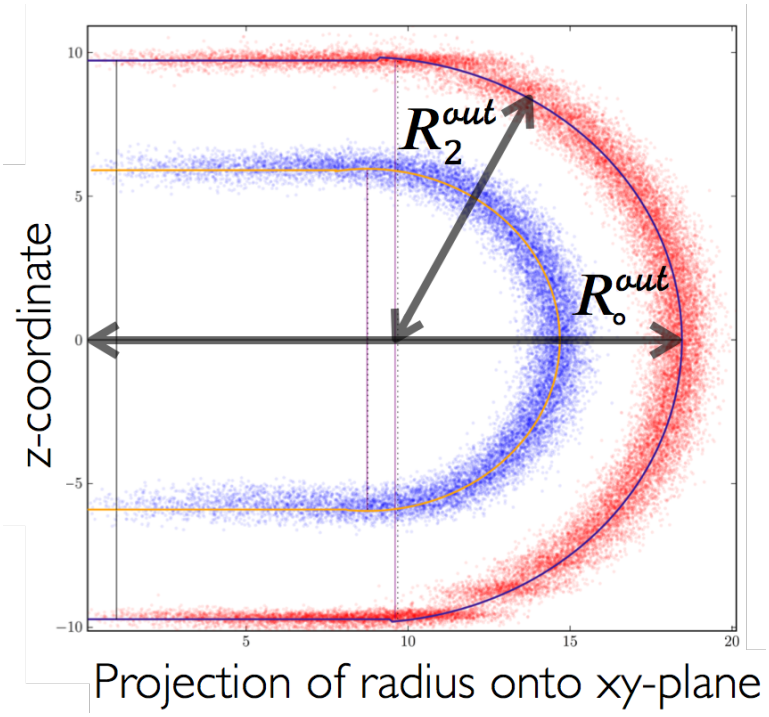

Figure 4.19: This figure shows how c_1 and c_2 are obtained by curve-fitting the free surface of the vesicle membrane (cf. Figures 4.18 and 4.23). When collapsed into (ρ, z)-coordinates, the vesicle profile looks like a horseshoe. So to relate this plot to Figure 4.18, think of the compressed vesicle as a "horseshoe of rotation". Inner lipid heads are shown in blue, and outer lipid heads are shown in red.

The tension calculated via (4.8) —using the ΔP data shown in Figure 4.15 and curvature as measured above— is plotted in Figure 4.22. As previously stated, one drawback to using Equation 4.8 is that this model assumes a true mathematical membrane with *vanishing thickness*. As can be seen from Figure 4.19 our bilayer has a significant thickness when compared to the size of the vesicle. For this reason we take as c_1, c_2 a weighted average of the inner and outer leaflets' corresponding curvatures, and treat the calculated tension as an *estimate* for the purpose of observing how the tension increases with squeezing force.

Schäfer approach

Schäfer begins with the balance of forces at the compressed vesicle's equator

$$2\pi R_\circ \gamma = \pi R_\circ^2 \Delta P - F. \tag{4.9}$$

which gives the surface tension as

$$\gamma_{Schäfer} = \frac{R_\circ}{2}\Delta P - \frac{F}{2\pi R_\circ}. \tag{4.10}$$

Then the substitution

$$\Delta P = \frac{F}{\pi R_i^{\,2}} \tag{4.11}$$

is included, giving

$$\gamma = \frac{R_\circ}{2}\left(\frac{F}{\pi R_i^{\,2}}\right) - \frac{F}{2\pi R_\circ}$$

$$\rightarrow \boxed{\gamma_{Schäfer} = \frac{F}{2\pi R_\circ}\left(\frac{R_\circ^2}{R_i^{\,2}} - 1\right)}. \tag{4.12}$$

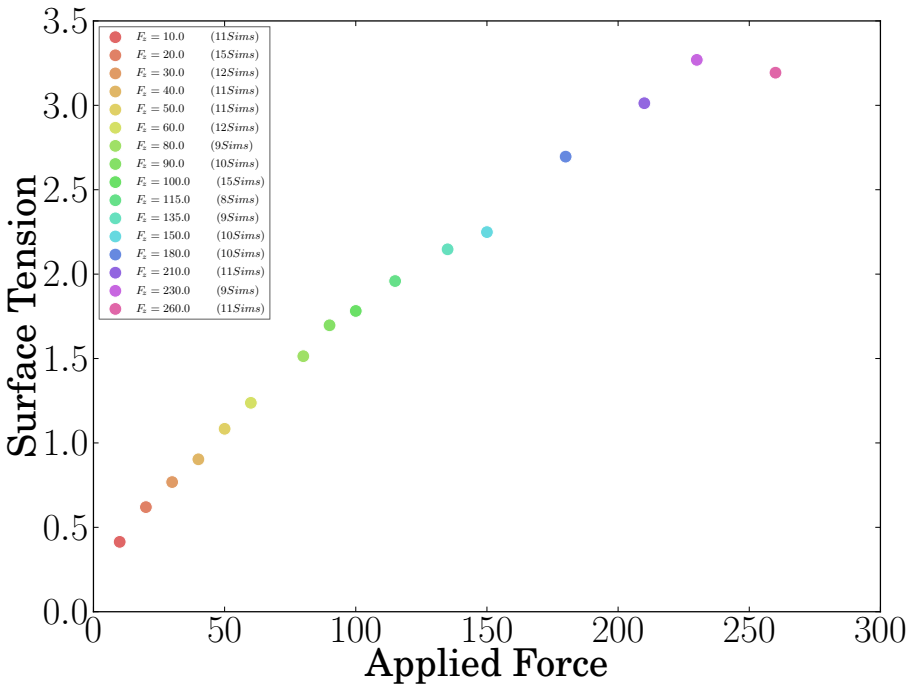

Figure 4.20: Surface tension calculated using the same expression as Schäfer et al. [34] obtained via the balance of forces at the vesicle equator (ignores bending). As can be seen in the high-force regime, this method may not be reliable for small vesicles.

But as we can see from Figure 4.20, this measure of tension at one point shows a *decrease* in tension with increasing force. At high forces this measure becomes erratic and appears to be unreliable.

Equation 4.11 is useful for large vesicles, where the membrane thickness can be ignored. In this case the contact area $A_c = \pi R_i^{\,2}$ is well defined. But for small vesicles —such as the one simulated in this work— with non-negligible membrane thickness, the contact area is not well defined.

Improvise: combine Schäfer and Young-Laplace

The expressions used by Schäfer relating ΔP, F and γ become problematic when the thickness of the vesicle membrane is non-negligible. This is because (among other reasons) said equations include the 'contact radius' R_i as an essential parameter. But for a thick membrane, R_i is poorly defined. Equation 4.8 and Equation 4.9 are both describing the same γ —let's equate the two.

$$\gamma_{Schäfer} = \gamma_{Young\text{-}Laplace} \tag{4.13}$$

$$\longrightarrow \Delta P \left(\frac{R_\circ}{2} - \frac{F}{2\pi R_\circ \Delta P} \right) = \Delta P \frac{R_\circ R_2}{R_\circ + R_2}$$

$$\longrightarrow \frac{F}{2\pi R_\circ \Delta P} = \frac{R_\circ}{2} - \frac{R_\circ R_2}{R_\circ + R_2}$$

$$= \pi R_\circ^2 \left(1 - \frac{2R_2}{R_\circ + R_2} \right)$$

$$\longrightarrow \frac{F}{\Delta P} = \pi R_\circ^2 \frac{R_\circ - R_2}{R_\circ + R_2}$$

$$\therefore \boxed{\Delta P = \frac{F}{\pi R_\circ^2} \left(\frac{R_\circ + R_2}{R_\circ - R_2} \right)}. \tag{4.14}$$

Plugging Equation 4.14 into Equation 4.10 gives

$$\gamma = \frac{\cancel{R_\circ}}{2} \left(\frac{F}{\pi \cancel{R_\circ^2}} \frac{R_\circ + R_2}{R_\circ - R_2} \right) - \frac{F}{2\pi R_\circ} \tag{4.15}$$

$$= \frac{F}{2\pi R_\circ} \left(\frac{R_\circ + R_2}{R_\circ - R_2} - 1 \right),$$

$$\therefore \boxed{\gamma_{\mathrm{YLS}} = \frac{F}{\pi R_\circ} \left(\frac{R_2}{R_\circ - R_2} \right)}. \tag{4.16}$$

The name γ_{YLS} was chosen to reflect the origin of this combined Young-Laplace-Schäfer expression for the tension.

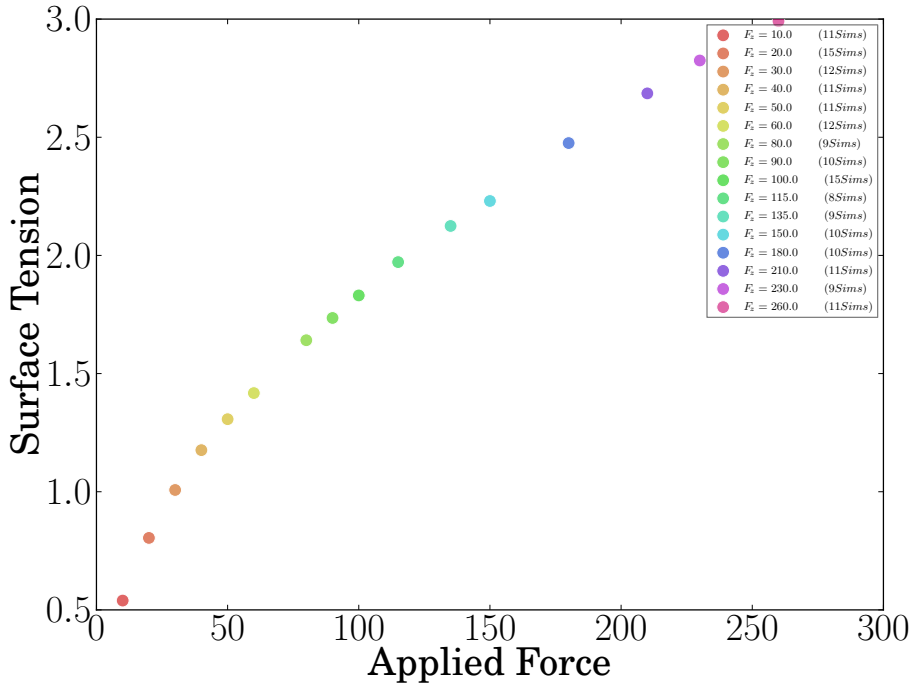

Figure 4.21: Tension calculated by combining theYoung-Laplace equation and the approach used by Schäfer et al. [34].

Since $\boldsymbol{R}_\circ$ and $\boldsymbol{R}_2$ can be measured much more accurately than $\boldsymbol{R}_i$, Equations 4.14 and 4.16 are much better behaved expressions for $\boldsymbol{\Delta P}$ and γ than Equations 4.11 and 4.12 respectively. We now have an alternate means of estimating the surface tension which does not require a direct measurement $\boldsymbol{\Delta P}$.

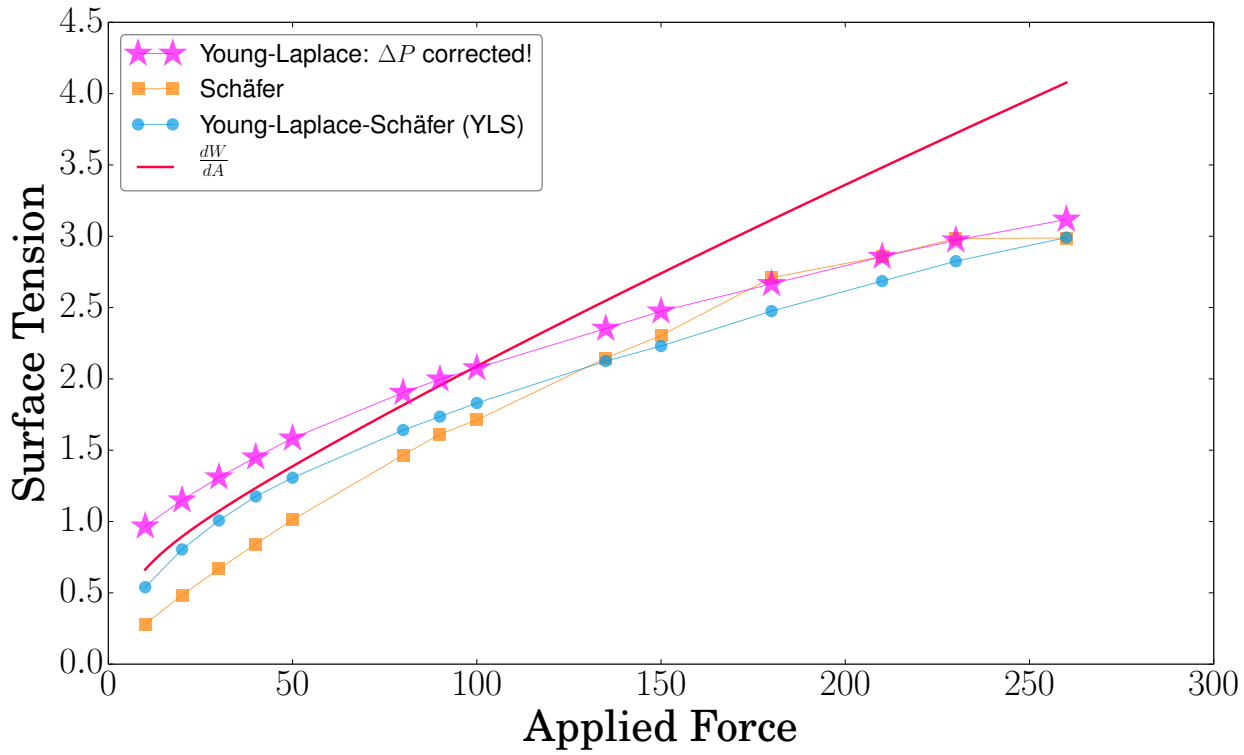

Figure 4.22: Comparing the three curvature-based measures of surface tension, the two most appealing are Schäfer and Young-Laplace-Schäfer —the latter being the best behaved. Additionally, the red 'reference line' is *analogous*[15] to $\frac{dW}{dA}$ as shown in Figure 4.17.

The trouble with γ_{YLS} is that it includes the force explicitly in the numerator but in the denominator it has the expression $(R_\circ - R_2)$ —which is unable to go zero as $F \longrightarrow 0$ (due to adhesion). Since the denominator doesn't go to zero at zero force, γ_{YLS} must go to zero as $F \longrightarrow 0$ —but the true tension is *nonzero* at zero force. Therefore, it seems reasonable to add a 'pre-stress' $\gamma_\circ$ to Equation 4.16 so that

$$\gamma_{\text{YLS}} = \frac{F}{\pi R_\circ}\left(\frac{R_2}{R_\circ - R_2}\right) + \gamma_\circ. \tag{4.17}$$

[15] Caution: Because the geometric measures of tension use a weighted average of the inner and outer surfaces, the $\frac{dW}{dA}$ shown here does as well. Hence, the reference line is *not* the same $\frac{dW}{dA}$ as shown elsewhere, but is merely analogous.

Bending correction $\Delta\gamma$ via shape equation

Since the vesicle is being flattened right up to the point of lysis, we can't assume that bending energy can be ignored. A complete expression for the shape equation of a membrane at equilibrium was derived by Zhong-can and Helfrich [38]:

$$\Delta P - 2\gamma H + \kappa(2H + c_\circ)(2H^2 - 2G - c_\circ H) + 2\kappa\Delta_s H = 0 \tag{4.18}$$

where Δ_s is the covariant surface Laplacian (Laplace-Beltrami operator on the surface of the vesicle)[16] given by

$$\Delta_s = \frac{1}{\sqrt{g}}\partial_i(g^{ij}\sqrt{g}\,\partial_j), \tag{4.19}$$

with g denoting the metric tensor. To find the components of g, we need to parametrize the vesicle's surface.

Free-surface parametrization

Equation 4.18 allows us to account for bending energy when calculating the surface tension. To use it we must first parametrize the vesicle surface (Figure 4.23), on which the metric components et cetera are calculated. The free surface of the compressed vesicle is described by the vector

$$\boxed{\vec{r} = \vec{r}(\phi, z) = \rho(z)\,\hat{e}_\rho(\phi) + z\,\hat{e}_z} \tag{4.20}$$

in cylindrical coordinates (ρ, ϕ, z) with

$$\boxed{\rho(z) = (R_\circ - R_2) + \sqrt{R_2^2 - z^2}}, \tag{4.21}$$

and

$$\hat{e}_\rho(\phi) = \cos\phi\,\hat{e}_x + \sin\phi\,\hat{e}_y. \tag{4.22}$$

[16] This operator is often denoted Δ or Δ_s. I'm using Δ_s because I don't want to write an equation containing two identical deltas (i.e. ΔP and $\Delta_s H$) having completely different meanings!

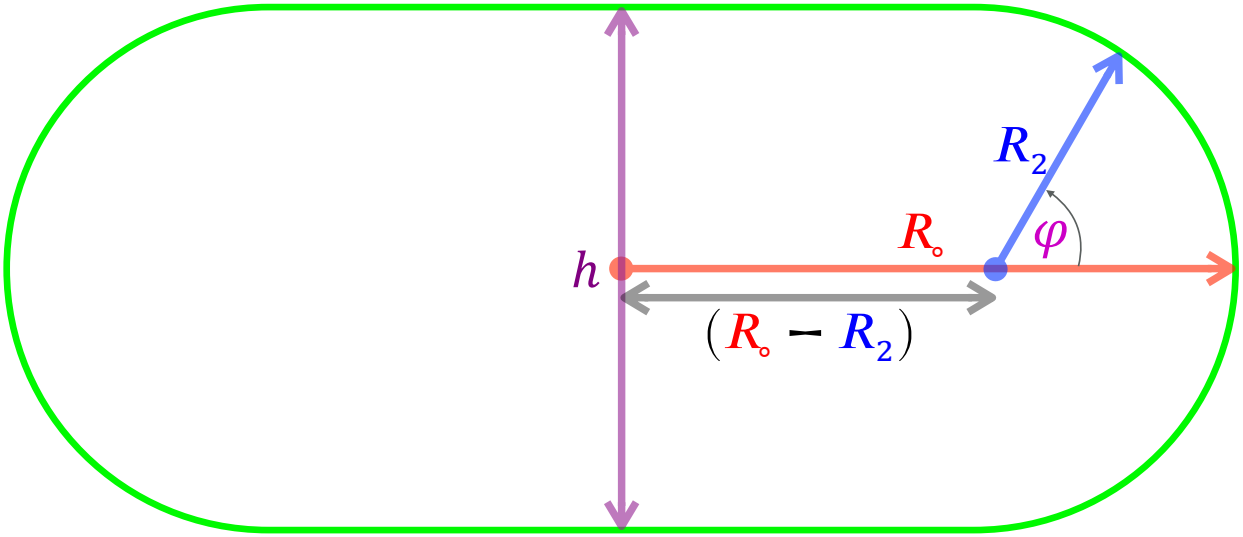

Figure 4.23: Parametrization of compressed vesicle (cf. Figures 4.18 and 4.19). In the context of projected area and membrane undulations, this profile delineates the 'mean surface' of the vesicle. The shape of the compressed vesicle is well approximated by a "filled torus" —a doughnut without a hole.One might also think of this shape as a "horseshoe of rotation". This approximation of the true surface is the same as used by [34, 37] and detailed in [12]. The variable names have been chosen to match those of [34] for ease of comparison. (Coordinates: In this figure, h lies along the z-axis. ϕ lies in the xy-plane.)

Explicit shape equation

Having parametrized the vesicle, we then carry out the general procedure described in [38] for computing (4.19). At the vesicle equator ($z = 0$), we have

$$\boxed{\Delta_s H\Big|_{z=0} = \frac{R_\circ - R_2}{2R_\circ^2 R_2^2}}. \tag{4.23}$$

Now plug (4.23) into (4.18) and with some algebra obtain

$$\gamma_{\text{Helfrich}} = \overbrace{\Delta P\left(\frac{R_\circ R_2}{R_\circ + R_2}\right)}^{\gamma_{\text{Young-Laplace}}} + \overbrace{\frac{\kappa}{R_\circ + R_2}\left\{\frac{(R_\circ - R_2)^3}{2R_\circ^2 R_2^2} - \frac{R_\circ - R_2}{R_\circ^2} + 2c_\circ - \frac{c_\circ^2}{2}(R_\circ + R_2)\right\}}^{\Delta\gamma}. \tag{4.24}$$

In other words

$$\boxed{\gamma_{\text{Helfrich}} = \gamma_{\text{Young-Laplace}} + \Delta\gamma}, \tag{4.25}$$

where $\Delta\gamma$ is the 'bending correction'.

Bending correction

So to account for bending, the correction to the Young-Laplace law is given by

$$\boxed{\Delta\gamma = \frac{\kappa}{R_\circ + R_2}\left\{\frac{(R_\circ - R_2)^3}{2R_\circ^2 R_2^2} - \frac{R_\circ - R_2}{R_\circ^2} + 2c_\circ - \frac{c_\circ^2}{2}(R_\circ + R_2)\right\}}. \tag{4.26}$$

At this point, Equation 4.26 contains two free parameters: the bending rigidity κ and spontaneous curvature $c_\circ$ of the bilayer. We use $\kappa \approx 4.0$, in accordance with that measured by Martin Bertrand[6] for flat bilayers composed of the same lipids, with the same interactions, and under conditions identical to our own. (This value is corroborated by our own measurements —see Figure 5.6) To get $c_\circ$, we measured the radius of the unabsorbed (i.e. free-floating) vesicle with all other initial conditions held equal —e.g. number of inner fluid particles, outer fluid density, et cetera.

We plot $\Delta\gamma$ and find that its influence is sufficiently small over the entire range of forces that the effects of bending to be ignored when estimating γ (see below).

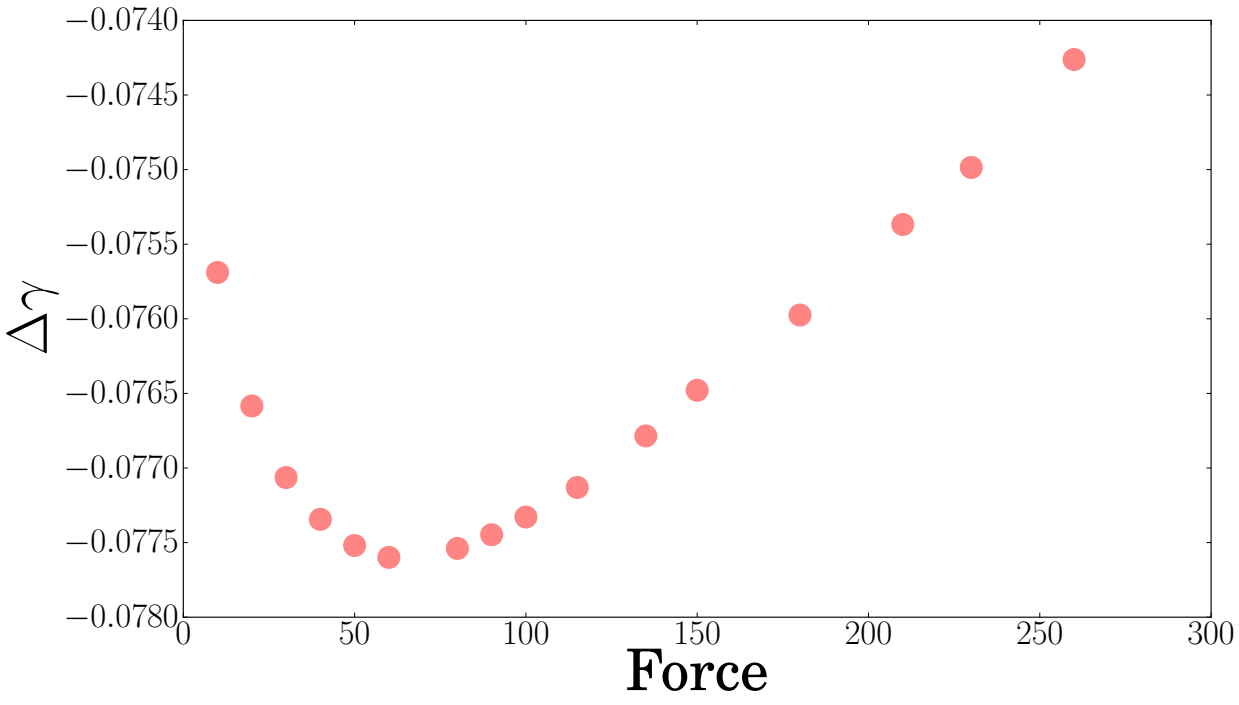

Figure 4.24: Here we estimate the contribution of bending energy, so that we may gauge its importance relative to the calculation of surface tension.

Notice in Figure 4.24 that the bending correction $\Delta\gamma$ varies negligibly (relative to the

surface tension) over the range of forces seen here. That is, since

$$\frac{\left|(\Delta\gamma)^{Max} - (\Delta\gamma)^{Min}\right| \lesssim 0.004\frac{\varepsilon}{\sigma^2}}{|\gamma^{Max} - \gamma^{Min}| \gtrsim 2\frac{\varepsilon}{\sigma^2}} \lesssim 0.002,^{17} \tag{4.27}$$

the bending correction varies so little that, for the purpose of estimating the surface tension, it can be treated as constant. And since the addition of a constant to the overall surface tension does nothing to change its functional dependence on the applied force, the *shape* of the curve in Figure 5.5 —$\alpha(F)$ plotted as a function of $\gamma(F)$— is *unaltered* by the bending correction $\Delta\gamma$.

[17] Interestingly, Schäfer et al.[34] obtain a factor of the same order of magnitude in their analysis of the importance of bending relative to surface tension.

Projected Area

The formula derived by Helfrich is concerned *not* with the actual surface area of the membrane, but rather the "projected area"[18]. In our case we use a fit to the vesicle profile, which yields an effective surface whose area is analogous to the projected area for closed vesicles.[19]

Projected area versus Force

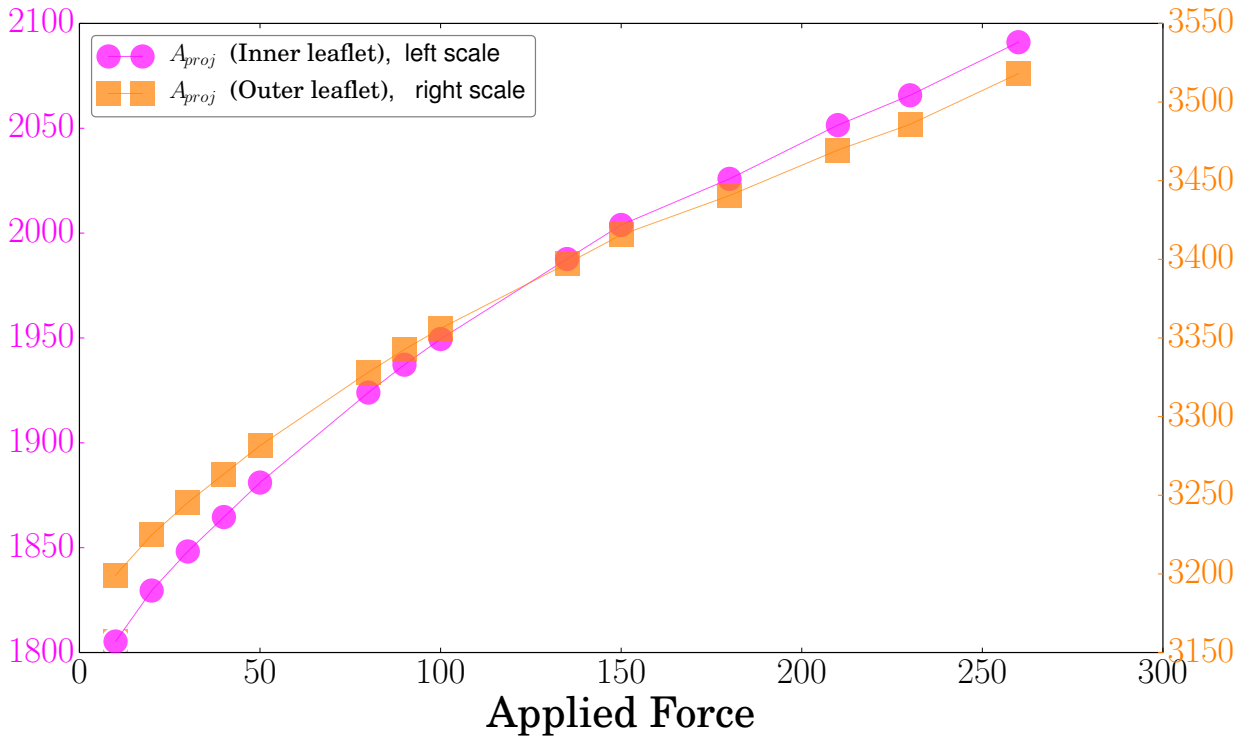

Figure 4.25: Projected Area (area of mean surface) versus force.

Calculating the projected area (above) relied on a 'parametrization algorithm' —an automated curve-fitting script designed to measure the vesicle's effective shape (mean surface).

[18] Distinction described in the Background section, under the heading Projected Area versus Triangulated Area

[19] The apparent surface area of a vesicle is the projected area of its membrane.

We want to check that the projected area calculation is accurate —and that the parametrization algorithm is returning honest measurements of the vesicle shape. One test is volume conservation.

Volume conservation

As mentioned in Projected Area versus Triangulated Area, the mean surface (or 'effective surface') associated with the membrane *must enclose the same volume as the vesicle itself.* To test for this condition Figure 4.26 compares the volume enclosed by the mean surface, to that enclosed by the actual membrane[20]. Significant disagreement[21] between the two volume measures would indicate a problem with A_{proj} —and the parametrization algorithm used to calculate it. This test can't *prove* that the parametrization algorithm (and A_{proj}) is accurate, but it is one way of screening for a problem.

Volume enclosed by mean surface

The first method uses the mean surface (given in Figures 4.23 and 4.18) associated with the lipid heads belonging to the inner leaflet of the vesicle. We integrate[22] the volume enclosed by the mean surface

$$V_{proj} = 2\int_0^{h/2} dz \int_0^{r(z)} dr \int_0^{2\pi} r\, d\phi, \tag{4.28}$$

$$= 2\int_0^{h/2} dz \int_0^{r(z)} 2\pi r\, dr = 2\int_0^{h/2} dz\, \pi r^2(z), \tag{4.29}$$

$$\rightarrow \boxed{V_{proj} = 2\pi(R_\circ - R_2)R_2^2 \arcsin\left(\frac{h}{2R_2}\right) + \pi h\left(\left((R_\circ - R_2)^2 + R_2^2 - \frac{h^2}{12} + (R_\circ - R_2)\sqrt{R_2^2 - \frac{h^2}{4}}\right)\right)}. \tag{4.30}$$

We plug the values of $R_\circ$, R_2, h returned by the curve fitting algorithm into Equation 4.30 in order to obtain the V_{proj} values shown in Figure 4.26.

[20] calculated by independent means
[21] There isn't any, they agree quite well.
[22] For details, see Volume Enclosed by Mean Surface in the Appendix.

Volume via density of vesicle contents

The second method is straightforward. Due to its irregular shape, it is not so easy to *directly* measure the volume of material contained within the simulated vesicle. However it is very easy to directly measure the *density* $\rho = \frac{N}{V}$ of the contents, and since the number of particles $N_{influid} \cong 4859$ sealed within the vesicle is known[23], we can simply calculate the inner volume as

$$V_\rho = \frac{N}{\rho}. \tag{4.31}$$

However we are calculating the volume of the vesicle contents to test our calculation of the mean surface (of the inner lipid heads). In *addition* to enclosing the vesicle contents this surface will, at any given moment, enclose some fraction of the volume occupied by the lipid heads. The fact that the particles have finite size does nothing to complicate the matter. The mean surface is calculated using the heads' centre of mass positions. Hence a lipid head lying exactly on the membrane has half of its volume enclosed by said surfrace, and so on. We assume that on the average, half of the inner leaflet's $\cong 1172$ lipid heads will be inside the mean surface. As a rough estimate, we treat the lipid heads as though their density is the same as that of the inner fluid.

Therefore, we use

$$N_{inside} = 4859 + \frac{1172}{2} = 5445 particles, \tag{4.32}$$

giving

$$\boxed{V_\rho = \frac{5445}{\rho_{influid}}}. \tag{4.33}$$

[23] At system setup, the particles originally pumped into the vesicle are of a different "particle type" than (although having interactions identical to) the outer fluid. Therefore we can count them. We can monitor the inner fluid particles to see whether any escape during the simulation (extremely rare except at the threshold of lysis tension).

Comparing the two —V_{proj} versus V_ρ

Happily, Figure 4.26 tells us that the two measures of enclosed volume agree quite nicely.

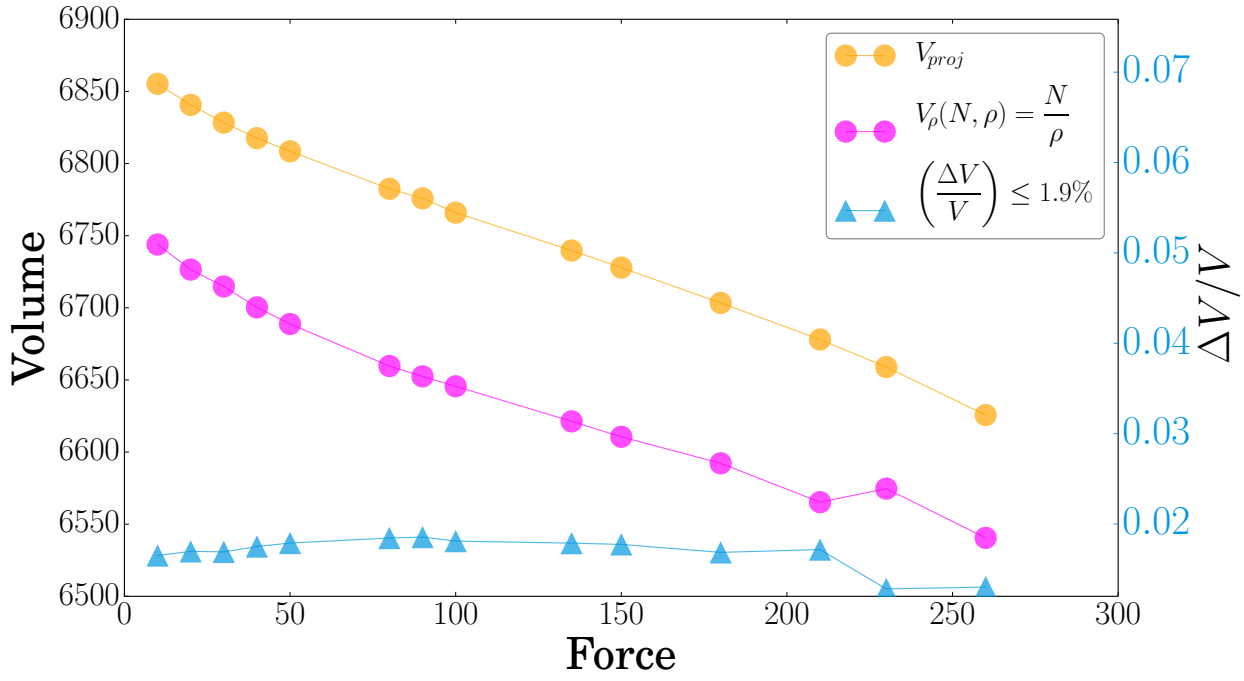

Figure 4.26: Comparing the volume enclosed by the vesicle to that enclosed by its mean surface. The two results agree to within 2%. This result is highly satisfactory —since *no volume constraint was implemented in the fitting software* which parametrized the mean surface, it suggests that Figure 4.23 is indeed a natural and honest description of the free membrane, and has yielded accurate measurements of the vesicle shape. The former is calculated via the number N_{inside} of particles inside the vesicle and the density ρ_{inside} of the vesicle contents. The latter is obtained by plugging surface-fit parameters (for the inner leaflet) into a volume integral of the theoretical vesicle shape.

It should be noted that the script which fits the mean surface was originally written only to find the mean curvature of the vesicle, and so was given lots of freedom in fitting other parameters. The script did *not* include a volume constraint. Nevertheless, the vesicle shape returned by the fitting script encloses a volume which agrees with that calculated by more direct means to within an accuracy of 2%!

Viscosity Calculation

The viscosity can be altered via the viscous drag parameter "*gamma*" in the DPD interaction[32]. In the following two figures, viscosity is calculated (via the indirect Einstein approach) for two Lennard-Jones fluids having different values of *gamma*. In Figure 4.27, we see the results of viscosity calculations for the standard (*gamma* = 1.0) LJ-fluid used in these simulations. Figure 4.28 shows the corresponding results for the high-viscosity LJ-fluid (*gamma* = 1000). For the high-viscosity LJ-fluid we measure $\eta \approx 205$, whereas for the standard LJ-fluid we obtain $\eta \approx 2.0$. Thus, a thousandfold increase in *gamma* produced a hundredfold increase in η.

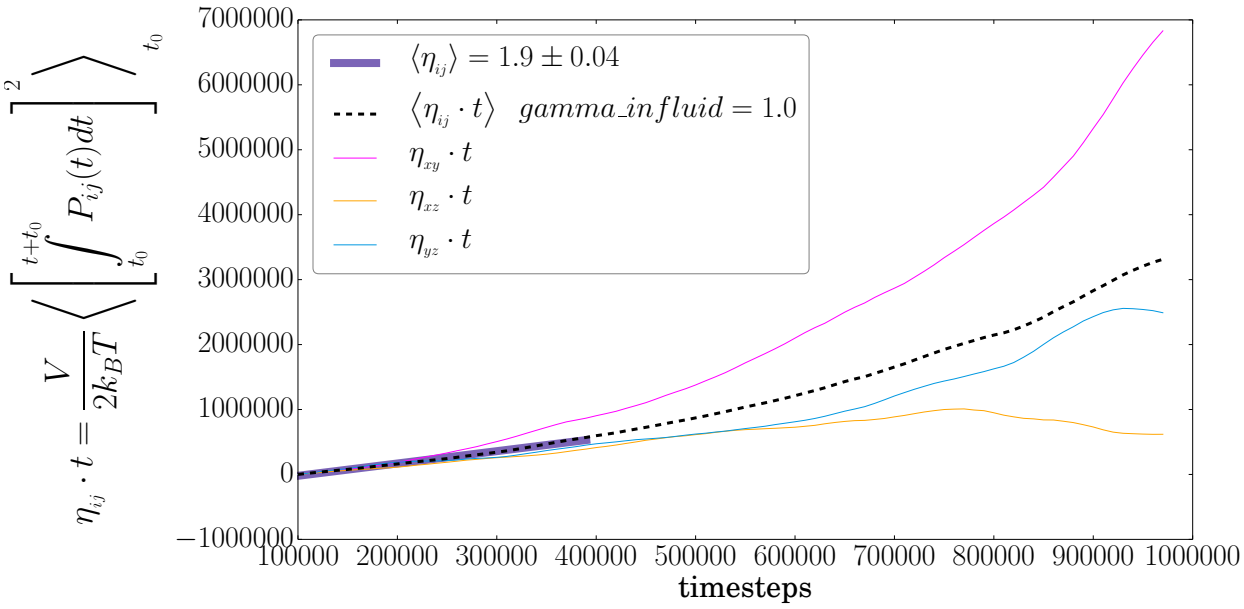

Figure 4.27: Viscosity calculation for standard Lennard-Jones fluid used in these simulations (*gamma* = 1). The slope yields a viscosity $\eta \cong 2.0$, which compares well with [29].

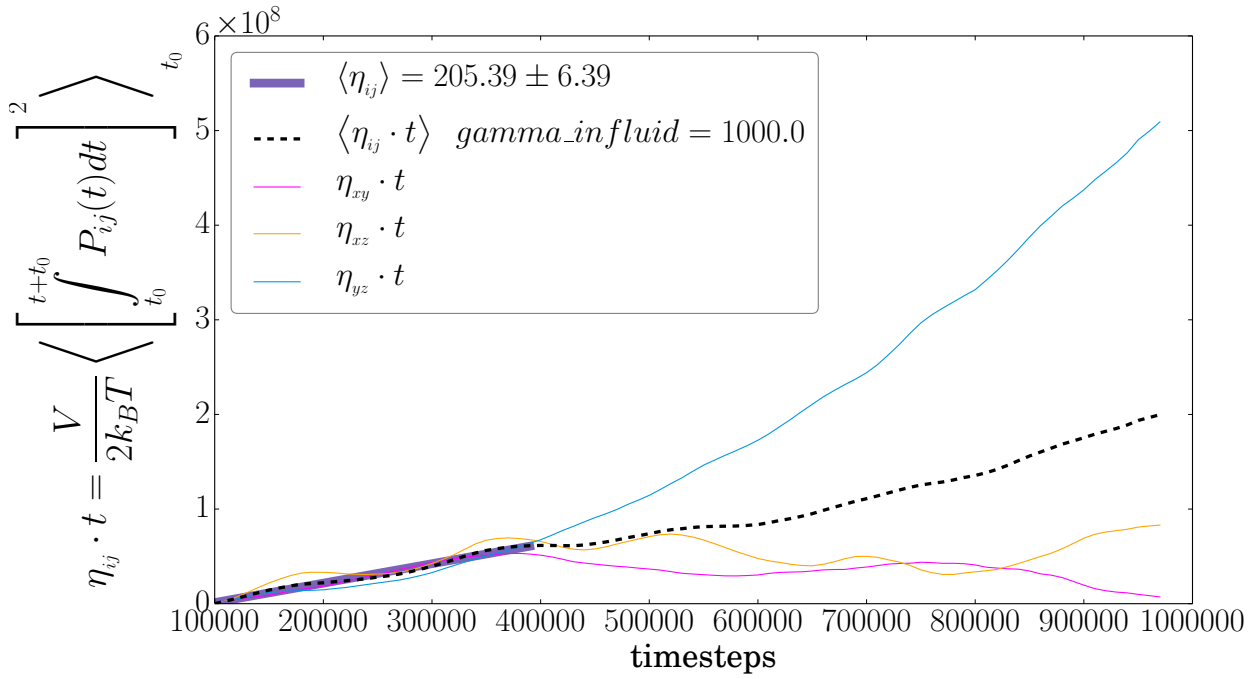

Figure 4.28: Viscosity calculation for Lennard-Jones fluid whose viscosity has been 'artificially' increased by setting the DPD drag coefficient ***gamma*** $= 10^3$. The slope yields a viscosity $\eta \approx 205$.

Density

Here we examine the phase-diagram of the Lennard-Jones fluid. Because the density of the vesicle contents increases significantly as a result of compression, we would like to check whether a phase transition (gas⟶liquid) could occur, as this could dramatically skew our results.

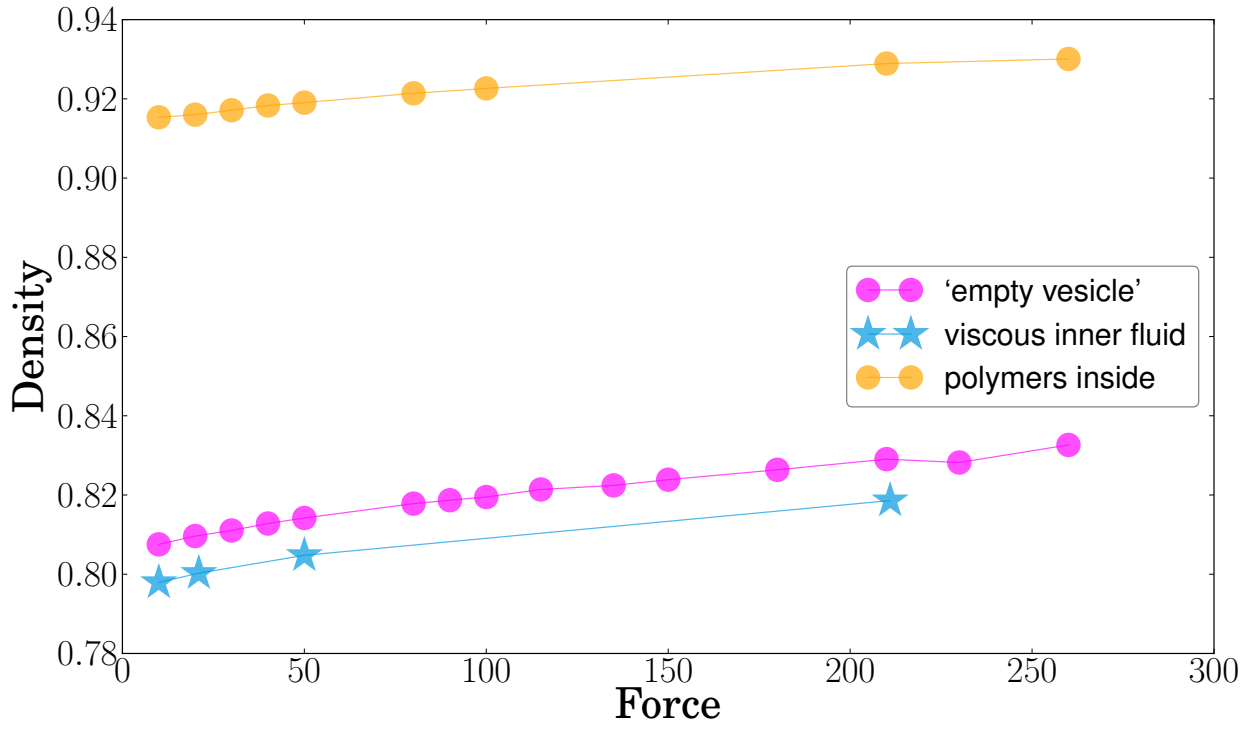

Figure 4.29: Density of vesicle contents versus force, for three different vesicles. The first vesicle (pink circles) contains standard Lennard-Jones fluid. The second vesicle (blue stars) has had the viscosity of its inner fluid 'artificially' increased via the DPD thermostat. The third vesicle contains $\gtrsim 85\%$ polymers (length=30 beads), the remaining portion of its contents being standard Lennard-Jones fluid. For the vesicle filled with polymers the density is similar to that of a LJ-liquid —see the blue line in Figure 4.30.

Referencing Figure 4.30 one can see from Figure 4.29 that for the ordinary and high-viscosity Lennard-Jones fluids used in our vesicles, the density remains well within the gas-phase regime. Even at high forces, neither of the two Lennard-Jones fluids approaches a phase transition.

One should note that Figure 4.30 applies to *LJ-fluid* (no polymers), whereas Figure 4.29 does include some data where a large fraction of the fluid particles have been bonded together into polymers.

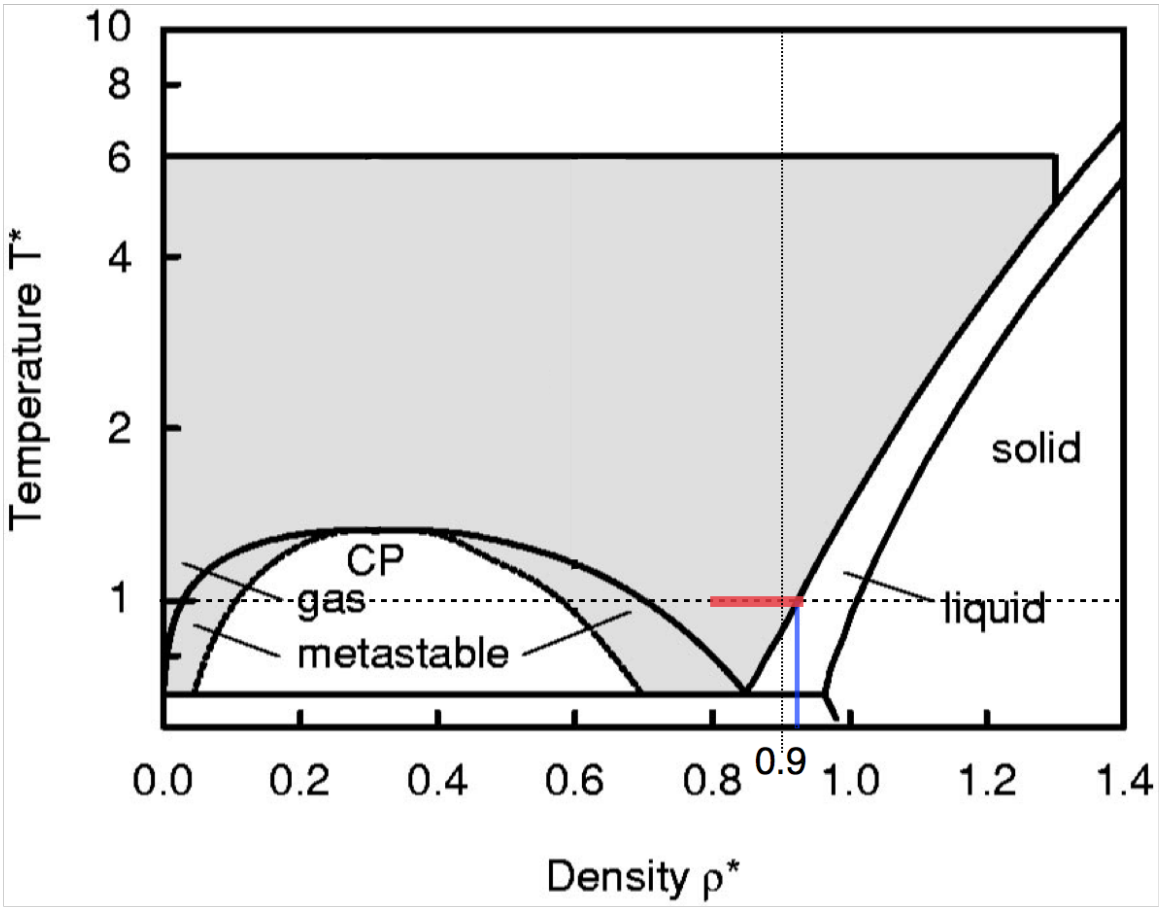

Figure 4.30: Phase diagram of the Lennard-Jones fluid (Adapted with permission[24] from [29]). The state of the Lennard-Jones fluid contained within the vesicle is confined to the region highlighted in red. The blue arrow indicates the boundary between the liquid and gaseous states at $T^{\star} = 1$.

[24] No additional uses are granted (such as derivative works or other editions).

Discussion

Relaxation Time τ

Relaxation time is Not constant!

In the following section we discuss the key result of these experiments, which is the observation that the relaxation time depends strongly on the applied stress —and that it does so in a manner consistent with an entropic spring. This result is predicted by the Helfrich model, which describes the effect of entropic undulations on the area expansion of membranes under tension. The functional dependence of the vesicle's relaxation time τ and of its projected area A_{proj} on the surface tension γ are well described (at least qualitatively) by the Helfrich model.

We begin with a simplified model of the vesicle, which serves as an analogy to the full-fledged Helfrich description of the vesicle. Then we derive τ as a function of the tension directly from the Helfrich model, and explore the additional information gleaned.

Accounting for $\tau(F)$

Undulations

Membranes behave as entropic springs. Thermal agitation pumps undulations in the vesicle membrane. When we press down on a vesicle we increase the membrane tension, which flattens out these undulations. This flattening of undulations by surface tension reduces the number of states[1] available to the vesicle, decreasing its entropy —just as pulling the ends of an entropic spring reduces the number of states[2] available to the spring. And just like the entropic spring, the squeezed vesicle exerts a temperature-driven restoring force pushing the

[1] In the case of the vesicle, available 'states' correspond to 'shapes' its undulating membrane (and contents) can assume. It is constantly fluctuating between different membrane shapes. Flattening the undulations literally reduces the number of shapes available to the vesicle.

[2] For the spring, 'states' correspond to spacial configurations of the beads.

system back toward greater entropy (undulations). This force is purely entropic, and must be countered in order to prevent the vesicle from reestablishing the suppressed undulations. If we increase the surface tension still beyond the point at which most undulations are suppressed, the membrane itself will begin to stretch. This stretching is called 'direct area expansion' while the flattening of undulations produces 'projected area expansion'. With direct area expansion, a different mechanism gains prominence in the generation of the restoring force —the hydrophobic effect. In this regime, further area expansion significantly increases the lipids' hydrophobic tails' exposure to the surrounding water.

There are then two distinct regimes to consider when studying the deformation of vesicles. In one regime the restoring force is primarily entropic and in the other, enthalpic.

The linear viscoelastic model describes the relaxation time as a viscosity divided by an elastic modulus

$$\tau \sim \frac{\eta}{K}. \tag{5.1}$$

To account for the force-dependence of τ, we redefine K as $K[F] = K(\alpha(F))$ in a form which incorporates the true entropic-enthalpic character of the relaxation process. As a result of the true modulus $K(\alpha)$ changing throughout the relaxation, an effective modulus $K_{effective}$ is observed which depends on the applied force. While it is true that the differential equations underlying the linear viscoelastic model assume a constant K, we take the view that for every probing force there is an *effective* K constant which one plugs in to give the correct τ. Indeed, for any particular force, the SLS model describes the relaxation quite accurately (Figure 4.7). That is to say, the force dependence of K does not mean that the SLS model which assumes constant K cannot accurately describe the relaxation.

Simplified model: two springs in series

One might wish naively to describe a vesicle as a pair of springs in series.

The 1^{st} spring k_1 represents the entropic component of the restoring force, and the second k_2 gives the restoring force associated with direct area expansion. But k_1 is special —it has a finite extensibility β so that for $F \geqslant F_1^{max} \equiv k_1\beta$, it behaves like a rigid rod of length β. For forces greater than F_1^{max}, only k_2 stretches.

We want to describe force-extension law in a manner analogous to Hooke's law. For $F \geqslant F_1^{max}$,

$$F = K \cdot (\Delta x_1 + \Delta x_2) = K\left(\beta + \frac{F}{k_2}\right)$$

$$\Longrightarrow K(F) = \frac{F}{\beta + \frac{F}{k_2}} = \frac{1}{\frac{\beta}{F} + \frac{1}{k_2}}.$$

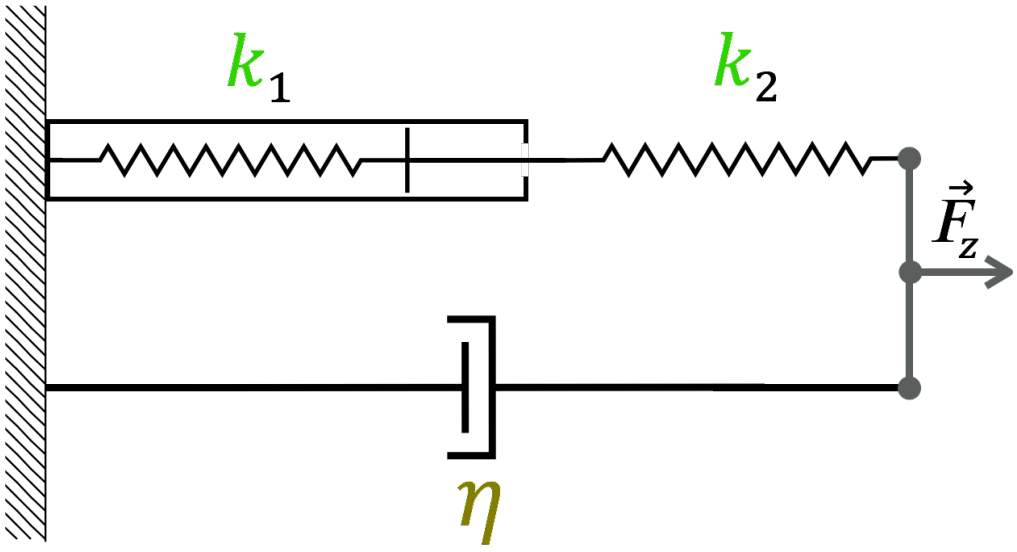

Figure 5.1: Two springs in series. The first spring has a finite extensibility above which it behaves like a rigid rod. The second spring is purely Hookean.

Plugging $K(F)$ into the SLS model then gives

$$\tau = \frac{\eta}{K(F)} = \eta \left(\frac{1}{\frac{\beta}{F} + \frac{1}{k_2}} \right)^{-1}$$

$$\therefore \boxed{\tau(F) = \eta \left(\frac{\beta}{F} + \frac{1}{k_2} \right)}. \tag{5.2}$$

It turns out that this model fits the simulation data beautifully (Figure 5.3). But unfortunately, it does not provide a theoretical basis for the actual value of β. And really, using constant β is just an approximation —valid so long as the vesicle is squashed by a force greater than the entropic restoring force in the regime where undulations have been mostly flattened.

For if $F < F_1^{max}$ then the 'two springs in series' model gives

$$F = K \cdot (\Delta x_1 + \Delta x_2)$$

$$= K \left(\frac{F}{k_1} + \frac{F}{k_2} \right)$$

$$\rightarrow K = \frac{F}{\frac{F}{k_1} + \frac{F}{k_2}}$$

$$= \frac{1}{\frac{1}{k_1} + \frac{1}{k_2}} = \left(\frac{k_2 + k_1}{k_1 k_2} \right)^{-1}$$

$$\Longrightarrow K(F < F_1^{max}) = \frac{k_1 k_2}{k_1 + k_2} = const.$$

and plugging into the SLS model then gives

$$\tau(F < F_1^{max}) = \eta \left(\frac{k_1 + k_2}{k_1 k_2} \right) = const. \tag{5.3}$$

But rather than settling abruptly to a constant at some threshold F_1^{max} Figure 5.3 suggests that τ continues to increase as $F \longrightarrow 0$. And in reality, there is no reason to expect an abrupt transition at low force, but rather a smooth increase in the relaxation time as $F \longrightarrow 0$. So Equation 5.3 is contradicted by the available data, and the simple 'two springs in series' model fails in the regime of vanishing F. Unfortunately, as the force vanishes so does α and with it the signal-to-noise ratio in the data, so that we can barely get a signal for $F \lesssim 10\text{F}$.

To improve the model, we need to replace β with an expression that varies smoothly from the entropic regime ($F \longrightarrow 0$) to the enthalpic (direct area expansion) regime ($F \gg 1$). We need to know the functional form of the entropic restoring force in order to connect this expression for τ to its thermodynamic origin.

In depth – Helfrich model

From linear viscoelasticity theory, we have

$$\tau \sim \frac{\eta}{K} \longleftarrow \textit{assumes } K \textit{ const.} \tag{5.4}$$

So to understand how the relaxation time varies with stress (surface tension), we need to know: which viscosity is η, and what modulus is K? The most physically appropriate viscosity is called the *dilatational-surface viscosity*[10, 33] η_d —the viscosity associated with stretching the membrane. We assume $\eta_d \cong \textit{constant}$[3] so that

$$\tau \propto \frac{1}{K}. \tag{5.5}$$

To obtain K we return to the heart of elasticity theory. Rewrite Hooke's law

$$\Delta(\textit{stress}) = K \times \Delta(\textit{strain}) \tag{5.6}$$

as

$$\Delta(\textit{strain}) = \frac{1}{K}\Delta(\textit{stress}). \tag{5.7}$$

This suggests a more general definition for the quantity K. For small $\Delta()$ we know that

$$\Delta(\textit{strain}) \cong \left(\frac{\partial(\textit{strain})}{\partial(\textit{stress})}\right)\Delta(\textit{stress}) \tag{5.8}$$

(becoming exact in the limit $\Delta(\textit{stress}) \to 0$), which leads us to redefine

$$\boxed{\frac{1}{K} \stackrel{\text{def}}{=} \frac{\partial(\textit{strain})}{\partial(\textit{stress})}}. \tag{5.9}$$

In the case of a stretching membrane $\textit{strain} = \alpha$ (i.e. area expansion) and $\textit{stress} = \gamma$ (i.e. surface tension), so that

$$\boxed{\frac{1}{K} = \frac{\partial\alpha}{\partial\gamma}} \tag{5.10}$$

defines the effective modulus K of the bilayer (in the vicinity of a specific value of γ).

[3] i.e. over the range of deformation studied here.

Now Equation 2.1 from Helfrich and Servuss [23],

$$\alpha(\gamma)\Big|_{\gamma>0} = \underbrace{\frac{k_B T}{8\pi\kappa}}_{\text{“}M\text{”}} \ln\left(\frac{1+\frac{\gamma A}{\zeta\kappa}}{\frac{A}{a}+\frac{\gamma A}{\zeta\kappa}}\right) + \frac{\gamma}{K_A}, \tag{5.11}$$

will take us home. (We have written $\zeta\kappa$ instead of $\pi^2\kappa$ in order to remove explicit assumptions about the bilayer shape.) To make things easier, we first expand Equation 5.11 as

$$\alpha\Big|_{\gamma>0} = \frac{\gamma}{K_A} + M\left\{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right) - \ln\left(\frac{A}{a}\left[1+\frac{\gamma a}{\zeta\kappa}\right]\right)\right\}. \tag{5.12}$$

The effective modulus is then

$$\frac{1}{K} = \overbrace{\frac{\partial}{\partial\gamma}\left(\left(\frac{\gamma}{K_A} + M\left\{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right) - \ln\left(1+\frac{\gamma a}{\zeta\kappa}\right) - \ln\left(\frac{A}{a}\right)\right\}\right)\right)}^{\left(\frac{\partial\alpha}{\partial\gamma}\right)_{\gamma>0}}, \tag{5.13}$$

$$= \frac{1}{K_A} + M\left\{\frac{1}{\left(1+\frac{\gamma A}{\zeta\kappa}\right)}\frac{\partial}{\partial\gamma}\left(1+\frac{\gamma A}{\zeta\kappa}\right) - \frac{1}{\left(1+\frac{\gamma a}{\zeta\kappa}\right)}\frac{\partial}{\partial\gamma}\left(1+\frac{\gamma a}{\zeta\kappa}\right)\right\},$$

$$= \frac{1}{K_A} + M\left\{\frac{\frac{A}{\zeta\kappa}}{1+\frac{\gamma A}{\zeta\kappa}} - \frac{\frac{a}{\zeta\kappa}}{1+\frac{\gamma a}{\zeta\kappa}}\right\},$$

$$\therefore \boxed{\frac{1}{K} = \frac{1}{K_A} + M\left\{\frac{1}{\frac{\zeta\kappa}{A}+\gamma} - \frac{1}{\frac{\zeta\kappa}{a}+\gamma}\right\}}. \tag{5.14}$$

Now, all simulations performed here occur in the regime $\gamma \ll \frac{\zeta\kappa}{a}$ and thus

$$\frac{1}{\frac{\zeta\kappa}{a}+\gamma} \cong \frac{a}{\zeta\kappa};$$

$$\rightarrow \frac{1}{K} \cong \frac{1}{K_A} + M\left\{\frac{1}{\frac{\zeta\kappa}{A}+\gamma} - \frac{a}{\zeta\kappa}\right\}. \tag{5.15}$$

Plugging in the measured area per lipid $a \approx 2.0$ and the fit result $\zeta\kappa \approx 20.0$, we take the

approximation

$$\boxed{\frac{1}{K} \cong \frac{1}{K_A} + \frac{M}{\frac{\zeta\kappa}{A} + \gamma}}. \tag{5.16}$$

Note that this approximation (5.16) is not particularly sensitive to the quantity $\zeta\kappa$ — it remains accurate to within a few % of (5.15) so long as $\zeta\kappa > 1$.[4]

Helfrich $\tau(\gamma)$

Returning to the relation between relaxation time, viscosity and elasticity (5.4), Equation 5.16 predicts (via the Helfrich model) a relaxation time

$$\boxed{\tau(\gamma) \sim \frac{\eta}{K} \cong \eta\left(\frac{1}{K_A} + \frac{M}{\frac{\zeta\kappa}{A} + \gamma}\right)}. \tag{5.17}$$

Note that this result differs from the relation (2.6) given by Dimova et al.[33] in some important ways. According to Equation 5.17, the relaxation time should approach a *finite* limit as the tension vanishes

$$\lim_{\gamma\to 0} \tau(\gamma) = \textit{constant} \quad \Big| \cong \eta\left(\frac{1}{K_A} + \frac{MA}{\zeta\kappa}\right). \tag{5.18}$$

And in the regime $\gamma \gg \frac{\zeta\kappa}{A}$ (i.e. larger tensions) we should observe

$$\tau(\gamma) \cong \eta\left\{\frac{M}{\gamma} + \left(\frac{1}{K_A} - M\frac{a}{\zeta\kappa}\right)\right\}, \tag{5.19}$$

so that the relaxation time decreases asymptotically

$$\boxed{\tau(\gamma) \propto \frac{\eta}{\gamma} + \textit{constant}}^{5} \tag{5.20}$$

with growing tension (c.f. Equation 5.14).

As it turns out, this is exactly the kind of behaviour we see! In Figure 5.2 we plot the relaxation time versus tension, and do a fit to $\tau(\gamma)$ using a function of the form

$$\tau \approx C_1 + \frac{C_2}{C_3 + \gamma}, \tag{5.21}$$

i.e. Equation 5.17. The only fit parameter we're really interested in is C_1. Based on the

[4] It is *extremely* unlikely that $\zeta\kappa \lesssim 1$.

[5] Notice that this agrees with the observation by Dimova et. al[33] that near lysis tension, $\tau \sim \frac{\eta}{\gamma}$.

Helfrich model as well as some dimensional analysis, we expect

$$C_1 \approx \frac{\eta_d}{K_A}, \tag{5.22}$$

where η_d is called the *dilatational*-surface viscosity. And since we have independently found $K_A \approx 8.4$,[6] we can thence extract the viscosity term $\eta_d \approx K_A C_1$.[7]

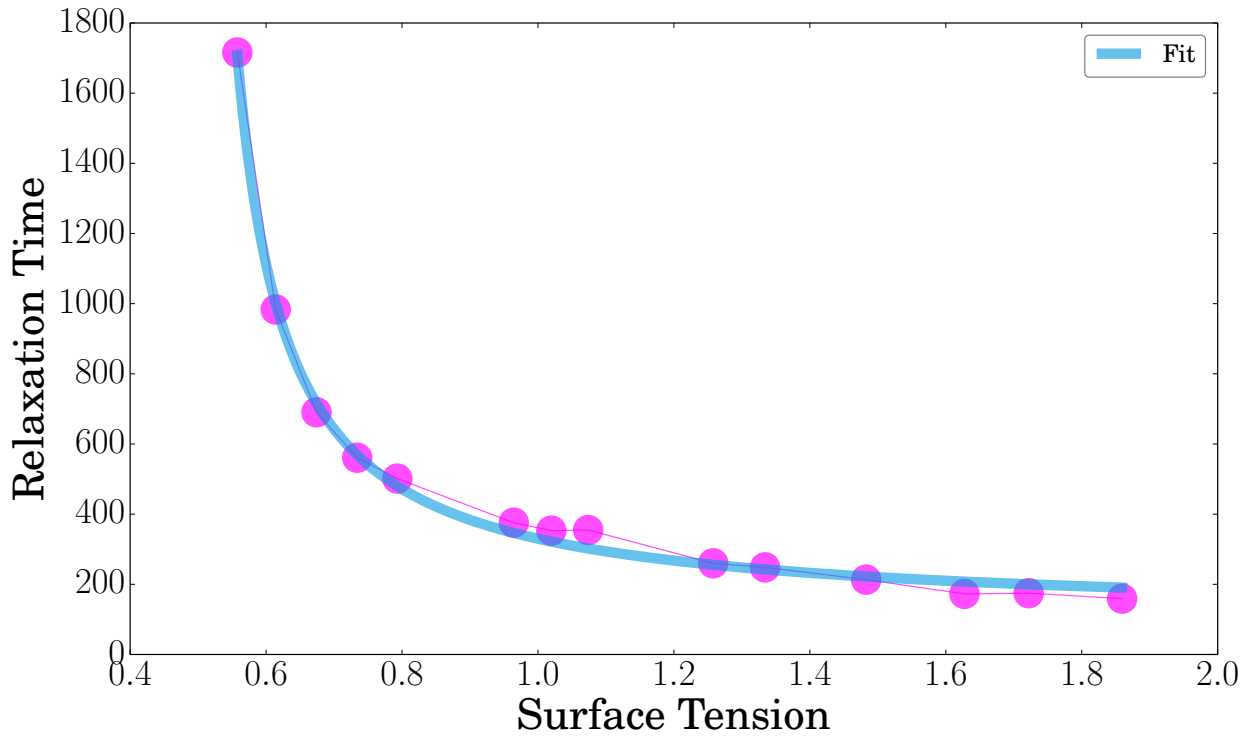

Figure 5.2: The Helfrich model leads to a correct qualitative description of $\tau(\gamma)$ —relaxation time versus surface tension. (The relaxation time is fit to a function having the form of Equation 5.17 —which was derived from the Helfrich model.) The fit (blue line) to $\tau(\gamma)$ also yields an estimate of the quantity (η_d/K_A).

The fit in Figure 5.2 estimates $\frac{\eta_d}{K_A} \cong 108 \pm 16\tau$, which corresponds to

$$\eta_d \approx 900\tau\frac{F}{\sigma} \approx 3.6 \times 10^{-11}\frac{Ns}{m}. \tag{5.23}$$

Interestingly, this viscosity is $\cong 3\times$ the value of η_s (*shear*-surface viscosity)[8] reported by den Otter et al.[8] for simulated DPPC bilayers. We expect *our* η_s to be smaller than that of

[6] See Figure 5.6

[7] cf. Equation 5.17.

[8] *dilatational*-surface viscosity (η_d) and *shear*-surface viscosity (η_s) have equivalent dimensions.

[8] since they used longer, two-tailed lipids. However for real lipid bilayers, the *dilatational*-surface viscosity η_d can be two orders of magnitude[10] larger than η_s. Given this fact, it is actually quite reasonable that our η_d should be larger than [8]'s η_s as well.

The parameter ζ

For planar bilayers[9] $\zeta = \pi^2$ and for quasispherical vesicles $\zeta = 24\pi$ —nearly an order of magnitude larger. Our vesicle starts out somewhat spherical, but becomes more and more oblate as it is flattened at higher forces. Therefore, in our case the parameter ζ in Equation 5.11 is in fact *not* a constant. While the Helfrich model succeeds as a qualitative description for both the relaxation time and area expansion (Figure 5.5), this may explain the quantitative inaccuracy of the fit-parameters.

Comparison to naive model

Note that Equation 5.17 (derived from the Helfrich model) agrees with the 'two springs in series' model, in that both predict a finite limit for τ at vanishing stress as well as the same asymptotic behaviour when the stress becomes large. However the Helfrich model describes τ as varying smoothly over the entire force range, whereas the 'two springs in series' model describes τ as hitting an abrupt plateau at $F < F_1^{max}$.

As a physical description the Helfrich model is superior, in that it provides the correct qualitative description of both the relaxation time as well as the area expansion —and reveals the basic connection between them. The form of $\alpha(\gamma)$ —and thence $\tau(\gamma)$— is derived from the equipartition of energy among thermally excited undulations in the bilayer. Thus the qualitative character of α and τ is grounded in the statistical mechanics of the membrane.

[9] (with periodic boundary conditions)

Effect of altering vesicle contents

It stands to reason that if the relaxation time τ of the vesicle is proportional to a viscosity term, then by increasing the viscosity of the vesicle contents, we ought to see a corresponding increase in the relaxation time. However, the inner fluid viscosity seems to play a minor role in changing the relaxation time. This is not surprising, since for vesicles as small as ours, the membrane and the contents are made up of similar numbers of particles.

Vesicle filled with polymers

One way to increase the viscosity of the vesicle contents, is to introduce long polymers within it. As can be seen in Figure 5.3, the presence of long polymers in high concentration produces the expected effect.

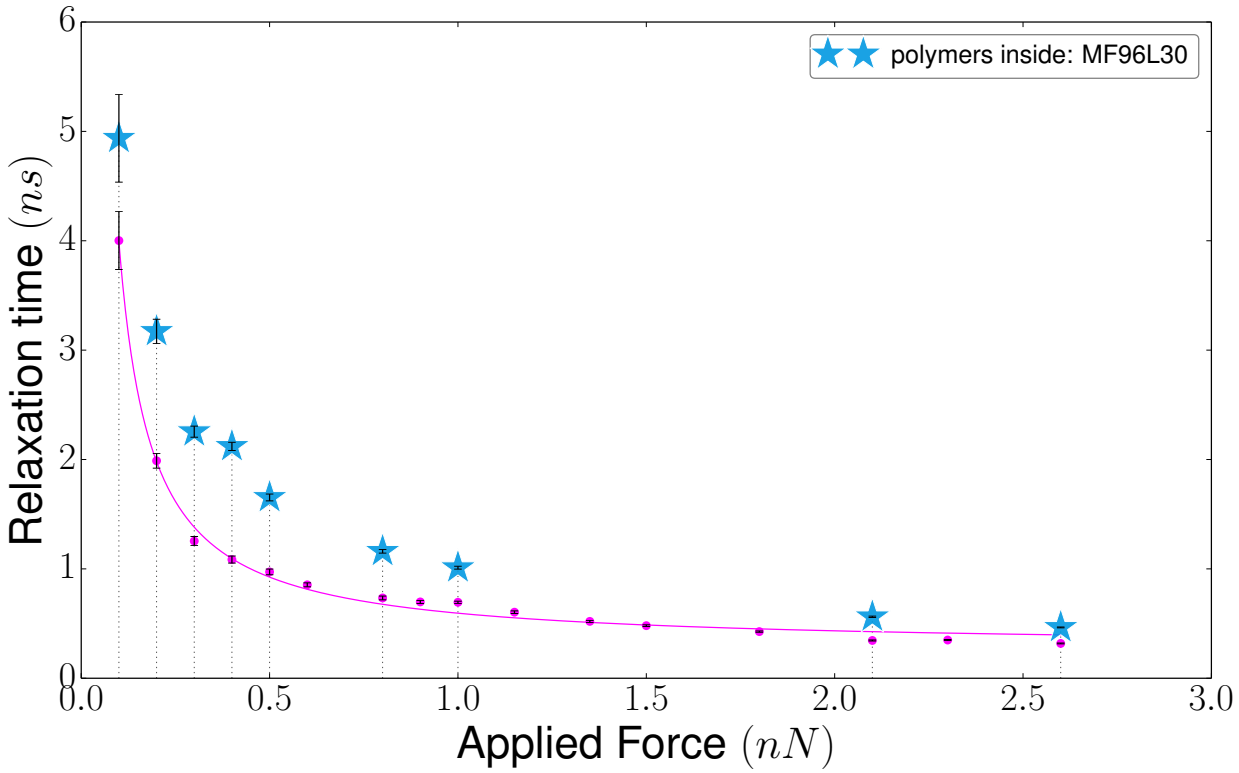

Figure 5.3: Comparing relaxation times for fluid-filled and polymer filled vesicles (polymer length = 30 beads). The vesicle filled with (an $\approx 85\%$ solution of) polymers shows markedly longer relaxation times than a fluid-filled vesicle over the same range of forces.

However, the reason for the observed increase in relaxation time may not be so straightforward. By polymerizing the inner fluid, we reduce to freedom of motion of the fluid particles which become incorporated into polymers. Therefore, polymerization of the inner fluid results in a decrease in the vesicle's internal pressure. To restore the internal pressure to that of the

vesicle filled with non-polymerized fluid, the vesicle must be "pumped up" —i.e. additional fluid particles must be added to the vesicle interior. Thus the addition of polymers leads to an *increase* in the density of the vesicle's contents —which surely would increase the viscosity in and of itself. There may also be interactions between the polymers and the bilayer —perhaps causing changes in the membrane's apparent elasticity by damping undulations.

(As an interesting side-note, the vesicle filled with polymers is apparently able to withstand *much* greater compression forces —despite containing more particles at higher density— than the 'empty' vesicle. The polymer-filled vesicle can withstand forces $\boldsymbol{F}_z$ in excess of $350\mathbf{F}$ whereas for the vesicle containing only fluid, lysis begins at $\boldsymbol{F}_z \gtrsim 270\mathbf{F}$. The explanation may be two-fold. Initially, the increased viscosity of the vesicle contents may dampen the impact of the AFM tip. The increased viscosity may also dampen local-density fluctuations in the membrane which lead to pore nucleation —making the vesicle more robust.)

Vesicle filled with high-viscosity fluid

Another way to increase the viscosity of the vesicle contents, is to 'artificially' increase the inner fluid viscosity —i.e. by altering the fluid particles' interactions. Instead of introducing polymers, we simply increase the viscosity by tuning the viscous drag coefficient in the DPD interaction (see Figures 4.27 and 4.28).

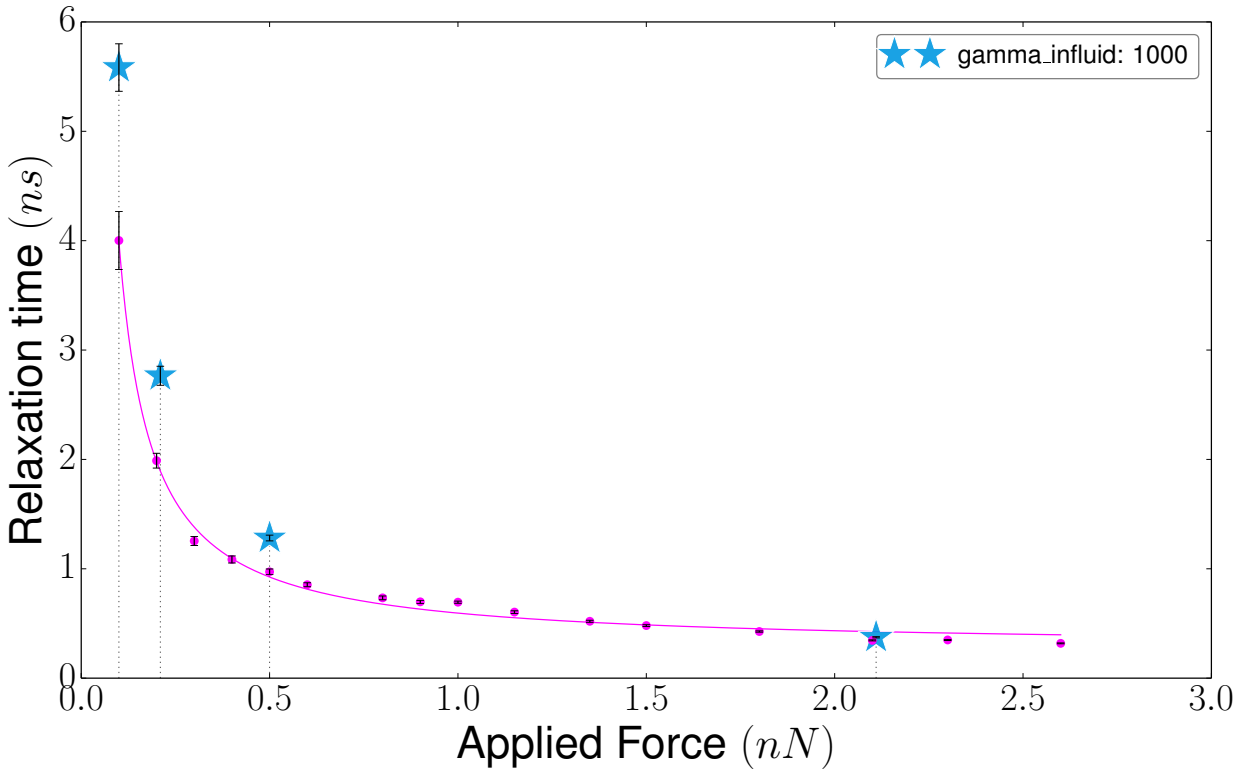

Figure 5.4: Increased relaxation time observed for vesicle filled with high-viscosity fluid (Blue stars). The viscosity of the inner fluid was increased via the viscous drag parameter in the DPD interaction. Despite increasing the viscosity $\approx 100\times$, the relaxation time increases only by $\approx 30\%$. Further, the increase is significant only at low forces.

The effect of increasing the inner fluid viscosity is rather small. In Figure 4.28, the vesicle with the longer relaxation times (blue stars) is filled with fluid having a viscosity $\approx 100\times$ that of the ordinary vesicle (pink circles). However, the two vesicles' relaxation times differ from one another by only $\approx 30\%$. Moreover, the effect vanishes at higher forces, which suggests that the membrane properties dominate the relaxation. This might further suggest, that the key to understanding $\tau(\gamma)$ lies not in the the vesicle's contents, but in the entropic elasticity of its membrane. As with the vesicle filled with polymers, the high-viscosity inner fluid may affect undulations (and other phenomena) in the bilayer.

Area Expansion α

Projected area versus tension: Helfrich model

The Helfrich model —which describes the area expansion α as a function of the tension γ— correctly predicts the observed dependence of the relaxation time $\tau(\gamma)$. We therefore would expect the projected area of the vesicle to behave, at least qualitatively, in accordance with the Helfrich model.

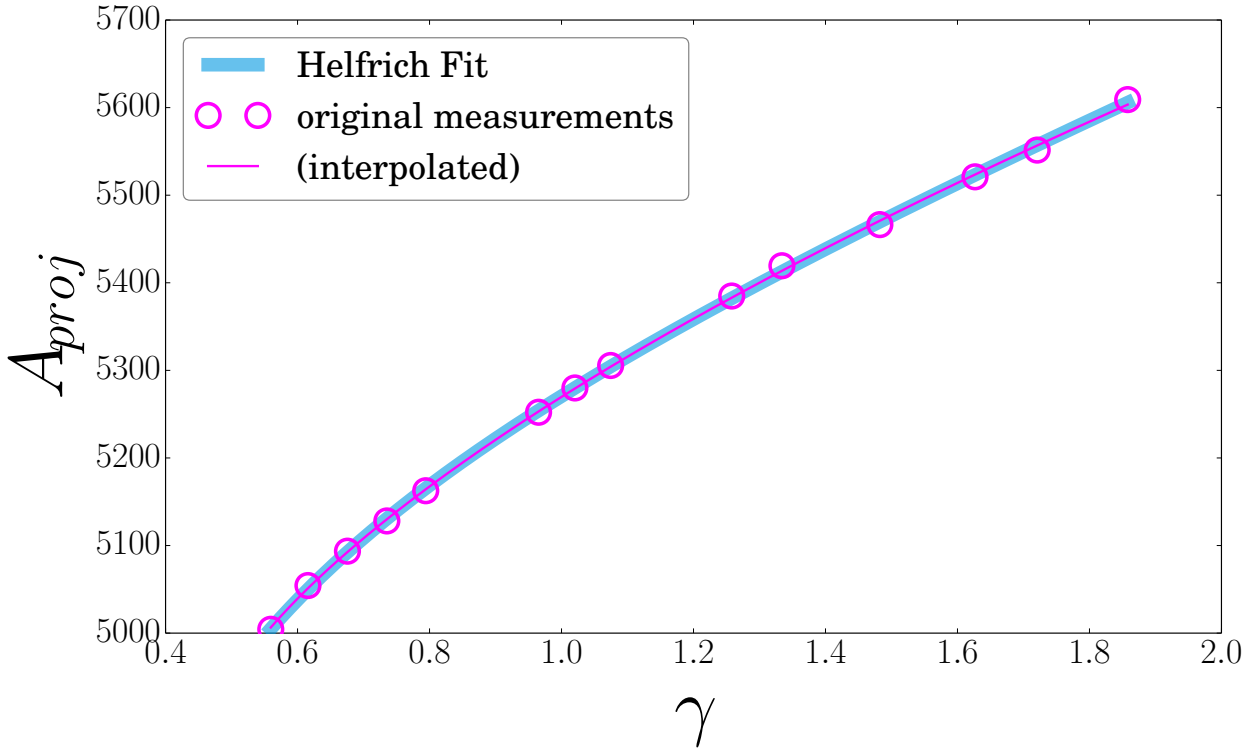

Figure 5.5: Fit to Helfrich model: Projected Area versus surface tension —wherein the tension (x–axis) has been estimated via the 'work approach' (see Figure 4.17, Equation 2.19).

For this reason, we plot in Figure 5.5 the area expansion as a function of the tension (calculated via the work approach – see Figure 4.17). A curve-fit of the Helfrich model to the projected-area data shows very good qualitative agreement. We produce the "qualitative fit" function

as follows:

$$\alpha(\gamma) \equiv \frac{A(\gamma) - A_\circ}{A_\circ} \tag{5.24}$$

$$\longrightarrow \boxed{A(\gamma) = A_\circ\Big(1 + \alpha(\gamma)\Big)}. \tag{5.25}$$

Plugging in $\alpha(\gamma)$ from the Helfrich model (Equation 2.1), this becomes

$$A(\gamma) = C_1\left(\left(1 + C_2\gamma + C_3\left\{\ln\Big(1 + C_4\gamma A\Big) - \ln(1 + C_4\gamma a) - \ln(N)\right\}\right)\right) + C_5, \tag{5.26}$$

where $C_1 \sim A_\circ$, $C_2 \sim K_A^{-1}$, $C_3 \sim \frac{k_B T}{8\pi\kappa}$, $C_4 \sim \frac{1}{\zeta\kappa}$, N is the total number of lipids in the bilayer, and a is the 'area per lipid', i.e.

$$a \equiv \frac{A}{N} = \frac{A^{inner\,leaflet} + A^{outer\,leaflet}}{3000 lipids}. \tag{5.27}$$

C_5 is an extra parameter that allows the fitting-software to boost the curve without affecting C_2 & C_3.

Equation 5.26 is somewhat of a permissive fitting function, it retains the qualitative character of the Helfrich model without imposing assumptions about vesicle shape which do not apply in our case.

Quantitative fit to linear regime

By reducing the number of free parameters in Equation 5.26, we are able to use the Helfrich model to obtain a quantitative estimate of K_A (see Figure 5.6). This is done by imposing the relation[15]

$$\kappa = K_A\left(\frac{\ell^2}{48}\right) \tag{5.28}$$

on the fitting function, where ℓ is the bilayer thickess. Bertrand et al.[6] have measured $\ell \cong 4.8\boldsymbol{\sigma}$ for a planar bilayer identical to that of our vesicle[10]. Therefore

$$\kappa \cong K_A\frac{(4.8)^2}{48} = K_A\frac{\cancel{(4.8)}(4.8)}{\cancel{4.8} \times 10} = 0.48K_A, \tag{5.29}$$

or more roughly $\kappa \approx \frac{1}{2}K_A$.

Imposing this condition on Equation 5.26 eliminates C_3 as a free parameter:

$$\kappa \approx 0.48 K_A \Longrightarrow C_3 = \frac{k_B T}{8\pi \underbrace{(0.48 K_A)}_{\kappa}} = \frac{k_B T}{3.84\pi K_A} = \frac{k_B T}{3.84\pi} C_2. \tag{5.30}$$

The fitting function thus becomes

$$A(\gamma) = C_1 \left(\left(1 + C_2 \left\{ \left[\gamma + \frac{k_B T}{3.84\pi} \left\{ ln\left(1 + C_4 \gamma A\right) - ln(1 + C_4 \gamma a) - ln(N) \right\} \right] \right\} \right) + C_5. \tag{5.31}$$

A fit of Equation 5.31 to the linear regime of $A(\gamma)$ is shown below.

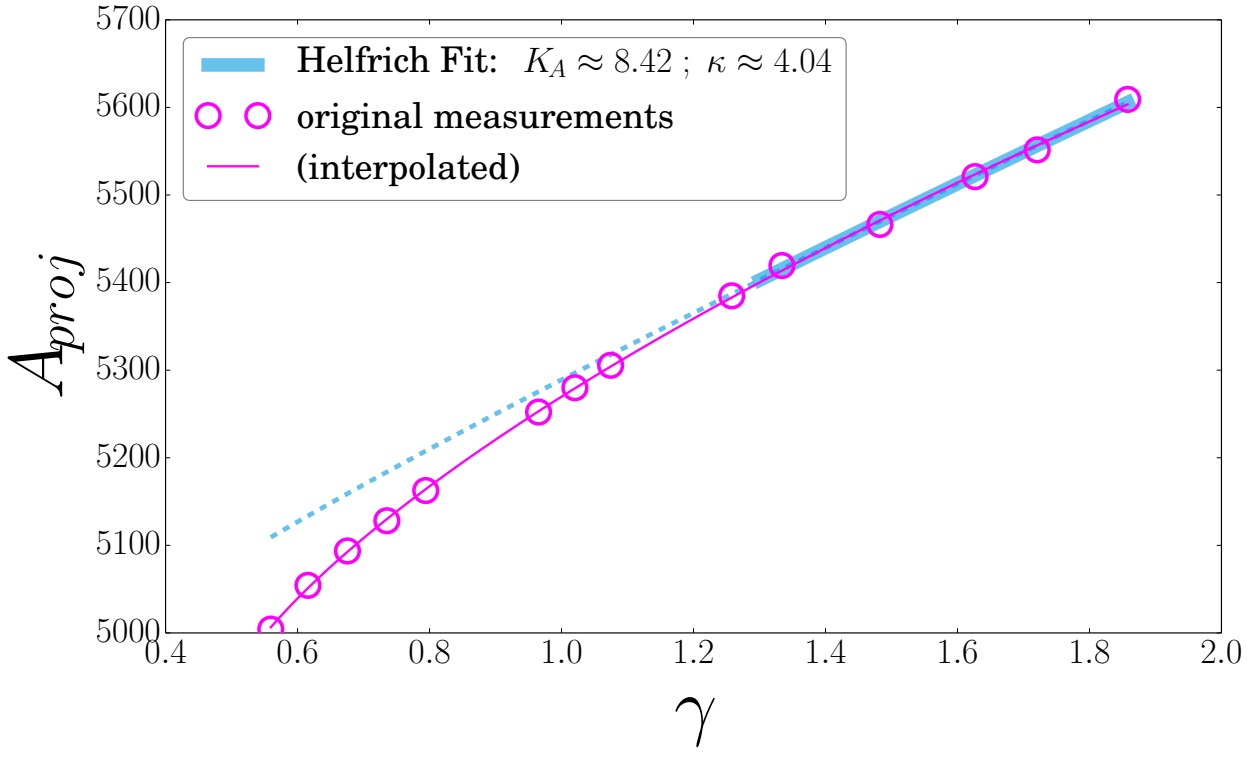

Figure 5.6: Estimating K_A & κ using the Helfrich model: quantitative fit to linear regime (solid blue line). The K_A returned by this fit agrees with that obtained by Bertrand et al.[6] (flat bilayers) to within 5%. (In its more constrained form, the Helfrich model does not simultaneously describe the entropic and linear regimes.)

The dashed line continues the fit function beyond the data selection window in which the fit was performed. We see that, in its more constrained form, the Helfrich model does not simultaneously describe the entropic and linear regimes of $A(\gamma)$. This is not surprising, since

[10] i.e. composed of identical lipids with identical interactions, immersed in identical solvent

our membrane is neither planar nor spherical, but takes on a variety of shapes over the range of forces studied here. As a result, the parameter ζ is not a constant, but instead $\zeta = \zeta(\gamma)$ —a complication which the Helfrich model does not account for. Hence, it is reasonable to just fit the high tension regime for K_A and not simultaneously expect a perfect fit to the entropic regime. Nonetheless, it is clear that our vesicle exhibits both entropic and linear regimes, and that the Helfrich model gives a complete qualitative description of both.

The K_A returned by the quantitative fit (solid blue line in Figure 5.6) agrees with that obtained by Bertrand et al.[6] to within 5%. This startlingly good agreement occurs without any suggestion to the curve-fitting software —prior to the fit, K_A is initialized at 1.0.

K_A has the dimensions of surface tension. Using the dimensionful conversions listed in Equation 3.6, we find that the conversion factor for the dimensions of surface tension is

$$\gamma \cong 0.02 \frac{\mathrm{N}}{\mathrm{m}}, \tag{5.32}$$

so that

$$K_A \longrightarrow K_A \gamma \approx 0.17 \frac{\mathrm{N}}{\mathrm{m}}. \tag{5.33}$$

Compression & Comparison to GUVs

Comparison to compression of giant vesicles (GUV data from Schäfer et al.[34])

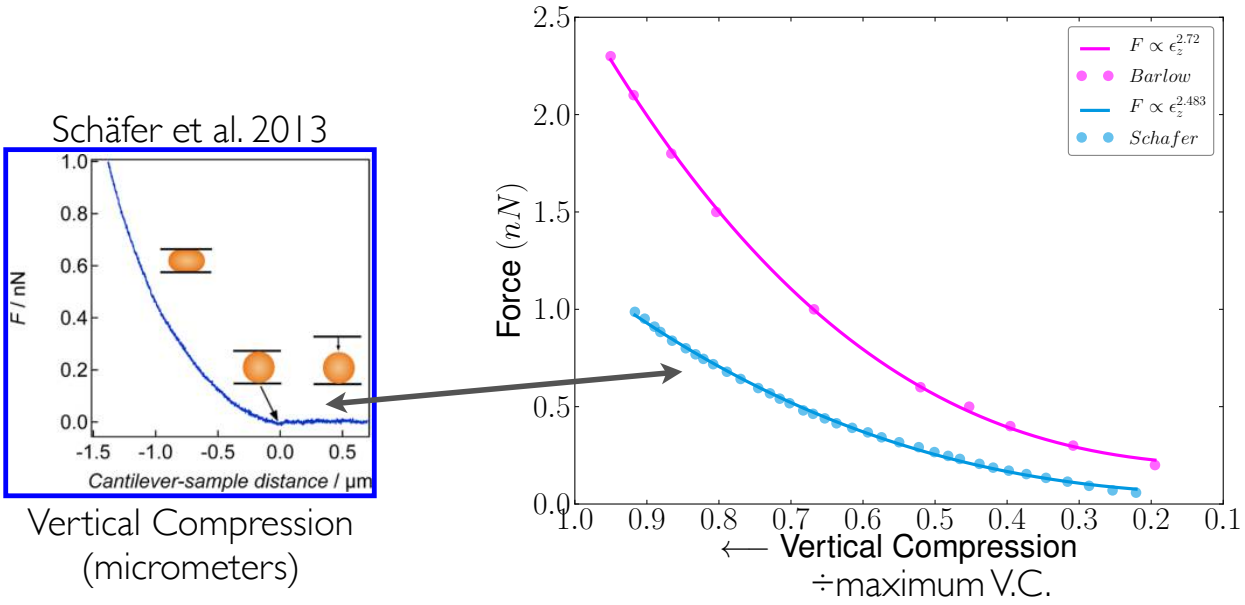

Figure 5.7: Force-compression curves for both real and simulated vesicles (left figure adapted with permission[11] from [34]). Right figure: The lower curve (blue data), transferred from Figure 4 of [34], shows the compression of a real "giant unilamellar vesicle" (GUV) at forces $\leqslant 1\text{n}\mathbf{N}$. The upper curve (pink data) shows the corresponding curve for our simulated vesicle (force values have been roughly converted to S.I. units).

A much better fit to the change in vesicle height in *both* simulation and GUV data is obtained with the function

$$Z_{fit}(F) = C_1 \mathit{ln}\Big(1 + C_2(F - C_4)\Big) + C_3(F - C_4),^{12} \tag{5.34}$$

which is used below in Figure 5.8. The qualitative similarity between simulation and experimental data is striking.

[11] No additional uses are granted (such as derivative works or other editions).

[12] Equation 5.34 has turned out to be useful as a general fit function for our purposes —it was based on the functional form of the Helfrich model.

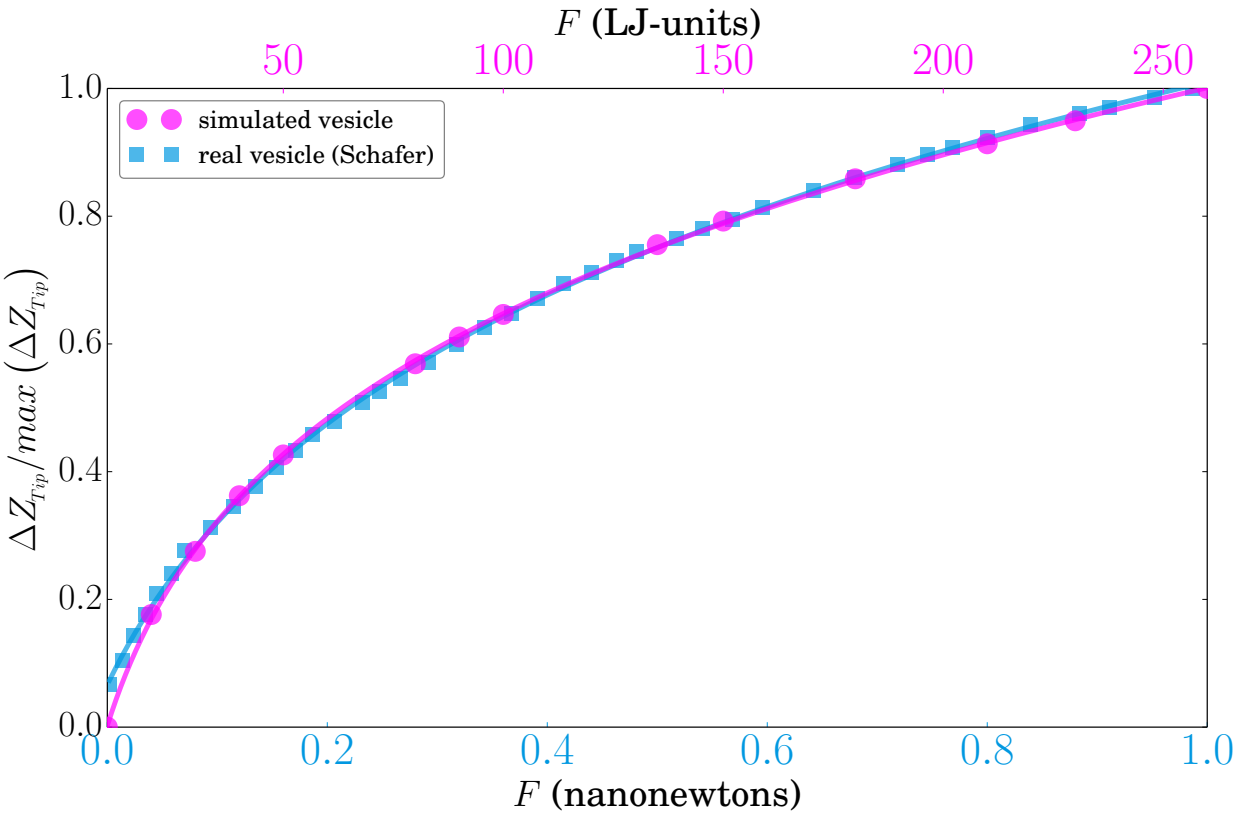

Figure 5.8: Vertical compression: Relative height variation as a function of applied force. The qualitative similarity between simulation and experimental data is striking. Both the giant vesicle data[34] and our simulation data are curve-fit using the same fitting function (Equation 5.34). Note that because our compression data stops at $F_z = 10.0\text{F}$ (whereas Schäfer's data stops much closer to $F_z = 0\text{nN}$) the comparison becomes problematic at low forces. Nevertheless we can see that the force-compression curves for real and simulated vesicles are very similar, and share the same functional form.

In Figure 5.8 the vertical compression is scaled as a fraction of the maximum compression which the (respective) vesicle can withstand. The scaled GUV[34] and simulation data overlap quite nicely. This happens in spite of (I) the immense difference in size and (II) the fact that our simulations use *compressible* fluid —real water is, for our purposes, incompressible. This suggests that it's the physical character of the undulating *membrane* —rather than the solvent— that determines the force-compression curve of a fluid-filled vesicle. The overlap means that the simulated bilayer vesicle, as a physical system, is practically indistinguishable from a real vesicle in this experiment.

Conclusion & Outlook

Central result: $\tau \neq const$

The central result of this work has been the observation that the vesicle's relaxation time τ is *not* constant (see Figure 5.2), but rather $\tau = \tau(\gamma)$ —the relaxation time is a function of the surface tension (mechanical stress). We have found that the relaxation time varies as the inverse of the tension, i.e. $\tau(\gamma) \propto \gamma^{-1} + const$, for large γ. This qualitative result follows from the Helfrich model, which describes the area expansion of bilayer membranes as a function of the tension. Further analysis in terms of the Helfrich model allowed us to estimate the membrane viscosity, via a curve fit to $\tau(\gamma)$. The estimated viscosity (*surface* viscosity of the membrane) compares very well with that observed for similar bilayers[8].

Vesicle contents

We were able to increase the relaxation time by increasing the viscosity of the vesicle's contents. This was done in two (separate) ways: (I) modifying the DPD viscous drag coefficient of the vesicle's inner fluid, and (II) filling the vesicle with polymers. In both cases, the effect was most significant at low forces. For the purpose of determining the importance of the contents' viscosity, we consider method (I) to be the most reliable since it introduced the least amount of confounding effects. While the *relative* increase in relaxation time was significant, the effect was minor when compared with the corresponding increase in the contents' viscosity. This suggests that the membrane dominates relaxation of small vesicles —which is completely reasonable given that the membrane/contents comprise similar numbers of particles.

Helfrich $\alpha(\gamma)$

Having used the Helfrich model as a physical basis for the observed $\tau(\gamma)$, we sought to test this explanation independently of the τ-data. The Helfrich model describes membrane *area expansion* $\alpha(\gamma)$ as the combined effect of flattening entropic undulations and direct stretching. If said model is the correct explanation of $\tau(\gamma)$, then $\alpha(\gamma)$ should exhibit the curvature at low tension and linearity at high tension which are characteristic of the Helfrich model. In Figure 5.5, this is exactly what we see. This lends credence to the postulate that it is the flattening of undulations which produces the nonlinearity in $\tau(\gamma)$ at low forces.

That is, if the relaxation of the membrane were purely elastic, then the curve in Figure 5.5 would be completely linear —$\alpha(\gamma)$ would be a straight line. But instead we see linearity only

at high tensions, with $\alpha(\gamma)$ becoming *curved* at low tension. This change in behaviour at low tension is precisely what one would expect to observe as a result of entropic undulations in the bilayer.

Since the Helfrich model accounts for how both relaxation time and area expansion vary with the tension, it seems very likely that it's an honest explanation.

Area compressibility modulus K_A

By fitting to the linear (high tension) regime of $\alpha(\gamma)$, we measured the area compressibility modulus at $K_A \cong 8.4\epsilon/\sigma^2$ —see Figure 5.6. This value compares well with previous simulations using similar lipids under similar conditions [6, 16, 15]. Converting our K_A into dimensionful units gives $K_A \approx 0.2\frac{N}{m}$, a value which compares well with experiments —e.g. AFM indentation of real bilayers[7], micropipette aspiration of real vesicles[11, 13].

fitting function

Interestingly, the functional form of the Helfrich model, when applied to other stress-strain curves (for which it was not intended) proves to be a very useful general fit function for various types of deformation. Examples include:[1] vertical compression, pressure difference (i.e. $P_{in} - P_{out}$), volume expansion etc.

Outlook

Really, we are interested in *cells*. We want to know how their mechanical properties result from their physical structure —and how sensitive said properties are to various structural changes. As a modest beginning we have studied the relaxation of a simple bilayer vesicle, and have already observed some interesting results. Due to entropic undulations in the membrane, the relaxation time is very sensitive to the tension —behaving linearly only as lysis tension is approached. Altering the viscosity of the vesicle contents will affect the relaxation time at low tension, but the effect diminishes as tension grows. This suggests that in small vesicles, it is the membrane viscosity which dominates relaxation at large stresses.

There is much to learn. Cells have cytoskeletons and organelles, and they live among other cells, supported by a network of external filaments called the extracellular matrix. The cytoskeleton can reorganize via chemical reactions triggered as a result of mechanical stimulus.

[1] Where the squeezing force plays the role of "strain" variable —though it is not in the units of strain.

Details such as these are topics for future work —much of which the simulation infrastructure created for this thesis can be adapted to do.

The many ways in which this work might be applied & expanded upon can be divided into (I) experiments on *real* —i.e. made of atoms— vesicles, cells & tissues, and (II) future simulations.

Experiments

The vesicles simulated here were soft and flexible, since they lacked both a membrane skeleton and cytoskeleton. Therefore, this work is most relevant to the mechanics of soft cells, such as stem cells and cancer cells. We expect that under parallel plate compression (and likely other techniques of deformation as well), cells and vesicles will exhibit relaxation times that vary greatly, depending on the magnitude of the applied stress. Anything which affects undulations in the membrane should thus affect the relaxation time.

If the relaxation time of real vesicles is found to exhibit the same characteristic tension-dependence observed here, then experimentalists studying cell dynamics will have a potential means of distinguishing relaxation of the cell membrane & flattening of undulations, from the relaxation of structures *within* the cell. It will also mean that, to draw valid comparisons among τ measurements from various experiments, one must either characterize entire $\tau(\textit{stress})$ curves, or ensure that the different measurements were performed in the high-stress regime, where τ flattens out.

And if (as we suspect) the membrane skeleton (cell cortex) suppresses undulations, then this work may provide a means of detecting changes in the cytoskeleton. For example, ischaemic neurons are believed to undergo degradation of the membrane skeleton. If (as a result of said degradation) the membrane undulates more strongly, then the effect of these undulations may show up in the entropic regime of the $\tau(\gamma)$, $\alpha(\gamma)$ curves —i.e. relaxation time / area expansion measurements could reveal pathology of the neuron. Because τ is most sensitive to undulations at *low forces*[2], only *gentle* perturbations of the cell would be needed to observe the relaxation time in the entropic regime —reducing the confounding influence of secondary effects.

While we have focused on the area expansion —which may be difficult to measure in vitro— as a means of observing relaxation in the membrane, experimentalists may prefer other measures of deformation, such as vertical compression of the vesicle et cetera. This should not be a hindrance, since the signature of the relaxation process will be present in a variety of different measurements.

[2] less squeezing force ⟶ less tension

Simulation

We consider future simulations which might expand on this work in the following way: beginning with the current system —a single vesicle made of 3-bead lipids in a box of Lennard-Jones fluid— what details might we add or change to produce the new simulations? From this perspective, many possibilities for future simulation research can be thought of in terms of four (by no means exclusive) categories:

- **(A)** Modification of surroundings (vesicle environment, substrate, adhesion site etc.),
- **(B)** Modification of bilayer,
- **(C)** Modification of vesicle contents,
- **(D)** New measurements.

One might draw an analogy to basis vectors $\{|A\rangle, |B\rangle, |C\rangle, |D\rangle\}$ for the space of possible simulations:

$$|\mathit{modification}\rangle = m_1 |A\rangle + m_2 |B\rangle + m_3 |C\rangle + m_4 |D\rangle .$$

Examples from each category include:

- **(A)** Adding an extracellular matrix, systems of many vesicles,
- **(B)** altering composition of bilayer[3], adding a membrane skeleton (cell cortex), larger vesicle,
- **(C)** adding a cytoskeleton (perhaps with contractile filaments), inclusion of organelles,
- **(D)** power spectrum of undulations (spherical harmonics), tapping mode AFM, etc.

As we move on to simulate more realistic cell prototypes, the successful inclusion of experimental data will be essential. Due to all of the existing AFM work on living cells, there are many experimental results waiting to be used for this purpose.

We look forward to testing these results in the laboratory, and if they hold true, expanding upon them in experimental and perhaps even medical applications. It is our hope that this work will support and inspire others, in the quest to understand & apply the physics of life. Their determination bodes well for the human condition.

[3] e.g. more detailed (specific) lipids, intra-membrane cholesterols (which can give rise to time-dependent spontaneous curvature & bending modulus of the membrane —results from the migration of free molecules among the leaflets in response to stress), lipid rafts, pores, channels, etc.

Bibliography

[1] Zeinab Al-Rekabi and Andrew E. Pelling. Cross talk between matrix elasticity and mechanical force regulates myoblast traction dynamics. *Phys. Biol.*, 10(6):066003, December 2013.

[2] Bruce Alberts, Alexander Johnson, Julian Lewis, Martin Raff, Keith Roberts, and Peter Walter. *Molecular Biology of the Cell*. Garland Science, 4th edition, 2002.

[3] Nina Amenta, Marshall Bern, and Manolis Kamvysselis. A new Voronoi-based surface reconstruction algorithm. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '98, pages 415–421, New York, NY, USA, 1998. ACM.

[4] C. Barbetta, A. Imparato, and J.-B. Fournier. On the surface tension of fluctuating quasi-spherical vesicles. *Eur. Phys. J. E*, 31(3):333–342, March 2010.

[5] Martin Bertrand. *Deformed soft matter under constraints*. PhD thesis, University of Ottawa, 2012.

[6] Martin Bertrand and Béla Joós. Extrusion of small vesicles through nanochannels: A model for experiments and molecular dynamics simulations. *Phys. Rev. E*, 85(5):051910, May 2012.

[7] Chinmay Das, Khizar H. Sheikh, Peter D. Olmsted, and Simon D. Connell. Nanoscale mechanical probing of supported lipid bilayers with atomic force microscopy. *Phys. Rev. E*, 82(4):041920, October 2010.

[8] W. K. den Otter and S. A. Shkulipa. Intermonolayer Friction and Surface Shear Viscosity of Lipid Bilayer Membranes. *Biophysical Journal*, 93(2):423–433, July 2007.

[9] Rumiana Dimova, Natalya Bezlyepkina, Marie Domange Jordö, Roland L. Knorr, Karin A. Riske, Margarita Staykova, Petia M. Vlahovska, Tetsuya Yamamoto, Peng Yang, and Reinhard Lipowsky. Vesicles in electric fields: Some novel aspects of membrane behavior. *Soft Matter*, 5(17):3201–3212, August 2009.

[10] Rumiana Dimova, Karin A. Riske, Said Aranda, Natalya Bezlyepkina, Roland L. Knorr, and Reinhard Lipowsky. Giant vesicles in electric fields. *Soft Matter*, 3(7):817–827, June 2007.

[11] E. Evans and W. Rawicz. Entropy-driven tension and bending elasticity in condensed-fluid membranes. *Phys. Rev. Lett.*, 64(17):2094–2097, April 1990.

[12] Evan A. Evans and Richard Skalak. *Mechanics and thermodynamics of biomembranes*. CRC Press, 1980.

[13] N. Fa, L. Lins, P. J. Courtoy, Y. Dufrêne, P. Van Der Smissen, R. Brasseur, D. Tyteca, and M. P. Mingeot-Leclercq. Decrease of elastic moduli of DOPC bilayers induced by a macrolide antibiotic, azithromycin. *Biochimica et Biophysica Acta (BBA) - Biomembranes*, 1768(7):1830–1838, July 2007.

[14] Deborah Kuchnir Fygenson, John F. Marko, and Albert Libchaber. Mechanics of Microtubule-Based Membrane Extension. *Phys. Rev. Lett.*, 79(22):4497–4500, December 1997.

[15] Rüdiger Goetz, Gerhard Gompper, and Reinhard Lipowsky. Mobility and Elasticity of Self-Assembled Membranes. *Phys. Rev. Lett.*, 82(1):221–224, January 1999.

[16] Rüdiger Goetz and Reinhard Lipowsky. Computer simulations of bilayer membranes: Self-assembly and interfacial tension. *J Chem Phys*, 108(17):7397–7409, May 1998.

[17] Louise Guolla, Martin Bertrand, Kristina Haase, and Andrew E. Pelling. Force transduction and strain dynamics in actin stress fibres in response to nanonewton forces. *J Cell Sci*, 125(3):603–613, February 2012.

[18] Kristina Haase and Andrew E. Pelling. Resiliency of the plasma membrane and actin cortex to large-scale deformation. *Cytoskeleton*, 70(9):494–514, September 2013.

[19] J. M. Haile. *Molecular Dynamics Simulation: Elementary Methods*. Wiley, June 1992.

[20] Alison Harman. *A Molecular Dynamics Simulation of Vesicle Deformation and Rupture in Confined Poiseuille Flow*. MSc Thesis, University of Ottawa, Ottawa, ON, 2013.

[21] E. Helfer, S. Harlepp, L. Bourdieu, J. Robert, F. C. MacKintosh, and D. Chatenay. Microrheology of Biopolymer-Membrane Complexes. *Phys. Rev. Lett.*, 85(2):457–460, July 2000.

[22] E. Helfer, S. Harlepp, L. Bourdieu, J. Robert, F. C. MacKintosh, and D. Chatenay. Viscoelastic properties of actin-coated membranes. *Phys. Rev. E*, 63(2):021904, January 2001.

[23] W. Helfrich and R.-M. Servuss. Undulations, steric interaction and cohesion of fluid membranes. *Il Nuovo Cimento D*, 3(1):137–151, January 1984.

[24] Alexandra L Hemsley, Diana Hernandez, Christopher Mason, AE Pelling, and FS Veraitch. Precisely delivered nanomechanical forces induce blebbing in undifferentiated mouse embryonic stem cells. *Cell Health and Cytoskeleton*, pages 23–34, January 2011.

[25] Dan D Lasic. Novel applications of liposomes. *Trends in Biotechnology*, 16(7):307–321, July 1998.

[26] C.T. Lim, E.H. Zhou, and S.T. Quek. Mechanical models for living cells—a review. *Journal of Biomechanics*, 39(2):195–216, 2006.

[27] F. London. The general theory of molecular forces. *Trans. Faraday Soc.*, 33(0):8b–26, January 1937.

[28] Siewert J. Marrink and Alan E. Mark. Molecular Dynamics Simulation of the Formation, Structure, and Dynamics of Small Phospholipid Vesicles. *J. Am. Chem. Soc.*, 125(49):15233–15242, December 2003.

[29] Karsten Meier, Arno Laesecke, and Stephan Kabelac. Transport coefficients of the Lennard-Jones model fluid. I. Viscosity. *The Journal of Chemical Physics*, 121(8):3671–3687, August 2004.

[30] Philip Nelson. *Biological Physics (Updated Edition)*. W. H. Freeman, 2008.

[31] Harald C. Ott, Thomas S. Matthiesen, Saik-Kia Goh, Lauren D. Black, Stefan M. Kren, Theoden I. Netoff, and Doris A. Taylor. Perfusion-decellularized matrix: using nature's platform to engineer a bioartificial heart. *Nat Med*, 14(2):213–221, February 2008.

[32] Carolyn L. Phillips, Joshua A. Anderson, and Sharon C. Glotzer. Pseudo-random number generation for Brownian Dynamics and Dissipative Particle Dynamics simulations on GPU devices. *Journal of Computational Physics*, 230(19):7191–7201, August 2011.

[33] Karin A. Riske and Rumiana Dimova. Electro-Deformation and Poration of Giant Vesicles Viewed with High Temporal Resolution. *Biophysical Journal*, 88(2):1143–1155, February 2005.

[34] Edith Schäfer, Torben-Tobias Kliesch, and Andreas Janshoff. Mechanical Properties of Giant Liposomes Compressed between Two Parallel Plates: Impact of Artificial Actin Shells. *Langmuir*, July 2013.

[35] Yaron R. Silberberg, Andrew E. Pelling, Gleb E. Yakubov, William R. Crum, David J. Hawkes, and Mike A. Horton. Mitochondrial displacements in response to nanomechanical forces. *J. Mol. Recognit.*, 21(1):30–36, January 2008.

[36] D. P. E. Smith. Limits of force microscopy. *Review of Scientific Instruments*, 66(5):3191–3195, May 1995.

[37] Mitsuki Yoneda. Tension at the Surface of Sea-Urchin Egg: A Critical Examination of Cole's Experiment. *J Exp Biol*, 41(4):893–906, December 1964.

[38] Ou-Yang Zhong-can and Wolfgang Helfrich. Bending energy of vesicle membranes: General expressions for the first, second, and third variation of the shape energy and applications to spheres and cylinders. *Phys. Rev. A*, 39(10):5280–5288, May 1989.

Open source software which made this research possible

HOOMD-blue simulation package: http://codeblue.umich.edu/hoomd-blue/

ESPResSo simulation package: http://www.icp.uni-stuttgart.de/~icp/ESPResSo

NumPy data analysis package: http://www.numpy.org

MDAnalysis data analysis package: http://code.google.com/p/mdanalysis/

matplotlib plotting package: http://matplotlib.sourceforge.net/

VMD molecular dynamics visualization software: http://www.ks.uiuc.edu/Research/vmd/

Appendix

Deriving Evans-Rawicz' α from Helfrich's α

To obtain the formula (Equation 2.5)

$$\underbrace{\alpha_{E.R.} \cong \overbrace{\frac{k_B T}{8\pi\kappa} ln\left(1 + \frac{\gamma A_{proj}}{C_o}\right)}^{entropic\ term} + \frac{\gamma}{K_A}}_{equation\ (2.5)}$$

for the area expansion α used by Evans and Rawicz[11], we begin with two expressions for the entropic (logarithmic) term derived by Helfrich and Servuss[1][23] for planar membranes. The first approximates the area expansion at zero tension

$$\left(\frac{\Delta A}{A}\right)_{\gamma=0} \cong -\frac{k_B T}{8\pi\kappa} ln\left(\frac{A}{a}\right), \tag{7.1}$$

and the second gives the entropic term for nonzero tension

$$\left(\frac{\Delta A}{A}\right)_{\gamma>0} = -\frac{k_B T}{8\pi\kappa} ln\left(\frac{\frac{\pi^2}{a} + \frac{\gamma}{\kappa}}{\frac{\pi^2}{A} + \frac{\gamma}{\kappa}}\right), \tag{7.2}$$

where γ is the tension, κ the bending modulus, A the 'projected' or 'effective' area of the membrane and a is the area per surfactant molecule.

Note that the π^2 factors appearing in Equation 7.2 are valid *only* for *planar membranes* (with periodic boundary conditions). For spherical vesicles that factor changes to 24π. **To avoid committing to a particular vesicle shape, we write ζ instead of π^2 from here on out.** Equation 2.5 is an approximation valid in the regime

$$\frac{\zeta\kappa}{A} \ll \gamma \ll \frac{\zeta\kappa}{a}.^{2} \tag{7.3}$$

[1] see equations #11 and #12 in [23]

First we rearrange Equation 7.2:

$$\underbrace{-\left(\frac{k_B T}{8\pi\kappa}\right)}_{\text{“}M\text{”}} \ln\left(\frac{\frac{\zeta}{a}+\frac{\gamma}{\kappa}}{\frac{\zeta}{A}+\frac{\gamma}{\kappa}}\right) = -M \ln\left(\frac{\cancel{\frac{\zeta}{a}}\frac{A}{\cancel{\zeta}}+\frac{\gamma}{\kappa}\frac{A}{\zeta}}{1+\frac{\gamma}{\kappa}\frac{A}{\zeta}}\right)$$

$$= -M \ln\left(\frac{\frac{A}{a}+\frac{\gamma A}{\zeta\kappa}}{1+\frac{\gamma A}{\zeta\kappa}}\right) = -M\left\{\ln\left(\frac{A}{a}+\frac{\gamma A}{\zeta\kappa}\right) - \ln\left(1+\frac{\gamma A}{\zeta\kappa}\right)\right\}$$

$$\rightarrow \boxed{\left(\frac{\Delta A}{A}\right)_{\gamma>0} = M\left\{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right) - \ln\left(\frac{A}{a}\left[1+\frac{\gamma a}{\zeta\kappa}\right]\right)\right\}}. \tag{7.4}$$

In their paper Evans and Rawicz state that they chose the tensionless state as their ‘reference state’, so we subtract Equation 7.1 from Equation 7.4

$$\underbrace{\alpha_{\mathcal{E}.\mathcal{R}.}}_{\text{entropic term}} = \left(\frac{\Delta A}{A}\right)_{\gamma>0} - \left(\frac{\Delta A}{A}\right)_{\gamma=0} \tag{7.5}$$

$$\cong M\left\{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right) - \ln\left(1+\frac{\gamma a}{\zeta\kappa}\right) - \ln\frac{A}{a}\right\} - \left(-M\ln\frac{A}{a}\right),$$

that is

$$\underbrace{\alpha_{\mathcal{E}.\mathcal{R}.}}_{\text{entropic term}} \cong M\left\{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right) - \underbrace{\ln\left(1+\frac{\gamma a}{\zeta\kappa}\right)}_{\text{“}G\text{”}}\right\}. \tag{7.6}$$

Next we rid (7.6) of G —using two approximations based on the conditions (7.3).

Since

$$\frac{\gamma a}{\zeta\kappa} \ll 1 \rightarrow \underbrace{\ln\left(1+\frac{\gamma a}{\zeta\kappa}\right) \cong \frac{\gamma a}{\zeta\kappa}}_{1^{st}\ approximation},$$

[2] Note that $C_\circ \equiv \zeta\kappa$ so that, with respect to Equation 2.5, the conditions (7.3) are equivalently written $\frac{C_\circ}{A} \ll \gamma \ll \frac{C_\circ}{a}$.

we have

$$\underbrace{\alpha_{\mathcal{E.R.}}}_{\text{entropic term}} \cong M\left\{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right)-\frac{\gamma a}{\zeta\kappa}\right\}.$$

And since $\frac{\gamma A}{\zeta\kappa} \gg 1 \gg \frac{\gamma a}{\zeta\kappa}$, we conclude

$$\underbrace{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right)-\frac{\gamma a}{\zeta\kappa} \cong \overbrace{\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right)}^{\text{The form we're after.}}}_{2^{nd}\ \text{approximation}},$$

and thereby obtain Equation 2.5;

$$\therefore \boxed{\underbrace{\alpha_{\mathcal{E.R.}}}_{\text{entropic term}} \cong \frac{k_B T}{8\pi\kappa}\ln\left(1+\frac{\gamma A}{\zeta\kappa}\right)}. \tag{7.7}$$

We write $\zeta\kappa$ in place of $\boldsymbol{C}_\circ$ (c.f. Equation 2.5) both for physical clarity and for ease of comparison with [23].

Choosing ζ

For quasispherical vesicles $\zeta = 24\pi$, so that Equation 2.5 is valid in the regime

$$\frac{24\pi\kappa}{A} \ll \gamma \ll \frac{24\pi\kappa}{a}. \tag{7.8}$$

In our simulations, this range corresponds to

$$\boxed{0.5\frac{\boldsymbol{\varepsilon}}{\boldsymbol{\sigma}^2} \ll \gamma \ll 150\frac{\boldsymbol{\varepsilon}}{\boldsymbol{\sigma}^2}}. \tag{7.9}$$

Note that in our case the vesicle is quasispherical only at low γ, becoming highly oblate for large γ. To fully account for this, ζ may itself need to be a function of the tension. Nonetheless, it is the direct area expansion term ($\frac{\gamma}{K_A}$) which dominates (2.5) in the high tension regime, so choosing $\zeta \cong 24\pi$ may be sufficient to capture the qualitative aspects of $\alpha(\gamma)$.

Volume Enclosed by Mean Surface

The first method uses the mean surface (given in Figures 4.23 and 4.18) associated with the lipid heads belonging to the inner leaflet of the vesicle. We integrate the volume enclosed by the mean surface

$$V_{proj} = 2\int_0^{h/2} dz \int_0^{r(z)} dr \int_0^{2\pi} r\, d\phi, \tag{7.10}$$

$$= 2\int_0^{h/2} dz \int_0^{r(z)} 2\pi r\, dr = 2\int_0^{h/2} dz\, \pi r^2(z),$$

$$= 2\pi \int_0^{h/2} dz \left\{ \left[\left((R_\circ - R_2) + R_2 \cos(\varphi) \right)^2 \right] \right\},$$

$$= 2\pi \left\{ \left[\int_0^{h/2} dz\, (R_\circ - R_2)^2 + \int_0^{h/2} dz\, R_2^2 \cos^2\varphi + \int_0^{h/2} dz\, 2(R_\circ - R_2) R_2 \cos\varphi \right] \right\}.$$

By inspection of Figure 4.23, and taking the z-axis as vertical[3],

$$\varphi = arcsin\left(\frac{z}{R_2}\right) \Longrightarrow \cos\varphi = \pm\sqrt{1 - \left(\frac{z}{R_2}\right)^2}. \tag{7.11}$$

From the same figure we can conclude that $-\frac{\pi}{2} \leqslant \varphi \leqslant \frac{\pi}{2}$, thus

$$\boxed{\cos\varphi = \sqrt{1 - \frac{z^2}{R_2^2}}}, \tag{7.12}$$

[3] i.e. h lies along the z-axis in Figure 4.23.

and likewise

$$z = R_2 sin\varphi \longrightarrow dz = R_2 cos\varphi \, d\varphi. \tag{7.13}$$

So that

$$V_{proj} = 2\pi \left\{ \left[\frac{h}{2}(R_\circ - R_2)^2 + R_2^2 \int_0^{h/2} dz \left\{ 1 - \left(\frac{z}{R_2} \right)^2 \right\} + 2(R_\circ - R_2) R_2^2 \int_{z=0}^{z=h/2} d\varphi \, cos^2\varphi \right] \right\},$$

$$= 2\pi \left\{ \left[\frac{h}{2} \left[(R_\circ - R_2)^2 + R_2^2 \right] - \int_0^{h/2} dz \, z^2 + (R_\circ - R_2) R_2^2 \int_{z=0}^{z=h/2} d\varphi \, \frac{1}{2} \{ 1 + cos(2\varphi) \} \right] \right\},$$

$$= 2\pi \left\{ \left[\frac{h}{2} \left[(R_\circ - R_2)^2 + R_2^2 \right] - \frac{1}{3} \left(\frac{h}{2} \right)^3 + (R_\circ - R_2) R_2^2 \left(\varphi + \frac{1}{2} sin(2\varphi) \right) \Big|_{z=0}^{z=h/2} \right] \right\}.$$

And since

$$\frac{1}{2} sin(2\varphi) = sin\varphi \, cos\varphi = \frac{h}{2R_2} \sqrt{1 - \frac{h^2}{4R_2^2}},$$

we have

$$V_{proj} = \pi h \left[(R_\circ - R_2)^2 + R_2^2 \right] - \pi \frac{h^3}{12} + 2\pi (R_\circ - R_2) R_2^2 \left(arcsin\left(\frac{h}{2R_2} \right) + \frac{h}{2R_2} \sqrt{1 - \frac{h^2}{4R_2^2}} \right),$$

$$\therefore \boxed{V_{proj} = 2\pi (R_\circ - R_2) R_2^2 arcsin\left(\frac{h}{2R_2} \right) + \pi h \left((R_\circ - R_2)^2 + R_2^2 - \frac{h^2}{12} + (R_\circ - R_2) \sqrt{R_2^2 - \frac{h^2}{4}} \right)}. \tag{7.14}$$

Finally, we plug the values of $R_\circ$, R_2, h returned by the curve fitting algorithm into Equation 7.14 in order to obtain the V_{proj} values shown in Figure 4.26.