McGill
University

Physics Department

Flexible polyelectrolytes: like-charged attraction, linear stability, & long-term structure

Author:
Adriano Ferrari

A thesis submitted to McGill University
in partial fulfillment of the requirements
of the degree of Doctor of Philosophy.

October 14, 2014

Acknowledgments

I thank my parents, Gustavo and Silvia, for all their loving support from half way across the world (and my brother Mauro, for his support from halfway across town).

Jess, thank you for helping me and encouraging me to finish, even though it meant less time for you, and for our growing companies. It will have been worth it, if it can mean more time for our growing *family*. I thank you, Zeno, for giving me the motivation for my final push. I am sorry to both of you for the time I lost to this. I will make it up to you.

I would like to thank Jorge Viñals for getting me started on this project. Finally, I would like to thank my supervisor, Martin Grant, for his guidance, and his unwavering patience, and for encouraging me when the going got tough with this thesis.

Abstract

Charged polymers are critical components of biological systems. DNA, F-actin, and bacteriophages are some examples of highly charged polymers. Because of their size, complexity, and of the long-range forces involved, full computer simulations can only be used for short times. Even standard theoretical approaches, such as mean-field theory, predict qualitatively wrong results.

We seek to determine how parallel pairs of flexible like-charged polymers behave. We use minimal models to represent the polymers, and determine three particular behaviors: 1) How and when the polymers exhibit like-charged attraction. 2) Whether deformations away from a straight line are stable or unstable. 3) If allowed to evolve freely, what long-term structures would form.

We represent our pair of infinite parallel charged polymers as a string of point charges linked with a spring potential. Counterions are represented by free point charges. With this system, we perform Monte Carlo simulations to find the following results:

1. Given two straight chains at fixed separation R, we reproduce the attractive interaction seen in other systems (real and simulated).

2. We find that the attraction decreases with temperature, lending support to the Wigner crystal model of this interaction.

3. We find a typical ground-state configuration, and calculate the interaction di-

rectly.

4. We find an expression that fits the force curve for a wide range of distances.

5. Using the force curve, we develop a simple approach to estimate the stability of the chains against small deformations.

6. Given wavelength deformations, we determine via simulations which modes decay and which grow, using three different methods. (linear stability analysis).

7. Initializing the system in the most unstable modes, we determine the structures formed in four cases: rigid, flexible-planar, flexible-free, and helical.

8. We find that flexible polyelectrolytes tend strongly towards a $1/f$ noise profile, when allowed to move freely at finite temperature.

In summary, we developed and studied a minimal model of flexible polyelectrolytes, and investigated many of the core features of this system.

Abrégé

Polymères chargés sont des éléments essentiels des systèmes biologiques. ADN, F-actine, et bactériophages sont des exemples de polymères fortement chargés . En raison de leur taille , de la complexité et des forces à longue portée impliqués , les simulations informatiques complets ne peuvent être utilisés pour de courtes durées. Même les approches théoriques standards (comme la théorie du champ moyen) prédisent des résultats qualitativement erronées.

Nous cherchons à determine comment les paires parallèles de polymères comme - chargées flexibles se comportent. Nous utilisons des modèles minimaux pour représenter les polymères, et de déterminer trois comportements particuliers: 1) comment et quand les polymères attirent (attraction comme chargée). 2) Si les déformations au départ de la ligne droite sont stables ou instables. 3) Si a permis d'évoluer librement, ce que les structures à long terme feraient.

Nous représentons notre paire de polymères chargés parallèles infini comme une chaîne de charges ponctuelles liées à un potentiel du ressort. Les contre-ions sont représentés par des charges ponctuelles libres. Avec ce système, nous effectuons des simulations de Monte Carlo pour trouver les résultats suivants:

1. donné deux chaînes droites à la séparation fixe R , nous reproduisons l' interaction attractive vu dans d'autres systèmes (réelles et simulées).

2. Nous constatons que l' attraction diminue avec la température , en appuyant sur le modèle de cristal de Wigner de cette interaction.

3. On retrouve une configuration typique de l'état fondamental , et nous calculons l'interaction directe.

4. Nous trouvons une expression qui correspond à la courbe de force pour une large gamme de distances.

5. En utilisant la courbe de force, nous développons une approche simple pour estimer la stabilité des chaînes contre de petites déformations.

6. Compte tenu des déformations de longueur d'onde , nous déterminons via des simulations qui modes décroissance et qui poussent , en utilisant trois méthodes différentes. (analyse de stabilité linéaire) .

7. L'initialisation du système dans les modes les plus instables , nous déterminons les structures formées dans quatre cas : rigide , flexible sur plane , flexible libre, et hélicodaux.

8. Nous constatons que les polyélectrolytes flexibles tendent fortement vers un profil de bruit à $1/f$, lorsqu'ils sont autorisés à se déplacer librement à température finie .

En résumé, nous avons élaboré et étudié un modèle minimal de polyélectrolytes flexibles, et étudié la plupart des fonctionnalités de base de ce système.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Problem

Bio-molecular or polymer systems are difficult to analyze theoretically. Even basic tenets of electrostatics such as "like-charges repel" can appear to be violated due to complicated inter-particle interactions and correlations. For instance, when a polyelectrolyte is placed in a solvent such as water, two chains of the same charge can attract each other[6].

This has a number of important implications, since many bio-polymers are polyelectrolytes: DNA, RNA, F-actin fibers, microtubules, and aggregating viruses, to name a few[7]. A great number of other polyelectrolytes (PEs) have various technological applications, and are used as thickeners, emulsifiers, or in water treatment. Medical applications of PEs include implant coatings and controlled drug release [1].

This attractive force, which has been observed in experiment[8] and simulations[9], and which has been attributed to numerous (sometimes competing) mechanisms[9][10][11], has been studied mostly for rigid rods[1], rigid finite-sized segments[12], rigid bundles[13], or completely free chains[14]. However, each of these studies was restricted in its scope because of the complexity of the system.

There has not yet been a study of a single pair of infinite, flexible polyelectrolytes, and the behaviors they exhibit. This is what we address in this work.

1.2 Minimal models

Much is said about the enormous and growing computational power available to researchers today. Even so, we can only perform full QM simulations of a few hundred atoms, for a few picoseconds[15]. The biological and chemical systems we are interested in are far outside the realm of what can be solved numerically for all situations of interest[16]. Barring a quantum leap in computer power, we need to turn to higher levels of abstraction to represent the molecules and polymers we wish to study. In most cases, researchers today use classical molecular dynamics (MD) simulations with phenomenological force fields. But in order to produce accurate models for simulation, we need to understand and account for all forces involved, even those that do not arise from mean-field theories and simple force fields.

We chose to perform simulations using as few elements as possible. To this end, we did not consider solvent effects (except by modifying the dielectric constant), we used periodic boundary conditions and only two parallel chains. All our particles are point charges, with repulsive potentials. Though we are studying this system with the intention to compare it to DNA or F-actin or other polyelectrolytes, our model polymer consists of nothing other than the charged particles of those polymers (i.e. no base-pairs, backbone sugars, etc.).

Though we could model a given polyelectrolyte system in far more detail than presented here, and even run it for long enough to obtain interesting results, we choose to restrict our attention to this model. By focusing on these core features, we are able to see more clearly what causes attraction, how the chains deform, and what structures they form.

1.3 Thesis Overview

This thesis is divided into two parts, "Background & Methods", and "Results & Discussion". The Methods section in Part I includes a description of our model system (section 3.1), and lay out the tools we will use to simulate polyelectrolytes in solution.

Part II is divided into three core chapters, covering the three stages of our analysis as described above. Namely, whether two parallel chains attract each other and how strongly they do so (Ch. 4), how this attraction initially affects the shape of the chains as they approach each other (Ch. 5), and what is the long-term (possibly non-linear) behavior of the chain-chain pairing (Ch. 6), including their strong tendency to form helical structures.

Part I

Background & Methods

Chapter 2

Theory & Review

2.1 Overview

Many systems, and most of those in biology, involve large molecules immersed in water. This is significant because water is polar, and weakens ionic bonds[17]. When the ionic bonds in a polymer are broken, the ions float free, leaving a charged surface behind. This is the mechanism by which DNA obtains its high charge[18]. This charge is necessary for function (for instance, in protein binding[19]). However, electrostatics is hard to calculate explicitly for any large system.

DNA is one important example of this. The phosphate groups on the DNA backbone dissociate in water, leaving behind a highly charged surface. The linear charge of DNA is $6e/nm$[20] (where e is the elementary charge). This charge is important for many functions (e.g. DNA binding proteins). DNA is far from the only polyelectrolyte system that has been studied extensively. For instance, F-actin and filamentous bacteriophages are both polyelectrolytes, with linear charges on the same order as that of DNA[21][22].

We will review the mean-field theory, Poisson-Boltzmann[23], that describes this situation. We will also present the Debye-Hükel linearization of this theory[24], and

how it leads to counterion condensation. We will also briefly outline two competing theories for how this condensation can lead to attraction between these highly charged surfaces[25][26][10].

2.2 Literature review

Until the mid-seventies, researchers believed that like-charged polymers and colloidal particles could only repel. This was supported by all the theories of the time, such as Poisson-Boltzmann [23], DLVO theory [27], and Manning condensation [28]. However, several experiments with polymers or colloids in electrolyte solutions showed attraction instead [29].

Manning condensation theory is a mean-field-theory that describes whether counterions condense onto charged rods or chains, and under what conditions (charge, temperature, and concentration) this takes place. Though it was more successful than other approaches at predicting the threshold for condensation in many experiments involving DNA[30], it still did not explain the like-charged attraction itself[31].

In 1980, Patey published a paper titled "The interaction of two spherical colloidal particles in electrolyte solution. An application of the hypernetted-chain approximation" [32], which had the results shown in Fig. 2.2.

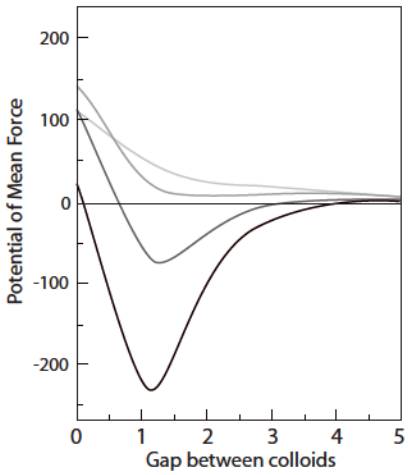

"It is immediately apparent that the [potential of mean force] is *not always repulsive* as we might have expected from previous work."
- G.N. Patey ([32], original emphasis).

Figure 2.1: The first clear sign that like-charged colloids might not repel.

Not only did his numerical results exhibit attraction, they also showed that the attraction between like-charged colloids in an electrolyte solution was due to electrostatics, and was not due to van-der-Waals or hydration forces. He went on to speculate how this force could cause some of the unexpected features found in previous experimental work. However, Patey had no explanation for this electrostatic attraction.

Over the next several years, Bloomfield [33] and R. Podgornik et al. [34] conducted experiments to attempt to determine the forces underlying DNA condensation. Meanwhile, J.X. Tang et al. [35] were addressing similar questions, but with *F*-actin and microtubules. All of these experimental papers led to the same conclusion: adding polyvalent salts leads to attraction between like-charged polymers[20].

A great number of experiments have been performed on polyelectrolyte systems, especially those present in biology. In the field of DNA experiments, the attractive interaction we are studying is termed DNA condensation, and has been examined in detail[36][37][38]. The core results of these studies is that attractive interactions by multivalent counterions is what makes it possible to compress the strongly charged molecule into a nucleus or virus[39][8].

Still there was no agreement on the physics behind this attractive interaction.

Various simple theoretical arguments (such as [40] or [41]) did show attraction, but they disagreed on a number of important features; they even disagreed on the underlying source of the attraction, and on whether the attractive force increased or decreased with temperature.

In 1997, Niels Grønbech-Jensen and his colleagues, perhaps spurred by the lack of direct investigations into the matter, simulated a truly simple model of a polymer pair [1], in the hopes of recreating on a computer the like-charged attraction seen in experiment and theory. They used two rigid line-charges to represent a polymer pair, several point charges to represent the counterions in the electrolyte solution, and Brownian dynamics to incorporate water and a heat bath at a specified temperature.

Despite the simplicity of this approach, the results reproduced the attractive force, and hinted at its source. In this thesis, we will extend the results of [1], by replacing their rigid line-charges with flexible chains of discrete particles[1], simulating much larger systems for longer times, and look at possible shape instabilities of the charged chains.

2.3 Polyelectrolytes

An electrolyte is simply a salt. In other words, when placed in solution, it will dissociate into positive and negative ions. A polyelectrolyte (PE) is a polymer that contains salts, so that when placed in solution, one charge-species of ions separates from the chain (becoming counterions), leaving a charged polymer behind.

[1]Manning published two versions of his counterion condensation theory. The first version [42] modeled the polymers as line-charges, while the later version [43] used discrete charges.

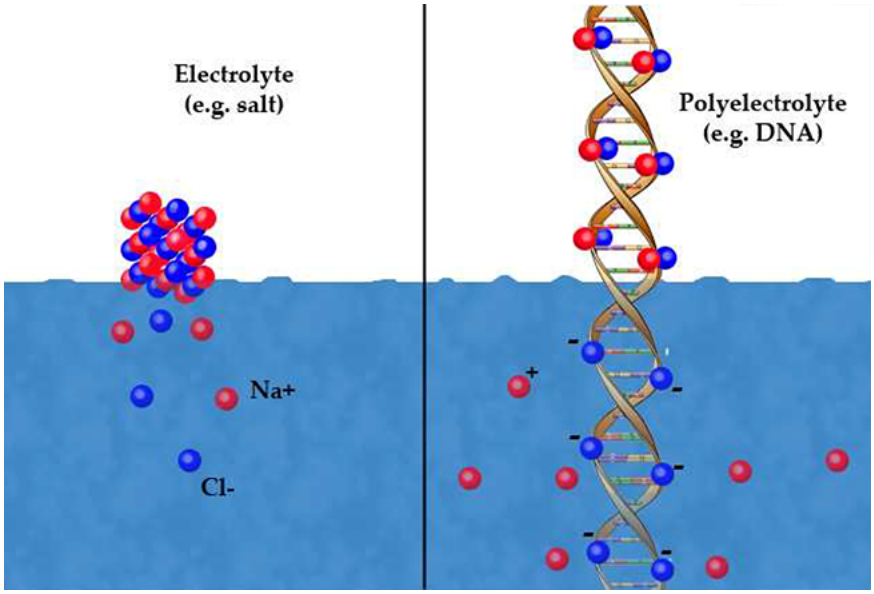

Figure 2.2: An electrolyte is a salt, dissociating into ions when in solution. In the case of DNA, the DNA chain becomes negatively charged, and therefore conductive, when dissociation occurs.

In water, it is the polarity of the H_2O molecules that breaks apart the ionic bonds in a given salt. For other solvents, a rough measure of a solvent's polarity is given by that solvent's dielectric constant, where dielectric constants larger than 15 are considered polar, in that they strongly shield the ionic bonds present in an electrolyte. As we'll see later, this dielectric shielding is enough to allow salts to dissociate at the appropriate temperatures, even when no explicit polar solvent molecules are used in our simulations.

Polyelectrolytes have a number of interesting properties. For instance, because the entire chain is of the same charge, different parts of the chain repel each other and the chain stretches out. This causes a solution of polyelectrolytes to become extremely viscous, since the extended chains are more likely to entangle [2].

In some situations, the counterions released from the chain (or other counterions introduced into the solution) remain loosely bound to the chain. This phenomenon is called counterion condensation. Because the counterions are of the opposite charge, the overall chain charge is diminished or neutralized, and chains no longer repel as

strongly. Thus, if salt is added to a PE solution, the chains return to their bunched or coiled arrangements. This process is at work with DNA condensation, the process by which DNA is wound tightly into a cell nucleus, virus, or even man-made gene delivery vehicles [3][4]. In the following chapters, we will review the mean-field description of the counterion-PE system, as well as the theory describing the condensation of the counterions onto the PE chains.

2.4 Poisson-Boltzmann Theory

As mentioned, there are many situations where a charged surface is immersed in a solution with ions present. We would like to know how these coions and counterions distribute themselves around the charged surface.

There are many possible variations to this situation. The solution can contain added salts, and therefore there will be more cations and anions in solution than those disassociated from the macromolecule[20]. These salts can be symmetric, or asymmetric, and the charged macromolecule can have regions of both positive and negative charge[44]. Fortunately, the simplest situation (a macromolecule of one charge, surrounded by neutralizing counterions and no added salts) is the most prevalent, and is applicable to DNA and other biopolymers in solution[20] [36] [13].

The mean-field approach is our starting point, and correctly predicts many of the important features of this system. Since our work consists of finding and explaining departures from the predictions of this theory, we will begin by reviewing the Poisson-Boltzmann equation.

We begin with the Poisson equation, which relates the electric potential $\psi(\vec{x})$ to the charge distribution $\rho(\vec{x})$. For a collection of ions, the charge density at any point is the sum of the individual ion number densities: $\rho(\vec{x}) = ez_+n_+(\vec{x}) + ez_-n_-(\vec{x})$. Here $n(\vec{x})$ is the number density of ions, and z is the ion valency.

The Poisson equation so far is:

$$\nabla^2\psi(\vec{x}) = -\frac{4\pi}{\epsilon_0\epsilon_r}\rho(\vec{x}) = -\frac{4\pi}{\epsilon_0\epsilon_W}\left(z_+n_+ + z_-n_-\right) , \tag{2.1}$$

where ϵ_W is the relative dielectric constant of water ($\epsilon_r \simeq 80$). Now, since the ions are free to move under the influence of the electric potential, then $n(\vec{x})$ is not fixed. In other words, the number density adjusts to the presence of the electric potential, and the equation remains self-consistent.

To find the distribution of ions, we incorporate the fact that, at thermal equilibrium, the ion density also obeys a Boltzmann distribution:

$$n_i = n_0^{(i)}\mathrm{e}^{ez_i\beta\psi} \tag{2.2}$$

Here n_0 is the reference density for which ψ is 0. This Boltzmann distribution is the mean-field result for the distribution of multiple ion species around a charged surface. For a single species of ion, combining the two equations, we obtain the Poisson-Boltzmann equation:

$$\nabla^2\psi(\vec{x}) = -\frac{4\pi}{\epsilon_0\epsilon_r}n_0\exp\left[-\ell_B\psi(\vec{x})\right] , \tag{2.3}$$

where we introduced the Bjerrum length $\ell_B = e^2/4\pi\epsilon_0\epsilon_r k_BT$. Solving this equation for various geometries is what yields the mean-field counterion distribution for that system.

In the case of a charged rod or cylinder, symmetry reduces the PB equation to a single variable, the radial distance r from the axis. In this case, the PB equation can be written as:

$$\frac{1}{r}\frac{\mathrm{d}}{\mathrm{d}r}\left(r\frac{\mathrm{d}\psi(r)}{\mathrm{d}r}\right) = \kappa^2\psi(r) , \tag{2.4}$$

where we introduce the Debye length $\kappa^{-1} = \sqrt{\epsilon_0\epsilon_r k_BT/Znq^2}$ (Z is the counterion

valency). We will use the PB equation in the next section, to derive the conditions for condensation.

2.5 Counterion condensation

When polyelectrolytes are placed in solution with ions of opposite charge, these counterions are attracted to the polyelectrolyte. If the attraction is stronger than the thermal energy, the counterions will condense onto the charged polymers. The conditions for this condensation were first derived by G. Manning and F.Oosawa, and is termed Manning-Oosawa condensation [42] [43].

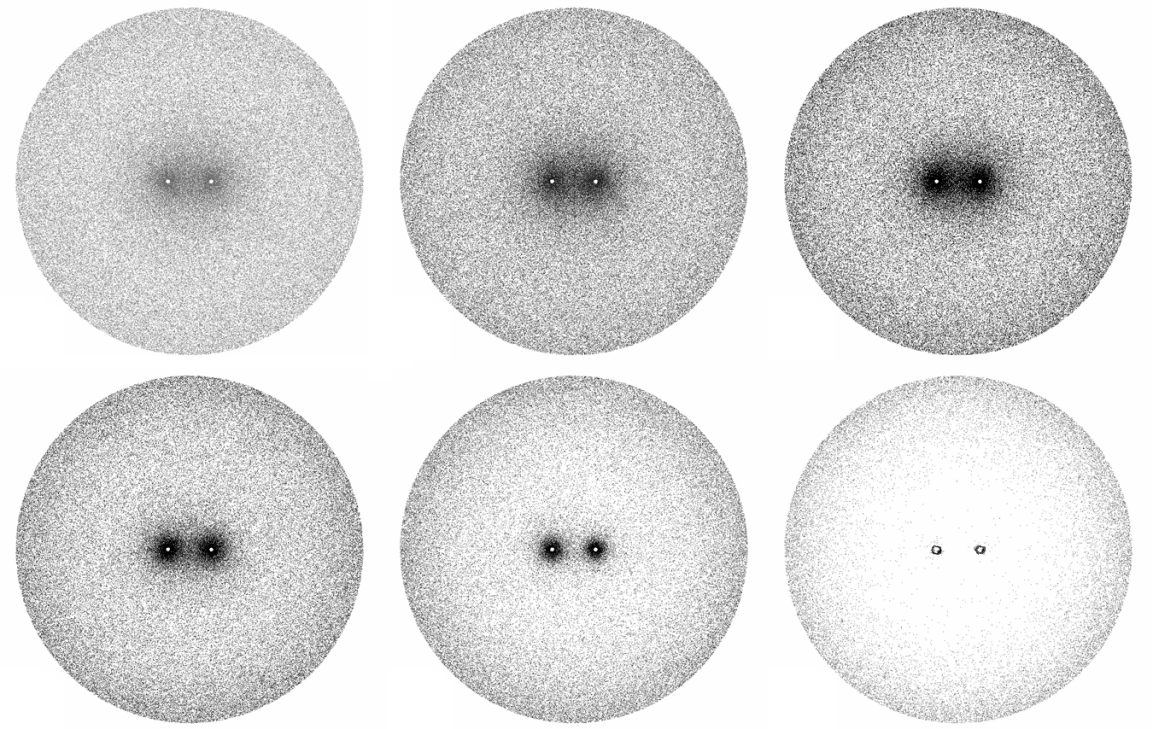

Figure 2.3: Simulations snapshots showing condensation in our system. Looking along the axes of two parallel polymers, here we see the counterion distribution as we lower the temperature. Lowering temperature, or increasing the counterion valency, leads to the counterion condensation phenomenon shown here. In the last figure, almost all counterions have condensed onto the polymers.

It is this condensation of counterions onto the polymer chain that produces all the

rich and complex behavior in our system, so we will begin by summarizing Manning-Oosawa condensation theory in the case of neutralizing counterions only. That is, there is no added salt, and all counterions are dissociated from the chain itself [5].

We treat the chain as a cylinder of radius a and linear charge density $\rho = e/b$, where e is the electron charge and b is the monomer spacing. We assume the cylinder is long enough so that we can ignore end effects. We will sketch the derivation of counterion condensation, based on [2]. For cylindrical symmetry, we saw in the previous section that we end up with the following equation:

$$\frac{1}{r}\frac{\mathrm{d}}{\mathrm{d}r}\left(r\frac{\mathrm{d}\psi(r)}{\mathrm{d}r}\right) = \kappa^2\psi(r)\,. \tag{2.5}$$

Here κ is the Debye-Hükel screening parameter introduced earlier, which has dimensions of inverse length. If we substitute $y = e\psi(r)/k_BT$ and $x = \kappa r$, we have:

$$\frac{\mathrm{d}^2y}{\mathrm{d}x^2} + \frac{1}{x}\frac{\mathrm{d}y}{\mathrm{d}x} - y = 0\,. \tag{2.6}$$

This is one of the Bessel equations, and the solution is given by:

$$y = \frac{2\xi K_0(\kappa r)}{\kappa a K_1(\kappa a)} \approx 2\xi K_0(x)\,, \tag{2.7}$$

where K_i is the modified, or hyperbolic, Bessel function of the second kind.

Here we introduced $\xi = (e^2/4\pi\epsilon_r\epsilon_0 b)/k_BT$, the Manning parameter, which is the ratio of the electrostatic energy of two charges at a distance b, to the thermal energy of the system.

For monovalent counterions, the charge density is:

$$\begin{aligned} q(r) &= -n_0 e\Big[\exp\left(+e\beta\psi(r)\right) - \exp\left(-e\beta\psi(r)\right)\Big] \\ &= -n_0 e\big[\exp\left(y\right) - \exp\left(-y\right)\big] \\ &= -n_0 e\big[\exp\left(2\xi K_0(\kappa r)\right) - \exp\left(-2\xi K_0(\kappa r)\right)\big]\,. \end{aligned} \tag{2.8}$$

Integrating this expression radially, for different values of ξ produces different behaviors. If $\xi < 1$, then the radial integral is well-behaved as $\kappa r \to 0$. If $\xi > 1$, then the integral diverges as $\kappa r \to 0$. In other words, there are an infinite number of counterions with a radius greater than the polyelectrolyte radius. To avoid this divergence, counterions condense onto the polyelectrolyte until the effective linear charge density reduces ξ to 1.

For our purposes, the key thing to note here is that $\xi = 1$ is the threshold for counterion condensation. If $\xi < 1$, then counterions condense onto the chains, and all manner of interesting behavior occurs[45][46][47]. Otherwise, we simply have two line-charges interacting with a slightly screened Coulomb potential.

2.6 Counterion mediated interactions

Once condensed onto the polyelectrolyte, counterions strongly influence the PE-PE interactions. There are two main theoretical approaches to studying how counterions affect these interactions. The first by Oosawa presents fluctuations of counterion density along the chain as the source of attractive interactions.

The other picture imagines that each counterion remains near a charge on the chain surface, and that counterions on the opposite chain will arrange themselves in a staggered or zipper configuration. This is the Wigner crystal model, and is in some sense a limit of the Oosawa model for wavelength going to b, the monomer

spacing. We will examine each in turn, very briefly. We then compare their attractive interactions, and the key characteristics of each.

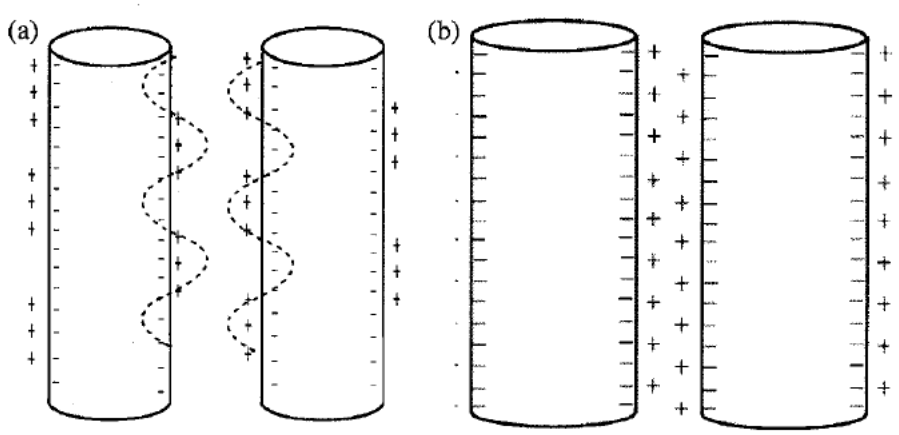

Figure 2.4: Oosawa (a) vs. Wigner crystal (b).[2]

2.6.1 Oosawa model

Oosawa tackled the problem of the distribution of counterions around a charged surface first by ignoring the discrete nature of the counterions, and working with charge density waves.[33] [35]

Oosawa's simple arguments lead to the following expression for the force f between two cylinders a distance R apart, in the limit of large R:

$$f(R) \approx k_b T \left(\frac{1}{Z^2 \ell_B R} - \frac{(Z\xi)^2}{1 + (Z\xi)^2} \frac{1}{R^2} \right) . \tag{2.9}$$

In this equation, the first term is due to the Poisson-Boltzmann repulsion between the rods. The second term, however, depends on the Manning parameter ξ, and describes the long-wavelength charge density fluctuations from Oosawa's model[1]. As we will see, though this expression yields attraction, it is not in line with our results, since this attraction is long-ranged, and increases with temperature.

2.6.2 Wigner crystal model

We have shown in section 2.5 that counterions will condense onto a charged polymer given certain conditions. In the limit $T = 0$, these condensed counterions will form a Wigner crystal. In this section, we will see how this crystal arrangement can lead to an attractive interaction, and that the persistence of the crystal structure correlations up to finite temperature is one model of like-charged attraction in polyelectrolytes.

We begin with the correlation energy of a single ion around a segment of the polyelectrolyte. The ion rests at a distance ρ from the PE center. If we approximate the PE segment as being a disk or radius ρ, the energy ε is given by,

$$\varepsilon = \frac{e^2}{\epsilon} qV(0) \tag{2.10}$$

$$\beta\varepsilon = \frac{e^2}{\epsilon k_B T} qV(0) \tag{2.11}$$

$$= \ell_B \, qV(0) \tag{2.12}$$

$$= \ell_B \, q(-2\pi\rho\sigma) \,, \tag{2.13}$$

where ℓ_B is the Bjerrum length, $V(0)$ the potential at the center, and σ the surface charge.

We now use the fact that the Wigner crystal is neutral to find that the surface charge of the PE in this segment must be equal to q. That is, $\sigma\pi\rho = q$. By introducing the fact that we only have one ion in this cell, we can say that the density of ions, n, is given by $n = 1/\pi\rho^2$. We substitute accordingly, and arrive at the correlation energy:

$$\beta\varepsilon = -2\ell_B q^2 \sqrt{\pi n} \,. \tag{2.14}$$

2.7 Like-charged attraction

We have shown that, in the Wigner crystal picture, the correlation energy of a single ion around a PE rod segment is $\beta\varepsilon = -2\ell_B q^2\sqrt{\pi n}$. If two such PE rods approach one another, the Wigner crystal lattices stagger due to the counterion-counterion repulsion. Therefore, a counterion on rod A will also be near an exposed region of rod B. This leads to a negative energy change, and therefore to attraction [48].

The Oosawa model, despite being a model of two continous line-charges, can also exhibit attraction. Using the Oosawa model of continuous charge density waves of wavelength λ_q, the pressure Π between two PEs is[49]

$$\beta\Pi = -\frac{\pi}{d^2}\left(\frac{\ell_B q^2\lambda_q}{1+\ell_B q^2\lambda_q}\right). \tag{2.15}$$

The Oosawa model demonstrates that attractive forces can exist at long-ranges as well. One important thing to note is that this pressure decreases with temperature. Combined with the discrete Wigner crystal model, we are ready to proceed to our own analysis of the like-charged attraction between two PEs, and to study how this interaction affects their structure and behavior when the rods are allowed to be flexible.

Chapter 3

General Methods

3.1 Overview

The system we will study consists of two chains of discrete charges, and several ions of opposite charge. These will all be enclosed in a periodic box; the number and charge of the counterions is chosen to make the unit cell neutral. The chains will both be parallel to the y-axis, and separated from each other by a distance R.

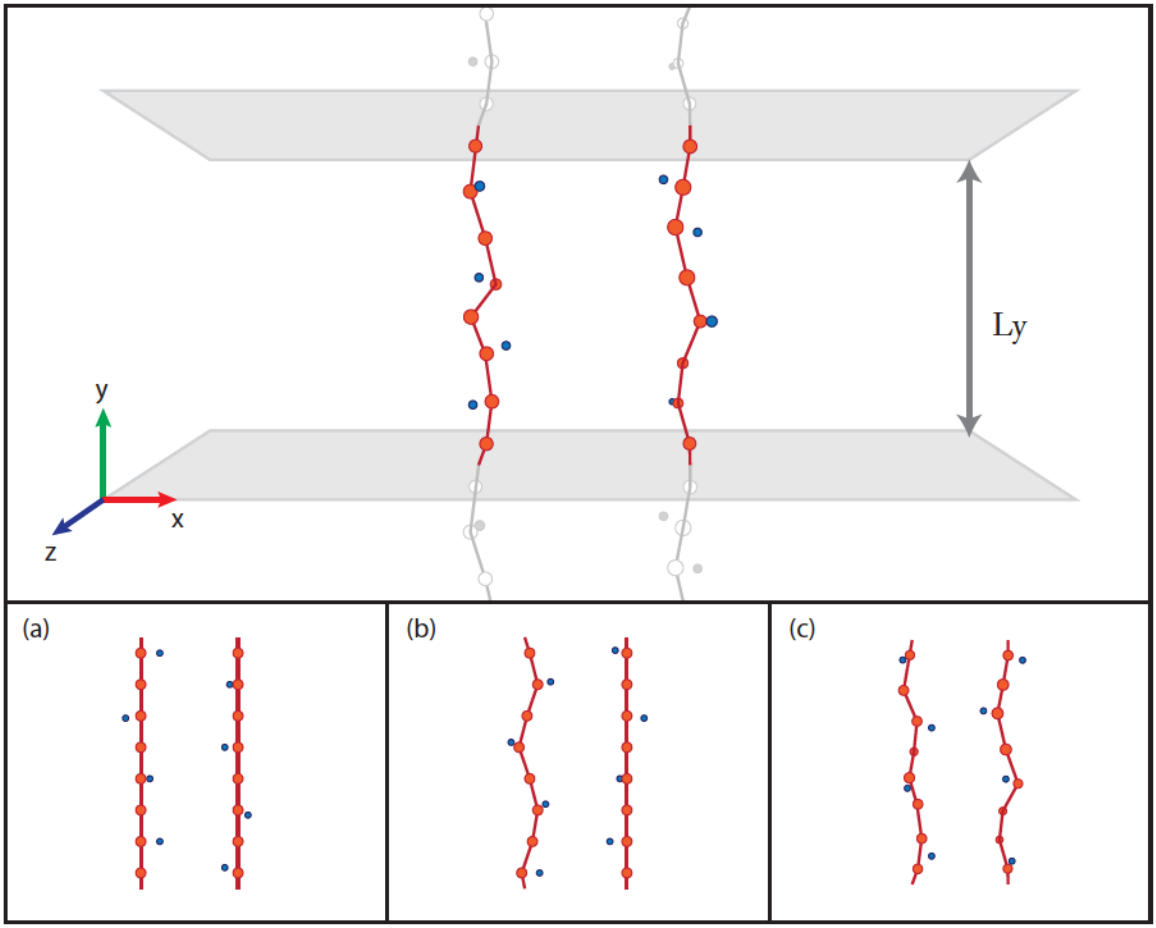

Figure 3.1: Our model system. N_ℓ chain particles are used to represent each chain. Periodic boundaries make these infinitely long. We place N counterions in the box, and allow them to move. The chains are either *a)* held fixed at a separation R , *b)* sinusoidally perturbed with wavelength $\lambda = L_y/n$, or *c)* allowed to move freely. These are discussed in chapters 4,5, and 6 respectively.

The charged monomers will be spaced a distance b apart from each other. The chains will be either straight or sinusoidally perturbed, depending on whether we wish to study the attractive force itself or its effect on perturbed chains. The counterions will start uniformly distributed in the box. They will then be subject to only two potentials: electrostatics, and repulsive Lennard-Jones. The first, electrostatics, is simply the Coulomb force due to the charged particles, and will cause the counterions to move towards the chain charges while keeping away from the like-charged counterions. The repulsive Lennard-Jones force gives all the particles a size, to prevent

them from overlapping each other.

Electrostatics alone would cause the counterions to bond directly to the chain particles, and not move from there. However, it's critical to remember the role of temperature in our simulations, through which both counterions and chain particles fluctuate.

Neighboring chain particles, or monomers, will be subject to a chain potential as well as the electrostatic and repulsive interactions. The monomers at the end of a chain will be attached, via this potential, to the periodic image of the beginning monomer.

3.2 Monte Carlo

3.2.1 Overview

The number of particles in our system ranges from 70 to 164 particles, each of which interacts with all others. Each counterion is free to move anywhere inside the simulation region, and the chains can be made free to translate and to fluctuate about the mean axis. The number of configurations in this system is extremely large. To explore such a large phase space, and to take full advantage of large-scale parallel computing, we decided to use Monte Carlo for our studies.

Every system can be characterized by a list of numbers, called the state of the system (list of particle positions, field values at lattice points, number of votes for each political candidate, etc). When there are far too many possible states, it becomes impossible to try them all to compare which state is most likely to actually occur. Moreover, most of these combinations are extremely unlikely (all the particles in one corner, for example).

Imagine that every state corresponds to a particular point on a 1 meter x 1 meter rubber grid. The system would be a ball rolling on this grid; wherever the ball is, the

system is in the state characterized by that point's list of numbers. The grid is not flat, but has peaks and valleys. Peaks correspond areas that are particularly unlikely (the ball/system will not stay there long). Valleys are places that are more likely, in that the ball/system will stay there longer.

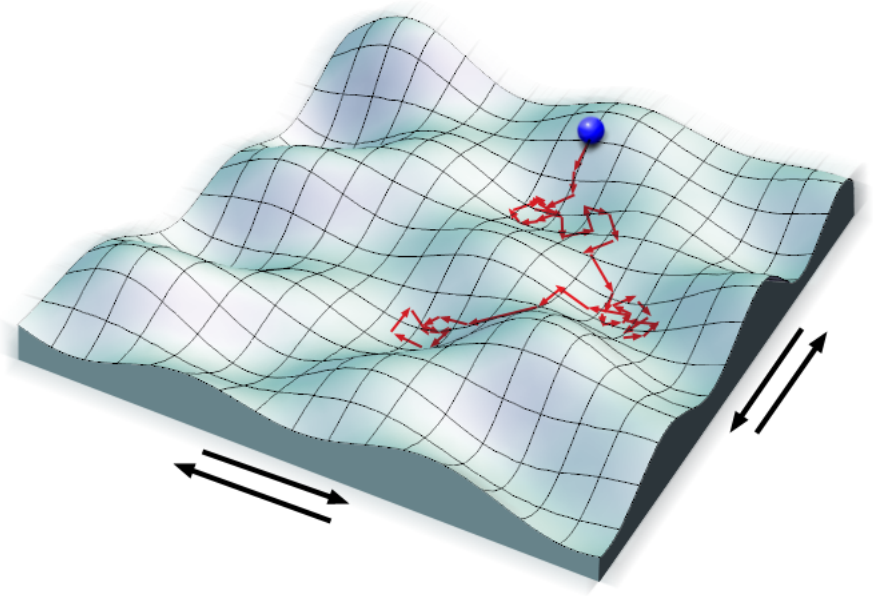

Figure 3.2: A physical analogy for Monte Carlo simulations. Each grid point is a list of numbers characterizing the state of a system. The ball is a marker for the current state of the entire system, and vibrating the sheet causes the marker to move from one state to a nearby state, staying more often close to valleys and away from peaks.

If we now shake the sheet and keep track of where the ball goes, we are performing a Monte Carlo simulation. The harder we shake, the higher the temperature we are using and the less likely that the ball will get stuck in a valley. If we shake too hard though, the shaking matters more than the peaks and valleys, and we are not getting much information about which areas are more likely. If we don't shake at all (zero temperature), the ball will simply roll towards the nearest minor dip and stay there even if there are far bigger valleys/more likely states elsewhere.

Since we often need to simulate a particular temperature, we need to make sure that this temperature will not move the ball/system too erratically, or conversely,

not move the system enough. At a fixed temperature, we can change the weight of the ball to achieve this balance. Keeping this rough analogy in mind, we will briefly review the theory of Monte Carlo simulations, and afterwards discuss the details of our implementation for this system.

3.2.2 Theory

In order to average a quantity $A(z)$ over a probability $p(z)$, we need to calculate

$$\langle A \rangle = \frac{\int A(z)\, p(z) \mathrm{d}z}{\int p(z) \mathrm{d}z}\,. \tag{3.1}$$

where the integral is over all possible values of z. If this cannot be done analytically, we can approximate the average by taking discrete samples:

$$\langle A \rangle \approx \frac{\sum_{i=0}^{N} A(z_i)\, p(z_i)}{\sum_{i=0}^{N} p(z_i)}\,. \tag{3.2}$$

In the limit of $N \to \infty$ this approximation converges on the correct average. However, for finite sampling, the rate of convergence depends very much on how the sample points are chosen and on the shape of the probability distribution $p(z)$ [50].

If $p(z)$ is practically zero everywhere except for a few sharp peaks, then any sample points that lie away from these peaks will not contribute anything to the average. In an ideal case, we could simply use $p(z)$ to pick our sample points. However, to do this, $p(z)$ must have an exact and invertible antiderivative. In practice, this is rarely possible (even for the simple Gaussian distribution), so we need a different way to choose sample points. Monte Carlo is one way to pick sample points in such a way that the sums in (3.2) converges quickly.

Since we are dealing with a physical system and want to average over equilibrium states, we need to ensure that we choose sample points in a way that is compatible

with physical equilibrium. To do so, we need to ensure that each step is time-reversible. That is, the probability of being in state x and going to state x', which we denote $T(x|x')$, is equal to the probability of being in state x' and going back to state x:

$$p(x)T(x|x') = p(x')T(x'|x)\,. \tag{3.3}$$

We know $p(x)$ up to a constant $\mathcal{Z}$ (which cancels out). Putting these known probabilities and rearranging yields the detailed balance condition, given the energies E and E' of the two states:

$$\begin{aligned} \frac{e^{-\beta E}}{\mathcal{Z}} T(x|x') &= \frac{e^{-\beta E'}}{\mathcal{Z}} T(x'|x) \\ \frac{T(x|x')}{T(x'|x)} &= e^{-\beta(E(x')-E(x))}\,. \end{aligned} \tag{3.4}$$

Any transition function that satisfies this criteria would work to move from one sample point to another. Following Metropolis [51], we make a reasonable choice: the move is always accepted if the energy is lowered (i.e. if the new state is more probable). Since the transition probability will be a function of x and x', we now have:

$$T(x|x') = \begin{cases} 1 & , \quad E(x') - E(x) < 0 \\ f(x,x') & , \quad E(x') - E(x) \geq 0\,. \end{cases} \tag{3.5}$$

Plugging this into the detailed balance condition and solving for f yields the Metropolis algorithm:

$$T(x|x') = \begin{cases} 1 & , \quad E(x') - E(x) < 0 \\ \exp(-\beta(E(x') - E(x))) & , \quad E(x') - E(x) \geq 0\,, \end{cases} \tag{3.6}$$

which can be rewritten as

$$T(x|x') = \min\left[1,\, \exp(-\beta\Delta E)\right] . \tag{3.7}$$

This then provides the procedure for choosing a new sample point in accordance with detailed balance, and thus moving through the configuration space of the system along equilibrium configurations. By taking measurements of the variable A at each step (whether the system has moved to a new configuration or not), we are performing an average over equilibrium configurations.

To summarize, if we have a system with a very large configuration space and we cannot calculate the antiderivative of the probability distribution, nor can we calculate the normalization factor, then we need to sample a number of discrete points. To move through the space in a way that satisfies equilibrium yet moves towards areas of high probability, we use the Metropolis algorithm.

3.2.3 Specifics

In our system, we have two kinds of particles: chain monomers and free counterions. We sometimes need to hold the chain particles fixed, while averaging over counterion configurations. At other times, we let both move freely. Because we sometimes want to average over many counterion configurations for each chain configuration, we also need to be able to adjust how much the chain particles move relative the free counterions. All of this can be accomplished if we choose to have two different step-sizes, one for each particle type.

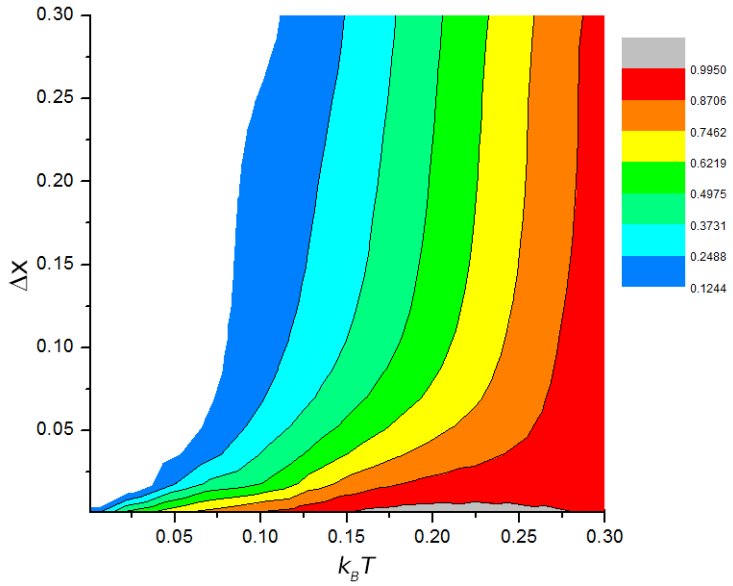

Figure 3.3: Here we see how step size Δx and temperature k_BT affect the acceptance rate for our simulations. For a given temperature, we can now choose a step-size that yields a close to optimal acceptance rate.

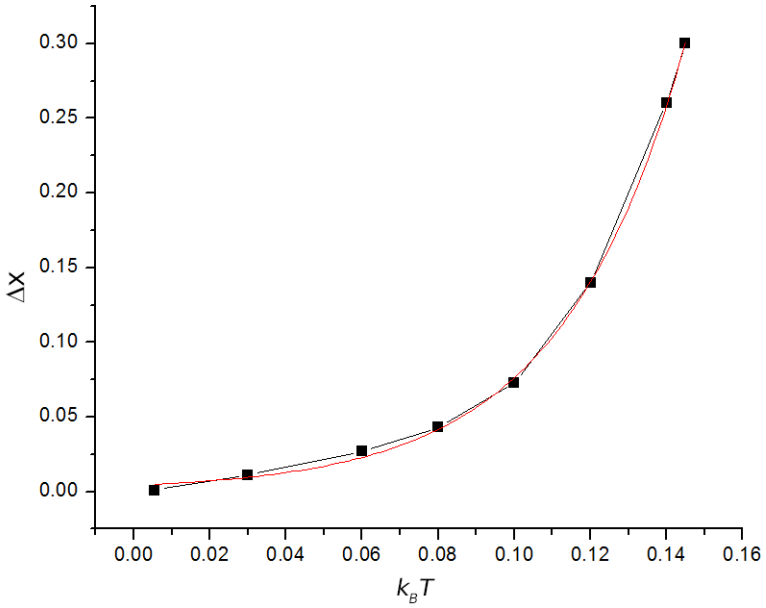

Figure 3.4: Fitting the previous figure allows us to find an analytical expression for optimal step-size given the temperature. We find an exponential relationship between Δx and k_BT.

The optimal MC acceptance rate for most systems is ~ 0.24 [52]. Because we varied the temperature in our simulations, we needed to find the step sizes that

would yield this optimal acceptance ratio for each temperature.

We also needed to have a different step size for counterions and for chain particles. This allowed us to tune each separately, as they had different potentials and therefore different acceptance rates for the same step-size and temperature. Having different step-sizes for chain and counterion particles also allowed us to easily fix the chains (setting chain step size $\Delta x_c = 0.0$), and to average over a great many counterion configurations for each chain configuration (by setting $\Delta x_c \ll \Delta_x$).

3.3 Periodic Boundary Conditions

We need to simulate infinite chains. To do this, we can make the boundaries of our simulation periodic. If we plan to orient our chains along the y-axis, we must set the periodic slab height L_y equal to some multiple[1] of b, the monomer spacing. The simulation box is $\sim 5L_y$ in the other two dimensions, which we found to be large enough for our purposes. We also prescribe that a counterion that passes through the top or bottom of the slab, comes out at the opposite side:

$$y_{wrapped} = \mathrm{mod}\left(y + 2L_y, L_y\right) , \tag{3.8}$$

where $\mathrm{mod}(a, b)$ is the remainder of a/b.

When computing the y distance between two particles i and j, the *image* of particle j might be closer to i than particle j itself. To restrict the range that y can take in our calculations, we need to take the distance to the *nearest image* (here *nint* is the nearest integer function):

$$y_{ij} = (y_i - y_j) - L_y \,\mathrm{nint}\left(\frac{y_i - y_j}{L_y}\right) . \tag{3.9}$$

[1]To avoid situations where every *nth* monomer has the same displacement, it is best to use a prime number multiple for the slab height. For example, box heights of $L_y = 41b$ or $L_y = 23b$ are better than $40b$ and $20b$.

In making the system periodic, we must be careful with the long-ranged Coulomb interaction[2]. We are faced with the problem of an infinite number of particles and an infinite number of interactions to consider. Since the Coulomb force gets weaker with distance, we could argue that we should only take the nearest N particles into account. However, because the Coulomb force decays relatively slowly (as opposed to, say, e^{-r}), we would still need a very large number of periodic images to get a reasonable estimate. In other words, the sum $\sum_{j=1}^{\infty} \frac{q_i q_j}{r_{ij}}$ converges slowly[53].

In 1921 Ewald published a simple trick [54] that made it more manageable (though it only applies to overall neutral systems). His trick was to split the interaction into long-range and short-range components. The short-range component considers only the particles in the simulation box. The long-range component is due to the periodic images. By summing the short-range component in real-space and the long-range component in reciprocal space, we have two rapidly converging sums to replace the original slowly converging sum. This technique has since been extended to other potentials and to other periodic boundary conditions. The case for Coulomb potential in our 1D slab geometry is called a Lekner potential[55], and is discussed in the following section.

3.4 Interaction Potentials

The evolution of our Monte Carlo simulation is determined entirely by the energy potentials we use and the initial conditions. There are no constraints on the system, other than those imposed through potentials. For instance, we do not force the particles to remain on the chain, we simply make it energetically favorable to do so. Because of this, it is important that we describe the potentials in detail.

[2]We would have to be careful with any infinite ranged interaction, but Coulomb is the only one in our system.

3.4.1 Lekner Potential

The dominant interaction in our simulations is the Coulomb potential. However, due to periodicity, we need to sum over not just the particle pairs, but also their periodic images (up to a predetermined cutoff). It is well known that this sum converges very slowly, and the standard procedure for periodicity in 3D is to use Ewald sums. Ewald sums break the sum into two components, one short-range and one long-range. The short-range sum is evaluated in real-space, while the long-range component in Fourier space. The extension of this result to periodicity along one dimension (our case) is called Lekner summation[56].

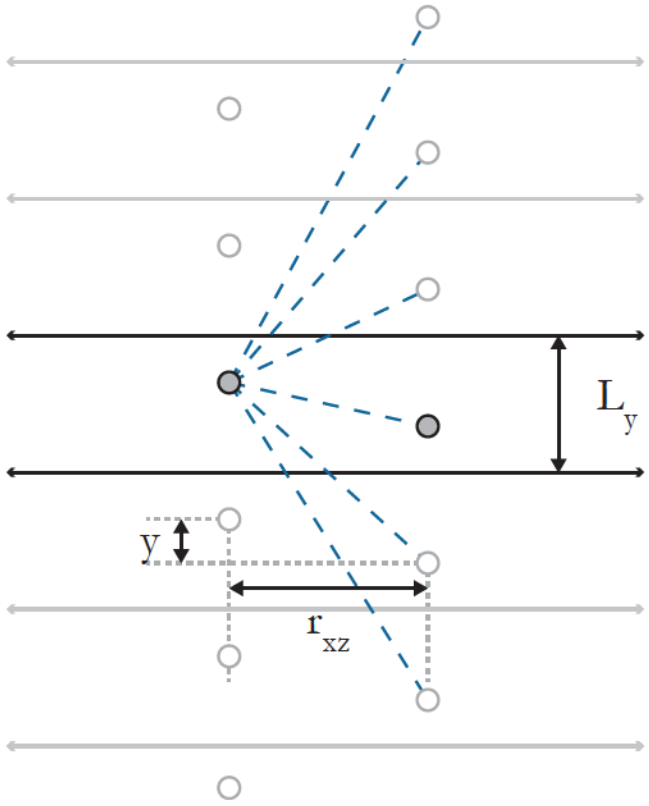

Figure 3.5: The coordinate system used. The slab is periodic in the y direction and infinite in the x and z directions. For the particles shown, y is the distance between their nearest images, and r_{xz} is their radial distance on the xz plane.

The potential energy between particles 1 and 2 is given by:

$$V_{12}(r_{xz}, z) = -2\frac{q_1 q_2}{L_y} \ln r_{xz} + 4\frac{q_1 q_2}{L_y} \sum_{n=1}^{\infty} \cos\left(2\pi n \frac{z}{L_y}\right) K_0\left(2\pi n \frac{r_{xz}}{L_y}\right), \qquad (3.10)$$

where $K_0(\cdot)$ is the modified Bessel function of the second kind, which is an exponentially decaying function[57].

3.4.2 Repulsive Lennard-Jones Potential

So far, the total force on a counterion is given by the Coulomb force (represented by the Lekner sum). With this alone, nothing prevents particles of opposite charge from coming arbitrarily close to one another. This is unphysical, because in a real system we would expect repulsion at short-range, no matter what charges the particles have, due to excluded volume effects[58]. We will use a repulsive Lennard-Jones model to capture the short-range repulsion, as is standard in studies such as these (e.g. [1], [14]).

The purely repulsive Lennard-Jones potential between two particles a distance r apart, is given by,

$$\mathcal{U}^{\mathrm{RLJ}} = \begin{cases} 4\epsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} + \frac{1}{4}\right] & , \quad r \leq 2^{1/6}\sigma \\ 0 & , \quad r > 2^{1/6}\sigma\,. \end{cases} \qquad (3.11)$$

Here ϵ is an energy, and it determines the strength of the interaction. The parameter σ is related to the effective size of the particle. This potential, and the resulting force, are continuous at $2^{1/6}\sigma$. The negative gradient of $\mathcal{U}_{RLJ}$ gives us the force:

$$\mathbf{F}^{\mathrm{RLJ}} = \begin{cases} 24\epsilon\sigma^6\left[\frac{2\sigma^6}{r^{13}} - \frac{1}{r^7}\right]\hat{\mathbf{r}} & , \quad r \leq 2^{1/6}\sigma \\ 0 & , \quad r > 2^{1/6}\sigma\,. \end{cases} \qquad (3.12)$$

To express this in Cartesian coordinates, use $\hat{\mathbf{r}} = \frac{1}{r}\left(x_{ij}\hat{\mathbf{x}} + y_{ij}\hat{\mathbf{y}} + z_{ij}\hat{\mathbf{z}}\right)$, $r = \sqrt{x_{ij}^2 + y_{ij}^2 + z_{ij}^2}$, and $x_{ij} \equiv x_i - x_j$. This force is calculated for all particle pairs (chain-ion, ion-ion, and chain-chain). Note that, due to periodicity, y values need to be wrapped to the nearest image.

3.4.3 Spring Potentials

We now turn to the bonded potentials that turn our collection of particles into a polymer chain. Our spring potential also serve to model the polymer's stiffness or persistence length, and as we shall see, our choices here affect the stability diagrams of Chapter 5.

As with many minimal polymer simulations, we need to ensure that the bonds between monomer units do not grow to unrealistic lengths. This is usually accomplished with the finitely extensible nonlinear elastic (FENE) potential[59]:

$$V_{\text{FENE}}(r) = -\frac{hL_{\text{max}}^2}{2}\ln\left(1 - \frac{r^2}{L_{\text{max}}^2}\right). \tag{3.13}$$

Here h is the spring constant. In all our simulations, we choose $h = 0.82nN/nm$ to approximate the spring constant of single-stranded DNA[60].

This potential has no minimum length, so it is important that the chain particles also have the repulsive Lennard-Jones potential (described in the previous section) to prevent the monomers from overlapping or colliding. The FENE potential, as implied by the name, can only extend to a maximum length of L_{max}, and the energy increases asymptotically as $r \to L_{max}$.

3.4.4 Helical symmetry-breaking potential

Our simple polymer model is built of charged particles connected in a spring. As is, there is no reason for a section of the polymer to prefer to bend in one direction

as opposed to the other. This leads to arbitrary fluctuations about the mean axis.

However, DNA has base-pairs along one side, and these bases form hydrogen bonds with their complements. These hydrogen bonds are weak relative the bare electrostatic charge of the backbone, but they do have one important consequence. Because the backbone bonds have an asymmetric preferred angle, the backbone tends to bend one way, and not the other[61].

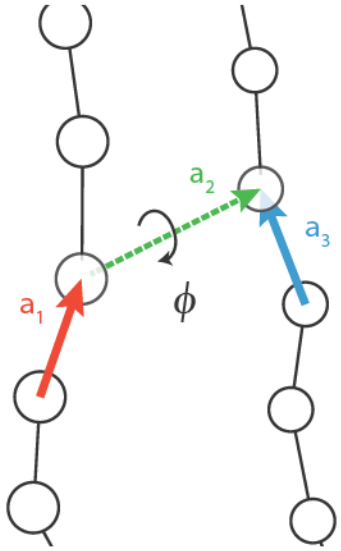

Figure 3.6: This is how we define the dihedral angle ϕ between corresponding particles on the two chains. The vectors a_1 and a_3 follow the bonded potential between consecutive particles on the same chain. There is no bonded potential along a_2.

We introduced a phenomenological potential between each monomer unit and the corresponding unit directly across from it. The form of the potential is as follows:

$$V_{bp}(r, \phi) = \epsilon_{bp}(\phi - \phi_0)^2 \,. \tag{3.14}$$

Here ϕ is the dihedral angle shown in Fig. 3.6 and is calculated from[3]:

$$\phi = \text{atan2}\left(|a_2| a_1 \cdot [a_3 \times a_2], [a_1 \times a_2] \cdot [a_3 \times a_2]\right) \,, \tag{3.15}$$

[3]Because our vector a_3 is reversed from the usual definition, our dihedral angle uses $a_3 \times a_2$ where other definitions would use $a_2 \times a_3$ instead.

where atan2(y, x) takes into account the signs of the arguments.

This potential will be absent from all our studies except for the very last section on "Flexible chains, symmetry-breaking". Even then, it will only be used with a very small energy (relative the other energy scales). We introduce it here for completeness.

3.5 Numerical Details

Our system has periodic boundary conditions, long-range forces with exponential singularities, and several other complicating factors. Because of this, we must take care with the numerical methods used. In particular, the Lekner sums are problematic. Because of the challenges in this system (periodic boundary conditions, long-range forces, short-range forces, risk of overflows/underflows, etc.), we will discuss a few of the techniques needed to obtain good behavior and results from this system.

3.5.1 Interpolation

Convergence problems

The Lekner potential is the most computationally intensive computation in each MC step. Moreover, though the sum converges quickly relative to the full Coulomb summation over images, it still does not converge quickly enough at small distances. Since we are often in a situation with small distances between ions and counterions, this is an important limitation. Worse, truncating the sum produces a false periodicity that strongly affects the results of the simulation (see Fig. 3.7).

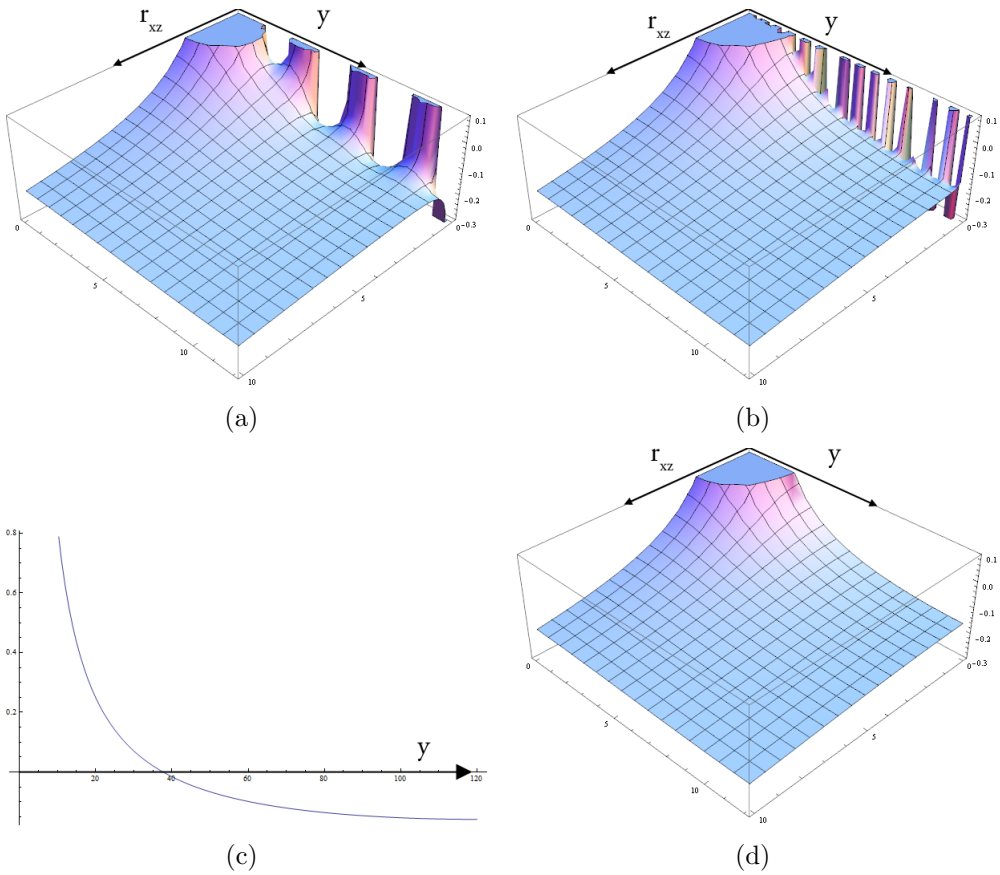

Figure 3.7: We see the need for care when using truncated Lekner sums. Subfigure (a) shows the Lekner potential when truncated at 7 terms, (b) when truncated at 30. Convergence is very poor as $r_{xz} \to 0$, as can be seen by the oscillations there. In (c) we see the analytical limit $r_{xz} \to 0$, and in (d) we apply this to the potential.

To remedy this situation, we have used a 2d interpolation (based on Numerical Recipes) for the calculation of the Lekner potential. This provides a significant speedup. Also, since we only need to calculate the interpolation table once, we can use a very large number of terms resulting in an increase in accuracy.

Convergence is much faster with Lekner sums than doing the direct Coulomb sum. However, there are still significant oscillations about the correct value when r_{xz} is small, as can be seen in figure 3.7. How close to 0 these oscillations occur depends on the number of terms kept in the infinite sum. For the points near $r_{xz} \sim 0$ where the oscillations start to appear, we transition smoothly to the analytical limit $r_{xz} \to 0$

instead

$$\lim_{r_{xz}\to 0}\left[-2\frac{q_1q_2}{L_y}\ln r_{xz}+4\frac{q_1q_2}{L_y}\sum_{n=1}^{\infty}\cos\left(2\pi n\frac{z}{L_y}\right)K_0\left(2\pi n\frac{r_{xz}}{L_y}\right)\right] \tag{3.16}$$

$$=\frac{q_1q_2}{12}\left[\gamma+\log\left(\frac{\pi}{24}\right)+\frac{\partial}{\partial n}\left[\mathrm{Li}_n(e^{-\mathrm{i}\pi y/12})+\mathrm{Li}_n(e^{\mathrm{i}\pi y/12})\right]_{n=0}\right]. \tag{3.17}$$

Here r_{xz}, z, L_y, and q_i are as before, K_0 is the zeroth order modified Bessel function of the second kind, and Li is the logarithmic integral, defined $\mathrm{Li}=\int_\mu^x \mathrm{d}t/\ln t$.

Parametrization

The Lekner sum is a function of the slab height L_y, which we need to be able to vary. Instead of calculating a new interpolation table for each value of L_y, we can instead parametrize the equation and create one interpolation table, applicable to all L_y. To do so, we need to define $\rho\equiv r_{xz}/L_y$ and $\zeta\equiv z/L_y$, and rewrite the Lekner function as follows:

$$V_{12}(r_{xz},z)=-2\frac{q_1q_2}{L_y}\ln r_{xz}+4\frac{q_1q_2}{L_y}\sum_{n=1}^{\infty}\cos\left(2\pi n\frac{z}{L_y}\right)K_0\left(2\pi n\frac{r_{xz}}{L_y}\right) \tag{3.18}$$

$$V_{12}(\rho L_y,\zeta L_y)=-2\frac{q_1q_2}{L_y}\ln(\rho L_y)+4\frac{q_1q_2}{L_y}\sum_{n=1}^{\infty}\cos\left(2\pi n\zeta\right)K_0\left(2\pi n\rho\right) \tag{3.19}$$

$$\hat{V}_{12}(\rho,\zeta)=\frac{2q_1q_2}{L_y}\Bigg[\underbrace{2\sum\cos\left(2\pi n\zeta\right)K_0\left(2\pi n\rho\right)-\ln\left(\rho\right)}_{I(\rho,\zeta)}-\ln\left(L_y\right)\Bigg] \tag{3.20}$$

$$\hat{V}_{12}(\rho,\zeta)=\frac{2q_1q_2}{L_y}\Big[I(\rho,\zeta)-\ln\left(L_y\right)\Big]. \tag{3.21}$$

Since it is free of L_y dependence, we can now generate a single interpolation table for $I(\rho,\zeta)$.[4]

[4]Another consideration: since the function increases exponentially as $r_{xz}\to 0$, we need to use a higher density of points in that region.

3.5.2 Equilibration

As with any Monte Carlo simulation, we need to perform several time-steps before gathering data in order to let the system forget the arbitrary initial conditions we imposed. A common approach to speed up the equilibration process, and to prevent the system from getting stuck in a local minima early in the equilibration stage, is to start at a high temperature and decrease it slowly towards the desired value in a simulated annealing. Though this did help, we found that the number of time-steps required at each temperature led to very long equilibration times.

Instead, we chose an alternative to this. We started the system as before (counterions randomly distributed near the chains), but all charges were scaled to $q' = q/100$. Since the interaction starts weak, counterions are free to move away from the chains. We then slowly ramped up the charge until its full value. In a sense, we were slowly introducing the electrostatic interaction. Also note that since all charges were scaled equally the system remains charge-neutral throughout this process.

3.6 Units

Our simulation units were chosen so that the parameters (lengths, energies) for a DNA-like system were of $\mathcal{O}(1)$. Our choices are as follows:

Table 3.1: Code Units

Code unit	SI	other
$\hat{E}$	$2.9390335 \times 10^{-21} J$	0.423 kcal/mol
$\hat{L}$	10^{-9}m	nm
$\hat{Q}$	$1.602 \times 10^{-21} C$	$1e$

Part II

Results & Discussion

Chapter 4

Like-charged Attraction

4.1 Overview

In our three-part study, our first goal is to find out if and when two charged polymers would attract each other. Once this is established, we can investigate how this attractive force affects the polymer's stability to deformation (Ch. 5), and what kinds of long-term structures result (Ch. 6). This chapter is dedicated to answering the first part of this study: determining how and when do our minimal model polyelectrolytes attract each other.

To study the attractive force between two straight like-charged polymers, we will examine three things. First, the force on the chains as a function of chain separation R, and how it compares to known results. Second, we want to determine how this force curve changes when we vary the parameters for counterion valence q_{CI}, or temperature T.

4.2 Setup

For all simulations in this chapter, we held the chains fixed at a separation R, while allowing the counterions to move freely. We used a parameter set that is similar to

the conditions of DNA in solution (within the constraints of our model). Our choice for the parameters for this simplified DNA-like system are summarized in the table below:

Table 4.1: Simulation Parameters

Parameter	Symbol	Real units	Code units
thermal energy	k_BT	300 K	$1.409\hat{E}$
monomer spacing	b	3.4 Å	0.34 nm
periodic height	L_y	108.8 Å	$32b = 10.88$nm
chain particle charge	q_{chain}	$-e$	-1.0
counterion charge	q_{CI}	$2e$	2.0
particles per chain	N_ℓ	32	-
number of counterions	N_{CI}	32	-

In section 4.3.3 we look at how the force curve changes when we vary temperature; so we run simulations for k_BT ranging from 0.001 to 1.75, while keeping all other parameters fixed at their default values.

4.3 Results

We now turn to one of the more unexpected features of this system: the fact that the negatively charged polymers attract each other. As is the case with DNA, we only see attraction when surrounded by divalent or polyvalent cations, and not with monovalent ions. As we will show, this is a result of the particular parameters of DNA in aqueous solution, and not a general result.

We will also examine how this attractive force is affected by various changes in the system parameters, such as counterion valency and temperature. From this, we are able to deduce that the dominant source of the like-charged attraction in our system is the Wigner crystals that form at zero temperature, and the persistence of those correlations to finite temperatures.

4.3.1 DNA-like system with divalent counterions

We begin by verifying that our simulations reproduce experimental fact: two DNA strands in solution with multivalent counterions will attract each other.

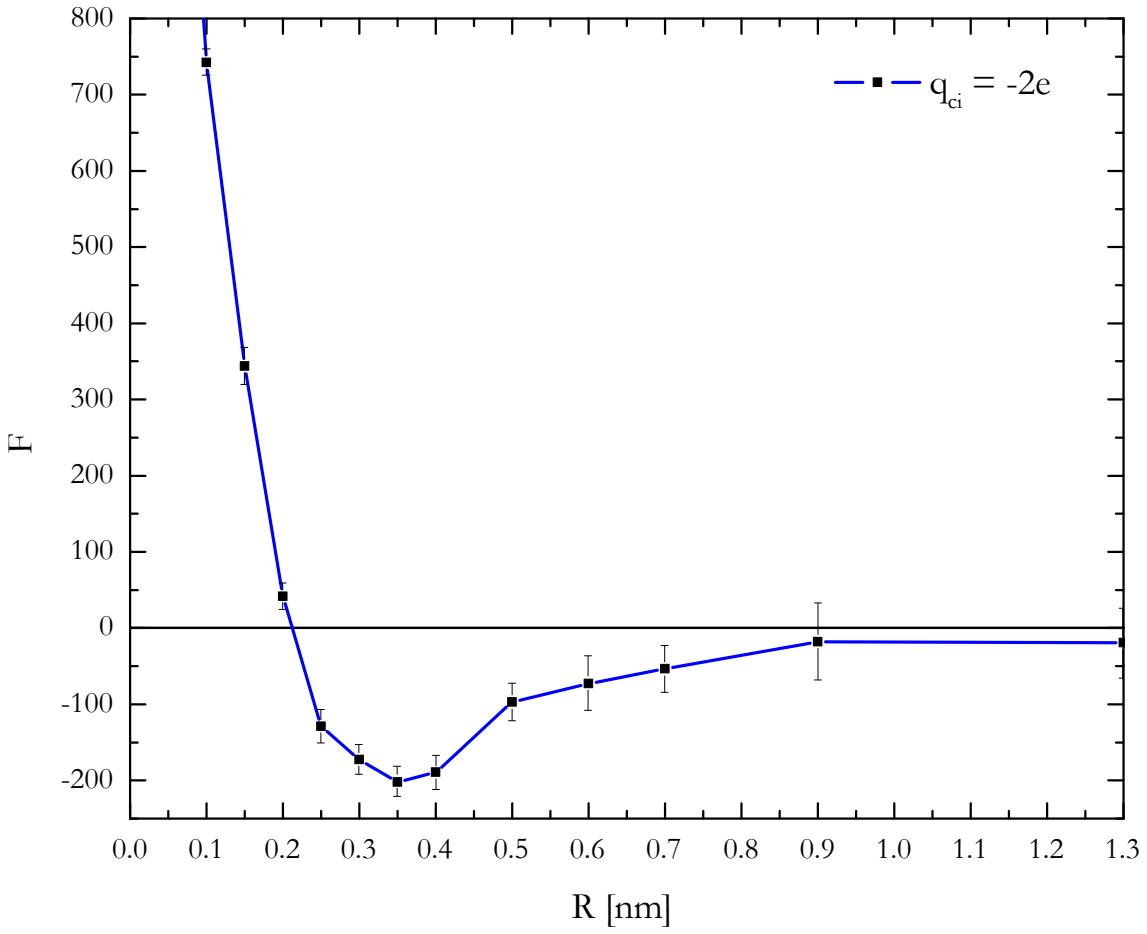

Figure 4.1: The force between two parallel charged polymers surrounded by counterions. We can see that the force is negative (attractive) when separation $R \gtrsim 0.2\,\text{nm}$.

Figure 4.1 shows attraction at $R \gtrsim 0.2\,\text{nm}$, on to infinity. We expect, of course, that the force should decay to zero at infinity. This attraction agrees, qualitatively and quantitatively, to that calculated by Gronbech-Jensen by other methods[1].

This figure tells us that the rods are very weakly attractive at long range, very strongly attractive at close range, and repel strongly once the counterions no longer fit between the chains. As we will see, the shape of this force curve will have implications

for the stability of the chains against deformation, once we move from these straight-chain/rigid studies to flexible ones.

4.3.2 Counterion valency

Now that we have seen that our system does exhibit an attractive interaction, we will see how this interaction is affected by varying one parameter, the counterion valency. Most experimental studies include some variation in the counterion charge, as it is a very easily accessible parameter; one simply needs to add different cations to the solution, and watch to see if DNA precipitates or not. Experiments such as these find that, for DNA, if we introduce multivalent counterions, they condense. If the solution contains only monovalent ions, there is no condensation.

We can perform similar studies by varying the charge of the counterions. Note that, in order to keep the charge of the chain monomers the same, and the overall system neutral (as required by the Ewald/Lekner approach to periodic electrostatics), we need to reduce the number of counterions accordingly.

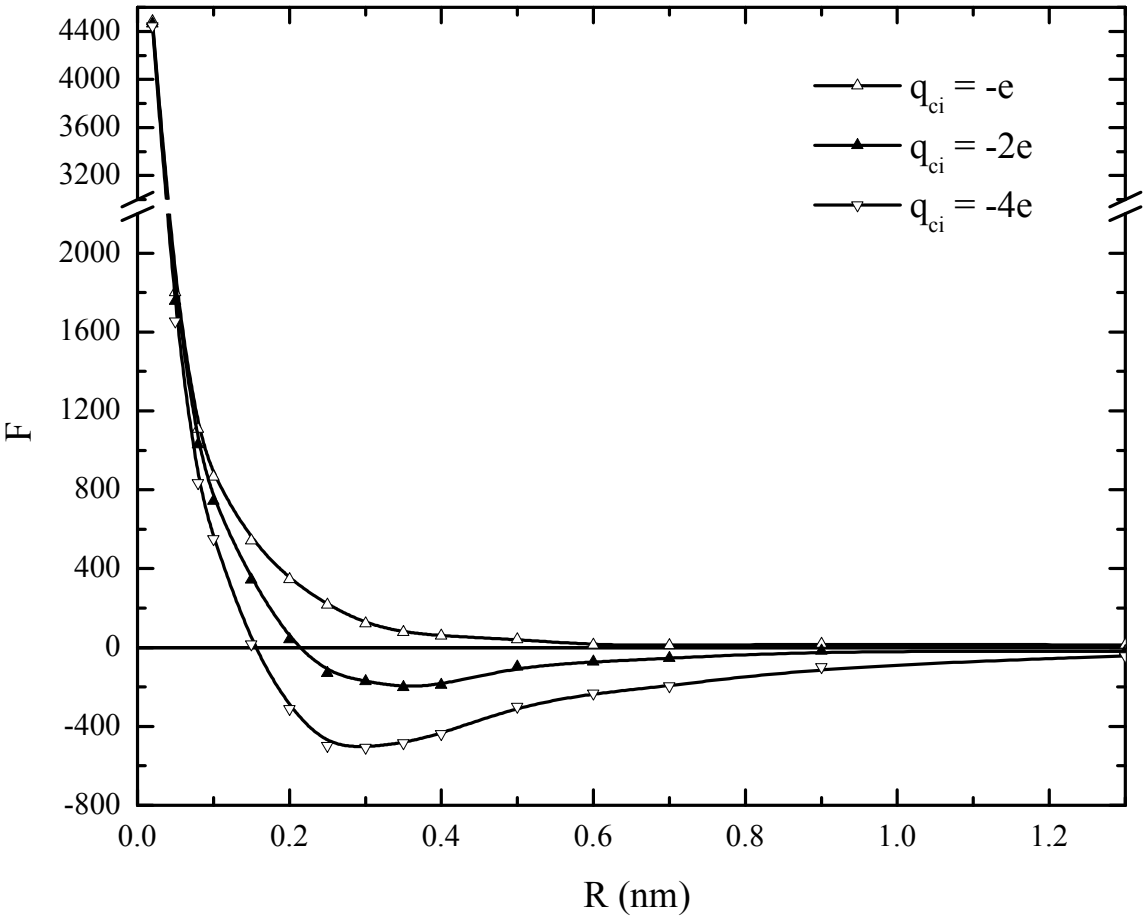

Figure 4.2: As expected from experiment, we see that monovalent counterions (open, upward triangles) do not exhibit attraction (that is, the force is never negative). For valencies 2 and 4, the attractive force increases, as expected from previous experimental and analytical predictions.

We see in Fig.4.2 that monovalent counterions are not sufficient to produce attraction, and that it requires divalent or higher[1]. We also see that the higher the valency, the stronger the attractive force. We note that with other parameter sets we can see attraction regardless of the valency of the counterions. However, it is a well-known experimental result that only multivalent counterions ($Z_{\mathrm{CI}} > 1$) induce attraction. This is a misleading and incorrect statement that is repeated often in the experimental literature (e.g. [33], [1], [62]). What they should say is: for these

[1]Throughout we refer to the counterion valencies, and keep the chain charged fixed. But it should be noted that what matters is the ratio of these two, not the absolute value of the charge.

given polymers (e.g. DNA), in this particular dielectric medium (water), at this temperature ($300\,K$), only multivalent counterions induce attraction. Though more cumbersome, this statement has the virtue of being correct (see [14] for more on this).

4.3.3 Temperature

The final parameter we varied was temperature. At very low temperatures, such as $k_BT = 0.001 \sim 200$mK, the electrostatic energy dominates over the thermal energy, and we always have attraction. As we increase the temperature, the strength of the attractive force gradually diminishes, until it disappears above a critical temperature $T \approx 0.79488 = 169$K. This temperature agrees very well with the critical temperature we expect from setting the Manning parameter $\xi = 1$, namely $T_c = 0.795$. Recall that the Manning parameter is the ratio of the electrostatic energy to the thermal energy.

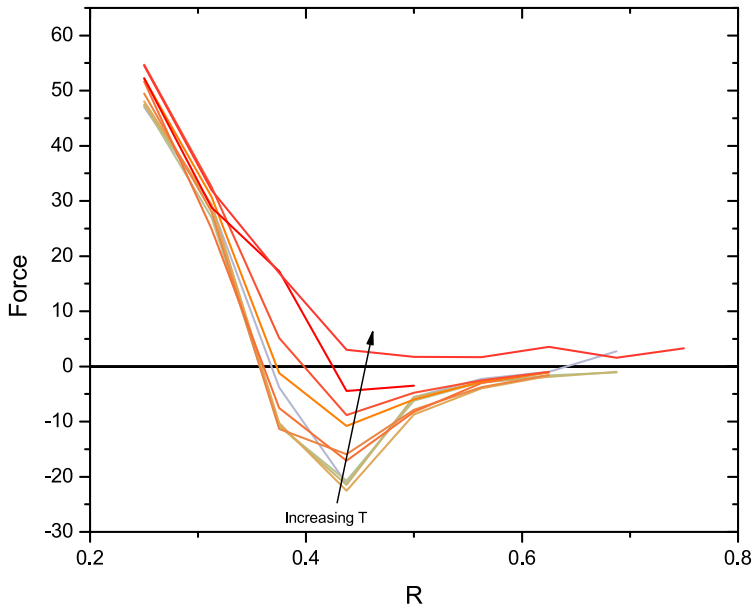

Figure 4.3: The force curves for increasing temperatures. As we can see, at higher temperatures, the curve is no longer attractive (negative) at any point.

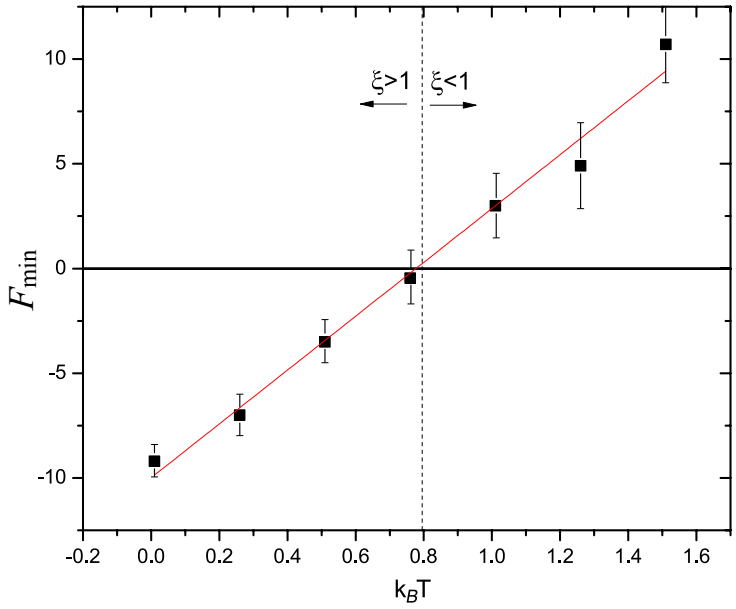

Figure 4.4: From the previous figure, we took the minimum value of the force, and plotted it as a function of temperature. We can see that the force begins attractive (negative) at $k_BT \sim 0$, and linearly increases (becomes less attractive) until the threshold at the Manning critical temperature of $T_c = 0.795$. Beyond this point, the force between the chains is increasingly repulsive.

These results demonstrate that the counterion induced attraction is strongest at zero temperature, but that the ground-state order (and associated attraction) survives into finite temperature. This suggests two important questions. First, we want to determine where the counterions are found at zero temperature (*i.e.* the ground-state configuration). Second, we would like to determine the cause of the linear decrease in the attractive force (Fig. 4.4).

Since we have run simulations at very low temperature, we already have the answer to the first question of where the counterions are located. At near-zero temperature, the counterions have minor fluctuations around their ground-state configuration, a snapshot of which is seen in the following figure.

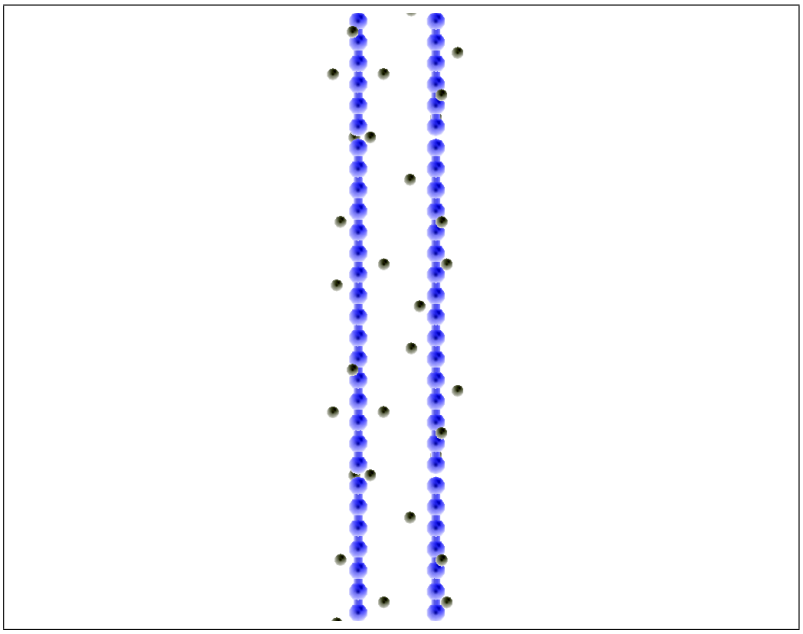

Figure 4.5: Here we have an initial ground-state configuration after equilibrating the counterions only (chains held fixed at distance $R = 13.6\mathring{A}$).

4.4 Discussion

We now have simulation results regarding the like-charged attraction between two chains. We have seen that, with our system, we can reproduce the core features of this interaction.

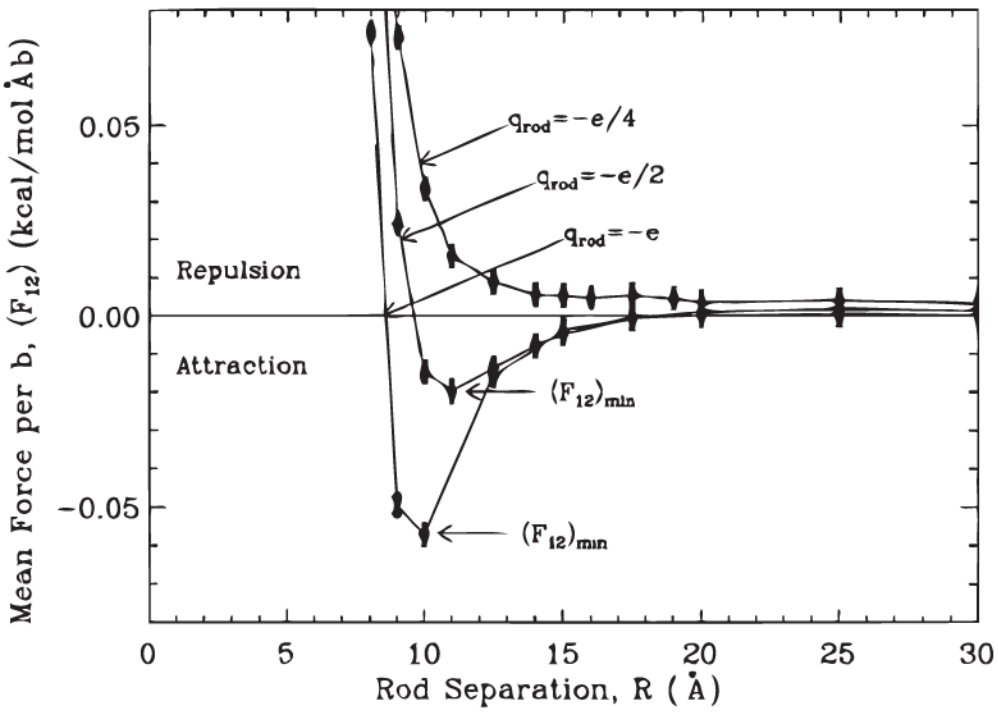

(a) b

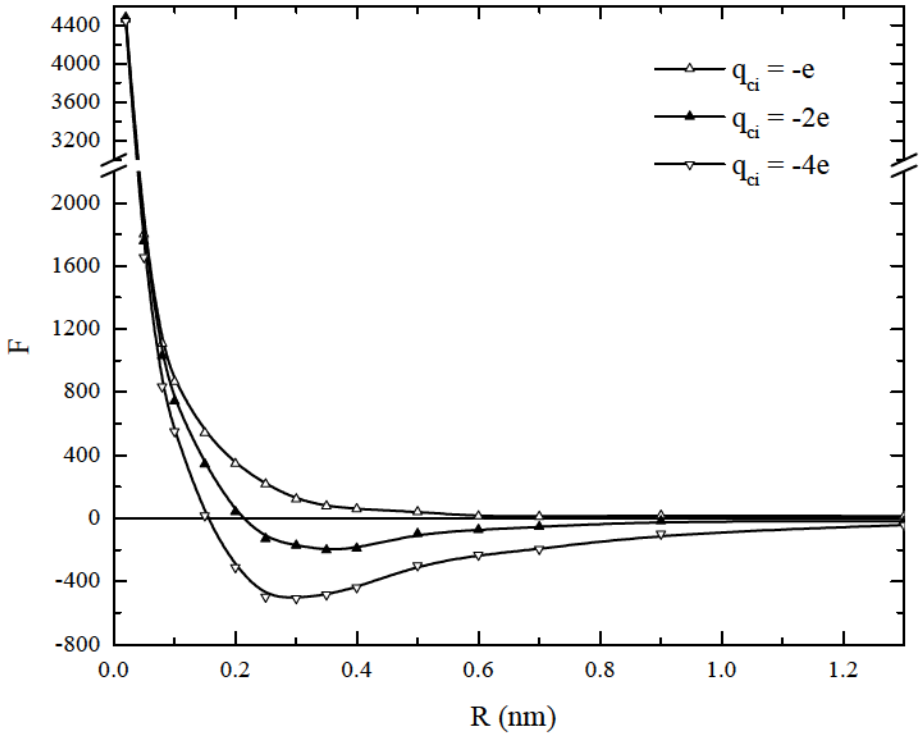

(b) b

Figure 4.6: A comparison between (a) Niels-Gronbech Jensen's results[1] and (b) our force curves. Note that we list the counterion charge, while they list the rod charge.

Moreover, since we have found the force to be caused by the ground-state Wigner crystal configuration, we can calculate the force on this particular configuration directly, using the Lekner sums and the appropriate particle positions.

4.4.1 Force curve fitting

In later chapters, we will need an analytic expression for the force curve, whether exact or an approximate fit. The force curve may seem similar to that generated by many long-range attraction/short-range repulsion potentials, such as the Lennard-Jones type forces, or generalizations thereof.

However, there is no simple power law, or combination of power laws, that fits the entire range of interest. In order to find the relevant ansatz for our fits, we incorporated a number of functions that are characteristic of this system. Namely, we incorporated exponentials, Bessel functions, and logarithms into *Mathematica*'s linear combination fitting function, to obtain:

$$F(R) = \left[AK_0(R)^2 + B\log(R)\right] \times \exp\left(-CR^2\right) - D\,. \tag{4.1}$$

The resulting function satisfies our requirements, and fits our curve with $R^2 = 0.983$ (parameters given below).

Table 4.2: Force curve fit parameters for DNA-like system.

A	503
B	968
C	2.82
D	5.04

A key difference between this function, and power-law functions, is the slow transition from attraction to repulsion. We will see later that the width of this well is what gives flexible polyelectrolytes the relatively large regions of instability that we

find.

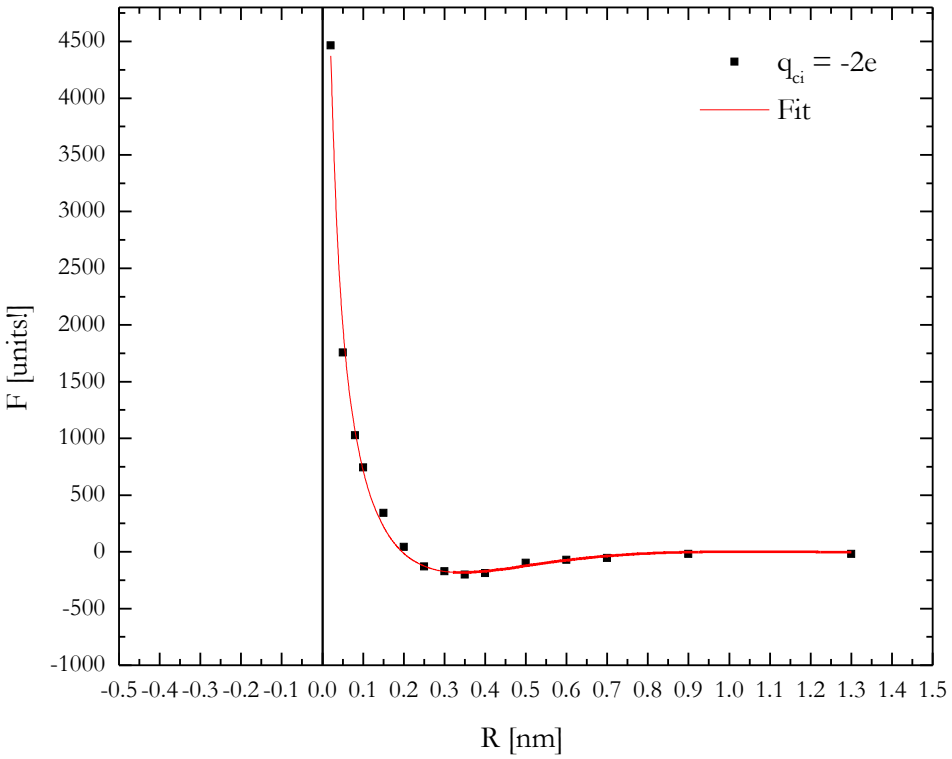

Figure 4.7: Though it seems a simple Lennard-Jones type power law should fit the force curve, this is not the case. Here we see our best fit, using the form $[AK_0(R)^2 + B\log(R)] \times \exp(-CR^2) - D$. Simpler forms cannot fit both the repulsive and attractive portions.

4.4.2 Direct computation

We have seen that at low-T, in a typical stationary configuration the counterions are arranged roughly as follows:

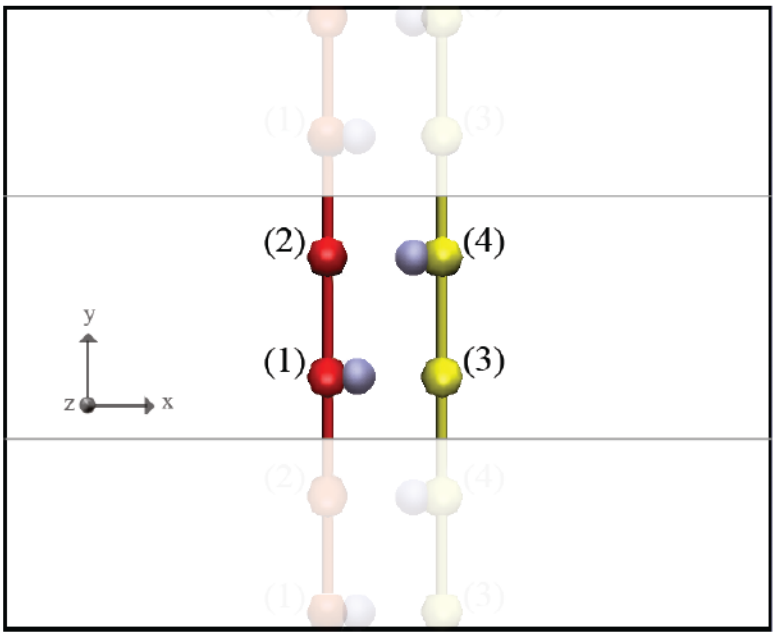

Figure 4.8: We will calculate the force between the chains due to this configuration. This configuration has half the inter-chain binding sites filled, and they alternate between occupied and empty along the chain.

Now that we have located the counterions in a typical configuration, we can use Lekner sums[55] to calculate the force on any one particle analytically. If we have periodicity along the y-axis, the Coulomb force component perpendicular to the y-axis, between two particles i and j, is given by[56]:

$$F_\rho = \frac{2q_iq_j}{L_y\rho} + \frac{8\pi q_iq_j}{L_y^2}\sum_{n=1}^{\infty} n\cos\left(\frac{2\pi n(y_j-y_i)}{L_y}\right) K_1\left(\frac{2\pi n\rho}{L_y}\right), \tag{4.2}$$

where $\rho = \sqrt{(x_j-x_i)^2+(z_j-z_i)^2}$, L_y is the height of the unit slab, and K_1 is the modified Bessel function of the second kind (first order).

To find the force on the entire right chain, we simply need to calculate the forces on particles (3) and (4) due to all other particles (both chains and counterions). The

force on i in the x-direction, due to all N particles is:

$$F_x^{(i)} = \sum_{j \neq i}^{N} \left[\frac{2q_i q_j}{L_y \rho} + \frac{8\pi q_i q_j}{L_y^2} \sum_{n=1}^{\infty} n \cos\left(\frac{2\pi n (y_j - y_i)}{L_y} \right) K_1\left(\frac{2\pi n \rho}{L_y} \right) \right] \frac{(xi - xj)}{\rho} \,. \tag{4.3}$$

If the counterion charge $q_{CI} = -2q_{chain}$, the system would be neutral and the force on the right chain would be $F_x^R = (F_x^{(3)} + F_x^{(4)})/2$. If the system is not neutral, we need to incorporate additional neutralizing charges. However, we cannot simply add a single compensating charge q_0 at the origin. In our simulations, any counterions not condensed onto the chains moved freely within the unit cell. To approximate this in our analytical calculations, we should add a uniform background charge of q_0 to a box of the same size as the simulation cell. Lekner sums with uniform background charges are highly non-trivial, so instead, we will create a uniform lattice of N' compensating charges q_0/N'. The force on the right chain is calculated as before, except we also include the N' neutralizing charges. This force is plotted in Fig. 4.9.

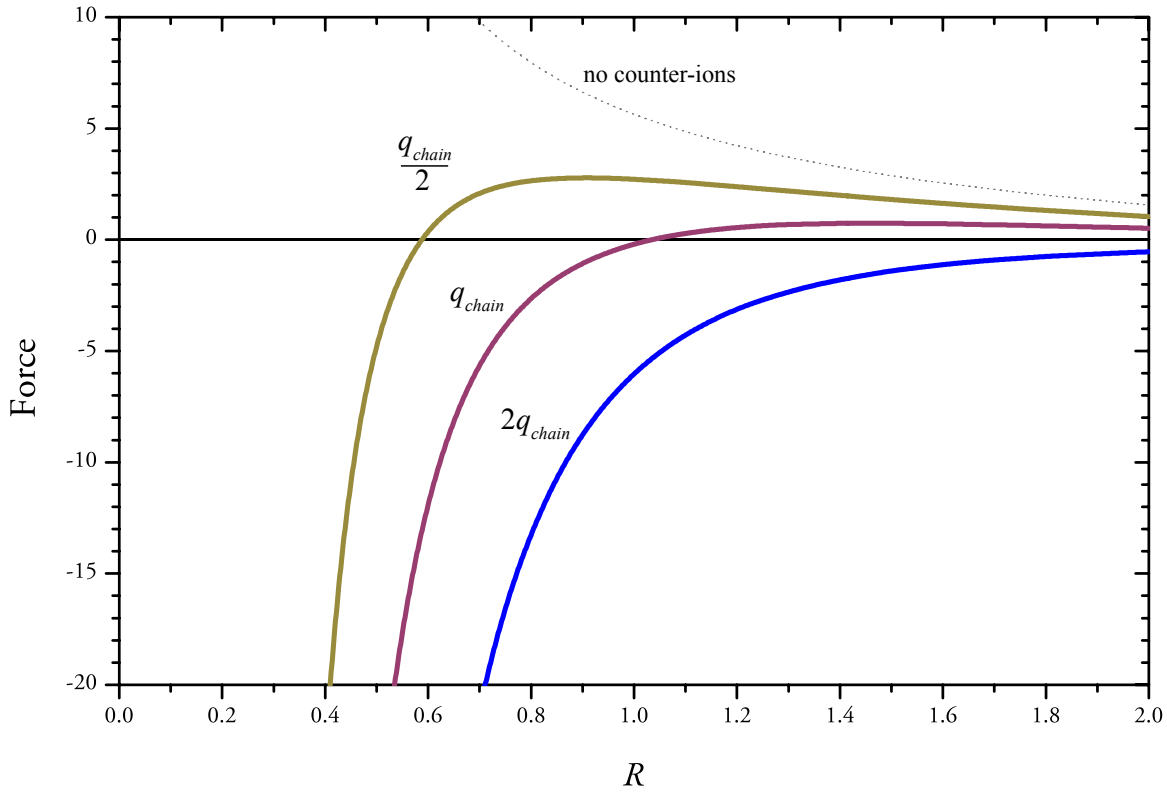

Figure 4.9: The analytical force curves for the Fig. 4.8 configuration. The curves are shown for different counterion charges q_{CI}. Note that we have not included any short-range repulsion effects. As we can see, all counterion valencies have an attractive (negative) region.

The configuration we have studied so far has only half of the possible binding sites filled. Yet, with only half the sites occupied by counterions of charge $q_{CI} = -\frac{q_{\text{chain}}}{2}$, the force becomes attractive at small separations. As we increase the counterion charge q_{CI}, the attraction gets stronger and longer-ranged.

We next consider the configuration shown in Figure 4.10. It has fully occupied counterion sites.

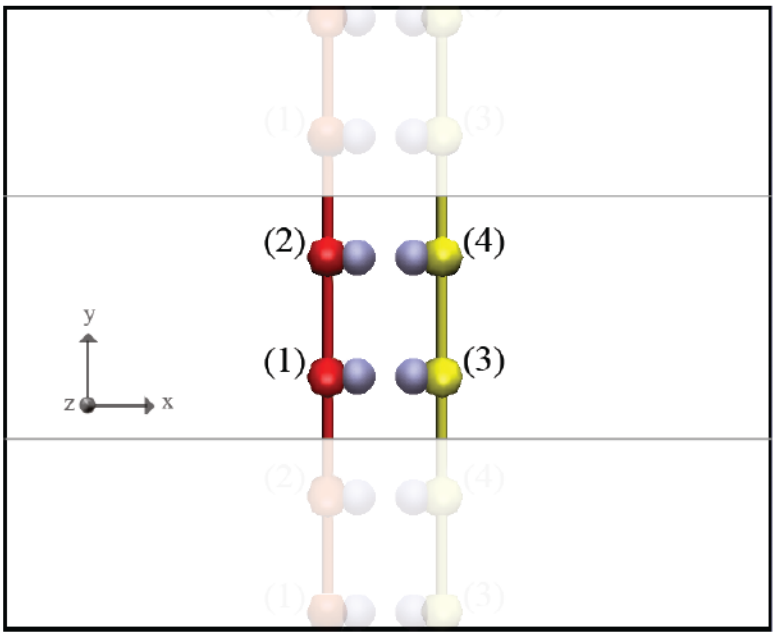

Figure 4.10: This configuration has all the inter-chain binding sites filled. The force is calculated for this configuration as well.

We use the same procedure as before to calculate the force on the right chain. Note that since this configuration has as many bound counterions as chain particles, the system is neutral when $q_{CI} = -q_{chain}$. For other charge ratios, we neutralize the system as before. The force curves are shown in Fig. 4.11.

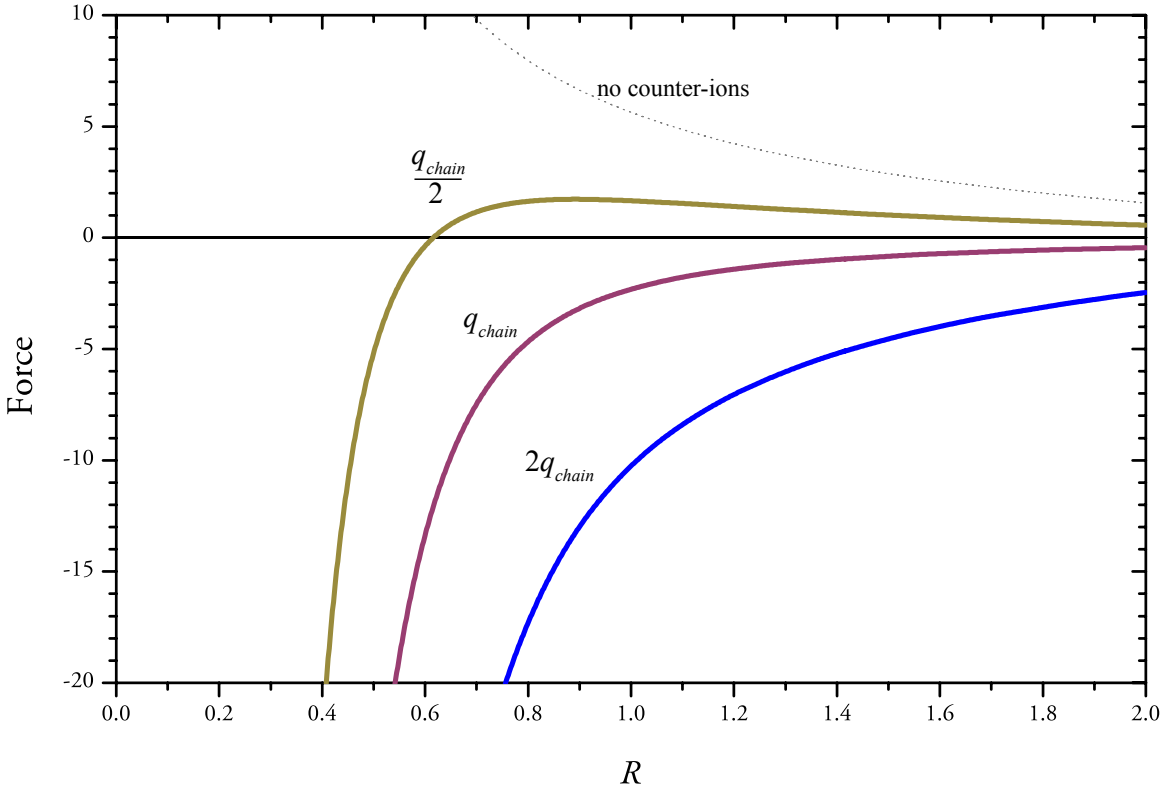

Figure 4.11: The analytical force curves for the Fig. 4.10 configuration. Once again, all three valencies exhibit an attractive (negative) region, and the highest valency, $2q_{chain}$ was attractive for all chain separations. Note again that we did not take into account and short-range repulsive forces.

Again the attraction is present (at short-range) even if the counterion charge is half the chain charge. If the counterion charge is at least equal to chain charge, the attraction is present for all R.

4.5 Summary

We have verified that attraction exists in our simulations. This is the essential first step for the following studies of how the chains initially deform on approach, and how they pair up and form structures. We have also determined the parameter ranges that are similar to a DNA-like system. This allows us to compare our minimal model of a polyelectrolyte to experiments of real systems, at least qualitatively. For instance, we

verified that for DNA-like parameters, only multivalent counterions induce attraction.

We have also performed the force analysis over a very wide range of temperatures, and found that the attraction is due to the zero-temperature Wigner crystals that form between the chains. We have used the Lekner sums to numerically calculate the force between two chains, and found that it is indeed attractive due to this arrangement.

We then found an analytical expression, eqn. 4.1, that fits the force curve, over the entire range of interest (that is, from 0 to where the force is negligible). This expression will be essential later, when we want to explore approximations that require an analytical expression for the force. Armed with these results of the force between two rigid polyelectrolytes, we can begin to explore how the attractive force affects flexible polyelectrolytes, something which has not been explored in detail.

Chapter 5

Stability Analysis

5.1 Overview

A straight chain models a perfectly inflexible polymer. This has some real-world analogs, but in biological systems, even stiff polymers (such as F-actin) have some degree of flexibility. This leads to an important question: whether the attractive interaction between two flexible like-charged polymers causes the polymers deformations to grow, leading to interesting structures, or to decay and cause the polymers to remain fully extended.

If they did deform and buckle in every possible scenario, then the entire discussion about how and why two polymers attract would be moot; no polymer pair would remain extended long enough for the attraction to have any effect. Fortunately, the polymers do not always buckle under this attractive force.

We will begin by reviewing linear stability analysis, and how to use it to determine whether the given chain configuration is stable, and if not, which sinusoidal mode we would expect to dominate in the linear regime. This will let us predict what kinds of structures we might see in a two-chain system, and by extension, in a solution of flexible charged polyelectrolytes.

Before performing extensive simulations of this system to answer these questions, we will try to predict, using a mean-field theory, what we might expect to see. Our approach uses the chain-chain force curve, a quantity that is relatively accessible experimentally, to make these predictions. Finally, we will explain the three numerical methods we used to determine the stability curves of our system, and compare their results to our mean-field predictions.

5.2 Motivation

Stability against deformation is an important part of the behavior of these polymers. We have already seen that a pair of PEs will attract each other, meaning that they will approach each other. We wish to determine whether they deform as they approach, or if they remain extended. We are also interested in the patterns and structures they might form, and how these are affected by the parameters of the system.

How pairs (and bundles, by extension) of polyelectrolytes behave as they approach each other is an important area of study. It affects how polyelectrolytes form packing structures, it affects DNA condensation[29], and can lead to novel structures such as F-actin liquid crystals[3]. The linear stability analysis of this chapter gives us a first-order approximation to answer these questions.

5.3 Theory

Linear stability analysis is a way to predict what large-scale periodic structures will form. The approach is to take the linear system, make a very small perturbation of a particular wavelength, and see if that perturbation will grow or if it will decay back to flat. The rate at which the perturbations grow or decay tell us which wavelength is likely to survive and grow to the point that it becomes visible at a larger scale.

5.3.1 Linear Stability Analysis

In this section we will present only the elements of linear stability analysis that are relevant to our particular system. First we must clarify what it means for something to be unstable. In the context of our system, a chain is said to be unstable against deformations when any deformation away from a straight line configuration grows over time.

Let us say we start with a straight chain that experiences a certain force, due to the counterions around it and to the other chain. Suppose that when we move one monomer to the right, it feels a force to the left. Then this small deformation away from linearity experiences a restoring force, so the deformation decays and the chain remains straight (i.e. that deformation is stable).

Now suppose that under different circumstances, when we move one monomer to the right, it feels a force to the right. Then this deformation would grow exponentially, and eventually the chain would deform drastically and possibly break. Or, once the deformation grows beyond a certain point, we are no longer in the linear regime (small amplitude), and non-linear effects may come into play.

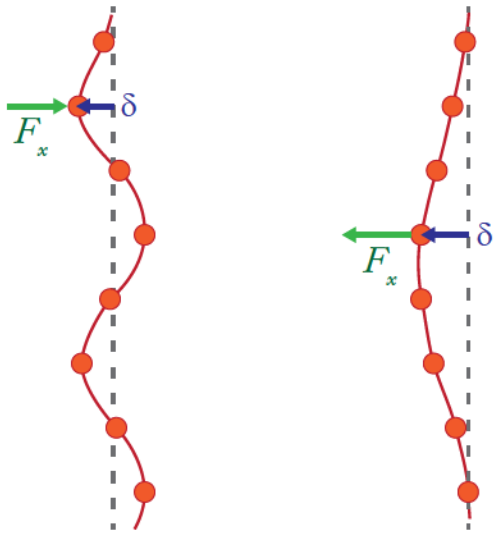

Figure 5.1: A perturbation can either grow or decay, depending on the exact nature of the forces on a given monomer. If the force is in the same direction as the displacement (right figure) it grows exponentially, otherwise it decays. Whether it grows or decays is a function of the wavelength.

How quickly a deformation grows is called the growth rate Γ. If it is positive, the deformation grows exponentially and the chain buckles; if negative, the deformation decays exponentially. For small deformations, it is expected that the growth rate of a deformation is directly proportional to the force on the deformed portion, F, and inversely related to the amplitude of distortion δ.

$$\Gamma = \frac{\Delta F}{\delta}. \tag{5.1}$$

The chains can be perturbed in many ways. We need to determine what kind of deformation we should use. Fortunately, for small perturbations, we can decompose any deformation away from a straight-line as a Fourier series of sinusoidal deformations. Each of these modes are independent in the linear regime, so we can examine them each in turn.

The wavelength of these sinusoids varies from $2b$ to L_y. The smallest mode is $2b$, due to the monomer spacing constraint; L_y is the longest mode allowed by our

periodic slab. Another way of looking at this: only an integer number of wavelengths is allowed in the y direction, in order to preserve periodicity ($\lambda = L_y/n$). Hence, our maximum wavelength of $L_y/1$. The Nyquist criterion for sampling frequency states that for our sample points spaced b units apart, the minimum wavelength is $2b$.

In summary, our goal in this chapter is to determine which scenarios cause the chains to buckle (become unstable), and which cause them to remain extended (stable). The tools we use to determine this are various implementations of linear stability analysis.

5.3.2 Mean-field Predictions

We can now be more concrete in explaining our approach. We will distort one chain sinusoidally with wavelength λ. For each particle i, we output the force $F_i(R)$ divided by that particle's displacement δ_i. Then the growth rate $\Gamma(R)$ is given by $\frac{1}{N_c}\sum_i (F_i(R) - F_{\text{straight}}(R))/\delta_i$, where $F_{\text{straight}}(R)$ is the force curve from the previous chapter. We will repeat the process for different wavelengths λ, and combine this information into stability diagrams (e.g. Fig. 5.8) that show where the growth rate $\Gamma(R, \lambda)$ is positive, and where is it negative.

Mean-field growth rate

If we have the force curve as a function of chain separation R, we can predict what stability diagram we might expect using the following procedure.

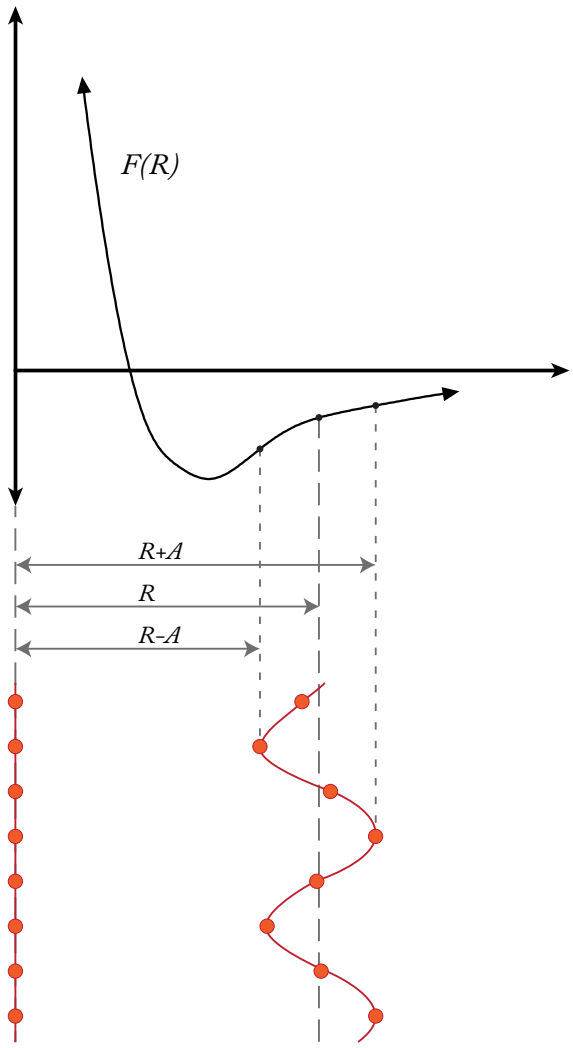

Figure 5.2: The force curve above is for a pair of straight chains a distance R from each other. If we deform one chain slightly (amplitude exaggerated here), then the force on a monomer i at a distance $R+\delta_i$ will be approximately equal to $F(R+\delta_i)/N_\ell$. We can average $\Delta F_i/\delta_i$ over all monomers to compute the stability of the chain against small perturbations.

Beginning with two straight chains, we deform one of them, as shown in Fig. 5.2. Now the force experienced by each monomer unit in the perturbed chain varies from $f(R-a_0)$ to $f(R+a_0)$. The total force is the average of all the forces felt by each monomer unit i at a displacement δ_i from the mean, as well as the restoring spring

force f_{sp}:

$$F_{tot}(R) = \frac{1}{N_l}\sum_{i=0}^{N_l} f(R+\delta_i) + f_{sp} \tag{5.2}$$

$$= \frac{1}{N_l}\sum_{i=0}^{N_l} f\left(R + a_0\sin(2\pi ik/N_l)\right) + f_{sp}\,. \tag{5.3}$$

As before, k is the wavenumber and N_l is the number of monomers on one chain. The spring force in the x direction on a monomer i is simply the sum of the spring force due to the monomers above and below it:

$$f_{sp} = f_{ch,x}(r_i, r_{i+1}) + f_{ch,x}(r_i, r_{i-1}) \tag{5.4}$$

$$= f_{ch,x}(R+\delta_i, R+\delta_{i+1}) + f_{ch,x}(R+\delta_i, R+\delta_{i-1})\,. \tag{5.5}$$

If we now plug in $\delta_i = a_0\sin(2\pi ik/N_l)$ and $f_{ch,x}(x_i, x_j) = h(x_j - x_i)/(1-(b^2 + (x_j-x_i)^2/L^2))$, and simplify, we get:

$$f_{sp} = a_0 hL^2\Bigg[\frac{\sin(2\pi ik/Nl) - \sin(2\pi(i-1)k/Nl)}{(b^2-1)L^2 + 4a_0^2\cos\left((2i-1)\pi k/Nl + \phi\right)^2\sin(\pi k/Nl)^2} + \frac{\sin(2\pi ik/Nl) - \sin(2\pi(i+1)k/Nl)}{(b^2-1)L^2 + 4a_0^2\cos\left((2i+1)\pi k/Nl\right)^2\sin(\pi k/Nl)^2}\Bigg]\,. \tag{5.6}$$

To get the growth rate, and therefore the stability diagram, we use $\Gamma = \Delta F/\delta$:

$$\Gamma = \frac{1}{N_l}\sum\frac{F_{tot}(R) - f(R)}{\delta_i}$$

$$= \frac{1}{N_l}\sum\frac{f\left(R + a_0\sin(2\pi ik/N_l)\right) + f_{sp} - f(R)}{a_0\sin(2\pi ik/N_l)}\,. \tag{5.7}$$

For any given force curve $f(R)$ (or approximation thereof), this equation yields the stability diagram for any free parameters we choose.

Force curve derivative

If we fix the wavenumber k at some number, then $\sin(2\pi ik/Nl) \sim \mathcal{O}(1) \equiv s$, and Γ becomes:

$$\Gamma = \frac{1}{N_l} \sum^{N_l} \frac{f\left(R + a_0 s\right) - f(R)}{a_0 s}\,. \tag{5.8}$$

The procedure above immediately suggests one possible simplifying idea: *in the limit $a_0 \to 0$, the growth rate should approach the derivative of the force curve.*

The k dependence is lost in this limit, but it is still a useful limit in that it gives us a great deal of qualitative information about the stability as a function of R or other parameters, based solely on the derivative of the force curve for that parameter set.

Lennard-Jones force

To simplify our qualitative explorations of the stability of our system, we will not use a fit or interpolation of the force curve found in Chapter 4. Instead, we use a Lennard-Jones type force to model the chain-chain interactions.

This allows us to predict qualitative behavior without getting sidetracked by the precise form of the force curve. As we shall see later, this qualitative behavior does indeed hold, even though the actual chain-chain force is different from Lennard-Jones.

The force we use in this section is given by:

$$F_{mf}(R) = -\frac{24\epsilon\sigma^6\left(R^6 - 2\sigma^6\right)}{R^{13}} \tag{5.9}$$

Using this force in the Γ equation yields, for each i:

$$\Gamma_i = \frac{\csc\left(\frac{2in\pi}{\mathrm{Nl}}\right)}{a_0}\left[f_{sp} + \frac{24\epsilon\sigma^6\left(R^6 - 2\sigma^6\right)}{R^{13}} - \frac{24\epsilon\sigma^6\left(-2\sigma^6 + \left(R + a_0\sin\left[\frac{2in\pi}{\mathrm{Nl}}\right]\right)^6\right)}{\left(R + a_0\sin\left[\frac{2in\pi}{\mathrm{Nl}}\right]\right)^{13}}\right]. \tag{5.10}$$

We can now use the mean-field $\Gamma = \sum \Gamma_i$ equation to plot stability curves for various parameters.

Chain force

Our first parameter variation is the restoring force h of the FENE chain potential. The amount of stretching a monomer experiences is directly related to the frequency of the deformation. The highest frequency mode corresponds to the greatest possible distance between neighboring monomers, since at this frequency the monomers zig-zag or alternate between $x_0 - a_0$ and $x_0 + a_0$. We therefore expect that the restoring force is strongest at high frequencies. Mean-field predictions of the stability, for various chain force strengths h, are shown below.

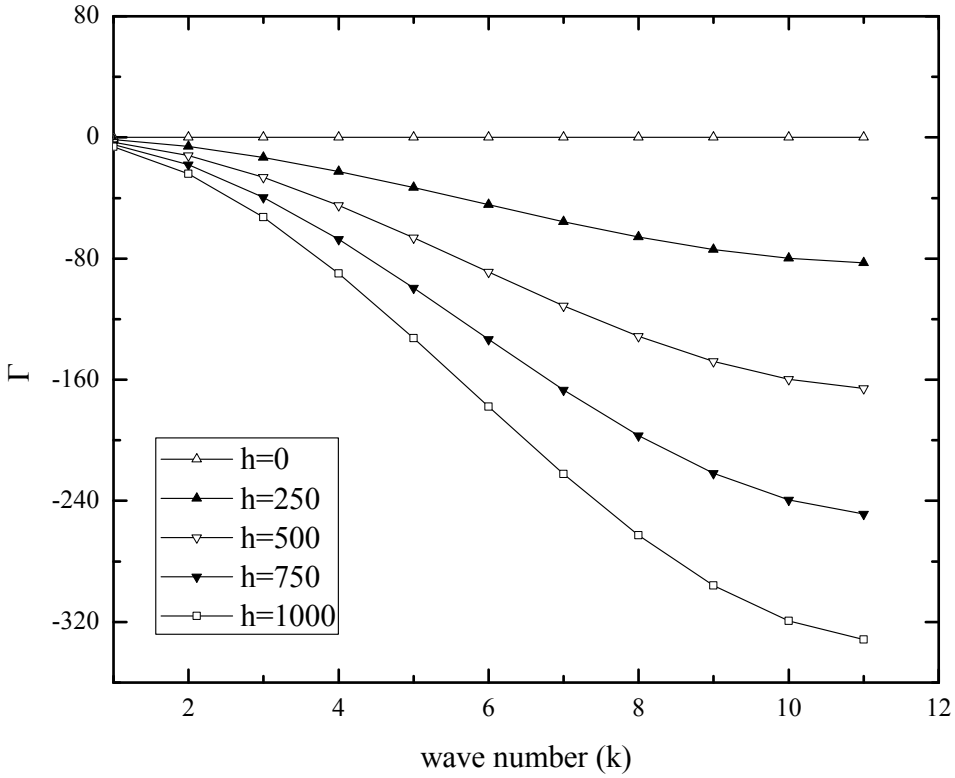

Figure 5.3: Here we see how varying the spring constant h results in more negative (more stable) high-frequency modes.

The fact that the higher frequency modes are most dampened (most negative),

and that the damping increases as the chain stiffness increases, are both intuitive results. However, the precise shape of these stability curves is not trivial, and we will use it as a basis of comparison with our simulation results.

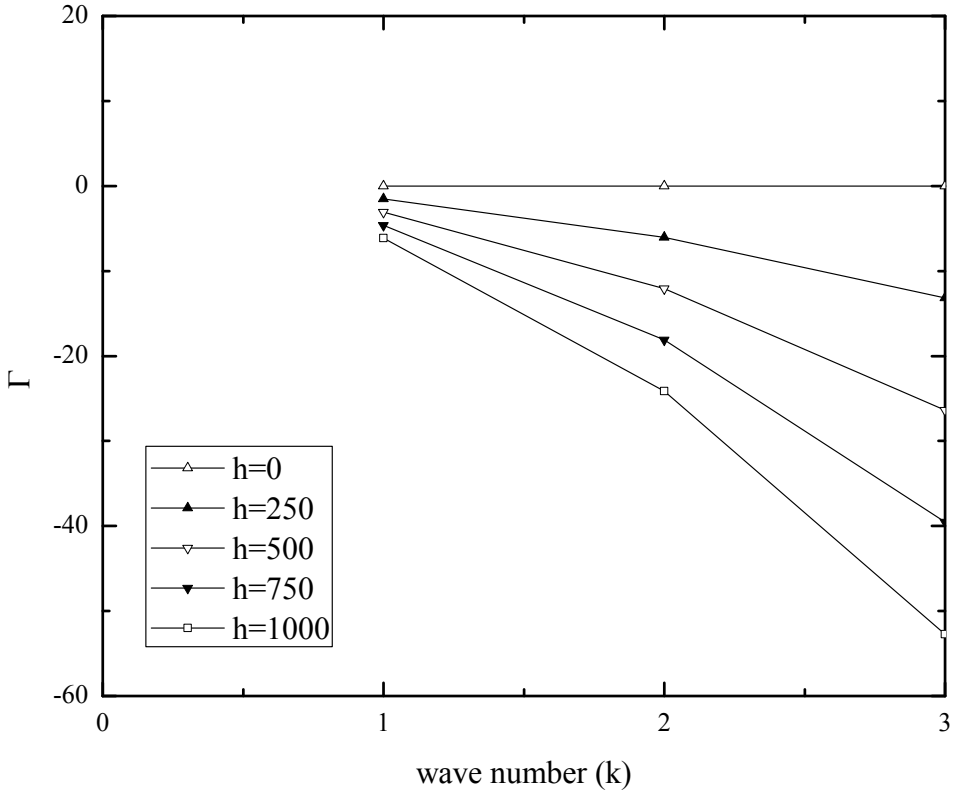

Figure 5.4: Zooming in to the first three wavenumbers, we can see that they are non-zero. We expect $\Gamma \to 0$ as $k \to 0$.

We have also seen that the stability curves approach 0 as the wave number approaches 0. Since we are removing the mean force before producing these diagrams, it is a requirement that the 0th mode be zero.

Interaction energy

Now we investigate how the interaction energy of the chain-chain force affects the stability/instability of the perturbations. Holding all other parameters fixed (e.g. chain-chain separation R, spring force h), we see in figure 5.5 that the effect of the interaction energy is to dampen all modes shown.

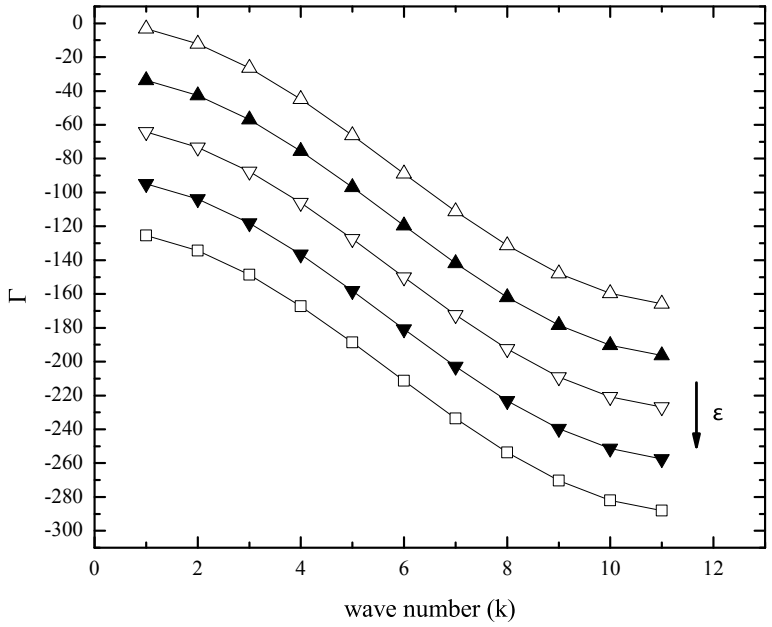

Figure 5.5: For the Lennard-Jones force example, we varied the energy parameter ϵ. The result of this, *for this particular chain separation* R, is to stabilize the perturbations.

However, as we shall see later, increasing the interaction energy does not always stabilize the deformations. More generally, a higher ϵ leads to stabilizing modes that are already stable, but also to further destabilizing unstable modes. In other words, ϵ controls the strength of the effect, whether positive or negative.

Chain separation

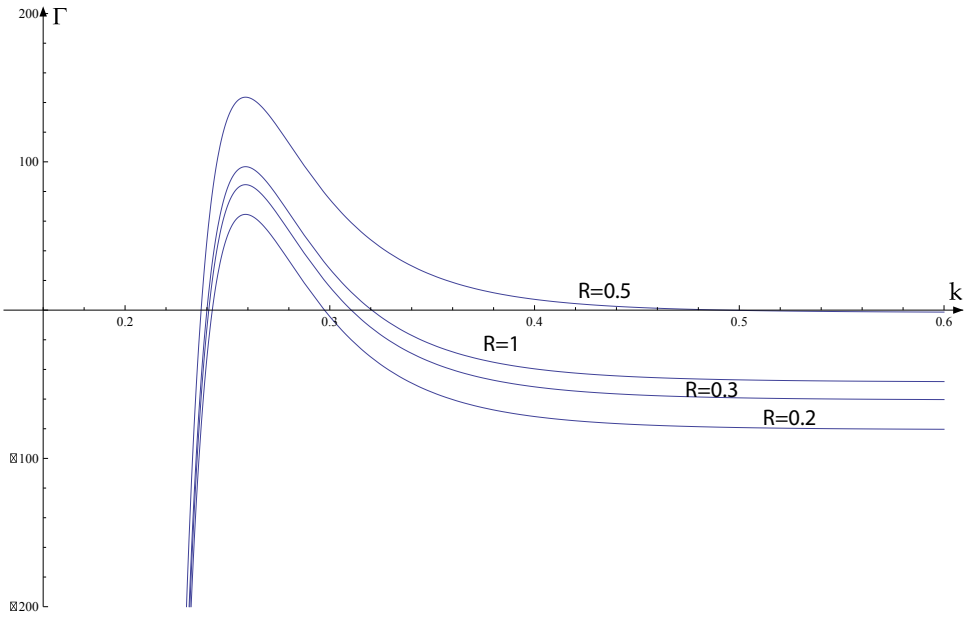

Figure 5.6: Here we hold ϵ and h fixed, and plot the growth rate for various chain separations R. We see that the growth rate needs to be a function of both k and R, since the sign of the growth rate depends on both of these variables.

Stability Diagrams

In the previous sections, we have seen that the linear stability of the chains depends on two variables, wavenumber k and separation R. All others are held fixed at their physiological defaults ($k_BT = 300K$, $h = 0.82nN/nm$, etc.). Since our main concern is the sign of the growth rate, we can summarize what we have learned about the stability of this system in a single diagram.

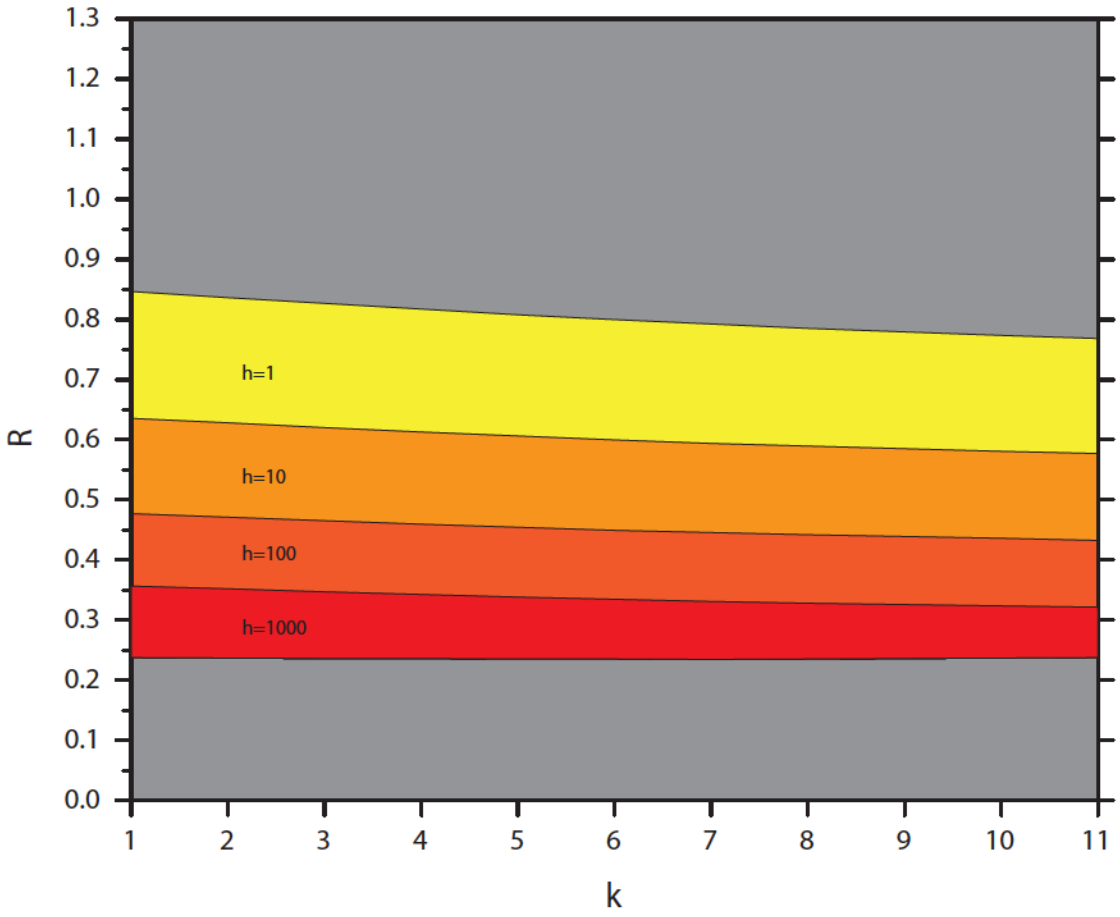

Figure 5.7: Stability diagram with Lennard-Jones potential. Chain separation goes up along the y-axis, wavenumber k along the x. In reds/oranges are the unstable regions. Note that each region contains those with stronger spring force as well. For example, $h = 1$ is a band that extends from $R \sim 0.8$ to 0.25.

This stability diagram shows that there is a band of instability within the attractive $R \gtrsim 0.2$ region. The distance that this band extends depends strongly on the stiffness of the chain. Based on this diagram, we expect that at large separations, chains would approach each other in a parallel and extended configuration. As they enter the instability band, the chains would begin to deform.

The above figure, as with the rest of this section so far, has been using the Lennard-Jones force curve to produce qualitative results. However, because we have found a good analytical fit to the force curve for these parameters, we can also use

that fit as our function in the growth rate equation (Eq. 5.7), to obtain a stability diagram and stability surface that is closer to our real system.

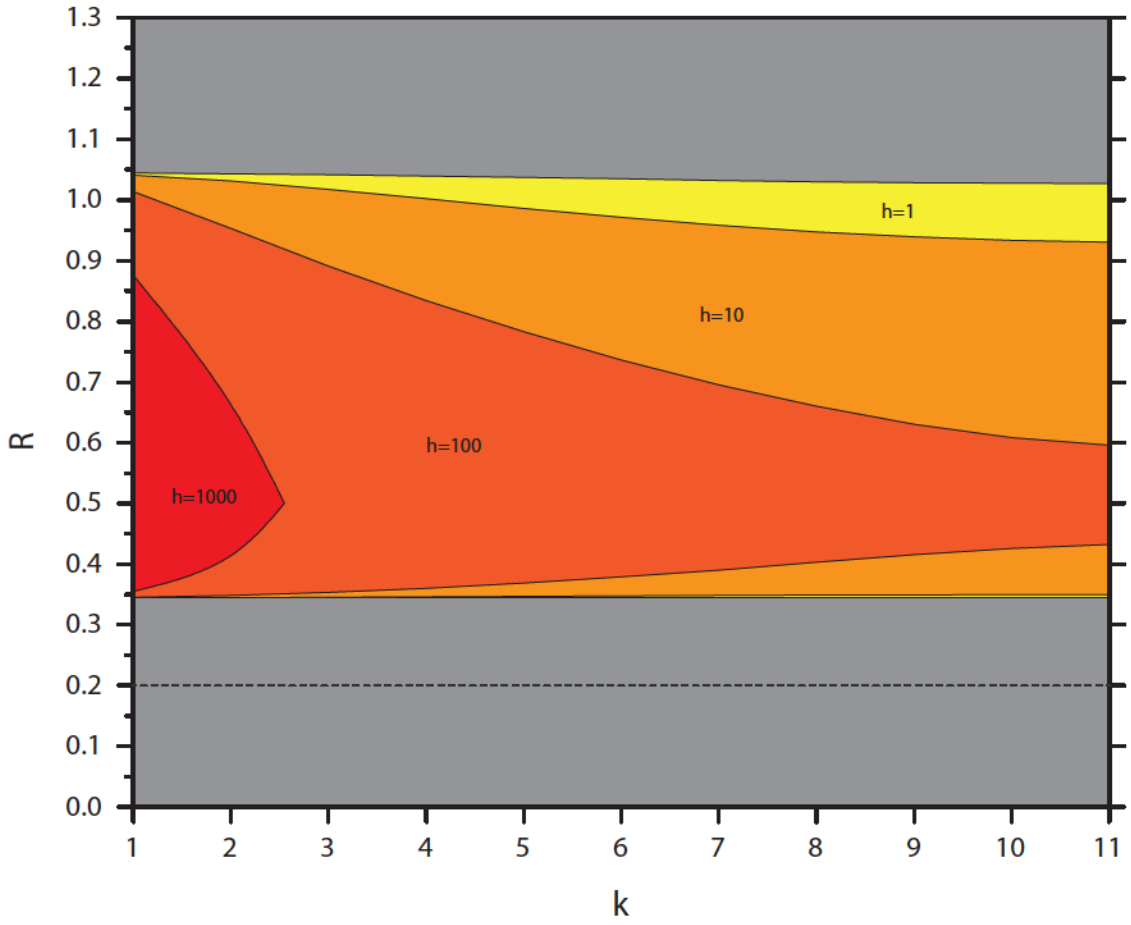

Figure 5.8: Stability diagram with force curve fitting function. The dashed line at $R \sim 0.2$ marks the beginning of the attractive region. Once again, each region contains those with stronger spring force as well.

It is important to note that we are still using the growth rate as approximated by the mean-field model. This approach does not take into account that the nearness of the rod should affect the counterion configuration and therefore the force and growth rate.

Still (as we will see), the general features are in line with our more rigorous results, and we can make a number of observations. It is clear that we are more likely to find unstable modes at low frequency than higher ones, and at low separation rather than

higher ones.

This tells us that the chains are stable while attracted towards each other, and become unstable near each other, beginning with long-wavelength oscillations. Whether this behavior persists beyond the small-amplitude linear regime is not something we can address with the current method.

5.3.3 Summary

With our mean-field model, and based solely on a qualitiative chain-chain interaction, we have gathered a number of predictions for how our simulations might behave. We have seen that the chain-chain separation R is an important parameter in determining the sign of the growth rate, and so we will investigate our system using various R's as initial conditions. We have also determined that the stability curves should be able to be approximated by the derivative of the force curve.

Using our predictions, we have also found hints as to where we might expect to find unstable modes. Namely, at long-wavelengths, with small chain-chain separations, and in situations where the attractive interaction is strong (e.g. multivalent counterions, and not monovalent ones). With this information, we are ready to explore the parameter space of our simulated system, and investigate the stability curves we find.

First, we will describe the three approaches we will use to generate stability curves from our simulations. Then, we will present the results of these simulations. And finally, in the discussion section, we will compare our results to the predictions made in this section.

5.4 Setup

Now that we have a qualitative expectation for the growth rate Γ as a function of the system parameters, we can attempt to establish the growth rate using computer simulations.

5.4.1 Static analysis

As we saw in section X, one way to find the growth rate is by averaging the force-over-displacement ($\Delta F/\delta$) over all particles on the chain. To do this with our simulations, we perturb the chains with a sinusoidal displacement of a particular wavelength (again, one that is L_y/n), and keep the chains fixed in space with this displacement. We then allow only the counterions to move. Every 10 Monte Carlo sweeps, we record and output the force felt by each chain particle.

For each timestep, we need to find out if the force on a given particle is stronger or weaker than the average force on the whole chain. If it is weaker, then when we subtract the mean undisplaced force, the ΔF will be negative. This procedure is identical to calculations of tidal forces: if the far side of the earth is pulled towards the moon with a weaker force than the near side, the effect once the center-of-mass force is removed is a residual outward force on both sides.

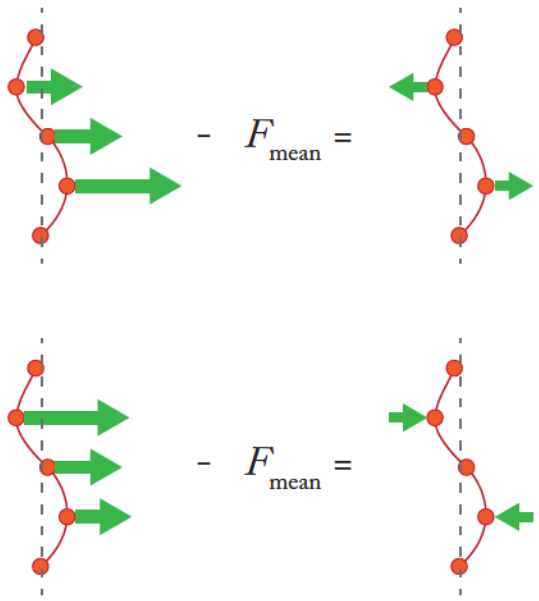

Figure 5.9: Even if the force on all chain monomers were in the same direction (say, attracted towards the other chain), the deformations may still be either stable or unstable. This depends on the force felt by the perturbed monomers after subtracting the mean force. This is analogous to tidal forces.

We then divide the force difference ΔF of each particle by that particle's displacement (which never changes, since the chain particles are held fixed). If we then average this quantity over time[1], we obtain an estimate of Γ, the growth rate, for that particular set of parameters.

This procedure, which we call static analysis, is much simpler than the upcoming dynamic analyses. It is also closer to the mean-field approach we described earlier. However, there is much that we cannot discover with this method (such as the point where the system leaves the linear regime), and so we later supplement this with dynamic alternatives.

[1]Note that the order of averaging is important. That is, $\langle F \rangle / \langle \delta \rangle \neq \langle F/\delta \rangle$. Averages over time and over particle number, however, are interchangeable.

5.4.2 Dynamic Analysis

Amplitude Tracking

An alternative approach to the static analysis described previously, is to introduce a sinusoidal perturbation of wavelength L_y/n, as before, and then allow the chain particles to move. We then track the root-mean-square amplitude as a function of time ($A_{rms}(t)$). If the amplitude grows, then this particular wave number n is unstable; if it decays, this n is stable. The rate of growth/decay can be compared across different n's.

```
For every 10th MC sweep:
    For each particle i on left chain:
        xL += x[i]/Nl;                //calculate avg x of left chain
    End

    For each particle i on left chain:
        sumsq += (x[i]-xL)*(x[i]-xL); //calculate sum of squares
    End

    rms = sqrt(2.0*sumsq/Nl);         //calculate rms

    Output rms;
End
```

This approach is not ideal. First, while a given mode may be decaying, a competing mode might be in the process of growing. At early times (linear regime), this other mode is negligible, but the exact point at which other modes start to wash out the signal of the current mode in question, is hard to determine.

Energy Tracking

We can also manually increase the amplitude from $0 \rightarrow \epsilon$, and track the potential energy of the system. If the potential energy decreases for a particular wavelength, then that mode is favored, and it will be unstable. If the energy increases, the selected mode is stable.

```
initChain(R,0); //initialize at flat, distance R

For every 10th MC sweep:
    For each pair of particles:
        totalE += energy(i,j); //calculate total interaction energy
    End

    Output totalE;

    initChain(R,amp+epsilon); //increment deformation amplitude
End
```

This method easily yields the sign of the growth rate. However, it is not immediately apparent how to go from the increase/decrease in energy into a growth rate. Because of this, and because Power Tracking (next section) yields better results, this method was not used.

Power Tracking

One final dynamic approach that is more accurate than amplitude tracking, is power tracking. As before, we start the chain with an initial small perturbation of a particular wavelength, and then allow both the chain particles and the counterions to move freely.

Instead of tracking the amplitude itself as a function of time, we take the Fourier transform of the position data, and analyze that. Because the system starts at one particular wavelength, the power of the Fourier transform will have a peak at the corresponding frequency, and be zero elsewhere. If the mode is unstable, that peak in power will decay exponentially. If it is stable, it would increase in power.

We initialize the system with a particular wavelength, and run n_{cp} random seeds of this. At each timestep, we output the coordinates. All coordinates, for all particles and all random seeds, are combined into a single results table. This results table is input into Matlab for analysis. The first step is to load the file, extract the x coordinates as a 1D array, and reshape it into an $N_\ell \times n_{ts} \times n_{cp} \times n_{ks}$ matrix.

```
function [gamma, ks, fxt] = autostab(filename, Nl, n_ts, n_cp, n_ks)

A = load(filename);
x= A(:,2);
clearvars A;

xtc = reshape(x,Nl,n_ts,n_cp,n_ks);
 // x coords in matrix of
 // Nl  : particle index
 // n_ts: timestep
 // n_cp: random seeds
 // n_ks: initial perturbation
```

We then find the mean x coordinate, so we can get the deformation $\delta = x - x_m$ away from the mean. We then obtain the power spectrum of δ, and then median of this over all random seeds.

```
xm = squeeze(mean(xtc,1));
xma = zeros(Nl,n_ts,n_cp,n_ks);
for i = 1:Nl
  xma(i,:,:,:) = xm(:,:,:);
end

fxtc = abs(fft(xtc-xma,[],1)); // power spectrum
fxt = squeeze(median(fxtc,3)); // median over seeds
```

We then extract the k values from the data (max peak at 0th timestep). We track the evolution of this peak over time, and fit to an exponential. The exponent gives us the growth rate we are seeking.

```
ks = [];
for i=1:n_ks
  m_index = find(fxt(:,1,i) == max(fxt(:,1,i)))-1;
  m_index = m_index(1);
  ks = [ks m_index]; // get k values
end
ks = ks(1,:)

gamma = nan*zeros(1,n_ks); // initialize gamma w/zeros
```

```
    for i=1:n_ks
      // for each k value, find a git to an exponential
      e = fit([0:(n_ts-1)]',fxt(ks(i)+1,:,i)','exp1');

// growth rate for that k value given by fit
gamma(i) = e.b;
    end

 // sort k values, add units
    [val ind] = sort(ks);
    ks = sort(ks);
    ks = 2.*pi.*ks./(Nl*0.34);
    gamma = gamma(ind);

    figure
    plot(ks,gamma)
```

As we will see, this method, though more difficult to implement, yields better results than amplitude tracking. It does not suffer from the washing out effect of amplitude tracking, since we can see each mode independently. Also, we can always track the height of the initial primary peak (the excited mode), and compare it to the other modes.

5.5 Results

Previously we described several methods we could use to determine the stability curves of this system. Here we will examine the results of each of these methods in turn. In the next section, we will discuss what these results imply for the long-term or large-scale structure of the polyelectrolyte systems, as well as compare these results to the predictions we made using the force-curve.

5.5.1 Static Growth Rate

In holding the chains fixed at a given amplitude, while outputting the average force felt by each monomer, we obtained a rough estimate of the growth rate Γ. In

all parameter sets used, the overall shape of the Γ surface remained the same. The only differences we found were in 1) the strength of the high-frequency damping, 2) the overall scale of the effect, 3) the zero-point of the surface.

Γ surface

Shown below is a Γ surface as a function of the perturbation wavenumber and the chain-chain separation. Despite the poor statistics and jagged results, we can see many of the features that persist when we obtain stability surfaces by more complex means. For instance, there is the flat region at large separations (right edge of diagram). The instability rises slowly and peaks at around $R \sim 0.4$nm, and then drops rapidly as the chains approach each other further.

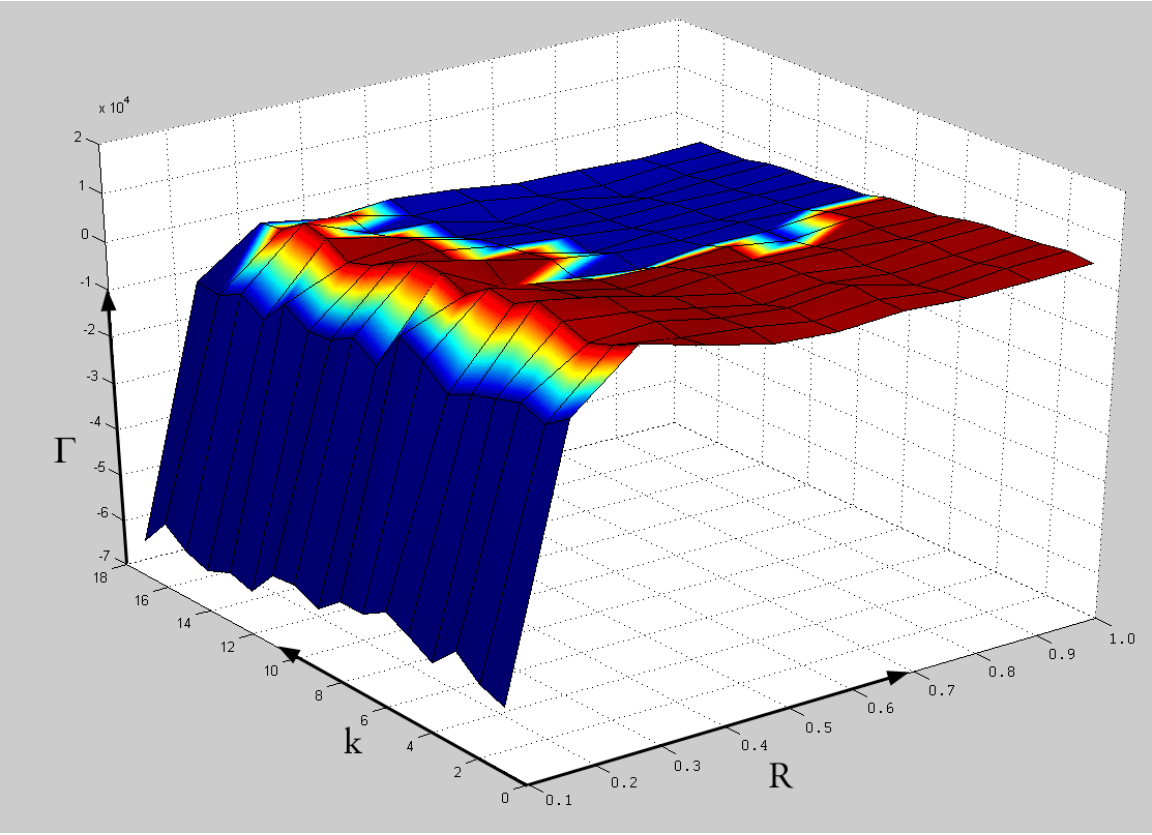

Figure 5.10: Here we see the stability surface of our system, based on the static analysis approach. Despite the poor statistics, the general shape is consistent with more involved methods.

The key point is that there is a positive region on the surface. These are the regions of interest, because it is here that the polyelectrolyte will tend to deviate from an extended configuration and form other (possibly non-linear) shapes. The fact that the positive regions are mostly in the intermediate-separation/long-wavelength corner imply that it is only when the chains are close that a long-wavelength structure begins to form.

5.5.2 Dynamic Growth Rate

We now present the results of our dynamic approaches, where we initialized the chains in sinusoidal configurations, and then let them evolve freely. Our aim for this section is to present all the results that lead up to our core results, the various stability curves and stability diagrams. We focus mostly on the data and the chain of reasoning. Discussion of these results is left to the next section, 5.6, where our focus will switch to the implications of these results.

Amplitude Tracking

The rms deviation from the average axis position, $A_{rms} \equiv \sqrt{\frac{1}{N_l}\sum (x_i - \langle x \rangle)^2}$, is an easy to calculate metric. Over a short time, this amplitude will tend to either grow or decay, depending on the wavelength. However, after long enough times, the amplitude will simply fluctuate about an average variance (one that is dependent on the temperature, the chain stiffness, amongst other parameters).

To be sure that we are tracking the growth/decay of only one mode, we were very conservative in our choice of cut-off time. As we can see, the growth or decay of the modes seem essentially linear, as opposed to exponential. What we are seeing is the first-order linearization of the exponential decay/growth.

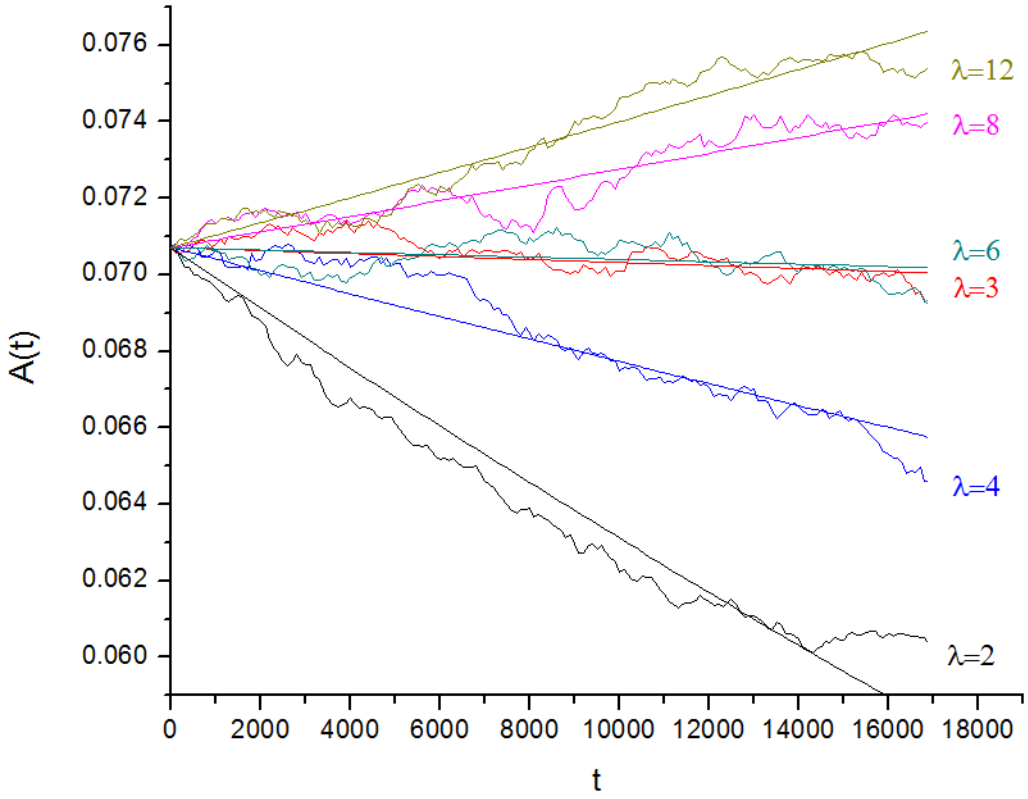

Figure 5.11: Here we plot the rms amplitude of a chain over time (MC time steps). The chains were perturbed with a small amplitude, and wavelengths $\lambda = 2, 3, 4, 6, 8, 12$. As we can see, some of these wavelengths lead to a growing amplitude, others to decay. The slope of the lines give us the growth rate Γ for that wavelength.

With the fits we obtain in the above figure, we can obtain the exponential rate constant (growth or decay) for each of the wavenumbers we simulated. Plotting these growth rates as a function of wavenumber yields our first simulation stability curve:

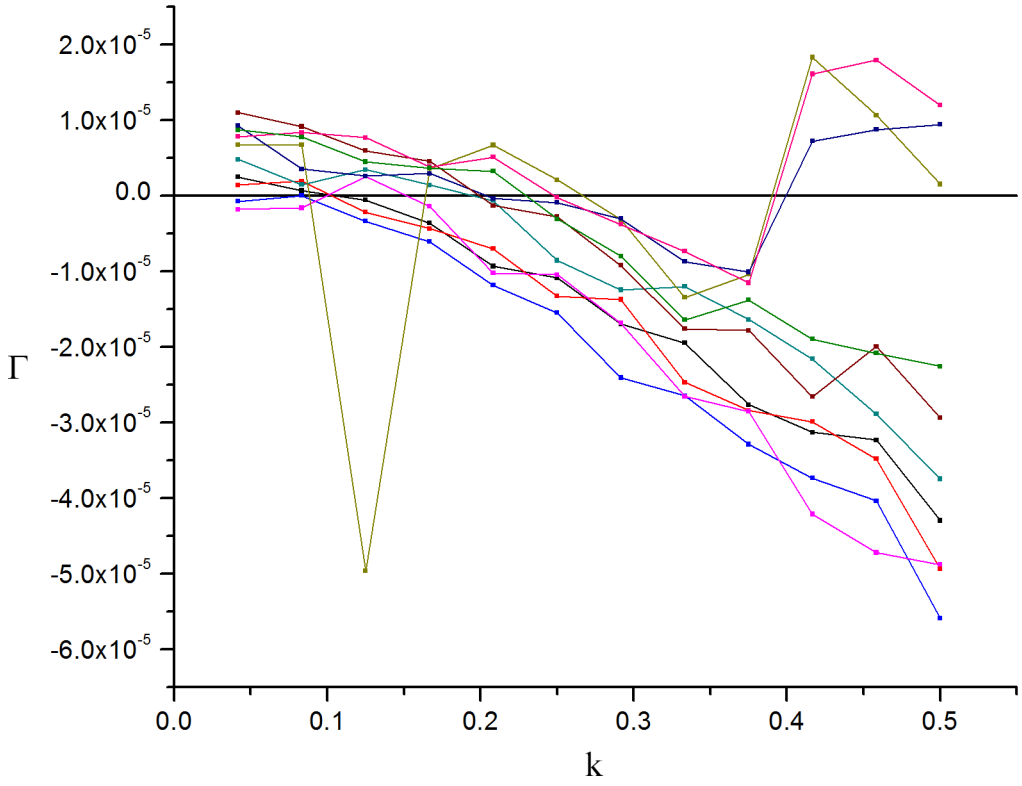

Figure 5.12: Stability curve, amplitude variation. Here we see several individual runs. The variance is large, because of the drawbacks of the amplitude tracking approach. Also, there are several clear outliers in these results.

Once again, we see the characteristic s-curve we have seen in other approaches. With this approach, we can also see there are a number of outliers, and that each run produces different results within a wide band. To determine the most likely behavior, we averaged the above curves (each of which was itself an average of several hundred random seeds). The resulting average is smooth, and represents the stability of our chains against deformation for this parameter set.

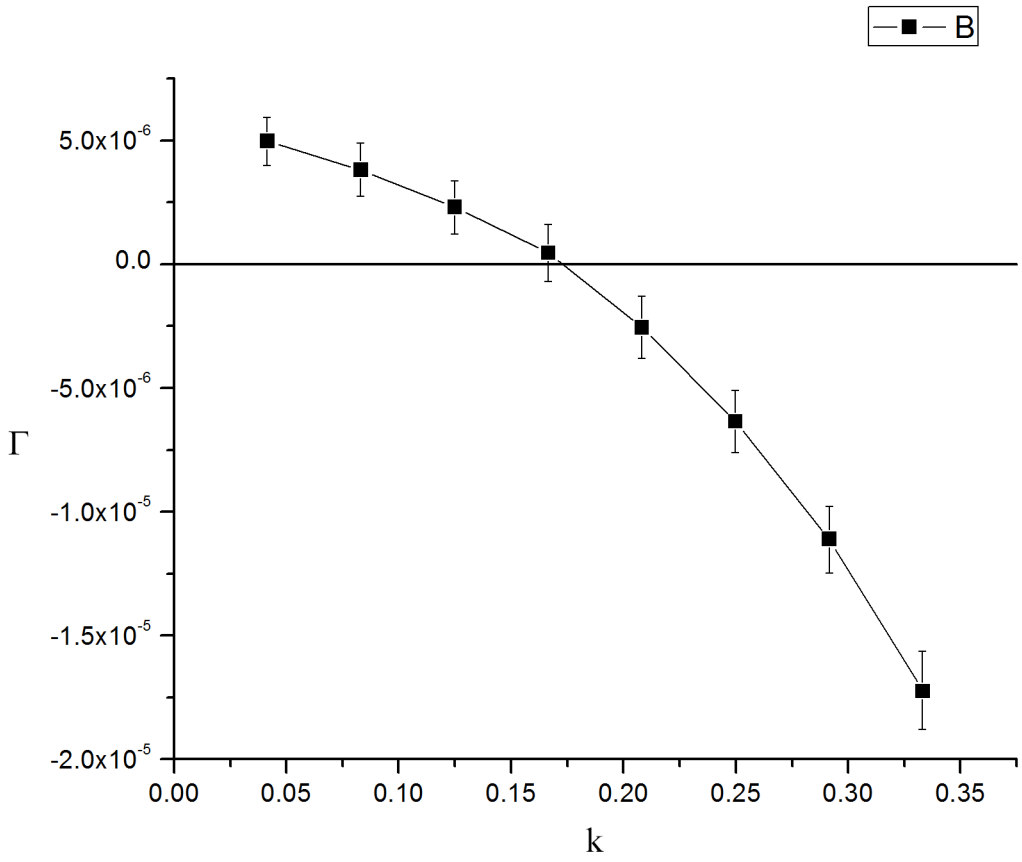

Figure 5.13: Averaging over a large number of initial configurations yields a very smooth growth rate curve. For this particular system ($R = 0.355$ and $h = 100$), we see that the transition between the stable (negative) and unstable (positive) regions occurs at $k = 0.172$, which corresponds to a wavelength of 36.5 nm, which is comparable to the persistence length at this stiffness. Note that, due to outliers and poor data in the high-frequency region, we have only plotted the first 8 points here, and are missing the upturned tail of the rest of the stability curve.

Once again, the shape of this is in line with what we expected from our mean-field approach. If we now take the same procedure and apply it to varying parameters (albeit with fewer statistics), we can see how parameters affect the result. In this case, we varied the chain stiffness h.

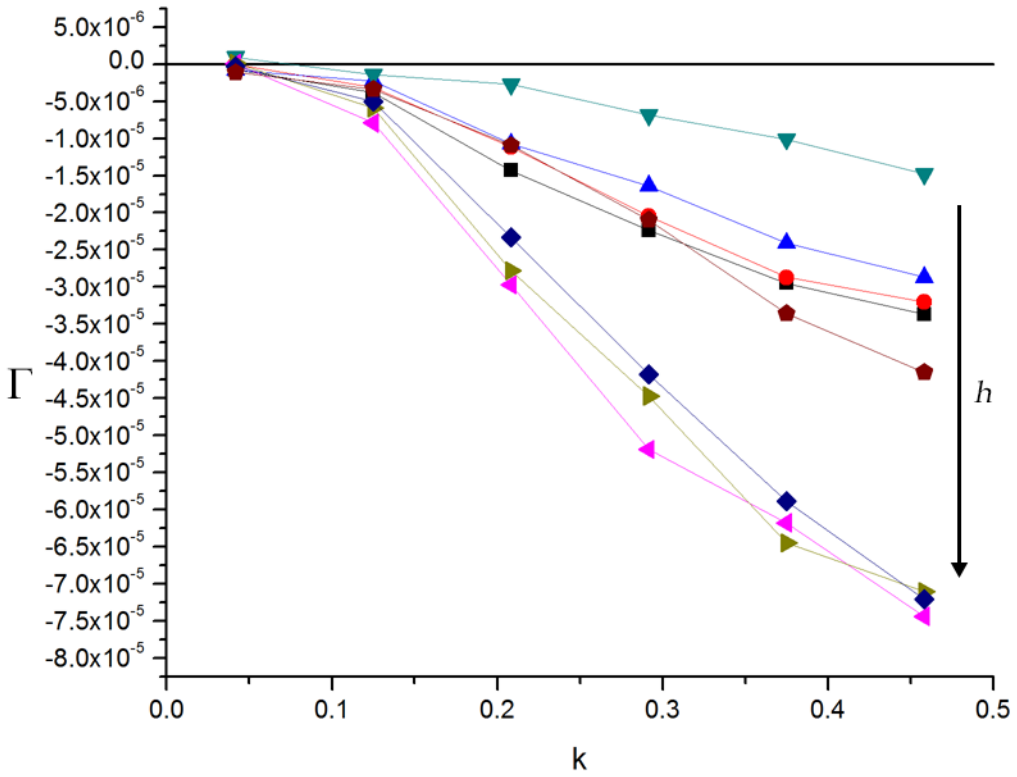

Figure 5.14: Growth rate as a function of wave number, for various spring constants. As expected intuitively, as the spring constant h increases, the more quickly any high-frequency perturbations will decay.

While these results are valid stability curves, the very large number of independent random seeds that were required to produce smooth averages led us to look for alternative ways of generating stability curves.

Energy Tracking

Here are the results from increasing the perturbation amplitude manually from 0 to ϵ over the course of T timesteps. The goal was to see whether the potential energy U increased or decreased as ϵ increased (i.e. $\mathrm{d}U/\mathrm{d}\epsilon$). While we did find that long-wavelength modes were stable and high-frequency modes unstable, two factors prevented us from proceeding towards a stability curve based on these potential energy measurements.

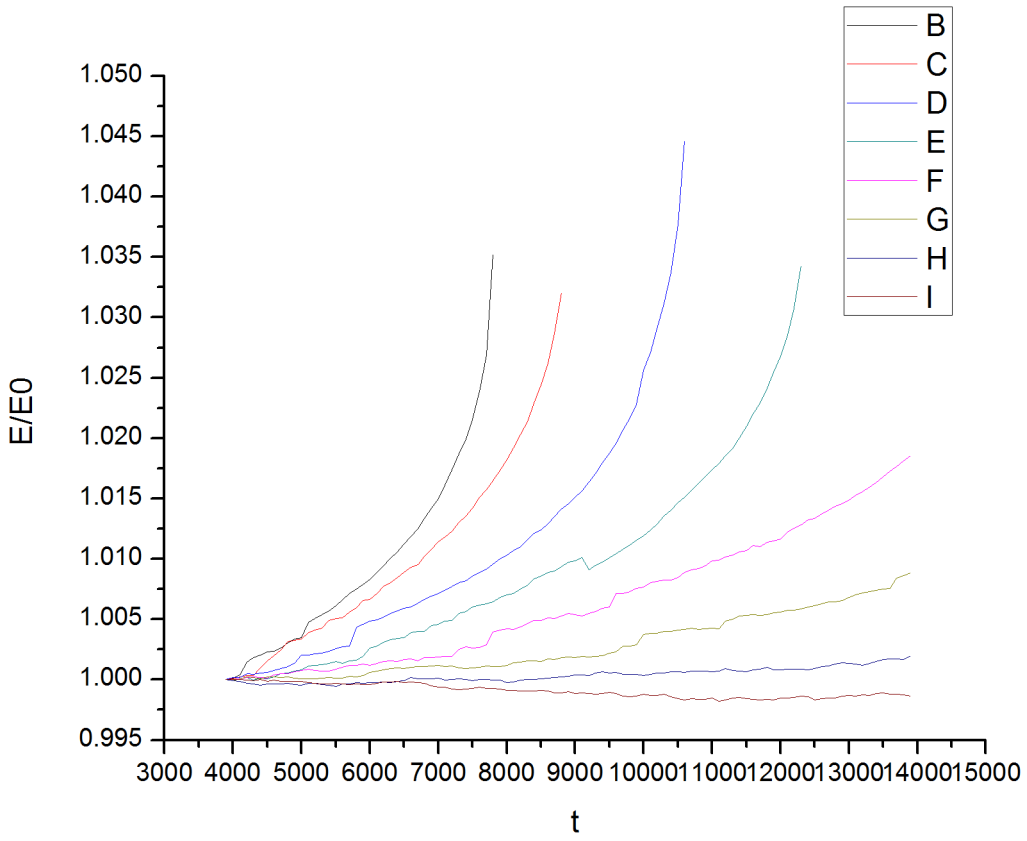

Figure 5.15: An alternative approach to determining the stability of a particular wavelength perturbation. Here we manually increase the perturbation amplitude, and plot the energy. If the energy increases, then the wavelength is unstable. If the energy decreases, the mode is stable.

First, the mapping of $\mathrm{d}U/\mathrm{d}\epsilon$ (a rate of change in potential energy vs amplitude) into a growth rate constant (change of force over displacement), is not trivial. And secondly, we needed to either increase the amplitude very slowly, or to wait for equilibration after each increase step. Both of these procedures resulted in long run-times for our simulations, making them prohibitively expensive in terms of CPU time.

Power Tracking

We started the chains with a small amplitude perturbation of wavelength λ, and allowed the chains to move freely. Here we see a fourier analysis of the displacement

from the mean, as a function of time. In other words, we have the spectrogram of the chain displacement.

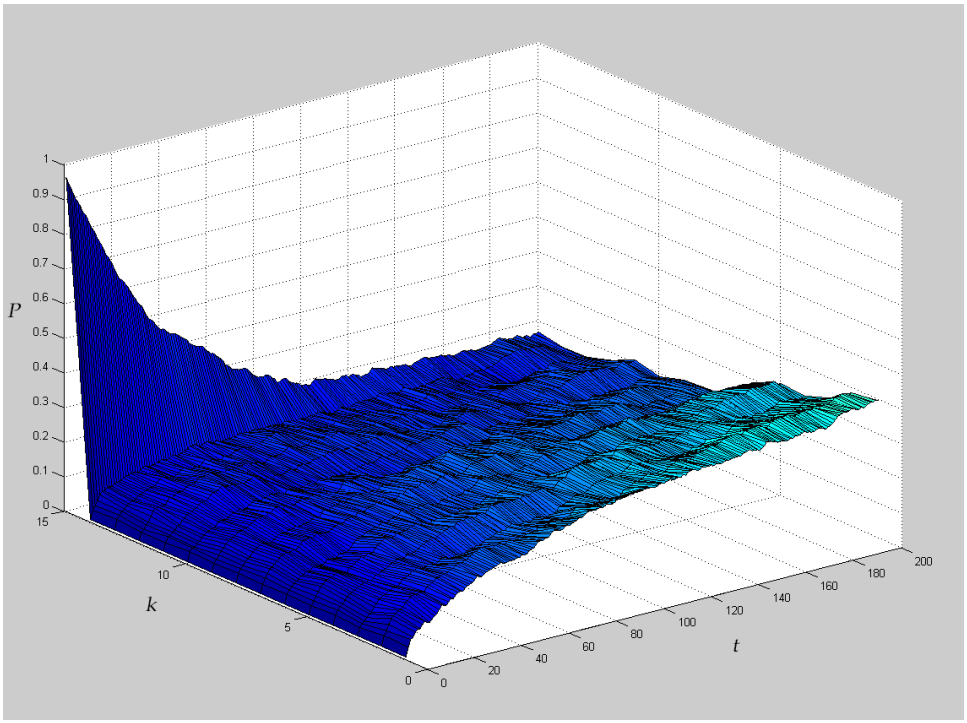

Figure 5.16: The spectrogram of the chain, perturbed at the highest possible frequency. We can see that the initial power spectrum (along bottom-left axis) shows a single peak at the maximum frequency, and zero power along other modes. As time advances (right-ward), this peak decays.

We now have a fourier power spectrum as a function of time. This tells us at a glance, which mode is dominant (if any) at any given time. We can also see that, even when we start with full power in one mode and none in any other, eventually some other modes begin to grow.

Each spectrogram we see here is actually the average of many hundreds of spectrograms, each using the same excitation mode, but using different random seeds for the initial configuration of the counterions and the initial perturbation phase.

In order to find the linear region, we plotted the ratio of the highest mode to the second highest. Since the highest mode at first will always be the initially excited mode, this tells us the ratio in power between the excited mode and the other modes.

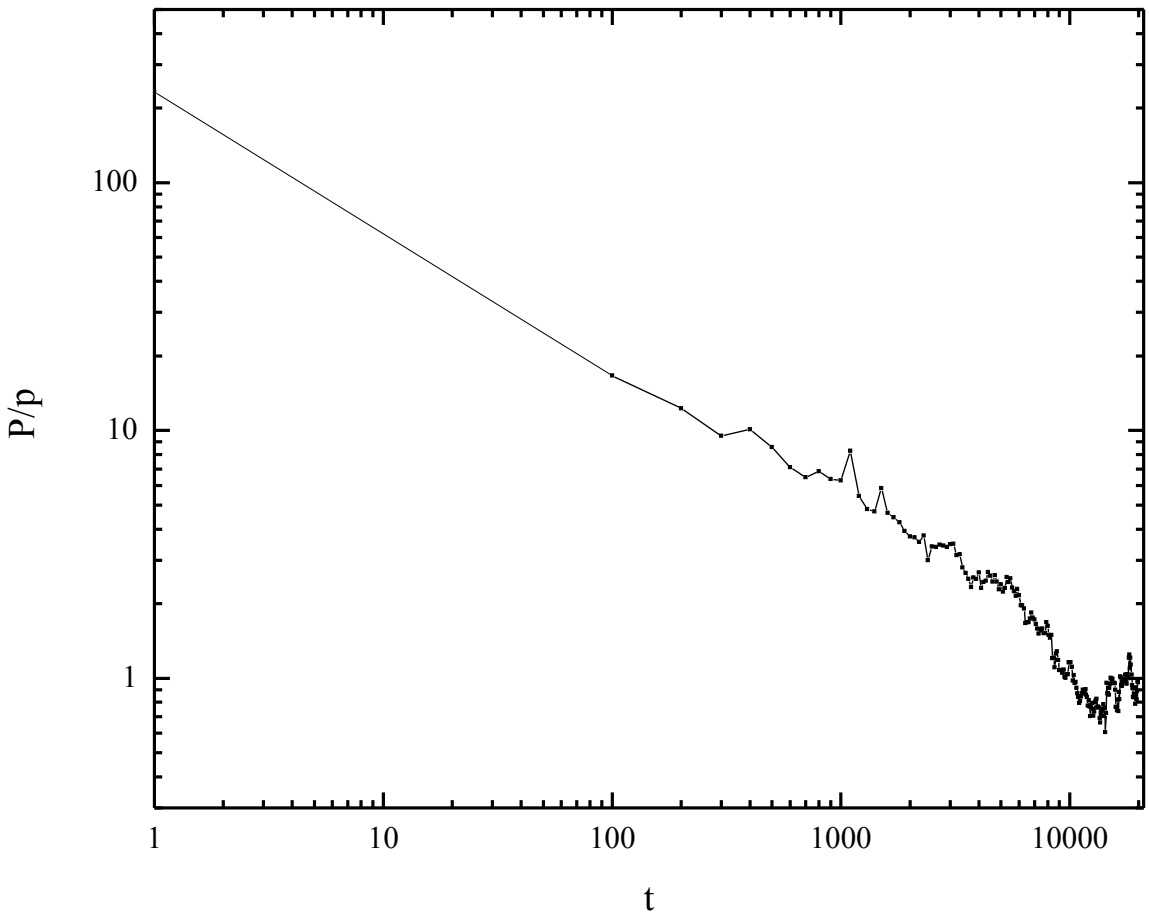

Figure 5.17: Ratio of initial excited mode to the second highest mode. We restrict our attention to $t < 500$.

Our cutoff is chosen so that the excited mode is still $10\times$ the next highest mode. This is what we define to be the linear regime for our purposes.

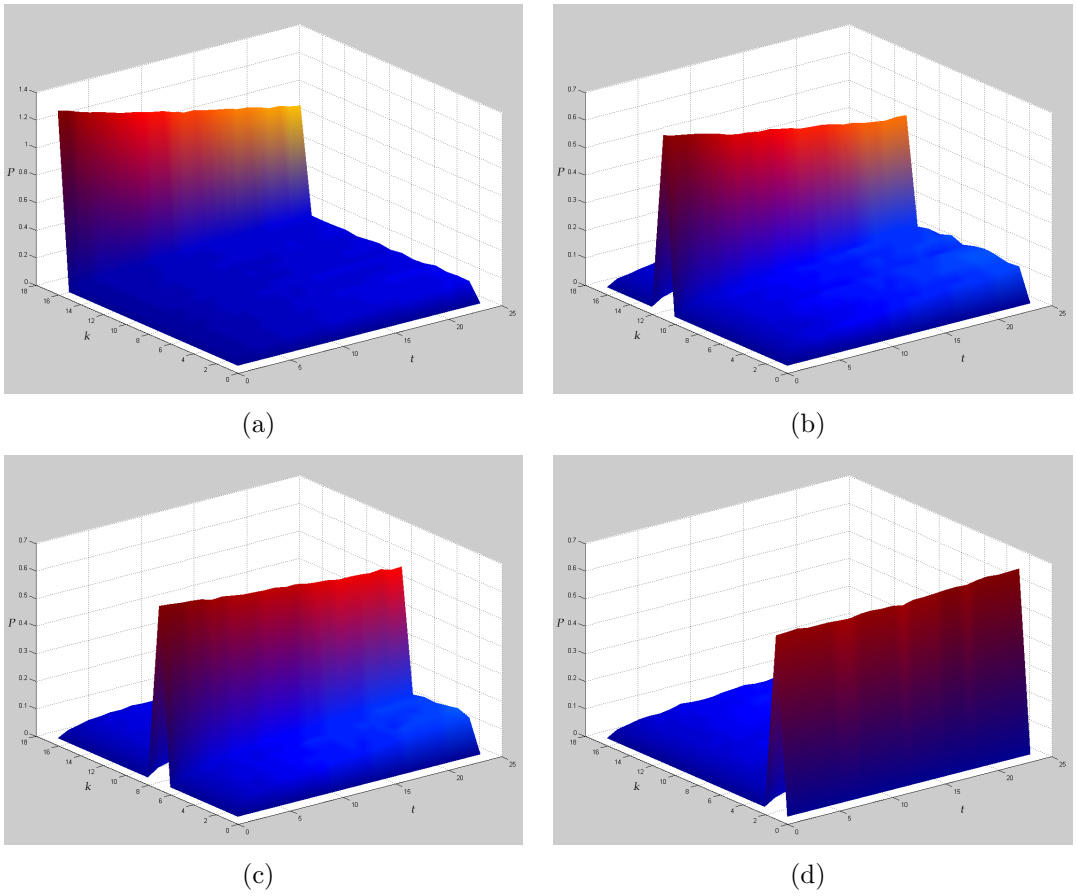

(a) (b) (c) (d)

Figure 5.18: Four example spectrograms, each with a different excited mode. Here we can see that the high frequency mode in (a) decays rapidly. Though not as clear, in figure (d) we have a low-frequency mode that is increasing.

For each spectrogram, we have a single excited mode that either grows or decays. Once again, it is the rate constant of that exponential growth/decay that we are after. For each spectrogram we isolate the dominant mode, and perform a fit to exponential, as seen in Fig. 5.19.

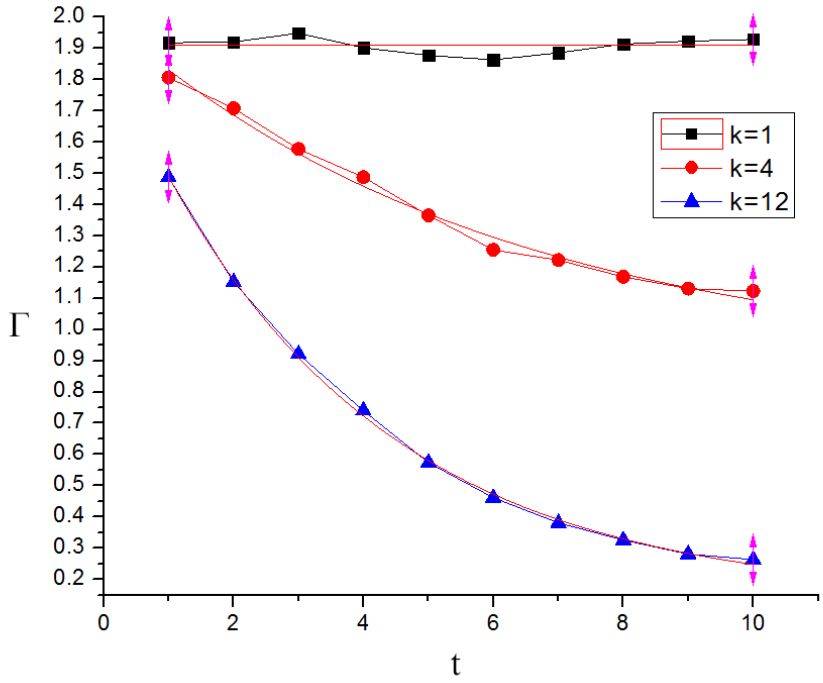

Figure 5.19: Growth or decay rate of for three example excited modes ($k = \{1, 4, 12\}$). An exponential fit is shown for each. The rate constant of this exponential is the growth rate Γ.

Performing this fit for all possible modes in the system, and plotting the rate constant as a function of either k or R yields the stability curves for those variables. We can see in the $\Gamma(k)$ results the characteristic curve due to the restoring chain force, and in the $\Gamma(R)$ results we see the influence of the force gradient.

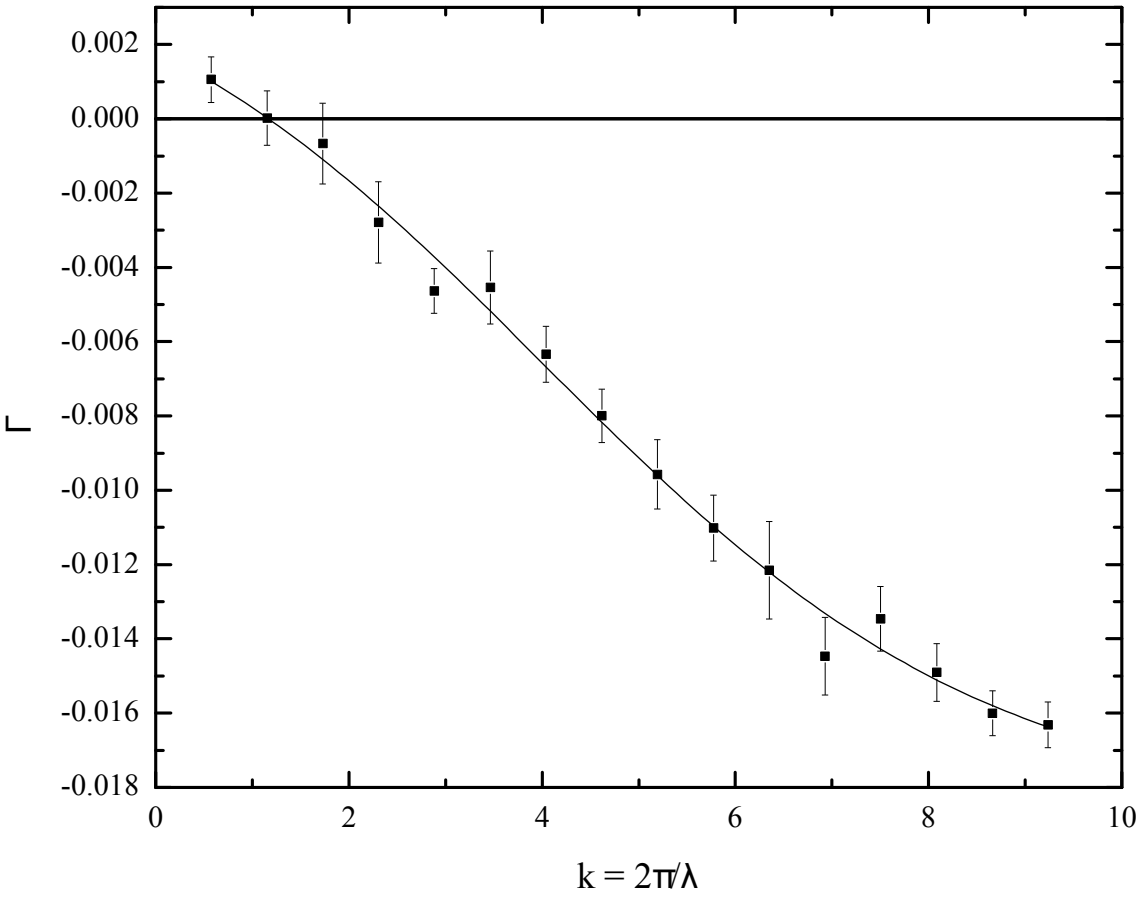

Figure 5.20: Here we have the growth rate as a function of wavenumber for fixed R. The s-curve shape is the same as seen in other methods.

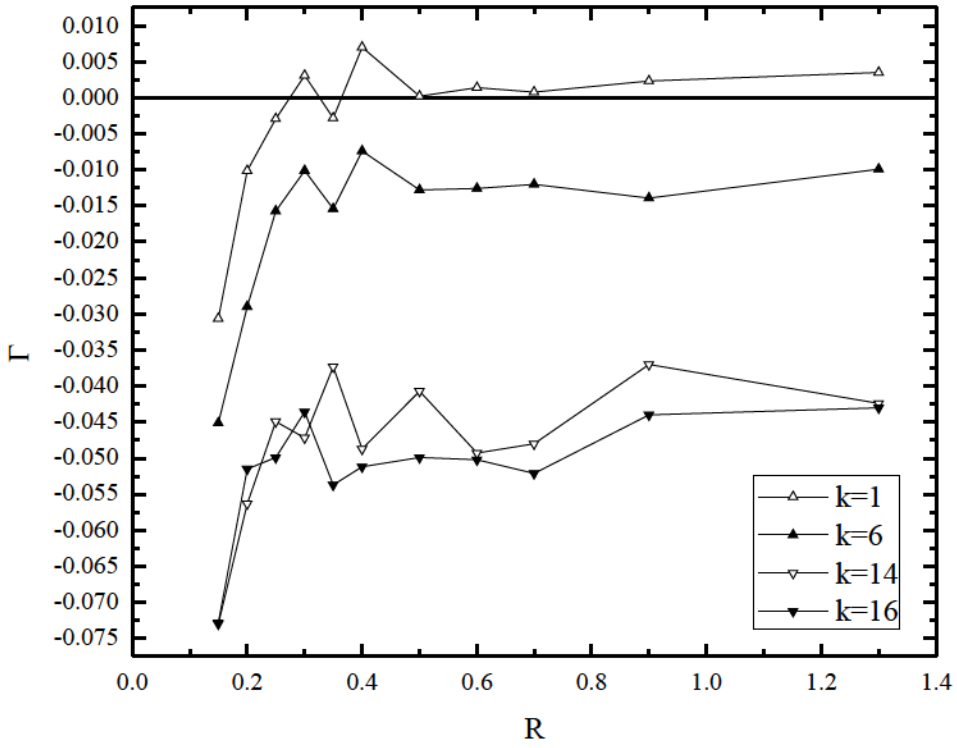

Figure 5.21: Growth rate as a function of R for four different wavenumbers. As expected, the higher wavenumbers are more stable (more negative). The peak at $R \sim 0.45$ means that it is at that separation that the polyeletrolytes are most likely to deform. For $k = 1$, there is a broad range over which the growth rate is positive and unstable.

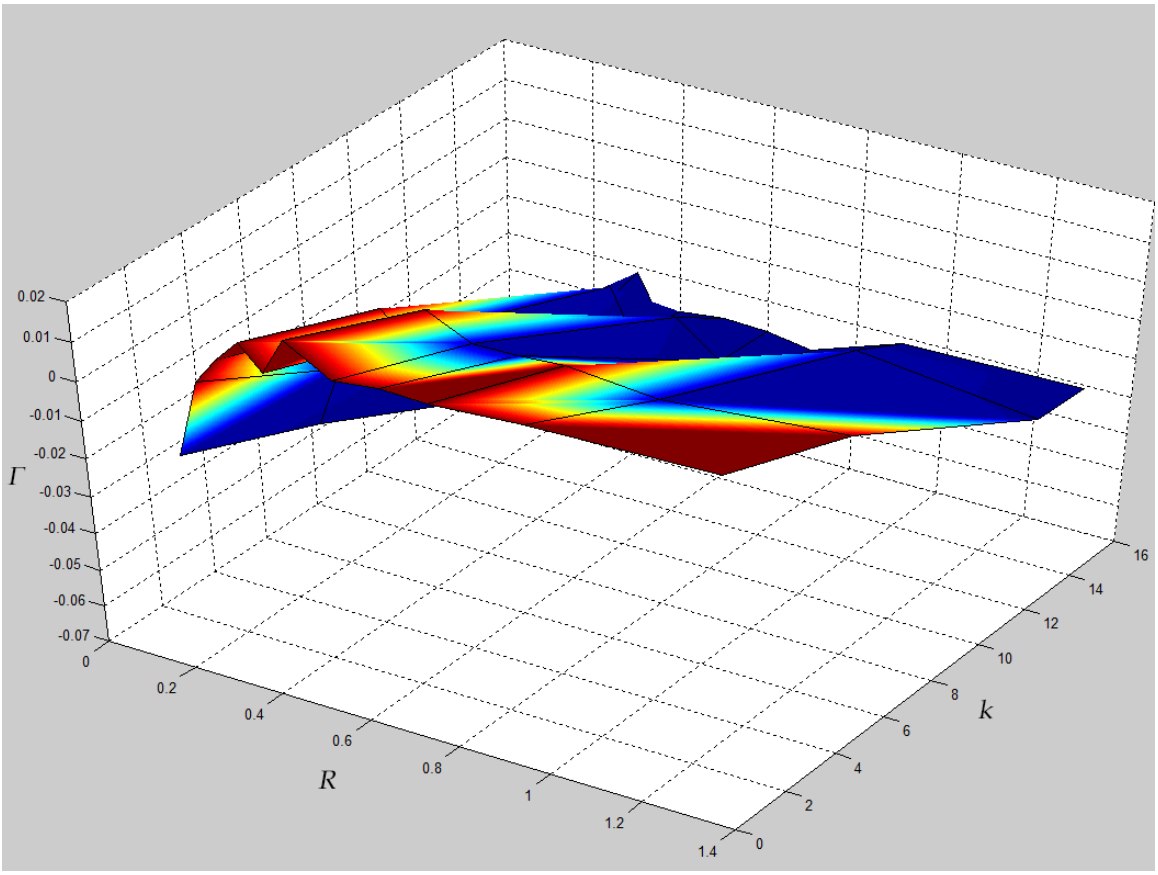

Figure 5.22: A stability surface generated by combining the four curves from the previous figure. Once again, red denotes instability, blue stability.

In these power tracking results we have seen some features that were missed with the other methods (but not, as we will see, with the force curve derivative). For instance, there is a persistent peak and valley pairing near the growth rate maxima at $R \sim 0.45$.

5.6 Discussion

Force curve derivative

In section 5.3.2, we claimed that $\mathrm{d}/\mathrm{d}R$ of the force curve was a good guide to the qualitative shape of the stability curve as a function of R.[2]

[2]Recall that the force-derivative approach discarded the wavenumber variations.

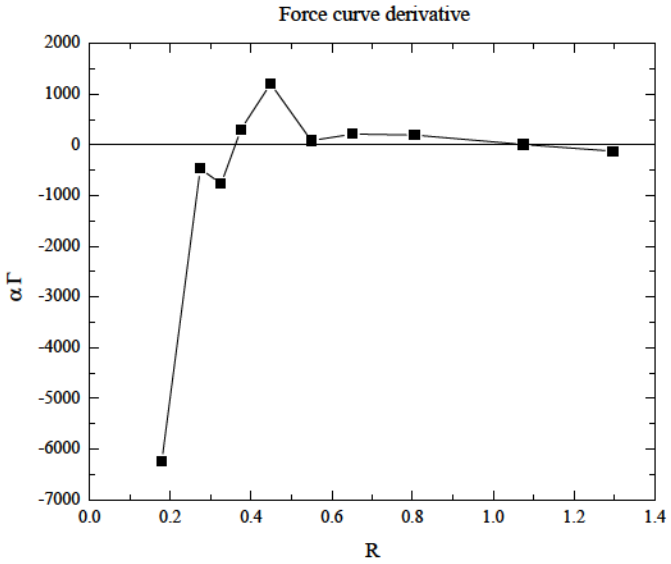

(a)

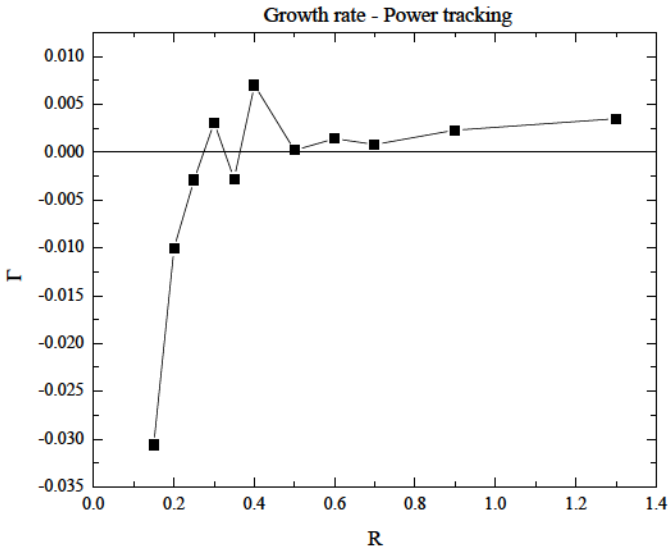

(b)

Figure 5.23: Here (a) is the numerical derivative of the force curve from Ch. 4, and (b) is the growth rate as a function of R, as found with our Power tracking results. The agreement is clear, even down to the presence of a sharp peak at $R \sim 0.45$. We however, do not have a concise explanation for this peak under either method. One possibility is that there is a semi-stable region at a radius just before the peak, which leads to a greater contrast with the unstable peak itself.

In Fig. 5.23 we can see that the agreement between the force curve derivative and the growth rate results from simulations is very good. Minor differences are apparent, however. For instance, the large-R behavior is not the same. We have not been able

to investigate whether these minor differences would persist if we had better statistics for both the force curve and the power tracking simulation.

5.6.1 Mean-field with real force

In our qualitative predictions, we had used a Lennard-Jones-type interaction between the chains to make our predictions. Here we will use the same procedure used to produce the stability diagram of section 5.3.2, but instead of using a L-J force or a fit, we will use an interpolation of the actual force curve from Chapter 4.

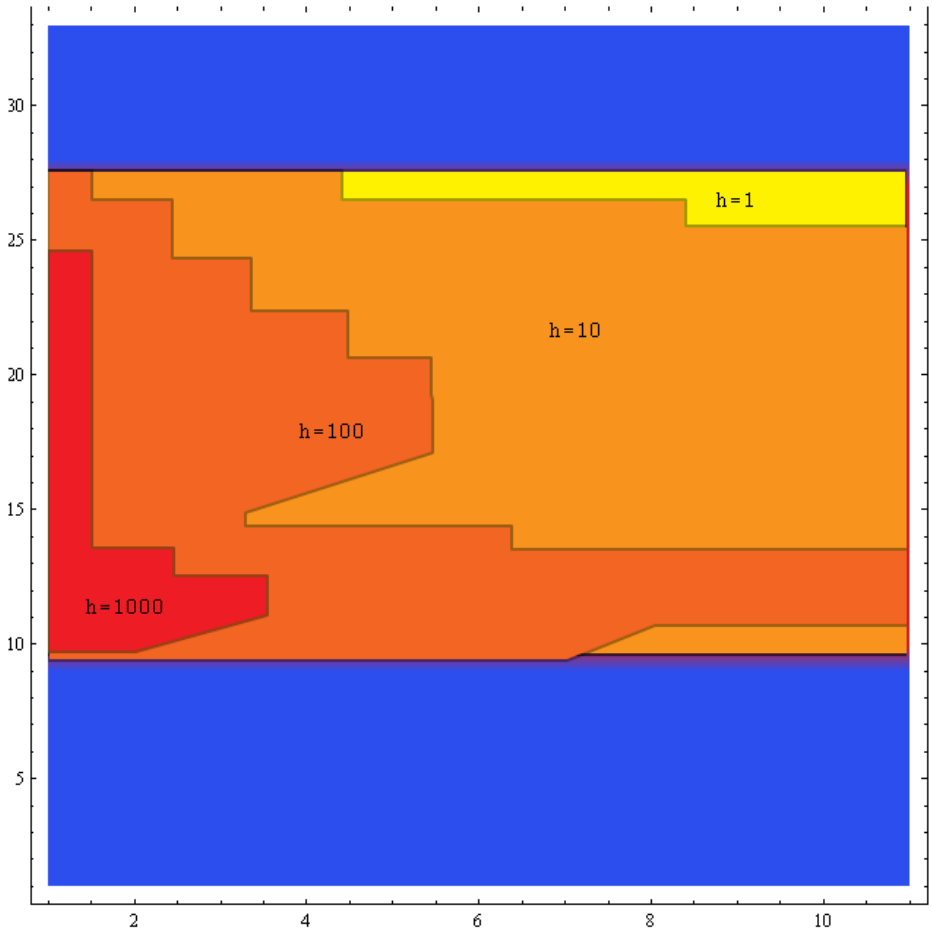

Figure 5.24: As before, the vertical axis is the separation R, and horizontally we have the wave vector k. Reds and oranges show the positive regions of the growth rate for various spring constants. The staircase pattern here is an artifact of the plots used and is not a feature of the system.

Because the stability diagram is related to the derivative of the force curve, the details of this stability diagram are sensitive to the exact interpolation method used on the force curve data (in this case, cubic B-spline).

Unlike other stability diagrams so far (using the Lennard-Jones force and the fit), this diagram shows regions where a chain enters a region of instability on approach, then leaves it, and then enters another one (see $h = 100$ region, at $k = 4$). This is due to a the peaks and valleys seen in the force curve derivative of the previous section (and in the stability curve result it is compared to).

5.6.2 Changing spring force

We have seen that our mean-field approach based on the force curve, predicts a strong k-dependence when the chain force is strong. This makes sense, since the higher frequency perturbations have a more rapid change in displacement from one monomer to another, and therefore a greater distance and greater restoring force. Though we expect, based on our mean-field results and on intuition, that the chain would be more stable for high frequency perturbations, the shape of the stability curve might have been very different from our predictions.

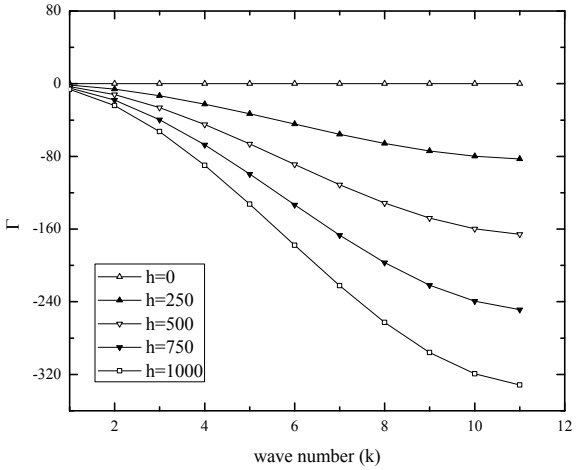

(a)

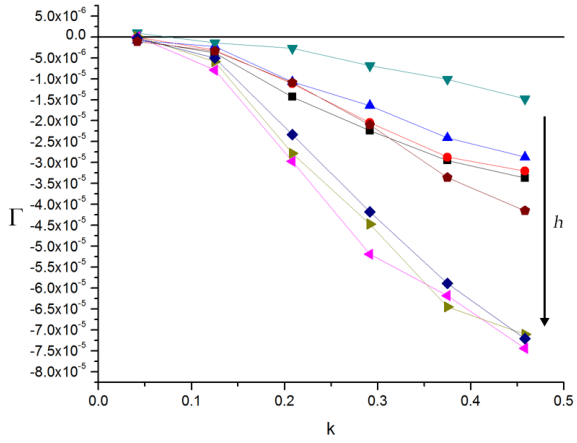

(b)

Figure 5.25: Comparison of our prediction (a), and simulation result (b). The growth rate as a function of k has a sigmoidal shape in both, and the strength of the high-frequency damping is related to the chain stiffness h.

As we can see from the comparison of our mean-field predictions, and of our simulation results, the k-dependence behavior is the same. This verifies that our mean-field approach is at least qualitatively correct. It also strengthens the argument that we should use low-frequency perturbations and low stiffnesses if we are trying to excite the chains to form interesting structures in the long-term.

5.7 Method comparison

In this chapter, we have attempted to determine under what conditions flexible polyelectrolytes are stable against small deformations. To determine this, we used various methods of increasing complexity. We review these here.

First, we perturbed the chain, held it fixed at that given amplitude, and calculated $\Gamma = \langle \Delta F/\delta \rangle$. From this we generated the Γ surface as a function of R and k, paying close attention to the sign of Γ. This method was coarse, and produced noisy results despite a large number of runs. Second we tracked the root-mean-square amplitude after exciting the system in a single mode. This procedure was straight-forward, and produced usable results. Finally, we performed a fourier transform of the chain's x-displacement, and tracked each excited mode's evolution over time. This method was most accurate.

We note that, if we simply want a qualitative sense of how a given system would behave, the force curve derivative would be the ideal approach. It requires no simulations of dynamics, and uses only the force curve to predict many of the features of the system's stability.

5.8 Summary

In this chapter, we have presented a simple way to predict many of the properties of a flexible polyelectrolyte's stability against deformation, by simply examining the force curve. We were able to predict how the stability of our system would behave when changing the chain stiffness, the interaction strength, and the chain-chain separation.

We found that the stability against deformation is a function of both the wavelength of deformation, and of the chain-chain separation. To summarize the dependence of the stability on these two variables, we produced stability diagrams for a

test potential of a Lennard-Jones form (Fig. 5.7), and for an analytical fit to our simulation force curve (Fig. 5.8). We also examined the stability surface of our simulation directly, using various methods, and found qualitative agreement with our predictions.

We found that one can even predict the behavior of flexible polyelectrolytes by simply differentiating the force curve. The derivative of the force curve was found to have a direct relationship to the growth rate, and even duplicated details such as peaks and valleys (Fig. 5.23). Overall, we were able to determine the stability of flexible polyelectrolytes against small deformations, using several methods, varying widely in approach and in how exact/approximate. All approaches agreed qualitatively, which gives us great confidence in results.

Armed with this information, we can now examine how our system behaves over the long-term, and how these initial linear-in-amplitude perturbations evolve over time. By using the results of this chapter, we can initialize they system in states that are more likely to produce interesting structures. That is the subject of the next chapter.

Chapter 6

Fluctuations & Structure

6.1 Overview

In the previous chapters, we established that two infinite parallel strings of point charges surrounded by multivalent counterions would experience an attraction. We also established that the force between them would lead to some sinusoidal modes which grow in amplitude. However, the analysis we used only tells us how the modes behave at very short times and small amplitudes.

In this chapter, we will see how the chain pairs evolve together over time, and whether or not they form any kind of structure. If they do, we wish to determine how that structure is affected by the parameters of the system. We also wish to determine which mode eventually wins and dominates the final structure. If there are several dominant modes, we are interested in the relationship between these excited modes.

6.2 Setup

We begin as before, with a random distribution of point charges to represent the counterions. During equilibration, the chains are initially held fixed and linear, a distance R apart from each other. Once equilibrated, we allow both the counterions

and the chains to move freely. Since we are interested in how flexible pairs of polyelectrolytes behave as they approach each other, we will begin by examining rigid rods, and then slowly introduce flexibility in stages and see how the chains react.

We have already seen that our minimal model can describe the like-charged attraction seen in real systems and in more complex simulations. Because we have already established where, in real parameter ranges, we are more likely to find unstable and non-linear behavior, we can greatly simplify our following analyses by restricting the thermal energy to ~ 0.

Our low temperature results form distinct Wigner crystals, reinforcing the support for this mechanism of attraction (versus Oosawa charge density fluctuations). Because lowering the temperature is equivalent to increasing the charge strength or lowering the solvent dielectric, we are essentially looking at the purest form of these interactions. Higher temperatures are smearing out these patterns, making them harder to detect.

6.3 Results

We begin our studies by simply letting the system evolve freely, at finite temperature, and examine the results. Once this has been done, we lower the temperature, implement various constraints, and examine the various ground-state structures that result under those constraints.

For each situation, we will first present a snapshot of the system in its final configuration, which is (or close to) the ground-state configuration for that parameter set. After we have gathered our results, we will compare them to what is know about similar systems, in the discussion section.

6.3.1 $1/f$ Noise

As part of our investigation of long-term structure, we initialized the system in a flat white noise profile, and watched it evolve over a long time. Though snapshots of the system at various time are not extremely revealing, a time series of the power spectrum is.

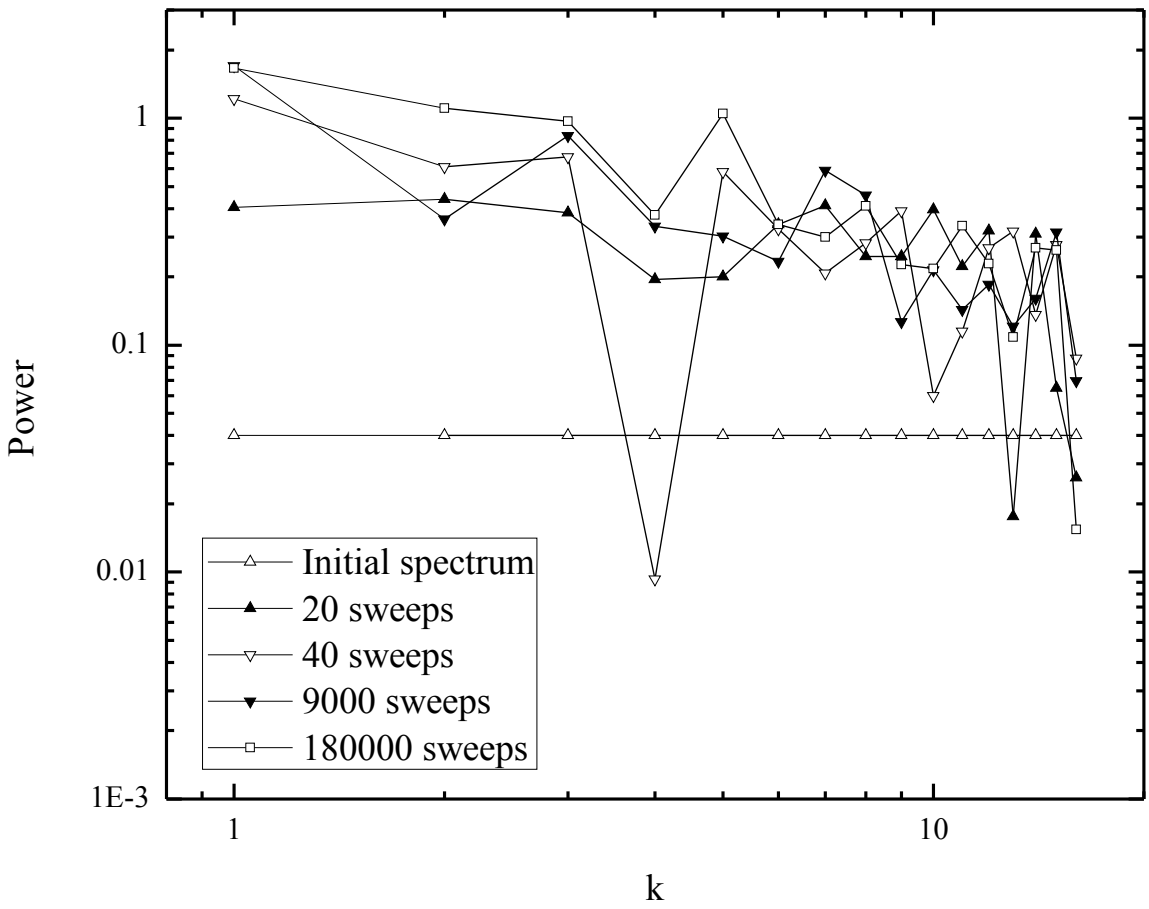

Figure 6.1: Individual snapshots of the power spectrum, for one particular random seed, at different times. The initial conditions are a white noise perturbation (flat line). As the system evolves, we can see from the snapshots that the initial white noise profile is lost after only 20 sweeps. We later collect the final power spectrum for thousands of starting configuration, and average them in Fig. 6.3

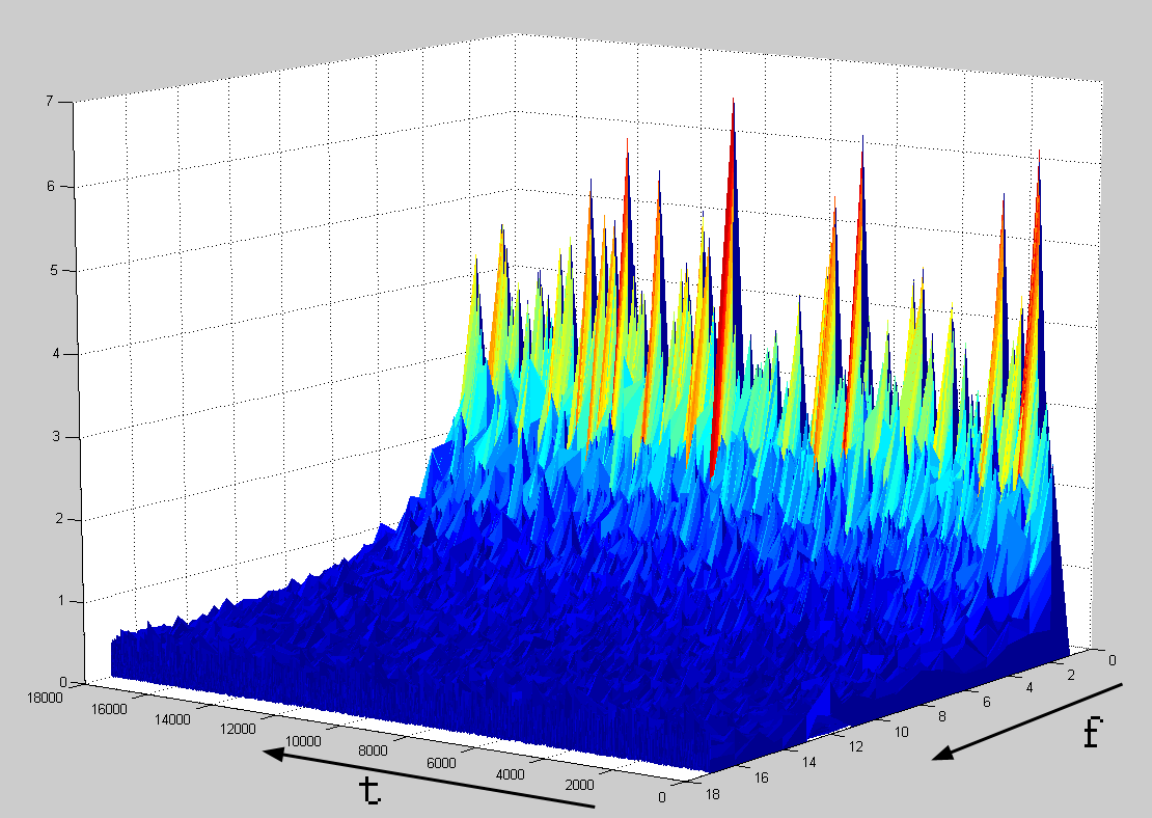

Figure 6.2: Power spectrum of the sinusoidal deformations under free evolution. At $t = 0$ the system begins with a flat (white noise) profile, but very quickly lower frequency modes begin to dominate (red peaks).

The power spectrum at different times in Fig. 6.1 show a great deal of variability. A spectrogram (Fig. 6.2) shows more clearly that there are no persistent patterns from one time to the next. This suggests that we can average the results over several MC sweeps. We have done this in Fig. 6.3, and found a result that was surprising.

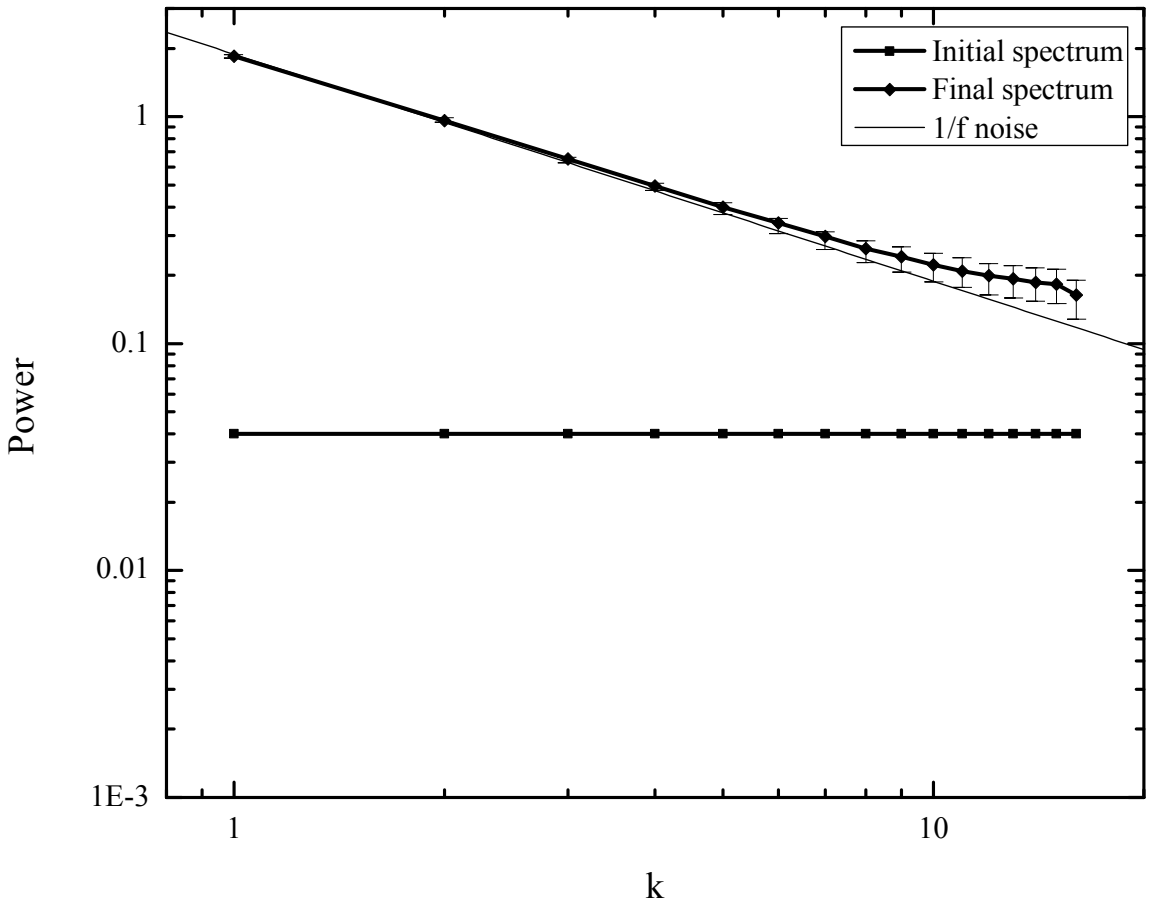

Figure 6.3: Here we see the initial white noise spectrum, and the result of averaging the power spectrum over time. The power spectrum is very clearly of the form $1/f$. The deviation from this at higher frequencies is common, due to greater variability in this region.

We found that the power spectrum tends very quickly, and very strongly, towards $1/f$ noise. We discuss the implications and possible causes of this in the discussion section. For now, we proceed with the original intention: to examine the structures that flexible polyelectrolytes form. To this end, we continue with the most constrained system, and examine its ground-state.

6.3.2 Rigid rods

We begin by examining the ground-state configuration of a pair of rigid chains, when fixed at a separation greater than the equilibrium chain-chain distance. In this

case, we enforce rigidity by simply increasing the chain force parameter h by several orders of magnitude.

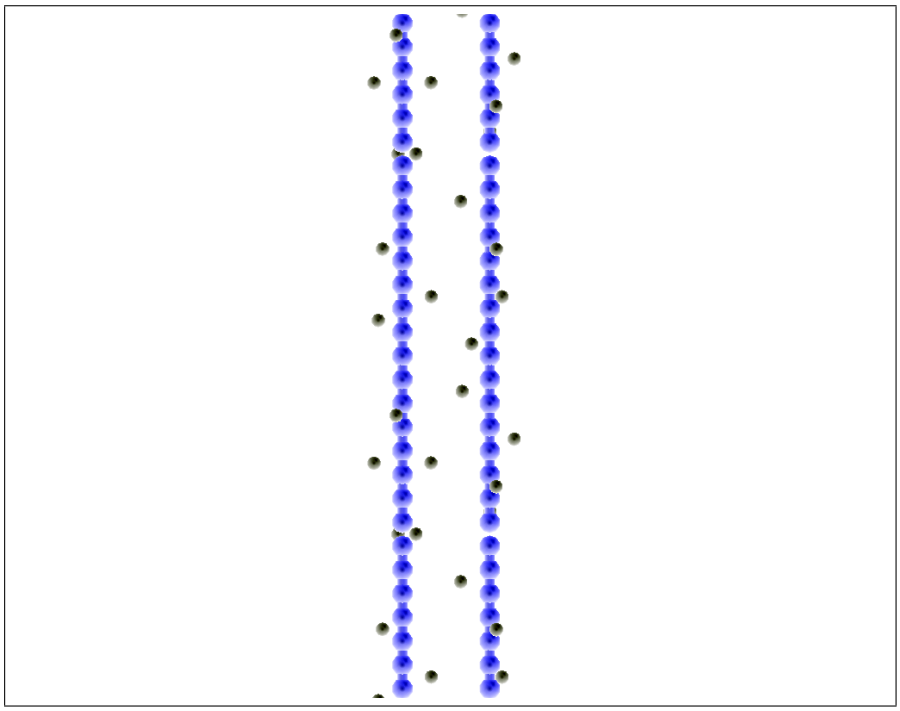

Figure 6.4: We held the rods (blue chained particles) fixed at distance $R = 1.5nm$, and allowed only the counterions (black particles) to move freely. Here we see a ground-state configuration for this system.

Here, the counterions have been allowed to move in order to find their ground-state configuration for this separation. Our starting configuration is similar to that of two widely separated chains. This suggests that, though the chains are being attracted to each other at large distances, their counterion arrangements remain at their independent configurations until the chains are very close to each other.

We then allow the chains to move freely as well. They immediately approach one another, and the counterions in between them stagger into a line. The end result is that there are twice as many counterions in between the rods as outside of them. Also to note, even though both rods and counterions were free to move in three dimensions, the final arrangement was entirely planar.

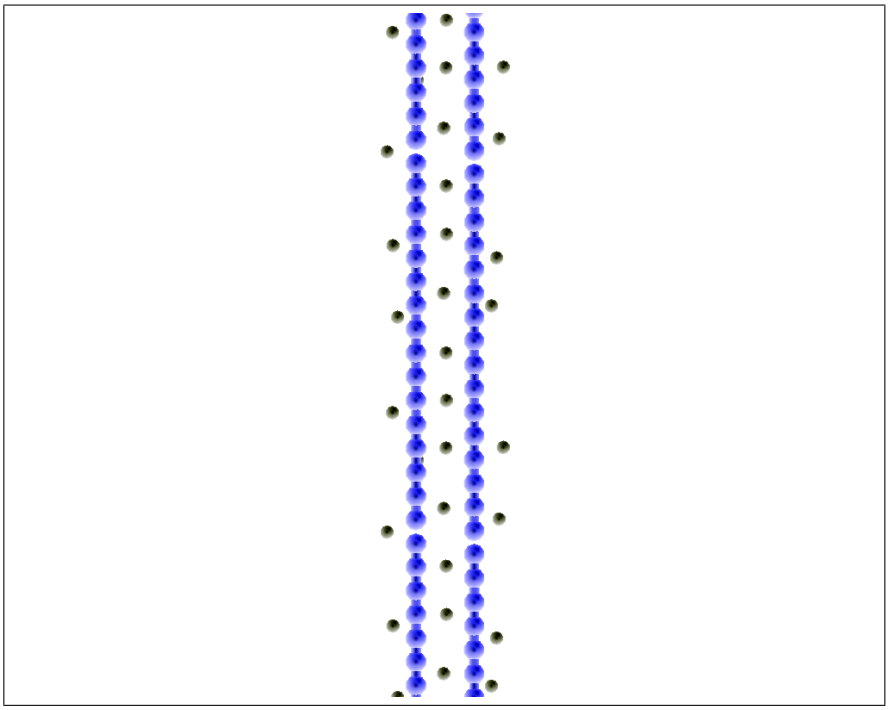

Figure 6.5: Same configuration as previous figure, after allowing the chains to move rigidly and laterally only. As we can see, they moved towards each other, and the counterions arranged themselves in a regular lattice, with twice as many between the chains as outside of them. This is known as a Wigner crystal [2] .

6.3.3 Flexible chains, restricted to plane

We now remove the rigidity constraint, but still allow our chains to move only on the xy-plane. This restriction has no physical analog, and is only a simplification so we can examine the structures that form. Counterions are free to move throughout 3D space.

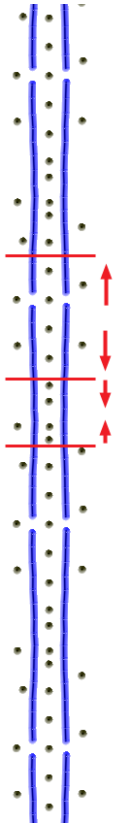

Figure 6.6: Here we can see that, when the chain is restricted to the xy-plane, the chains form mostly long-wavelength perturbations. A key thing to note is that the counterions also form charge density waves of the same wavelength.

We find that, for moderate values of the spring constant, the chains remain planar but undergo long-wavelength fluctuations (in this case, of the same size as the periodic height L_y). Once again, this seems to be related to the tendency of the system towards $1/f$ noise, as can be seen by the first few modes of the power spectrum in Fig. 6.7.

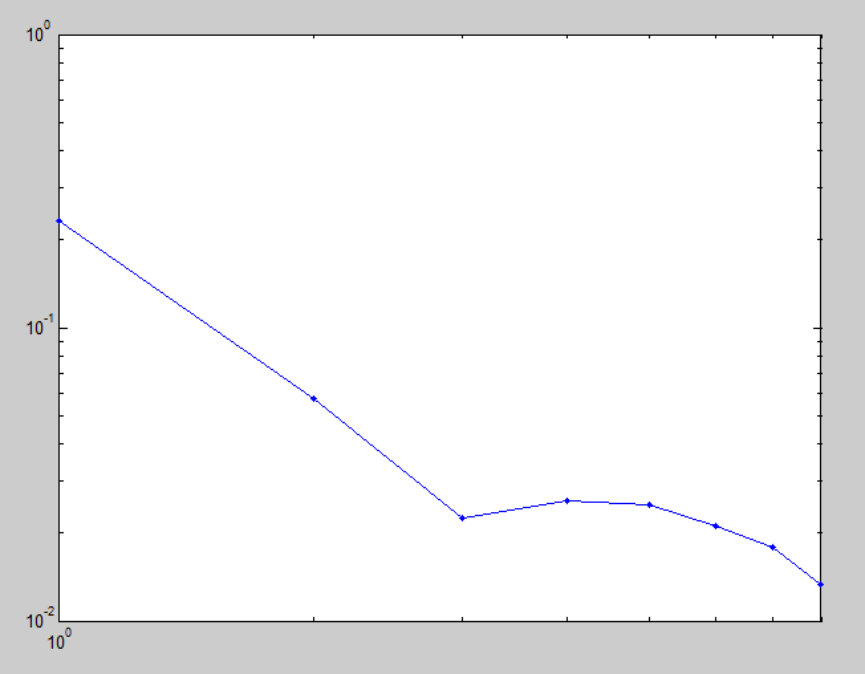

Figure 6.7: The power spectrum for the flexible planar ground state. The two longest-wavelength modes are dominant, and the first three have a $1/f$ slope.

Even though we have been mostly interested in somewhat stiff polymer parameters (to model DNA, F-actin and the like), we have also produced results for the ground-states of far more flexible chains. The structure formed when the chain particles are restricted to the xy-plane, can be seen in Fig. 6.8.

Figure 6.8: When restricted to the xy-plane, but given a very weak spring constant, the polyelectrolyte chains form this structure. It is essentially a 2D ionic crystal, with weak bonds holding all negative chain particles together.

6.3.4 Flexible chains, free

If we remove the xy-plane restriction, similar results are obtained. As we will see, this tendency towards coplanarity is a general feature of this system, at least for the parameter sets we use.

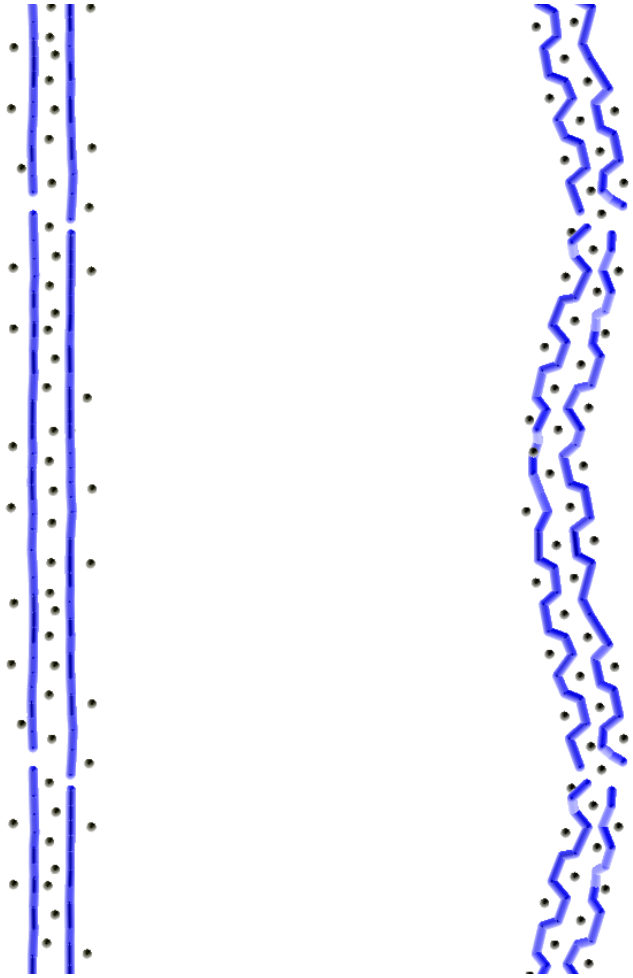

(a) (b)

Figure 6.9: Our low-temperature results for flexible and free polyelectrolytes, for two different spring constants.

One thing to note, however, is that when allowed to move off-plane, the loose runs exhibit a very clear long-wavelength mode which, unlike in the rigid rod case, is in phase for both chains. The counterions and chain particles in this case are still arranged in a 2D ionic crystal.

6.3.5 Flexible chains, symmetry-breaking

In the previous sections we have progressively removed the artificial constraints from the system, until the chains were allowed to move freely under the three main po-

tentials (electrostatics, repulsive core, and chain forces). We now investigate whether the introduction of a very weak potential to prefer (locally) one sense of torsion over another would be all that's required to produce a helical configuration.

We ran simulations using free chains with the symmetry breaking potential:

$$V_{bp}(r, \phi) = \epsilon_{bp}(\phi - \phi_0)^2 \,,$$

where again ϕ is the dihedral angle between corresponding monomers on opposing chains. What we found is that, even when the base-pair energy was weak when compared to the electrostatic energy ($\epsilon_{ES}/\epsilon_{bp} \sim 1/100$), the chains would coil into helices.

Figure 6.10: The final stable configuration of a chain pair, when allowed to move freely, with a weak symmetry-breaking force. The average pitch for this configuration is $2.4nm$, which is reasonably close to DNA's $3.4nm$, given the simplicity of the model.

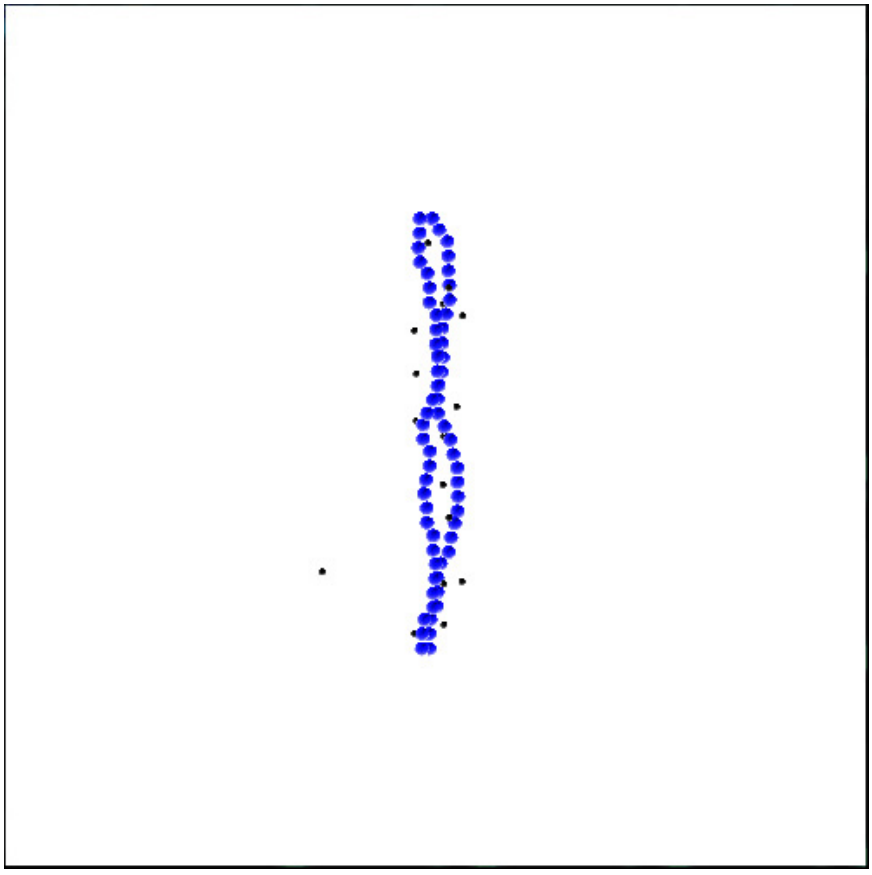

Figure 6.11: Here we have the final configuration for a stiffer chain. The resulting pitch is $8.8nm$, which is $\sim 40\%$ as coiled as DNA. We were unable to use parameters for stiffer chains, such as F-actin, due to the size of the simulation required (F-actin has a pitch of $70nm$ [3], meaning we would only have one turn per $70nm/0.34nm = 200$ monomers).

We then examined the degree of helicity as a function of the symmetry breaking potential e_{bp}. We wished to determine whether helicity increased linearly as the potential was increased, or if there was a threshold value, below which helicity was 0, and above it finite. To do this, we simply calculated the angle between the two matching monomers on each chain, and summed the difference of this angle along the entire length of the chains. If the sum was positive, that indicated a right-handed twisting.

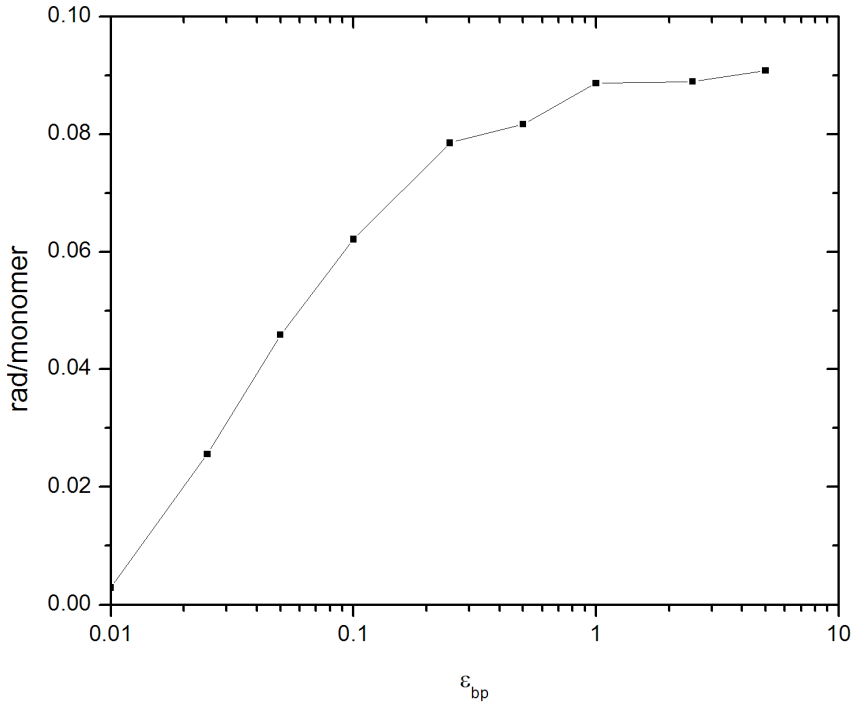

Figure 6.12: Amount of helicity as a function of the symmetry-breaking potential e_{bp}. We can see that even at near-zero values, the potential was enough to produce a net positive helicity. As the potential was increased, the total helicity for the chain quickly approached a saturation value at $\sim 0.09\,\mathrm{rad}$ (the average $\Delta\theta$ per monomer).

What we found is that helicity begins to appear as soon as any amount of symmetry breaking is introduced, within the precision of our study. The amount of helicity over the whole chain then rapidly reaches the saturation value for that given preferred angle. The implications of this will be addressed in the discussion section.

6.4 Discussion

6.4.1 $1/f$ Noise

We have seen that, whether we start our system with a single excited mode (stable or unstable), with white noise, or even flat, the chains quickly begin to fluctuate with a $1/f$ power spectrum. This type of noise (also called *flicker noise*, *long-term memory*, or *pink noise*) has been found in economics, behavioral psychology, atmospheric

physics, electronic signals, city populations, and the neural activity of the brain. No single simple explanation exists for its ubiquity. A power law spectrum also implies that the perturbations of this system are scale-free. In other words, when we increase or decrease the scale of the system (the periodic height L_y), our noise profile should be identical. We have verified that this is indeed the case, with simulations at twice and half the size of our main results.

Based on the Bak-Tang-Weisenfeld sandpile model of self-organized criticality[63], we can attempt to find a phenomenology to explain the $1/f$ scale-invariant fluctuations in our system. While other models that aren't based on self-organized criticality might apply, we found that the BTW sandpile model had the closest phenomenological mapping to our system. Based on the sandpile model, three things are needed: a medium, a random disturbance, a critical threshold of stability, and a restoration after reaching criticality. For the sandpile model itself, the following are the elements needed for self-organized criticality, and $1/f$ noise.

1. The sand pile as a medium.
2. Random addition of sand grains to surface sites.
3. A critical threshold of instability for each site.
4. A way for the disturbing sand grains in an unstable site to propagate to its neighbours.
5. An eventual restoration of the medium.

In our system, we have the following:

1. The chain as a medium.
2. Random addition of counterions to "sites" along the chain.

3. A critical threshold of stability for the whole system (as seen by stability diagrams). Since random local addition of counterions changes the system's average R & k, this addition can also move the system past the stability threshold

4. When the local pertubation is unstable and grows, the deformation propagates to the site's neighbours.

5. The chain-chain attraction, and the spring restoring force, will eventually bring restore the system.

Though we have only phenomenological similarities between the two models, there are parallels for every element. As with other systems with self-organized criticality[64], ours is robust to fine-tuning: all initial conditions and parameters resulted in $1/f$ noise spectrum eventually. A theory of the self-organized criticality of this system would be of interest. It should be noted, however, that this $1/f$ noise has only been observed so far to extend over one decade, and that a more in-depth study of this phenomenon should first establish whether this noise profile extends over a larger region.

6.4.2 Structures

Given that most studies on flexible polyelectrolytes use finite and free segments, our results in this chapter are novel. We have seen that various features from line-charge studies persist even when the constraints are removed, such as the collinearity of the counterions both inside and outside the chain, the coplanarity of the entire system.

If we couple the fact that the counterions and chain particles are entirely coplanar, and that they are always fully condensed if the Manning parameter is $\xi < 1$, then we might be able to perform these Monte Carlo simulations entirely in 2D, while using the same potentials (albeit with $z = 0$). This would greatly reduce the phase space of the system, and speed up the simulations.

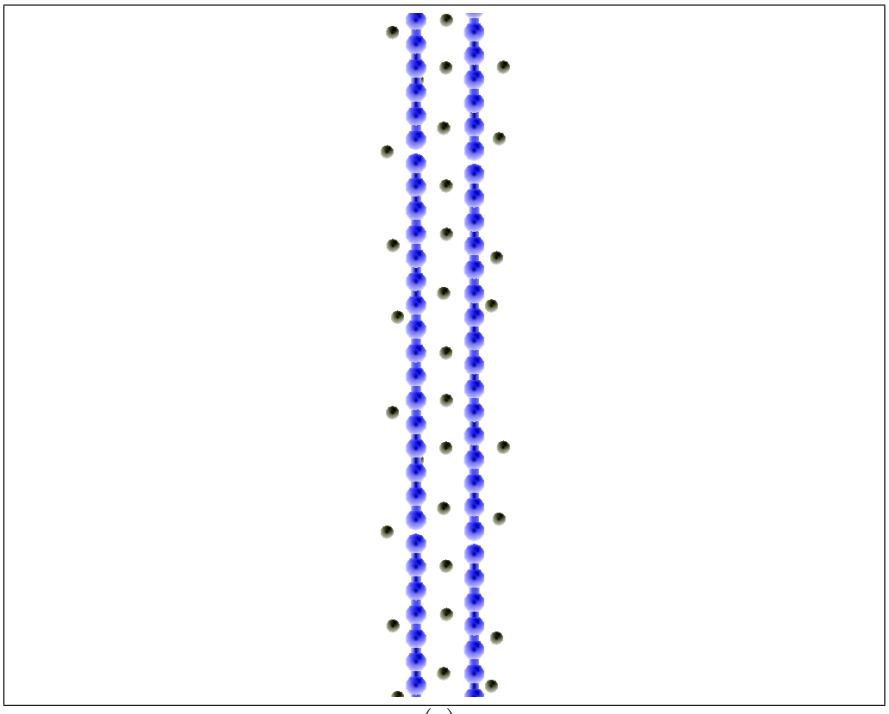

(a)

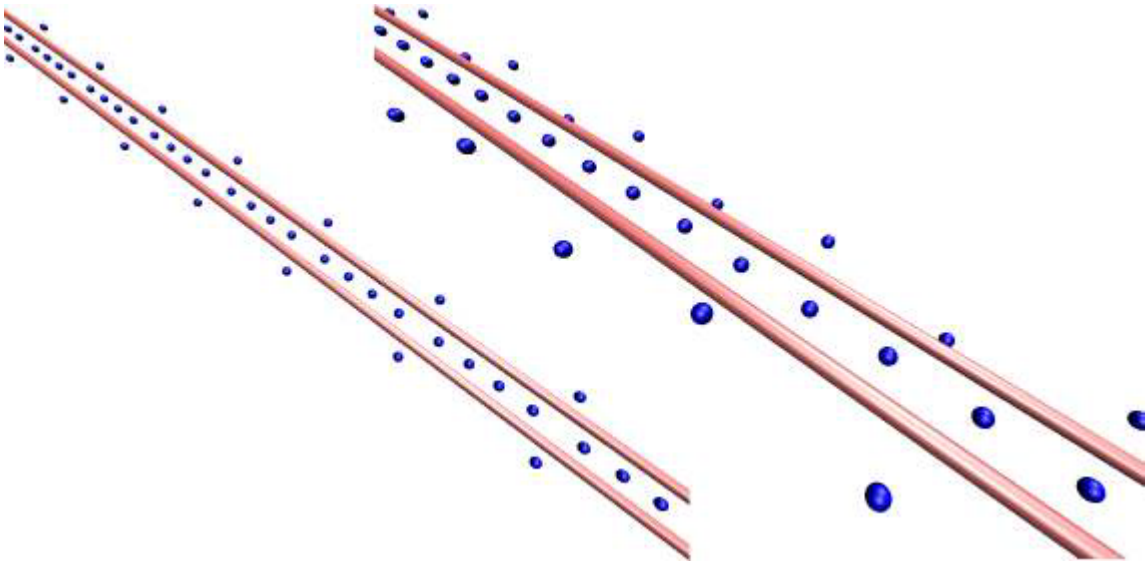

(b)

Figure 6.13: Results from (a) our simulations, and (b) from [4]. We can see that the ground state configuration for rigid lines is the same in both cases, despite using different simulation methods.

As can be seen by the comparison Fig. 6.14, our results are comparable with studies by [4]. They used a different simulation package (ESPResSo, whereas our was as described earlier), different periodic potentials (their MMM1D vs. our Lekner sums),

different repulsive potentials (Weeks-Chandler-Anderson vs. Repulsive Lennard-Jones). This further supports the robustness of these results, since they do not depend on the details of the simulation methods used.

The main thing to note with our flexible chain simulations is that the power spectrum is still partially $1/f$ even at these low temperatures, and that the perturbations result in a corresponding wave of charge density.

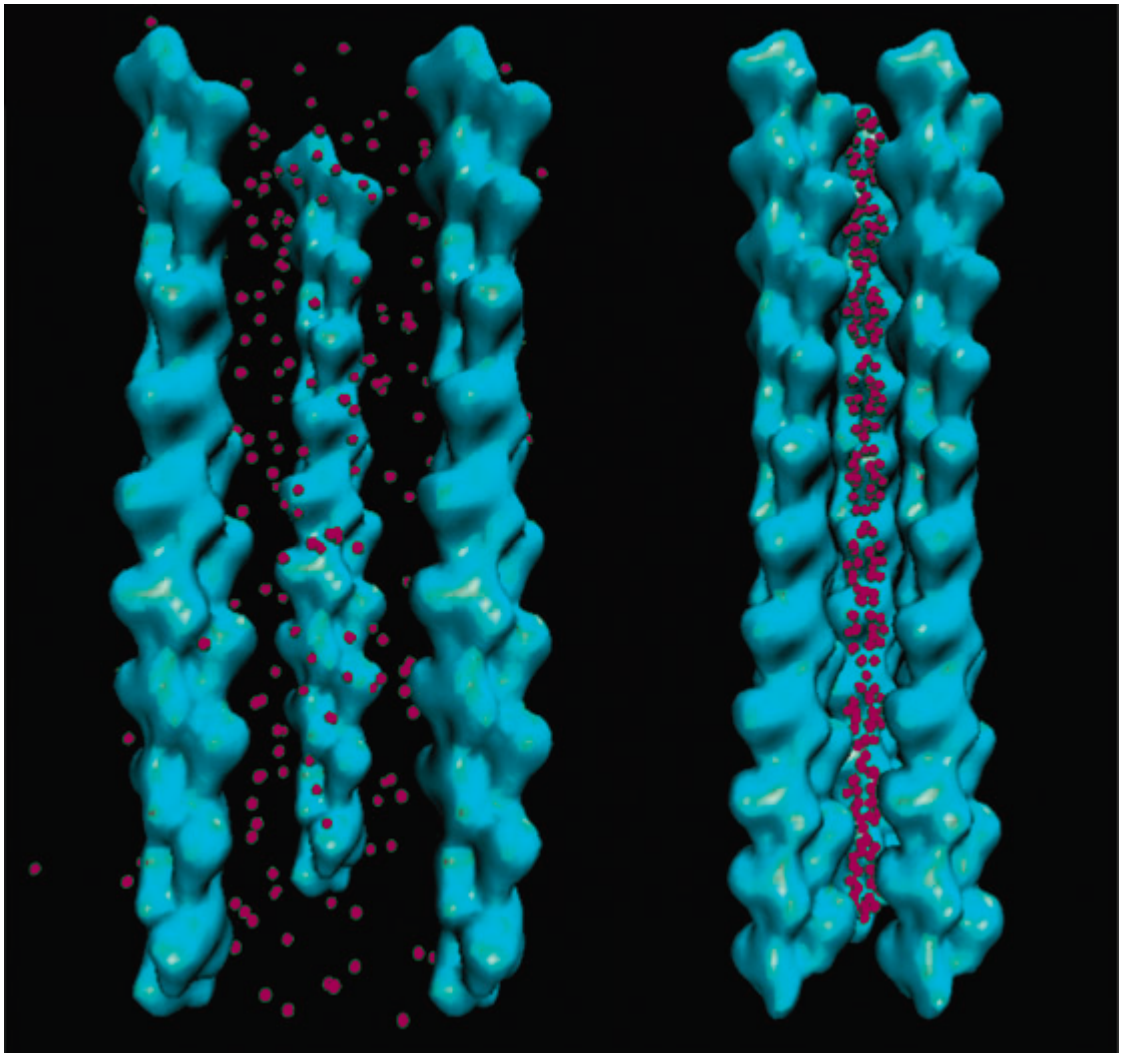

Figure 6.14: Results from [5]. Here we have (a) uncondensed F-actin and (b) condensed F-actin bundles. The counterions are seen to form a charge density wave, which in turn couples with the F-actin helix. This has been verified with X-ray measurements as well.

This kind of charge density wave has been seen wherever there is any topological

variation along the chain axis. For instance, F-actin has grooves that, when pairs of F-actin attract, cause the counterions between them to form waves. From our other studies it is clear that these waves are not required to create like-charged attraction.

The ground-state structures we found depended greatly on the constraints on the system, but they all shared some general features. First, unless one direction was made energetically favorable (using the symmetry-breaking potential), the chains and counterions remained essentially coplanar even when allowed to be completely free.

We also note that even in these zero-temperature limits, the final configuration exhibits some degree of $1/f$ noise. This suggests that this is a fundamental feature of the system and the energies involved, and not only a feature of higher temperature fluctuations.

As for helicity, we note that the potential used is ~ 100 times weaker than the other potentials, yet still had a strong effect overall. Also it should be noted there were instances of torsion and helicity even when there was no symmetry-breaking potential.

We found that, for the parameters in the DNA-like system, the amount of helicity in the final structure was very sensitive to the symmetry breaking potential. Even when $\epsilon_{ES}/\epsilon_{bp}$ ranged from 1/1000 to 1/10 there was a significant tendency towards forming helices, as can be shown in figure 6.15.

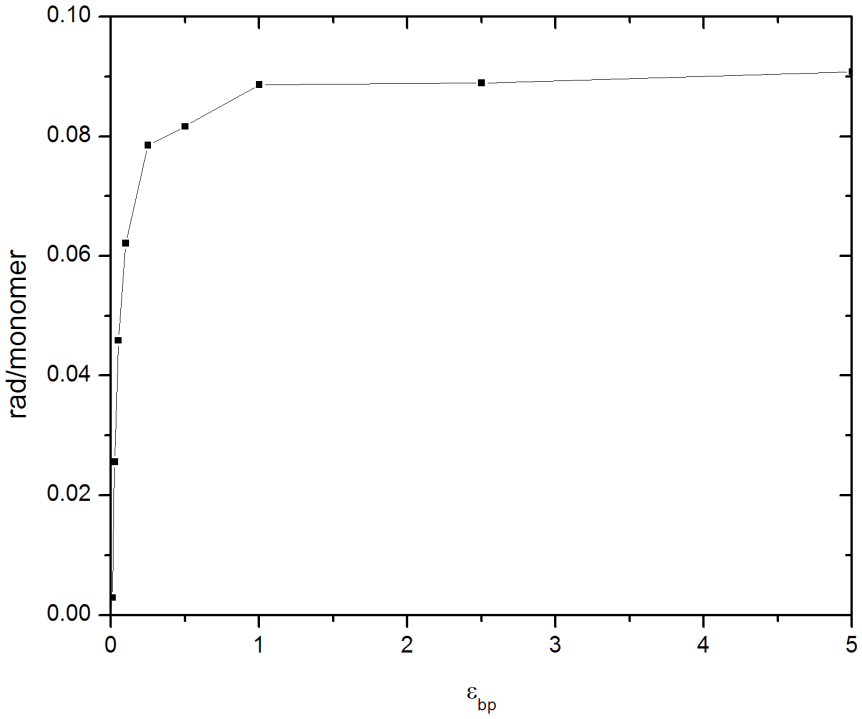

Figure 6.15: The tendency towards helicity starts immediately upon the addition of a symmetry-breaking potential, and quickly rises to a saturation point (at 0.09 rad per monomer) as the strength of this parameter is increased.

While our parameters attempt to model a DNA-like system, our model is still too minimal and coarse to make general statements about the tendency of real DNA strands to form helices, and whether the counterion-induced like-charged attraction has a significant role to play in this process. Nevertheless, the strength of the tendency towards helicity is surprising, and merits further study.

Chapter 7

Conclusions

7.1 Further work

There is much more that could be done, even without extending our model system. For instance, we have not explored the time-dependent dynamics of the chains when free or investigated how they fluctuate. In focusing on the attractive force, the stability to deformation, and the final configurations, we have not considered the rich behavior that a time-dependent study would yield.

A closer link to experimental results, by determining thresholds of stability for real finite polyelectrolytes for instance, would help us to use this model to lend support (or detract from) theories regarding the origin of certain behaviors and interactions (such as the like-charged attraction or DNA's ability to condense).

If we were to extend our current model slightly, we could investigate a great number of other behaviors. For instance, we could incorporate a hydrogen-bonding interaction between adjacent monomers of opposite chains, to simulate the base-pair bonding of DNA.

Another aspect that has not been addressed is end-effects. We have not considered how our results would change if we turn our infinite chains into finite chain segments.

Also, experiments with adjustable "dumbell" counterions have shown that counterion size can have an important effect on the interaction between polyelectrolytes.

Whether the chains are initialized with a single excited mode, white noise, or flat, there seems to be a strong tendency towards a 1/f noise profile. We have a phenomenological explanation based on the sandpile model, but we have not investigated this in detail (either analytically or numerically). We have not determine whether the noise is a result of the periodicity of the system, and whether it is truly scale-invariant.

7.2 Summary

In this work, we have successfully examined the behavior of a minimal model of two infinite, flexible polyelectrolytes. With such a complex problem space (infinite range electrostatics, large number of potential configurations), it is noteworthy that we were able to reproduce many of the behaviours seen in real systems.

For instance, we successfully reproduced the like-charged attraction seen in other works (both experimental and simulated). We extended this by finding an approximate analytical expression for the force curves, based on modified Bessel functions of the second-kind (which arose from the discrete particle-particle electrostatic interactions in a 1D periodic slab).

We then used the force curve to predict how two parallel polyelectrolytes would respond to deformations. Since most previous studies examining the like-charged attraction between polyelectrolytes used rigid rods, this was an interesting area of examination. We found that the polyelectrolytes were mostly stable outside a cut-off separation, and that the deformations were linearly unstable (at long-wavelengths) as the chains approached each other. This agreed with predictions based on a simple mean-field force curve approach.

Finally, we used the information about the force and the linear stability to initialize the system in states that were more likely to lead to interesting structures. We examined the ground-state structures under several constraints: rigid rods, planar flexibility, full flexibility, and bond-angle preference.

We found that the general behavior of each step was a predecessor for those following. Step zero is counterion condensation. Step one is "approach", where the rods close the gap and approach their equilibrium separation. Step two is "zippering", where the counterions surrounding each individual chain would begin to arrange themselves in a staggered line between the chains. If the chains are then allowed to flex, they would now begin Step three, "deformation". The wavelength of this deformation was dependent on the stiffness of the chain force, as expected from the stability analysis.

We found that even if the rods were not restricted to moving within the xy-plane, they would remain on that plane unless a symmetry breaking potential was introduced. On adding this potential, we found that the chains would wrap around each other in a helix, and that the helix would undulate with a characteristic wavelength. We also found that the tendency towards helicity did not require the symmetry breaking potential to be above a certain threshold.

Finally, we found that pairs of infinite flexible polyelectrolytes in solution with neutralizing counterions exhibit fluctuations of $1/f$ noise. This result was robust to all parameters used. We briefly presented a possible parallel to the sandpile model of self-organized criticality.

Though we used only point-charges, interacting via electrostatics, a chain potential, and a repulsive core, our model has given us information on how a pair of long flexible polyelectrolytes behave as they approach each other.

In summary, we gain valuable insight with models such as these. They allow us to determine the minimal set of features required to reproduce the behavior we are

interested in. This reduction of the problem to its core elements is in itself a result, as we are exposing the mechanisms that make the most significant impact on overall behavior and structure.

Bibliography

[1] N. Grønbech-Jensen, R.J. Mashl, R.F. Bruinsma, and W.M. Gelbart. Counterion-induced attraction between rigid polyelectrolytes. *Phys. Rev. Lett.*, 78(12):2477–2480, Mar 1997.

[2] Qi Wen. Exploring the mechanism of like-charge attraction in polyelectrolyte solutions. 2007.

[3] J. Ks, H. Strey, J.X. Tang, D. Finger, R. Ezzell, E. Sackmann, and P.A. Janmey. F-actin, a model polymer for semiflexible chains in dilute, semidilute, and liquid crystalline solutions. *Biophysical Journal*, 70(2):609 – 625, 1996.

[4] A. Arnold and C. Holm. Interactions of like-charged rods at low temperatures: Analytical theory vs. simulations. *The European Physical Journal E*, 27(1):21–29, 2008.

[5] Gerard CL Wong and Lois Pollack. Electrostatics of strongly charged biological polymers: ion-mediated interactions and self-organization in nucleic acids and proteins. *Annual review of physical chemistry*, 61:171–189, 2010.

[6] Mark J. Stevens. Bundle binding in polyelectrolyte solutions. *Phys. Rev. Lett.*, 82:101–104, Jan 1999.

[7] Jay X. Tang, Tadanao Ito, Terence Tao, Peter Traub, and Paul A. Janmey. Opposite effects of electrostatics and steric exclusion on bundle formation by f-actin and other filamentous polyelectrolytes. *Biochemistry*, 36(41):12600–12607, 1997.

[8] E. Allahyarov, G. Gompper, and H. Lowen. Dna condensation and redissolution: interaction between overcharged dna molecules. *Journal of Physics: Condensed Matter*, 17(20):S1827–S1840, 2005.

[9] E. Allahyarov, I. D'Amico, and H. Löwen. Attraction between like-charged macroions by coulomb depletion. *Phys. Rev. Lett.*, 81:1334–1337, Aug 1998.

[10] T. E. Angelini, H. Liang, W. Wriggers, and G. C. L. Wong. Like-charge attraction between polyelectrolytes induced by counterion charge density waves. *Proceedings of the National Academy of Sciences of the United States of America*, 100(15):8634–8637, 2003.

[11] A. Naji and R.R. Netz. Attraction of like-charged macroions in the strong-coupling limit. *The European Physical Journal E*, 13(1):43–59, 2004.

[12] Rebecca M Nyquist, Bae-Yeun Ha, and Andrea J Liu. Counterion condensation in solutions of rigid polyelectrolytes. *Macromolecules*, 32(10):3481–3487, 1999.

[13] J. X. Tang and P. A. Janmey. The polyelectrolyte nature of f-actin and the mechanism of actin bundle formation. *J. Biol. Chem.*, 271, 1996.

[14] S. Liu and M. Muthukumar. Langevin dynamics simulation of counterion distribution around isolated flexible polyelectrolyte chains. *The Journal of Chemical Physics*, 116(22):9975–9982, 2002.

[15] Silvio a Beccara, Pietro Faccioli, Marcello Sega, Francesco Pederiva, Giovanni Garberoglio, and Henri Orland. Dominant folding pathways of a peptide chain from ab initio quantum-mechanical simulations. *The Journal of Chemical Physics*, 134(2):–, 2011.

[16] Martin Karplus and J Andrew McCammon. Molecular dynamics simulations of biomolecules. *Nature Structural & Molecular Biology*, 9(9):646–652, 2002.

[17] J.M. Berg, J.L. Tymoczko, and L. Stryer. *Biochemistry, Fifth Edition.* W.H. Freeman, 2002.

[18] Andrey V Dobrynin and Michael Rubinstein. Theory of polyelectrolytes in solutions and at surfaces. *Progress in Polymer Science*, 30(11):1049–1118, 2005.

[19] J. Faraudo and A. Travesset. Phosphatidic Acid Domains in Membranes: Effect of Divalent Counterions. *Biophys. J.*, 92(8):2806–2818, 2007.

[20] W. M. Gelbart, R. F. Bruinsma, P. A. Pincus, and V. A. Parsegian. Dna-inspired electrostatics. *Physics Today*, 53(9), 2000.

[21] T. E. Angelini, H. Liang, W. Wriggers, and G. C. L. Wong. Like-charge attraction between polyelectrolytes induced by counterion charge density waves. *Proc. Natl. Acad. Sci. USA*, 100(15), 2003.

[22] J. X. Tang, P. A. Janmey, A. Lyubartsev, and L. Nordenskiold. Metal ion-induced lateral aggregation of filamentous viruses fd and m13. *Biophys. J.*, 83(1), 2002.

[23] M. Fixman. The poisson–boltzmann equation and its application to polyelectrolytes. *The Journal of Chemical Physics*, 70(11):4995–5005, 1979.

[24] P. Debye and E. Hückel. Debye-Hückel theory of electrolytes. *Phys. Z*, 24:185, 1923.

[25] B. I. Shklovskii. Wigner crystal model of counterion induced bundle formation of rodlike polyelectrolytes. *Phys. Rev. Lett.*, 82, 1999.

[26] I. Rouzina and V. A. Bloomfield. Macroion attraction due to electrostatic correlation between screening counterions. 1. mobile surface-adsorbed ions and diffuse ion cloud. *J. Phys. Chem.*, 100, 1996.

[27] E. J. Verwey and J. T. G. Overbeek. *Theory of the Stability of Lyophobic Colloids*. Elsevier, Amsterdam, 1948.

[28] D. Stigter. Evaluation of the counterion condensation theory of polyelectrolytes. *Biophys. J.*, 69(2):380–388, 1995.

[29] R.W. Wilson and V.A. Bloomfield. Counterion-induced condensation of deoxyribonucleic acid. a light-scattering study. *Biochemistry*, 18(11):2192–2196, 1979.

[30] Victor A Bloomfield et al. Dna condensation by multivalent cations. *Biopolymers*, 44(3):269–282, 1997.

[31] J. J. Arenzon, Y. Levin, and J. F. Stilck. The mean-field theory for attraction between like-charged macromolecules. *Physica A*, 283(1):1–5, 2000.

[32] G. N. Patey. The interaction of two spherical colloidal particles in electrolyte solution. an application of the hypernetted-chain approximation. *The Journal of Chemical Physics*, 72(10):5763–5771, 1980.

[33] V. A. Bloomfield. Condensation of dna by multivalent cations: Considerations on mechanism. *Biopolymers*, 31(13):1471–1481, 1991.

[34] R. Podgornik, D.C. Rau, and V.A. Parsegian. Parametrization of direct and soft steric-undulatory forces between DNA double helical polyelectrolytes in solutions of several different anions and cations. *Biophys. J.*, 66(4):962–971, 1994.

[35] J. X. Tang, S. Wong, P. Tran, and P. Janmey. Counterion induced bundle formation of rodlike polyelectrolytes. *Ber. Bunsenges. Phys. Chem.*, 100(6):796–806, 1996.

[36] P. G. Arscott, C. Ma, J. R. Wenner, and V. A. Bloomfield. Dna condensation by cobalt heaamine(iii) in alcohol-water mixtures: Dielectric constant and other solvent effects. *Biopolymers*, 36, 1995.

[37] E. Raspaud, M. O. De la Cruz, J. L. Sikorav, and F. Livolant. Precipitation of dna by polyamines: A polyelectrolyte behavior. *Biophys. J.*, 74, 1998.

[38] J. Pelta, F. Livolant, and J.-L. Sikorav. Dna aggregation induced by polyamines and cobalthexamine. *J. Biol. Chem.*, 271, 1996.

[39] V. A. Bloomfield. Dna condensation. *Curr. Opin. Struct. Biol.*, 6, 1996.

[40] B.-Y. Ha and A. J. Liu. Effect of non-pairwise-additive interactions on bundles of rodlike polyelectrolytes. *Phys. Rev. Lett.*, 81(5):1011–1014, 1998.

[41] F. Oosawa. Interaction between parallel rodlike macroions. *Biopolymers*, 6(11):1633–1647, 1968.

[42] G. S. Manning. Limiting laws and counterion condensation in polyelectrolyte solutions i. colligative properties. *The Journal of Chemical Physics*, 51(3):924–933, 1969.

[43] G. S. Manning. Limiting laws and counterion condensation in polyelectrolyte solutions iv. the approach to the limit and the extraordinary stability of the charge fraction. *Biophysical Chemistry*, 7(2):95–102, 1977.

[44] Markus Deserno, Christian Holm, and Sylvio May. Fraction of condensed counterions around a charged rod: comparison of poisson-boltzmann theory and computer simulations. *Macromolecules*, 33(1):199–206, 2000.

[45] O. A. Saleh, D. B. McIntosh, P. Pincus, and N. Ribeck. Nonlinear low-force elasticity of single-stranded dna molecules. *Phys. Rev. Lett.*, 102:068301, Feb 2009.

[46] Mark J. Stevens, Dustin B. McIntosh, and Omar A. Saleh. Simulations of stretching a strong, flexible polyelectrolyte. *Macromolecules*, 45(14):5757–5765, 2012.

[47] Ngo Minh Toan and D. Thirumalai. On the origin of the unusual behavior in the stretching of single-stranded dna. *The Journal of Chemical Physics*, 136(23):–, 2012.

[48] J. J. Arenzon, J. F. Stilck, and Y. Levin. Simple model for attraction between like-charged polyions. *The European Physical Journal B - Condensed Matter and Complex Systems*, 12:79–82, October 1999.

[49] Sho Asakura and Fumio Oosawa. On interaction between two bodies immersed in a solution of macromolecules. *The Journal of Chemical Physics*, 22(7):1255–1256, 1954.

[50] K. Binder. *Monte Carlo and Molecular Dynamics Simulations in Polymer Science*. Oxford University Press, 1995.

[51] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6), 1953.

[52] Gareth O Roberts, Andrew Gelman, Walter R Gilks, et al. Weak convergence and optimal scaling of random walk metropolis algorithms. *The annals of applied probability*, 7(1):110–120, 1997.

[53] Wikipedia. Ewald summation — wikipedia, the free encyclopedia, 2013. [Online; accessed 22-November-2013].

[54] P.P. Ewald. Die berechnung optischer und elektrostatischer gitterpotentiale. *Annalen der Physik*, 369(3):253–287, 1921.

[55] J. Lekner. Summation of coulomb fields in computer-simulated disordered systems. *Physica A*, 176, 1991.

[56] N. Grønbech-Jensen, G. Hummer, and K.M. Beardmore. Lekner summation of coulomb interactions in partially periodic systems. *Molecular Physics*, 92(5):941–946, 1997.

[57] Milton Abramowitz. *Handbook of mathematical functions : with formulas, graphs, and mathematical tables*. Dover Publications, New York, 1970.

[58] Paul J Flory. Statistical mechanics of dilute polymer solutions. *The Journal of Chemical Physics*, 17(12):1347–1348, 1949.

[59] Kurt Kremer and Gary S Grest. Dynamics of entangled linear polymer melts: A molecular-dynamics simulation. *The Journal of Chemical Physics*, 92:5057, 1990.

[60] Simona Cocco, Jie Yan, Jean-Francois Léger, Didier Chatenay, and John F Marko. Overstretching and force-driven strand separation of double-helix dna. *Physical Review E*, 70(1):011910, 2004.

[61] Wikipedia. Dna — wikipedia, the free encyclopedia, 2013. [Online; accessed 22-November-2013].

[62] B.-Y. Ha and A. J. Liu. Counterion-mediated attraction between two like-charged rods. *Phys. Rev. Lett.*, 79(7):1289–1292, Aug 1997.

[63] Per Bak, Chao Tang, Kurt Wiesenfeld, et al. Self-organized criticality: An explanation of 1/f noise. *Physical Review Letters*, 59(4):381–384, 1987.

[64] Jean M Carlson and John Doyle. Complexity and robustness. *Proceedings of the National Academy of Sciences of the United States of America*, 99(Suppl 1):2538–2545, 2002.