Nucleation: theory and applications to protein solutions and colloidal suspensions

2007 J. Phys.: Condens. Matter 19 033101

(http://iopscience.iop.org/0953-8984/19/3/033101)

INSTITUTE OF PHYSICS PUBLISHING JOURNAL OF PHYSICS: CONDENSED MATTER

J. Phys.: Condens. Matter **19** (2007) 033101 (28pp) doi:10.1088/0953-8984/19/3/033101

TOPICAL REVIEW

Nucleation: theory and applications to protein solutions and colloidal suspensions

Richard P Sear

Department of Physics, University of Surrey, Guildford, Surrey GU2 7XH, UK

E-mail: r.sear@surrey.ac.uk

Received 2 November 2006, in final form 21 November 2006
Published 3 January 2007
Online at stacks.iop.org/JPhysCM/19/033101

Abstract

The status of our understanding of nucleation is reviewed. Both general aspects of nucleation and those specific to protein solutions and colloidal suspensions are considered. We conclude that although we have what we believe is a quite good understanding of homogeneous nucleation in one simple system, hard-sphere colloids, in other systems there are still basic questions that are unanswered. For example, for even the most studied protein, lysozyme, there is an ongoing debate about whether the observed nucleation is homogeneous or heterogeneous. We review theoretical and simulation work on both homogeneous and heterogeneous nucleation. As heterogeneous nucleation appears to be much more common than homogeneous nucleation, and as earlier reviews have tended to focus more on homogeneous nucleation, we place particular emphasis on heterogeneous nucleation.

(Some figures in this article are in colour only in the electronic version)

Contents

1. Introduction 2
2. Introduction to nucleation 3
   2.1. Heterogeneous nucleation 5
   2.2. The nucleation theorem 9
   2.3. Heterogeneous nucleation with a distribution of barrier heights 9
   2.4. Nucleation at surface phase transitions 10
   2.5. Nucleation at phase transitions in porous media 11
   2.6. Nucleation on soluble impurities 13
3. Computer simulation 14
   3.1. Simulation techniques 15
4. Nucleation in protein solutions 16
   4.1. Proteins other than lysozyme 18

5. Nucleation when there is more than one phase transformation possible 18
5.1. Theory of homogeneous nucleation near a second first-order phase transition 20
5.2. Theory of heterogeneous nucleation near a second first-order phase transition 22
6. Nucleation in colloidal suspensions 23
6.1. Hard spheres 23
6.2. Charged spheres 24
6.3. Non-spherical particles 24
7. Conclusion 25
Acknowledgments 26
References 26

1. Introduction

We see phase transitions all around us: rainfall and the freezing of water are just two examples. These and most other transitions are first-order transitions, and therefore they occur via nucleation and then growth of the nucleus into a macroscopic rain droplet, ice crystal, etc. Here we review work on nucleation: the process which starts off a first-order phase transformation. Nucleation is an activated process: the growing nucleus of the new phase must overcome a barrier. It is the nucleus at the top of this barrier that determines the nucleation rate as the rate decreases exponentially as the barrier height increases. The nucleus is still microscopic at the top of the barrier and thus it is in almost all systems too small to be observed in experiment. However, the nucleus can and has been studied in detail via computer simulation. Another experimental difficulty that has slowed progress is due to the fact that nucleation is almost always heterogeneous, i.e., it occurs with the nuclei in contact with surfaces. These surfaces are typically of impurities that are impossible to either eliminate or characterize. As an example, consider the nucleation of rain droplets. This occurs on nanoparticles in Earth's atmosphere[1]. Dusek *et al* [2] studied the nucleation of rain droplets. They observed nucleation due to the nanoparticles under conditions such that the barrier to homogeneous nucleation is estimated to be of order $10^6\,kT$. This is of course large enough to totally eliminate homogeneous nucleation, but under these conditions heterogeneous nucleation is rapid.

Work on homogeneous nucleation in simple liquids, and in water, is already well served by reviews; see Debenedetti's excellent book [3] and the reviews by Oxtoby [4, 5]. Therefore, this review will devote more attention to heterogeneous nucleation than earlier reviews have done. Heterogeneous nucleation is not only much more common than homogeneous nucleation, it is also less well understood, so there is a greater need for work to understand it. This review will also examine experimental work on protein solutions and colloidal suspensions. Experimental work on other systems will not be reviewed. However, see [6] for a recent review on nucleation in metals. Condensed matter physicists in general and those who study protein solutions and colloidal suspensions in particular, have considered mostly homogeneous nucleation. However, materials and atmospheric scientists are more used to dealing with heterogeneous nucleation. We will refer to the most relevant work from the atmospheric and materials science literature below.

This review is structured as follows. The next section is an introduction to the basic theory of nucleation. This includes nucleation at surface transitions, in porous media, and in the presence of soluble impurities. Section 3 considers computer simulation, the most powerful technique we have for understanding nucleation. Section 4 reviews work on the crystallization of proteins. Most of this work is on a single protein, lysozyme. Lysozyme has a fluid–fluid

[1] The situation is similar on Mars [1].

transition within the coexistence region of its solution-to-crystal transition. This has inspired considerable theoretical work on understanding nucleation of one phase transition near to or within a second phase transition. We review this work in section 5. Following this we consider work on nucleation in colloidal suspensions. The final section is a conclusion.

2. Introduction to nucleation

Nucleation is an activated process; it relies on a rare event, the formation of the critical nucleus. This makes the timescale for nucleation much larger than the characteristic timescale of the microscopic dynamics of the system. The critical nucleus has a high free energy which means it has a very low probability of forming within a time of order the characteristic timescale of the microscopic dynamics. The basic physics of nucleation is best illustrated with the help of the simplest theory: classical nucleation theory (CNT). CNT is described in a number of excellent books and reviews, e.g., [3–5, 7] so we will not review it here in any detail, but we will go through the terms in the CNT expression and discuss the underlying physical ideas. For homogeneous nucleation, CNT expresses the rate per unit volume as the product of an exponential factor and a pre-exponential factor or prefactor

$$\text{rate} = \rho Z j \exp\left(-\Delta F^*/kT\right). \tag{1}$$

The exponential factor is $\exp(-\Delta F^*/kT)$, where ΔF^* is the free energy cost of creating the critical nucleus, the nucleus at the top of the barrier. By definition the probability of an event occurring is proportional to the exponential of minus the free energy cost of the event over the thermal energy kT. Thus the exponential factor comes from the, very low, probability of forming a nucleus at the top of the barrier.

The prefactor is the product of the three terms: the number density of molecules ρ, the rate at which molecules attach to the nucleus causing it to grow, j, and the Zeldovich factor Z. The number density of molecules ρ is essentially the number of possible nucleation sites per unit volume, as for homogeneous nucleation the nucleus can form around any one of the molecules present. An upper bound on the rate at which molecules attach to the nucleus causing it to grow, j, is provided by the diffusion-limited flux onto the nucleus. If the critical nucleus is approximated by a sphere of radius R^* then this upper bound is $j \simeq \rho D R^*$, where D is the diffusion constant for the molecules. The growth of the nucleus may however be an activated process, in which case j can be orders of magnitude lower than the diffusion-limited value. Studies of the growth of large protein crystals (the nuclei are too small to study) have found that there the growth is in indeed an activated process [8], and can be a factor of 10^6 slower than the diffusion-limited growth rate. If this also applies to nuclei, j will be a factor of 10^6 smaller than the diffusion-limited value.

The final factor is the Zeldovich factor Z. Essentially the Zeldovich factor is there as the probability that a nucleus at the top of the barrier will go on to form a crystal is less than one. Therefore the rate at which a nucleus actually crosses the barrier and grows into a new phase is Zj not j. The Zeldovich factor Z is approximately equal to $[\partial^2(\Delta F/kT)/\partial n^2]^{1/2}$ evaluated at the top of the barrier. Here ΔF is the free energy as a function of the number of molecules in the nucleus n. The Zeldovich factor is dimensionless. Within CNT, the second derivative of the free energy at the top of the barrier scales as $kT/(n^*)^{4/3}$, where n^* is the excess number of molecules in the critical nucleus. Thus within CNT $Z \simeq 1/(n^*)^{2/3}$, and so for nuclei with $n^* = 10$–100 molecules Z is of order 0.1 or 0.01. Finally, note that the exponential factor typically varies much more rapidly with supersaturation than the prefactor, and so the prefactor is often taken to be a constant.

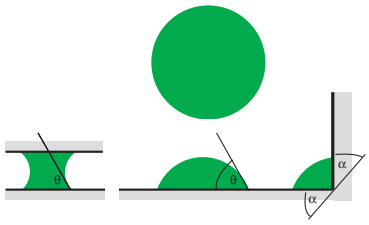

Figure 1. Schematic diagrams of critical nuclei in the bulk, on a flat surface, in a 90° corner or wedge, and in a slit pore. The first three are nuclei of a new bulk phase and are at the same value of the supersaturation. The nucleus in the pore (on the left) is at a lower chemical potential. The transition in the pore is shifted by the interaction with the walls of the pore to lower values of the chemical potential. The contact angle θ is indicated on the nuclei on the flat surface and in the pore. We specify wedge angles via the angle α, which here is 45°. The internal angle, here 90°, equals $180° - 2\alpha$. Note that the radius of curvature of all the critical nuclei except the one in the pore is the same as it is set by the supersaturation via (4). The nucleus for homogeneous nucleation is a sphere, whilst those for heterogeneous nucleation on flat surfaces and in wedges are parts of spheres of the same radius. The schematic diagrams are two-dimensional cross-sections through three-dimensional nuclei.

Having considered the prefactor, let us look at the free-energy barrier. CNT treats the nucleus as if it were a macroscopic phase. If we restrict ourselves to the nucleation of one fluid inside the bulk of another phase, then the nucleus is spherical and its free energy has just two terms: a bulk term and a surface term. See figure 1 for a schematic diagram of such a nucleus. If the nucleus has a radius R then the bulk term is the free-energy change involved in creating a sphere of radius R of the new phase. The surface term is the free-energy cost of the interface at the surface of this sphere. Thus the free energy is

$$\Delta F = -\frac{4\pi}{3} R^3 \rho_n \Delta\mu + 4\pi R^2 \gamma, \tag{2}$$

where $\Delta\mu$ is the chemical potential of the phase the nucleus is forming in minus the chemical potential of the phase nucleating, and γ is the interfacial tension. The number density of the nucleating phase is ρ_n. For nucleation from a dilute solution or suspension, the solution can be treated as an ideal gas and then $\Delta\mu/kT = \ln(\rho_b/\rho_{\mathrm{co}}) = s$, where ρ_b and ρ_{co} are the number densities in the bulk phase in which nucleation is occurring, and at coexistence, respectively. This equation defines s.

The free energy at the top of the barrier, $\Delta F^*_{\mathrm{HOMO}}$, is easily found by setting the derivative of ΔF to zero. Then we have

$$\Delta F^*_{\mathrm{HOMO}} = \frac{16\pi}{3} \frac{\gamma^3}{(\rho_n \Delta\mu)^2}. \tag{3}$$

This occurs for a critical nucleus of radius

$$R^* = 2\gamma/(\rho_n \Delta\mu). \tag{4}$$

If the nucleation barrier ΔF^* is many kT then homogeneous nucleation will be slow. However, it has long been appreciated [3, 9] that nucleation can occur not just in the bulk but with the nucleus in contact with a surface. This is heterogeneous nucleation, and often the interaction between the surface and the nucleus dramatically reduces the free energy of the nucleus. Then the rate near a surface can be many orders of magnitude higher than in the bulk, resulting in heterogeneous nucleation being dominant and homogeneous nucleation being irrelevant.

2.1. Heterogeneous nucleation

Heterogeneous nucleation is also an activated process and so if all nucleation sites on the impurities have identical nucleation barriers then the functional form of the nucleation rate is the same as for homogeneous nucleation. However, the prefactor contains not the number density of molecules ρ but the number density of sites for heterogeneous nucleation, ρ_I, the number density of impurities times the number of places a critical nucleus can form on each impurity. Thus for heterogeneous nucleation with a single nucleation barrier we have

$$\text{rate} = \rho_I Z j \exp(-\Delta F^*/kT), \tag{5}$$

where ΔF^* is the barrier for heterogeneous nucleation, which will depend not only the supersaturation but also on the nature of the surface of the impurity. As impurities are presumably rather variable, in fact we would expect the barrier to vary from nucleation site to nucleation site. We will consider this possibility below, but for the moment we will continue to assume that heterogeneous nucleation is characterized by a single barrier height. This would be the case if the impurity surfaces are uniform.

The simplest surface is a smooth infinite plane with a uniform surface. It is straightforward to extend CNT to nucleation on a smooth plane, and this was done by Volmer in the 1920s [3, 10]. The result is that the barrier is

$$\Delta F^*_{\text{FLAT}} = \Delta F^*_{\text{HOMO}} f(\theta), \tag{6}$$

i.e., it differs from the barrier to homogeneous nucleation only by the presence of a function only of θ. The angle θ is the angle the interface between the nucleus and the bulk phase makes with the surface. It is called the contact angle, and a nucleus on a flat surface is shown in figure 1. The contact angle θ is determined by the interactions between the surface and the molecules in the nucleus. Attractions between the surface and the molecules that are stronger than those between the molecules in the nucleus will lead to a small angle θ as the nucleus spreads into a thin droplet to maximize its contact area with the surface. However, if the surface tends to repel the molecules then the nucleus is pushed away from the surface, resulting in a contact angle $\theta > 90^\circ$.

The angle θ and hence the nucleation rate depends on three interfacial tensions, not one as is the case for homogeneous nucleation. The three interfacial tensions are γ, γ_{bs} and γ_{ns}, which are the tensions of the interfaces between the nucleus and the bulk phase, the bulk phase and the surface, and the nucleus and the surface, respectively. The contact angle θ is related to the three interfacial tensions by Young's equation: $\gamma_{\text{ns}} + \gamma \cos(\theta) = \gamma_{\text{bs}}$.

The function $f(\theta)$ is given by

$$f(\theta) = (1/2) - (3/4)\cos(\theta) + (1/4)\cos^3(\theta). \tag{7}$$

Its two limits are $\theta = 0$, where $f(0) = 0$, and $\theta = 180^\circ$, where $f(180^\circ) = 1$. As $\theta \to 180^\circ$ the nucleus detaches itself from the surface; it becomes a complete sphere not in contact with the surface and so the barriers for heterogeneous and homogeneous nucleation are the same. In this limit there is what is called drying. In the other limit we have wetting. See [11] for an introduction to wetting and drying phenomena. As $\theta \to 0$, $f \to 0$, i.e., the nucleation barrier disappears. This is because when $\theta = 0$ then there is a thermodynamic driving force for forming a thick wetting layer of the new phase even at coexistence ($\Delta\mu = 0$). For $\Delta\mu < 0$ the free energy of a nucleus that looks like a flat layer is a monotonically decreasing function of thickness, and so there is no barrier. There is however a caveat to the statement that there is no barrier to nucleation when $\theta \to 0$. This is that there may be a nucleation barrier to the formation of a wetting layer. The formation of a wetting layer may itself involve a first-order surface phase transition and if so there will be a nucleation barrier to this transition; see

section 2.4. This barrier may result in the surface not being at equilibrium, i.e., not having the wetting layer, in which case there may be a barrier to be overcome before the new bulk phase can appear. The behaviour that underlies this barrier is not captured within the simple CNT for the nucleation of a new bulk phase that we are considering.

In the middle of θ's range, $\theta = 90°$, the surface does not attract the nucleating phase in preference to the bulk phase, i.e., the interfacial tensions γ_{ns} and γ_{bs} are equal. Here we have that $f(90°) = 1/2$. The barrier at the surface is half that in the bulk and so the rate at the surface can be many orders of magnitude higher than in the bulk. The basic physics here is that at the surface there is already an interface with an interfacial free energy cost and so this free-energy cost has already been paid and does not contribute to the free energy of the nucleus. It is essentially the same physics as that which attracts colloidal particles to liquid–liquid interfaces creating Pickering emulsions [12].

So far we have only considered nucleation on a perfectly smooth surface. As nucleation can occur on dirt or any other impurity present, there is no reason to expect the surface to be perfectly smooth on the lengthscale R^*. To illustrate the effect of surface geometry we examine nucleation in wedges. A wedge is shown on the right in figure 1. It is formed from two perfectly smooth planes each at an angle α to a common plane. These planes meet along a line and the internal angle is $180° - 2\alpha$. The internal angle is the angle that defines the volume where the nucleus forms. A wedge with $\alpha = 45°$ and so a 90° internal angle is shown in figure 1. In a wedge the nucleus contacts two surfaces not one as on a flat surface, and so its free energy is lower. Sholl and Fletcher calculated the CNT nucleation barrier in a wedge and found it to have the form [13]

$$\Delta F^*_{\mathrm{WEDGE}} = \Delta F^*_{\mathrm{HOMO}} f_{\mathrm{W}}(\theta, \alpha), \tag{8}$$

i.e., it differs from the barrier to homogeneous nucleation only by the presence of a function of θ and α. The form of f_{W} is not as simple as that of f so we do not reproduce it here. However, its behaviour is simple. For $\alpha = 0$, it reduces to $f(\theta)$. For $\alpha > 0$, $f_{\mathrm{W}}(\theta, \alpha)$ is always less than $f(\theta)$ for all θ, although the difference tends to zero as θ tends to 180° and the nucleus is pushed off a flat surface and out of a wedge. Thus nucleation is always faster in a wedge than on a smooth surface. The most notable feature of f_{W} is that $f_{\mathrm{W}}(\theta, \alpha) \to 0$ as $\theta \to \alpha$, i.e., the nucleation barrier tends to zero as θ approaches α from above, or as α approaches θ from below. For $\theta \leqslant \alpha$ there is no nucleation barrier. So for a narrow wedge, α quite large; even if we are far from wetting there is no nucleation barrier. This should be contrasted with the case for a flat surface where $f \to 0$ as $\theta \to 0$, which is at wetting.

The absence of a barrier for $\theta \leqslant \alpha$ can be understood from elementary thermodynamics. This tells us that at coexistence a wedge with an angle $\theta \leqslant \alpha$ is filled with a macroscopic amount of the liquid phase [14]. If such wedges are present then if we start in the vapour phase and increase the chemical potential up to that at coexistence then when we have reached coexistence the wedges will have filled with liquid. Thus even at a supersaturation of zero, the liquid phase is present. The nucleation barrier is abolished at all supersaturations. This is true even for surfaces with substantial values of θ, provided α is large enough. Thus if we wish to abolish the nucleation barrier we can either use a flat surface which the nucleating phase wets or create narrow wedges with $\alpha \geqslant \theta$.

It should be borne in mind that this is all for a large wedge, with a mouth that is large with respect to the radius of the critical nucleus. Then the nucleation barrier is overcome before the nucleus grows out of the wedge. If the wedge is small, then although filling the wedge may involve a small or non-existent barrier there may be a barrier to nucleation of the new phase out of the mouth of the wedge and into the bulk phase. This barrier has been observed in the studies of nucleation at rectangular pores by Page and Sear [15]. The critical nucleus for nucleation

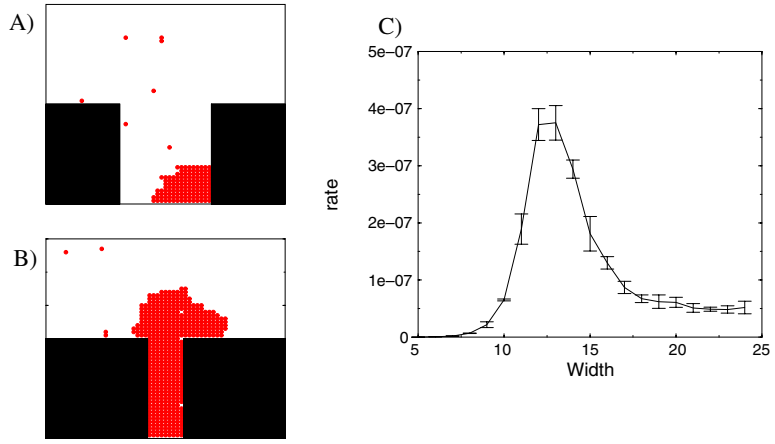

Figure 2. (A) and (B) are snapshots of critical nuclei for nucleation in a pore and out of a pore, respectively. (C) is the rate of nucleation as a function of pore width in lattice sites. The rate is for a single pore and is per simulation cycle. A cycle is one attempted spin flip per spin. The error bars are the standard deviations of the rates obtained in four independent runs. The model is the two-dimensional Ising model with parameters $J/kT = 0.8$ and $h/kT = 0.05$ [15]. The systems are 60 by 60 lattice sites in size and the pores are 30 sites deep and have fixed spins along both sides and the bottom. In (A) and (B) the fixed spins are shown in black, while the up spins are mid-grey (red). The sites with down spins are left white. It is the up-spin phase that is nucleating from the down-spin phase. The pore in (A) is 24 lattice sites across while that in (B) is 9 sites across. (A) and (B) are from figure 1 of [15] and (C) is from figure 3 of [15].

out of a narrow pore is shown in figure 2(B). Pores are microscopic holes or indentations in the surface.

Page and Sear [15] studied nucleation at models of microscopic pores and found two nucleation barriers whose heights varied in opposite directions when the width of the pore was varied. The critical nuclei for the two barriers are shown in figures 2(A) and (B). The fact that one rate increases while the other decreases as the pore width increases results in the overall nucleation rate having a maximum at one value of the pore width. The nucleation rate is shown in figure 2(C). The results are from simulation and the model is the two-dimensional Ising model. Also, Valencia and Lipowsky [16] employed CNT to study nucleation on patterned surfaces, surfaces with small patches on them. They also found two barriers, although there the intermediate state, i.e., the state after one barrier but before the other, had a higher free energy than the initial state. In contrast, for the pores studied by Page and Sear the intermediate state, a filled pore, had a lower free energy than the initial state. Also, Granasy *et al* [17] have presented qualitative results showing nuclei growing on surfaces with a variety of geometries such as steps and groves. These were obtained using their phase-field technique.

Having considered heterogeneous nucleation on flat surfaces, and in wedges and pores, we can compare heterogeneous and homogeneous nucleation. In figure 3 we have plotted barriers as a function of θ for homogeneous nucleation, heterogeneous nucleation on a flat surface, and heterogeneous nucleation in a 90° wedge. We have used a log scale and the nucleation rate varies exponentially with ΔF^*; thus the ratio between the highest and lowest rates plotted is exp(1000), a truly enormous number. At flat surfaces with small contact angles θ, or in narrow wedges with $\theta - \alpha$ small, the barrier to nucleation can be small or negligible even under conditions where the barrier to homogeneous nucleation is huge.

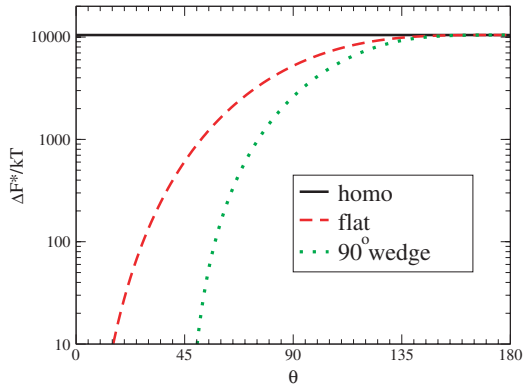

Figure 3. Classical nucleation theory predictions for the nucleation barrier ΔF^* as a function of the contact angle θ between the nucleating phase and a surface. The solid line is the barrier for homogeneous nucleation, the dashed red curve is for heterogeneous nucleation on a smooth flat surface, and the dotted green curve is for heterogeneous nucleation in a right-angled wedge, $\alpha = 45°$. We took $\gamma = kT/a^2$, $\rho_n = 1/a^3$ and a supersaturation $\Delta\mu/kT = 0.04$. Here a is the size of the molecules which provides the microscopic lengthscale but whose value does not affect the barrier height.

If such enormous reductions in the barrier seem unlikely consider the recent work of Dusek *et al* [2] on the nucleation of water droplets (rain) on small (tens of nanometres) particles in the Earth's atmosphere. They observe rapid nucleation on the particles at supersaturations of less than $s = 0.01$. Now, the vapour–liquid interfacial tension of water is known: it is $\gamma = 7 \times 10^{-2}$ J m^{-2} [18]; and the number density of water $\rho_n = 4 \times 10^{28}$ m^{-3}. Thus, the CNT estimate for the nucleation barrier at $s = 0.01$ is $\Delta F^*_{\mathrm{HOMO}} \simeq 10^6\ kT$. Here heterogeneous nucleation of liquid water is occurring at supersaturations such that it is inconceivable that homogeneous nucleation could occur.

The potential of surfaces, especially non-planar surfaces, to dramatically lower the nucleation barrier is presumably the reason why heterogeneous nucleation is much more common than homogeneous nucleation. It is also interesting to note that the limit of supersaturation of the atmosphere is set by the maximum size of the nanoparticles present. Dusek *et al* found a strong dependence of the nucleation rate on the size of the particle the nucleation is occurring on. It is likely that the supersaturation at which nucleation occurs is often determined by the size or type of impurities present. In the experiments of Dusek *et al* the dependence on the maximum size of the particles present is easy to understand. Fletcher used CNT [19] to study heterogeneous nucleation not on infinite surfaces but finite spheres; see [20] for more recent work[2]. He found that if the radius r of the particle is much larger than the radius of the critical nucleus then the barrier is close to that on a flat surface but that as r becomes comparable to R^* the barrier rises steeply. From (4) and using parameter values for water we find that at $s = 0.01$ the critical radius for the nucleation of liquid water is approximately 50 nm. This is comparable to the sizes of the nanoparticles upon which the nucleation is occurring and so we would expect the nucleation barrier to be very sensitive to particle size, which is just what Dusek *et al* found [2].

[2] References [19, 20] consider the nucleation of fluid phases. A computer simulation study [21] of the nucleation of a crystalline phase on a particle found that crystalline nuclei behaved rather differently from fluid nuclei. The crystalline lattice is strained by the need to bend around the curved particle and this strain energy appeared to increase rapidly with the size of the nucleus. This not only affected the nucleation barrier but caused the nucleus to detach from the particle once it had reached a particular size.

2.2. The nucleation theorem

If, as is almost always the case in experiment, the critical nucleus cannot be directly observed, then we cannot directly obtain information about the nucleus. So, we do not know how big it is, what is its structure, etc. However, if we have the rate as a function of chemical potential, then we can use the nucleation theorem to estimate perhaps the most basic property of the nucleus—the number of molecules it contains.

The nucleation theorem [5, 22, 23], relates the rate of change of the free-energy barrier to the excess number of molecules in the nucleus n^*. It is

$$\frac{\partial \Delta F^*}{\partial \mu} = -n^*. \tag{9}$$

It applies to both homogeneous and heterogeneous nucleation. The excess number of molecules in the critical nucleus is, by definition, the difference between the number of molecules in a volume that contains the nucleus and the number of molecules there would be in the same volume without the nucleus. Now, within CNT the derivative of the log of the rate with respect to $s = \Delta\mu/kT$ is

$$\frac{\partial \ln(\text{rate})}{\partial s} = \frac{\partial \ln(\rho_I Z j)}{\partial s} - \frac{\partial \Delta F^*/kT}{\partial s} \tag{10}$$

for heterogeneous nucleation. If we assume that the variation with chemical potential is dominated by the term coming from the barrier then we can neglect the first term on the right-hand side. The second term is simply n^*. Therefore we have

$$\frac{\partial \ln(\text{rate})}{\partial s} \simeq n^*. \tag{11}$$

As we have neglected the prefactor and this is only difference between the CNT expressions for homogeneous and heterogeneous nucleation, this equation also applies to homogeneous nucleation. Using (11), if in an experiment the rate is measured over a range of conditions and the variation of the chemical potential is known then an estimate of the excess number of molecules in the critical nucleus can be obtained.

In order to analyse experimental data, an equation of the CNT form, i.e.,

$$\text{rate} = A \exp(-B/s^2), \tag{12}$$

with A and B constants, can be fitted to the data. See for example [24, 25]. Then, applying the nucleation theorem and assuming (11) holds, we have that the excess number of particles in the critical nucleus is

$$n^* = 2B/s^3. \tag{13}$$

Thus if B is determined for the data the number of molecules in the critical nucleus is known.

Now, if the nucleation is homogeneous then there will be a single barrier height and the form of (12) is reasonable. However, if the nucleation is heterogeneous it will be occurring on impurities and there is no reason to expect the surfaces of the impurities to be completely uniform; the surface chemistry and geometry is expected to vary from point to point. Then instead of there being a single nucleation barrier there will be a range of barriers, with different barriers at different points on the surfaces of the impurities.

2.3. Heterogeneous nucleation with a distribution of barrier heights

If most of the nuclei that form are crossing barriers with the same or almost the same height and CNT is a reasonable description of the physics of the nucleation, then the nucleation rate is given by (5), with a free-energy barrier that scales with s as $1/s^2$. The rate then has the

functional form of (12). However, if for example the nucleation is occurring on surfaces with a wide range of contact angles θ, then the $1/s^2$ scaling of the individual barriers remains but instead of there being one value for B there is a distribution of Bs. This is clear from (6) for the nucleation barrier at a surface. The scaling with supersaturation of individual barriers is the same as that for homogeneous nucleation but each barrier height is a function of θ. Surfaces with smaller contact angles θ will have smaller values of B.

As the impurities are uncharacterized we do not know the values of the Bs. In this case, the author has proposed modelling the Bs with random variables [26–28]. The procedure is to assume that Bs at the different sites are independent random variables with a Gaussian probability distribution of mean $\bar{B}$ and standard deviation σ. The mean and standard deviation are then the parameters of the model. Thus if $p(B)$ is the probability that a nucleation site has a barrier B/s^2 then the total nucleation rate is given by

$$\text{rate} = A \int \mathrm{d}B\, p(B) \exp(-B/s^2), \tag{14}$$

where we have assumed that the prefactor A does not vary from impurity to impurity. For the Gaussian distribution of Bs we obtain

$$\text{rate} = A \exp(-\bar{B}/s^2 + \sigma^2/2s^4). \tag{15}$$

A distribution of barrier heights always results in a rate larger than if all the barriers were of the mean height. This expression is valid so long as $\bar{B} \gg \sigma$ and results in a different, flatter, functional form from that predicted by CNT with a single barrier height. The rate varies more slowly with supersaturation than when there is only a single barrier.

Experimental data have also been analysed [24] by using the nucleation theorem (11) to obtain an estimate of the number of molecules in the critical nucleus. If there is a range of barrier heights there will also be a range of critical-nucleus sizes. With our Gaussian distribution of Bs we have that the rate is given by (15), and so, using the nucleation theorem (11), we have that

$$\frac{\partial \ln(\text{rate})}{\partial s} \simeq \langle n^* \rangle \simeq \frac{2\bar{B}}{s^3} - \frac{2\sigma^2}{s^5} \tag{16}$$

which defines the average of the excess number of molecules in the critical nuclei, $\langle n^* \rangle$. The average is taken weighting critical nuclei by their contribution to the overall rate, $\propto \exp(-B/s^2)$ [28]. The $\simeq$ signs are there as we neglecting the effect of any variation of the prefactor with supersaturation. In equation (16) the second term comes from the distribution of barrier heights and has the opposite sign to the first term. Hence the distribution of barrier heights tends to reduce the variation of the apparent size of the critical nucleus with supersaturation. Thus a distribution of barrier heights is a possible explanation of the finding of Galkin and Vekilov [24] that the size of the nucleus in lysozyme was essentially independent of supersaturation.

2.4. Nucleation at surface phase transitions

So far we have considered only nucleation at a bulk transition. Phase transitions can also occur at surfaces; see [11] for a recent review. If a surface transition is first order then it will start with nucleation. In this section we will briefly consider the possible nucleation behaviour at surface transitions on a large surface immersed in a bulk phase, which may itself undergo a bulk phase transition.

Near and at coexistence for a bulk vapour–liquid transition there may be a surface transition from a state in which the surface is covered with a thin layer of the liquid to one in which

it is covered by a thick layer of the liquid. This is a first-order transition. If it occurs at bulk coexistence it is called a wetting transition and if it occurs off coexistence it is called a prewetting transition. Although there are no quantitative experimental studies of nucleation rates in transitions of this type, there have been studies of superheating and supercooling at wetting transitions [29, 30]. There has also been a significant amount of work on the mechanism of dewetting, which is the process whereby a thick film that covers a surface breaks up into droplets, or evaporates if it is volatile. This starts with either nucleation or an alternative mechanism called spinodal dewetting; see for example [31, 32]. For further work see the recent review of dewetting by Thiele [33]. The nucleation at wetting transitions has been studied theoretically [11, 30, 34, 35].

Of course all nucleation at surface phase transitions occurs at surfaces. However, we can still preserve the distinction between homogeneous and heterogeneous nucleation for surface transitions if we define homogeneous nucleation as being nucleation that occurs on a smooth perfect surface of a substrate while heterogeneous nucleation occurs on defects such as impurities, dust particles, etc. Experimental work has shown nucleation occurring on dust particles [31] but the theoretical work has been on what we have just defined as homogeneous nucleation. There is therefore a need for theoretical and/or computer simulation studies of heterogeneous nucleation of dewetting and other surface transitions.

2.5. Nucleation at phase transitions in porous media

Porous media are solids that are like Swiss cheese: they are full of holes. These holes or pores can be uniform in size and shape, as in zeolites where they are essentially unusually large voids in a crystalline lattice. They can also have a range of sizes and shapes, in which case we have disordered porous media. Phase transitions can occur inside these pores. Pores that are from say a few molecules across to a few hundred molecules across modify but do not destroy phase transitions such as the vapour–liquid transition. See the review of Gelb *et al* [36] and the volume edited by Lu and Zhao [37] for introductions to phase transitions in porous media. Notably, considerable hysteresis is observed for vapour–liquid transitions in porous media, suggesting that the dynamics of the phase transition is hindered by free-energy barriers.

The literature on phase-transition hysteresis in porous media is extensive and we do not have space to review it all here. We will content ourselves by outlining some of the basic physics that is believed to underlie this hysteresis. We will see that it is easy to imagine that the dynamics of pore filling and emptying may be complex, especially for disordered porous media with a range of pore sizes connected by channels with a range of widths. Having looked at this we will review some of the recent work that explicitly calculates nucleation barriers in simple confined systems, e.g., inside cylinders.

The most recent work in this area has focused on the balance between pores filling with a new phase due to nucleation inside the pore itself and filling due to a neighbouring pore being already filled and the new phase nucleating from that neighbouring pore. See the work of Monson and co-workers [38–40] and the work of the group of Neimark [41, 42]. This work uses theory and simulation but does not explicitly calculate nucleation rates. To get a better understanding of the competition between pore filling via nucleation inside the pore and filling via nucleation from an adjacent pore, let us consider the simple system of two connected pores in figure 4.

The two pores in figure 4 are of comparable size and so at equilibrium will fill with liquid at approximately the same value of the chemical potential. However, nucleation barriers need to be overcome. First, let us consider nucleation inside each of these pores. On the left-hand side of the top pore there is a narrow wedge, a wedge with a large value of α. In the section

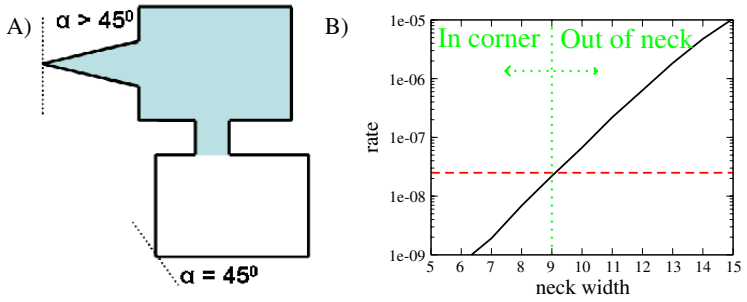

Figure 4. (A) A schematic diagram of a simple model of a porous medium. Porous media consist of interconnected networks of pores. Our model consists of only two, interconnected, pores. The top pore is shown filled with the liquid phase, which is shown in light blue. Corners are defined by their values of α. This is shown explicitly for a right-angled corner and for a corner in a narrow wedge. (B) A plot of nucleation rates as functions of neck width. The black curve is for nucleation out of a neck, and the red dashed curve is for nucleation in a corner. The model is the two-dimensional Ising model, and the contact angle between the nucleating phase and the walls $\theta = 90°$. The rates are from [15]. The vertical dotted green line separates the region where nucleation out of the neck is faster and so dominates from the region where nucleation in the corner is faster and so it dominates.

on heterogeneous nucleation we found that the nucleation barrier tends to zero as $\theta - \alpha$ tends to zero and so wedges with large values of α have small nucleation barriers; see [13]. Thus on increasing the supersaturation nucleation will occur in the narrow wedge with $\alpha > 45°$ before it occurs in the corners with $\alpha = 45°$. The bottom pore only has corners with $\alpha = 45°$, so nucleation in this pore will occur at larger supersaturations than in the top pore.

Thus the first point to note is that nucleation can occur in different pores at different supersaturations, if these pores have different shapes. Now consider what happens when in the top pore nucleation has occurred and so it is filled with the liquid phase while the bottom pore is still filled with the vapour. Of course this bottom pore can still fill via nucleation inside it, but it can also fill via nucleation out of the neck between the filled top pore and the bottom pore. Nucleation out of a neck is shown in figure 2(B). Thus for the pore at the bottom there are two competing mechanisms for filling. The pore will of course fill via the faster process. If the pores are two dimensional and contain Ising spins then we can use the simulation data of [15] to obtain quantitative values for the rates of both processes. For the parameter values of [15], the nucleation rate out of a neck is faster than nucleation in a 90° corner if the neck is more than nine spins wide. Thus if the neck is more than nine spins across it will fill from the pore it is connected to, whereas if the neck is narrower it will fill via nucleation in a corner of the pore itself. This is shown in figure 4(B) where the nucleation rates are plotted as functions of the neck width. To the left of the dotted green line the neck is sufficiently narrow that nucleation in the corner of the pore dominates, while to the right nucleation out of the neck dominates.

In parallel to the work on understanding the balance between pores filling from adjacent pores or via nucleation inside them, there have been a number of papers [43–48] dedicated to analysing the nucleus inside the confined environment of a pore. On a flat surface the nucleus is part of a sphere, but in pores the nucleus may adopt other, more complex shapes. For example, in a narrow slit pore the nucleus may span the pore. This is shown in figure 1 on the left. See [43–47] for studies of the nucleus shape and barrier height for nucleation in slit pores. Slit pores are the space between two parallel planes a microscopic distance apart. Nucleation in cylindrical pores has also been studied [45, 46, 49]. Slits and cylinders are the two simplest

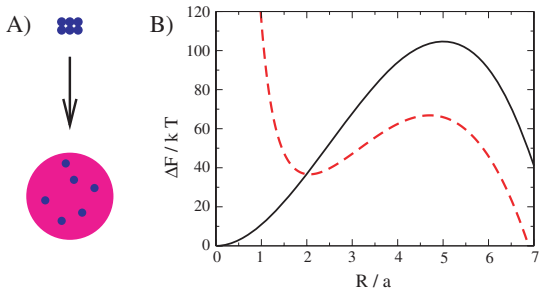

Figure 5. (A) A schematic diagram of nucleation occurring on a cluster of molecules. These attract each other strongly in the bulk phase, and so form a compact cluster there, but are highly soluble in the nucleating phase. At the top is shown the tightly bound cluster of six impurity molecules in the bulk phase. They are shown as small dark grey (blue) circles. Below is shown a nucleus (mid-grey (magenta)) of the new phase, with the six impurity molecules diffusing inside it. (B) A plot of the free energy of a nucleus as a function of its radius R. The solid black curve is for homogeneous nucleation within the standard CNT (2). The dashed red curve is for nucleation on $n = 10$ soluble impurity molecules. The free energy plotted is that of (17) but with ΔF_{SOL} subtracted off. The parameter values are $\rho_n = 1/a^3$, $\gamma = kT/a^2$, $B_2 = 4a^3$ and $\Delta\mu/kT = 0.4$.

geometries for a pore. This is all theoretical work; it is not clear how a nucleus in a porous medium can be observed in an experiment.

2.6. Nucleation on soluble impurities

We have considered heterogeneous nucleation on surfaces that are flat and surfaces that have some curvature, e.g., have corners. However, all these surfaces have been inert solids in the sense that although they interact with the nucleus the surfaces themselves are not changed by this interaction. There is no reason to expect that all impurities are inert in this sense: in particular, they may be soluble in the nucleating phase. The effect of soluble impurities on the nucleation of water droplets was studied by Köhler between the wars [50]. He considered small particles of sodium chloride acting to promote the nucleation of water. Sodium chloride is of course highly soluble in water. Since then atmospheric scientists have made a number of studies of this phenomenon [51]. Also, Talanquer and Oxtoby [52], and Reiss and co-workers [53, 54] have studied nucleation on soluble particles; see also [55].

As a complete treatment of nucleation in the presence of soluble impurities is complex [53–55], we will restrict ourselves to outlining a simple illustrative theory for the special case of an impurity that consists only of a handful of soluble molecules. See [53–55] for more detailed treatments. Schematic diagrams of the impurity before and after the nucleus forms around it are shown in figure 5(A). In general the situation will be more complex than this; for example, the impurity might consist of an insoluble solid surface with soluble impurity molecules on top of it. If a nucleus has formed around a small cluster of n molecules, then the CNT free energy acquires extra terms due to these molecules. Equation (2) becomes

$$\Delta F = -\tfrac{4}{3}\pi R^3 \rho_n \Delta\mu + 4\pi R^2\gamma + \Delta F_{\mathrm{SOL}} + nkT \ln[n/(4\pi R^3/3)] + n[n/(4\pi R^3/3)]B_2 + \cdots. \tag{17}$$

The last three terms are the first three terms in a density or virial expansion for the free-energy change on going from an aggregate of n impurity molecules to n molecules in solution inside the nucleus. These n molecules are assumed to be highly soluble in the nucleating phase and so they diffuse freely inside the nucleus. We also assume them to be involatile in the nucleating

phase, so we work at fixed number of molecules n in the nucleus. This implies that the free-energy cost of one of these molecules evaporating from the nucleus must be large enough to prevent them leaving the nucleus during the timescale of the nucleation.

The first additional term in (17), ΔF_{SOL}, is the density-independent part of the free-energy change in going from a compact impurity particle of n molecules to the n molecules free in solution in the nucleus. As it is independent of the radius R of the nucleus we will find that it does not affect the height of the free-energy barrier. However, of course the free-energy change of dissolving the impurities must be negative. The second term is the ideal gas term for the n molecules. The third term is the lowest-order term from the interactions in between the n molecules; it is the second virial coefficient term [56]. The second virial coefficient is B_2, which is positive if the molecules repel each other, which they will if they are highly soluble in the nucleating phase. There are also, third, fourth, etc virial terms, but we neglect them.

We have plotted a free energy ΔF for a nucleus with $n = 10$ molecules dissolved in it in figure 5(B). The most striking feature of the plot is that the free energy now has a minimum at a small value of R, call it R_{M}. This is close to $2a$ in figure 5(B). A droplet at a minimum in the free energy is of course stable. Without the soluble impurities the free energy at small R is dominated by the surface term. This is a monotonically increasing function of R, and so there is no minimum. However, if the nucleus contains impurities their contribution to the free energy decreases with increasing radius. Thus their contribution competes with the surface tension and there is a minimum in the free energy as a function of radius. This minimum is at the radius where the forces from the surface tension and the impurities balance. For radii less than R_{M}, the free energy actually decreases as R increases and so nuclei of radius R_{M} will form spontaneously. Their formation does not require the crossing of a barrier.

However, in figure 5(B) we see that between $R = R_{\mathrm{M}}$ and the large values of R where the free energy decreases monotonically, there is still a free-energy barrier. This must be overcome in order to reach the large values of R where the monotonically decreasing free energy provides a driving force for the growth of the nucleus into a new bulk phase. The height of this barrier is $\Delta F^* - \Delta F(R = R_{\mathrm{M}})$. Because it is difference between these two free energies the ΔF_{SOL} terms cancel and so the nucleation barrier does not contain ΔF_{SOL}. Note that the impurities have significantly reduced the nucleation barrier. The free-energy barrier for nucleation is $105\,kT$ without the impurities and $30\,kT$ with them. This corresponds to a speed up by a factor of exp(75).

3. Computer simulation

It is the free energy of the nucleus at the top of the barrier that determines the rate. Here the nucleus is at a maximum in the free energy and so the nucleus will only remain there for a time of order j^{-1}. As only one nucleus is required to obtain a new bulk phase and as it persists for such a small time, studying it in experiment is typically not possible. Exceptions are colloidal suspensions such as those studied by Gasser *et al* [57], where the particles are larger than the wavelength of light and so can be seen in a microscope, and where the timescale j^{-1} is of order 10 s [57]. In protein solutions the nuclei are much smaller than the wavelength of light and so cannot be detected via light microscopy. There has been a study using atomic force microscopy of clusters of protein molecules in a crystallizing solution [58], but it is not clear that the clusters observed were nuclei near the top of the barrier. Thus except in colloidal suspensions there have been no clear observations of the critical nucleus in experiment. From a theoretical perspective, nuclei are an awkward size: they typically contain somewhere between tens and thousands of molecules. It is easier to deal with one or two, or an infinite number (the thermodynamic limit) than tens or thousands of molecules. This has left computer simulation as our primary

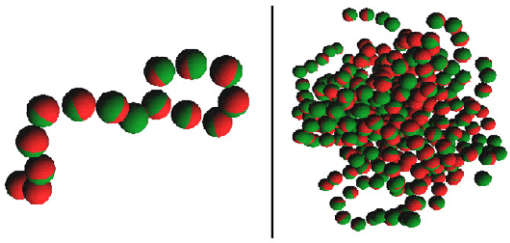

Figure 6. Simulation snapshots of nuclei of the liquid phase for the Stockmayer model of polar fluids. They are from the work of ten Wolde *et al* [59]. The Stockmayer potential is a Lennard-Jones potential plus a permanent dipole. The snapshot on the left is a small nucleus (much smaller than the critical nucleus) showing that small nuclei are chain-like. The snapshot on the right is of a much larger nucleus that is now quite globular but still has 'hair' of chains of the particles on its surface. This unusual structure is not taken into account by CNT and could only have been found via simulation.

technique for understanding nucleation, although density-functional theory calculations have also been important.

Computer simulation of nucleation can require specialized techniques, but its key advantage is that it is exact, unlike density-functional theory, and crucially it allows us to see and to measure precisely the properties of the critical nucleus. Frenkel and co-workers in particular have obtained many beautiful and insightful results on the structure of nuclei. For example, they studied the nucleation of a model particle with a large dipole moment [59, 60]. Dipoles tend to align nose-to-tail into chains and these chains perturb the structure of the nucleus, making it a little hairy. Computer simulations not only showed the structure of the nucleus for the first time but related this to CNT's underestimate of the nucleation barrier. Two snapshots from the work of ten Wolde *et al* [59] are shown in figure 6.

The specialized techniques are required due to the often huge timescales involved. In figure 2 there are rates of 10^{-7} per simulation cycle. Therefore even a rough estimate of the rate via simple direct simulations would require say 100 simulations each of approximately 10^7 cycles. This would require a prohibitive amount of computer time, and rates many orders of magnitude lower are often encountered.

3.1. Simulation techniques

The best established technique for the computer simulation of low nucleation rates is umbrella sampling. See the book of Frenkel and Smit [61] for an introduction. The umbrella-sampling technique involves splitting the rate into a product of a prefactor and $\exp(-\Delta F^*/kT)$, as in CNT. The free-energy barrier is determined by defining an order parameter; the value of this order parameter must be small in the phase the nucleation is occurring in and large in the nucleating phase (or vice versa) so that it distinguishes between the two phases. Then an artificial potential that is a function of this order parameter is imposed on the system. This potential is used to push the system to the top of the barrier and as this is done the work done to push the nucleus to the top of the barrier is calculated. This work done is ΔF^*. The prefactor is then found by simulating the nucleus near the top of the barrier. For detailed descriptions of this technique applied to a specific system, hard spheres, see [62, 63]. For a recent discussion of the role of order parameters, see [64].

Rates have also been calculated via transition path sampling techniques [65], which, as the name suggests, calculate rates by sampling the paths between the initial state, the phase before a

nucleus forms, and a final state, which is a nucleus that is a long way past the top of the barrier. These techniques have been pioneered by Chandler and co-workers. See [66] for an application by this group to nucleation in the Ising model. Recently, techniques called transition interface sampling have been developed by Bolhuis and co-workers, from transition path sampling techniques [67, 68]. They have been further developed by Allen and co-workers [69–71]. As their name suggests, these techniques sample transition paths by chopping the paths into short segments with a series of surfaces or interfaces in order-parameter space. The nucleation rate of sodium chloride crystals from the melt has been calculated [72] via both umbrella-sampling and transition-interface-sampling techniques [69–71]. The two techniques gave consistent rates.

4. Nucleation in protein solutions

Unfortunately, quantitative studies of the nucleation of protein crystals are rare. This is despite the importance of protein crystallization for protein structure determination [73]. Lysozyme is the only protein for which there have been quantitative studies of the nucleation rate of the crystalline phase. These have been performed by a number of groups [24, 25, 74–77]. The results obtained by different groups do not agree with each other [78–80]. Lysozyme is itself an atypical protein in the sense that it crystallizes very readily. Biologists working to crystallize proteins in order to determine their structure are typically working with proteins that are very hard to crystallize; the easy-to-crystallize proteins were crystallized years or decades ago. Thus it is not clear how much of what we learn about lysozyme will be applicable to other proteins [81]. Proteins are of course a very diverse set of molecules.

Most of this section will be devoted to a review of the experimental results for lysozyme, although we will briefly consider the, largely qualitative, work on the nucleation of other proteins. Our analysis of the data for lysozyme follows that of [28]. We will start with the pioneering work of Galkin and Vekilov [24, 74, 75]. They studied the nucleation of lysozyme in small droplets. They fitted their rates as a function of supersaturation to an equation of the CNT form, (12), and obtained [24]

$$\text{rate} = 10^9 \exp(-65/s^2)\ \text{m}^{-3}\ \text{s}^{-1} \tag{18}$$

with $s = \ln(\rho/\rho_{\text{co}})$ the supersaturation $\Delta\mu$ divided by kT, assuming the solution is ideal. The nucleation occurred at concentrations around 50 mg ml^{-1}, which correspond to volume fractions around 5%. Equation (18) is equation (5) of [24], converted to SI units. From the factor in the exponent they obtain a crystal–solution interfacial tension of $\gamma \simeq 1$ mJ m^{-2}. Taking lysozyme to be 3 nm across, a surface tension of kT per area of 3×3 nm^2, is 0.4 mJ m^{-2}. Thus the interfacial tension is about $2.5kT$ per protein molecule at the surface. This is very reasonable. We expect that the free-energy strength of the bonds that hold the protein molecules together in the protein crystal to be a few kT.

Other groups have also made quantitative measurements of the nucleation rate of lysozyme [25, 76, 77] and the rates measured by different groups with different methods do not agree. This has created some controversy [78–80]. It was suggested that the higher rates include nucleation that is heterogeneous even though the claim was that the experiments measured the rate of homogeneous nucleation.

In their experiments, Galkin and Vekilov observed a few nuclei that seem to form almost instantaneously in their experiments. They ascribed these nuclei to heterogeneous nucleation. Having discounted these nuclei they counted the number of the nuclei that formed over longer periods of time and used this number to obtain the rate of (18). They attributed this nucleation rate to homogeneous nucleation. Now, within CNT, the prefactor, which Galkin and Vekilov estimate to be equal to 10^9 m^{-3} s^{-1}, is equal to ρZj; see (1). The Zeldovich

factor Z is approximately the curvature of the free energy at the top of the barrier divided by kT. Hence it is dimensionless and typically of order 0.01–0.1 [3, 28]. For a dilute protein solution, $\rho = 10^{24}$ m^{-3}. This leaves the growth rate j. The growth of critical nuclei has not been measured but the growth of large crystals has been [8]. These rates are sensitive to supersaturation but can be significantly less than 1 nm s^{-1} and are rarely much more. This corresponds to an attachment flux of less than one molecule per lattice site per second. However, even if we set the growth rate j for the entire nucleus to be 0.1 and $Z = 0.01$, we still obtain a prefactor of 10^{21} m^{-3} s^{-1}. This is 12 orders of magnitude larger than that in Galkin and Vekilov's fit [24]. The data of Galkin and Vekilov, and the similar data of Bhamidi *et al* [25] seem incompatible with homogeneous nucleation as described by CNT. Either CNT is failing or the nucleation is heterogeneous.

In heterogeneous nucleation the prefactor is $\rho_I Z j$, see (5), where ρ_I is the number density of sites at which heterogeneous nucleation can occur. Thus if we assume that the density of impurities ρ_I is approximately 12 orders of magnitude lower than ρ we can obtain the value of the prefactor in Galkin and Vekilov's fit [24]. Thus the nucleation data [24, 25, 74, 75] are consistent with the nucleation being heterogeneous.

We have seen that the small size of the prefactor implies nucleation on impurities present at number densities approximately 10^{12} times lower than the number density of protein molecules, and we lack the ability to detect impurities at this concentration. However, there are impurities that are known to be present that occur at much higher concentrations. The protein lysozyme tends to dimerize and these dimers are impurities. The concentration of these dimers has been quoted as close to 1% in commercial lysozyme [82]. Burke *et al* [82] have studied the effect of dimer impurities on nucleation by systematically varying their concentration. They find that increasing the concentration of dimer from close to zero to 1% of the lysozyme concentration increases the nucleation rate significantly. This is consistent with the results of recent simulations of a simple model [83], the Ising model. These simulations found that impurities only two lattice sites long can significantly decrease the barrier to nucleation.

Finally, Galkin and Vekilov [24] also used the nucleation theorem [22, 23] to estimate the size of the critical nucleus. They obtained a size of close to ten molecules, which was essentially constant over the range of supersaturations they were able to cover. This is consistent with heterogeneous nucleation with a wide range of nucleation barriers; see section 2.3 and [28].

One interesting approach to distinguishing between homogeneous and heterogeneous nucleation is via the use of what is often called either the Kolmogorov–Johnson–Mehl–Avrami (KJMA) approximation or Avrami's law [84, 85]. This is based on the idea that *if* the nucleation is homogeneous and only a small fraction of the initial phase has transformed into the equilibrium phase that is nucleating, then the fraction of the volume occupied by the equilibrium phase should vary as t^4. The fourth power of time is due to a third power of time from the growing volume of domains whose linear dimension is increasing linearly with time, plus an extra power of time due to the number of domains increasing linearly with time. Thus the t^4 dependence is a consequence of a constant rate of domain growth, and a constant domain nucleation rate. Both should be constant so long as the supersaturation is constant, if the nucleation is homogeneous. If, however, the nucleation is heterogeneous and occurs on a few nucleation sites where the barrier is unusually low, then these sites may be used up early on, resulting in the nucleation rate dropping even at constant supersaturation. Then the fraction of the volume occupied by the new phase will vary with time more slowly than t^4; we expect the dependence to be closer to t^3. Using a phase-field model Castro and co-workers [84, 85] have studied this phenomenon and observe precisely this effect. This may prove to be a useful way of inferring the origin of the nucleation in protein solutions without needing to observe the nucleation itself.

4.1. Proteins other than lysozyme

The more quantitative work by physical scientists on the nucleation of protein crystals has concentrated on lysozyme. There is quantitative data for lysozyme [24, 25, 75] but only qualitative studies have been performed for other proteins [58, 86–88]. This is despite the fact that nucleation is essential to crystallization, and protein crystals are required by structural biologists for structure determination via x-ray diffraction [73, 89]. The work by the scientists who work to crystallize proteins has typically been highly empirical [73] and has focused on developing techniques that improve the probability of obtaining the desired large crystals. For example, over the last 20 years there has been work on developing nucleants: substances that are added to the protein solutions to induce nucleation [73, 90, 91].

Nucleants work by providing surfaces where the barrier to heterogeneous nucleation is low. The most effective nucleants are disordered porous media [73, 91]. As we saw in the sections on heterogeneous nucleation and on nucleation in porous media, nucleation can be much faster in a pore than on a flat surface. This is presumably the explanation for the fact that porous media are better at inducing nucleation than are smooth flat surfaces, although we have no direct experimental evidence of nucleation at these pores. The pores in the porous media have sizes comparable to what we expect the size of the critical nucleus to be, which is around 10 nm. The fact that disordered porous media are more effective than ordered porous media is interesting. A possible explanation follows from the fact that the nucleation barrier is very sensitive to the size of the pore, see figure 2, and it is likely that it is also highly sensitive to its shape. Then by chance a few of the pores of a sample of disordered media may have just the right size and shape to dramatically lower the barrier to nucleation. In an ordered porous medium, as all the pores are the same it is highly unlikely that this one size and shape will be optimal. See [15, 26, 27, 92] for work on this problem. It may be that disorder is generically useful for promoting rare events, for example protein misfolding [93] or even evolution [92].

5. Nucleation when there is more than one phase transformation possible

In solutions of some proteins, for example lysozyme [94], there is a fluid–fluid transition within the fluid–crystal transition, i.e., there is a second phase transition in addition to the phase transition whose dynamics we are interested in. Thus at sufficiently large supersaturations a lysozyme solution may have a higher chemical potential than two other phases. One of these is the crystal and the other is a protein solution with a higher concentration, a protein-rich solution. A calculated phase diagram that is qualitatively the same as that of lysozyme [94] is shown in figure 7. It was calculated using the type of simple model of proteins proposed in [95]. Consider a protein solution that is to the right of the solution–crystal coexistence curve (the solid black curve) *and* within the fluid–fluid phase transition (the dashed green curve). Two different phases could nucleate: the crystal and the protein-rich solution. The obvious question then is: Which one appears first? To answer this question it is necessary to know the two nucleation rates. For example, CNT predicts that the barrier height scales as $\gamma^3/(\Delta\mu)^2$. Now the crystal phase has the lower chemical potential which tends to make the nucleation barrier for crystallization smaller than the one for the fluid–fluid transition, but if the interfacial tension between the protein-rich and protein-poor solutions is much smaller than that between the solution and the crystal then CNT would predict that the protein-rich phase would nucleate first.

Also, if the crystalline phase forms then we are at equilibrium and the system will stop evolving; however, if the protein-rich phase forms and coexists with protein-poor solution then we are not yet at equilibrium and the crystalline phase may nucleate in either or both

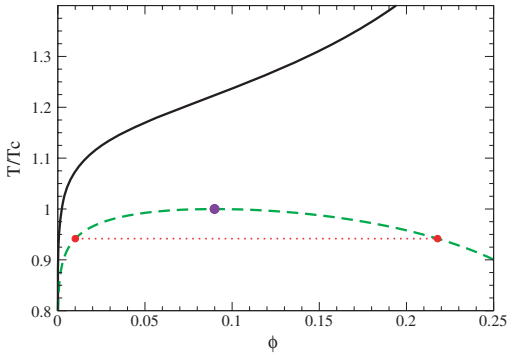

Figure 7. Theoretical phase diagram of a simple model of a protein, in the temperature–concentration plane. T is the temperature, T_c is the temperature of the critical point, and ϕ is the volume fraction of the protein. The solid black curve is the volume fraction of the dilute protein solution that coexists with the crystalline phase. The volume fraction of the crystal is not shown; it is above the maximum volume fraction shown. The dashed green curve is the coexistence curve for the transition between a protein-poor solution and protein-rich solution. The critical point is marked by a larger (purple) circle. A tie line connecting two coexisting fluid phases is marked by a dotted red line and the densities of these two phases are indicated by smaller (red) circles.

the coexisting phases. At coexistence the two fluids will have the same chemical potentials and hence the same supersaturations, but presumably the interfacial tension between the crystal and the protein-rich solution will be lower than that between the crystal and the dilute protein solution. Then for homogeneous nucleation CNT would predict that the crystal would nucleate in the protein-rich solution.

Galkin and Vekilov [96] measured the nucleation rate of lysozyme crystals in the vicinity of the transition between the protein-poor and protein-rich phases. This was at solution conditions that are different from those in the work discussed in the previous section [24]. They found that as they cooled the solution, thus increasing its supersaturation, the nucleation rate went through a maximum. This is very different from the usual rapid and monotonic increase of the rate with increasing supersaturation. For a number of different solution conditions the maximum was consistently close to the coexistence curve, i.e., it occurred very near where the solution crossed the line (on the protein-poor side) into the coexistence region. This is clearly seen in figure 8, which is reproduced from [96]. The vertical dotted lines mark the position of the coexistence curve and the maxima in the curves all lie very close to these dotted lines. The obvious inference is that the fluid–fluid transition is directly affecting the nucleation, perhaps because the protein-rich phase is nucleating first. Vekilov and co-workers have developed and studied models of the effect of the fluid–fluid transition on the nucleation of the crystalline phase of lysozyme [97–99]. The study of Pan *et al* [98] in particular employed a two-step nucleation process to explain the maximum in the nucleation rate as a function of temperature. Indeed they were able to reproduce the experimental behaviour.

Lysozyme is not alone in having a fluid–fluid transition. They appear frequently during attempts to crystallize proteins; see the work of McPherson and co-workers [88] for examples. Unfortunately, there have been no quantitative studies of the effect of the fluid–fluid transition on the nucleation of crystals of proteins other than lysozyme. However, there have been observations that in some systems the appearance of droplets of one fluid phase coexisting with another is associated with the nucleation of the crystal phase [88]. Also, the crystals may appear at the interface between the fluid coexisting faces, suggesting heterogeneous nucleation of the crystalline phase at a fluid–fluid interface.

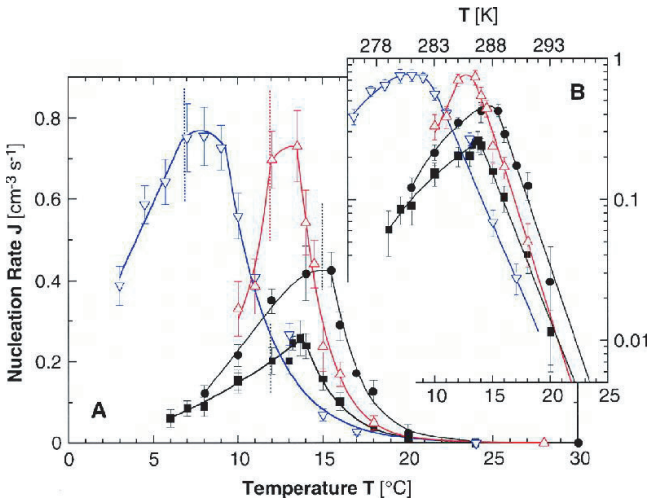

Figure 8. The rate of nucleation of lysozyme crystals, J, as a function of temperature T, at pH 4.5 in a 50 mM sodium acetate buffer with 4% (wt/vol) NaCl. The black squares are for a lysozyme concentration of 50 mg ml^{-1}, and with no additives; the black circles are for a concentration 80 mg ml^{-1}, no additives; the down-pointing triangles are for a concentration 50 mg ml^{-1}, and with 5% (vol/vol) glycerol; the up-pointing triangles are for 50 mg ml^{-1}, with 0.2% (wt/vol) PEG 5000. The vertical dotted lines indicate the respective temperatures of fluid–fluid separation, and the curves are just guides for the eye. The inset is the same data but on a semilogarithmic scale. This is figure 1 of [96].

Having outlined the problem of nucleation in the presence of an additional phase transition, we will consider the theoretical work on understanding this problem. We will look at homogeneous and heterogeneous nucleation in turn.

5.1. Theory of homogeneous nucleation near a second first-order phase transition

First-order fluid–fluid transitions end at a critical point, see figure 7, where there are large lengthscale critical fluctuations [100, 101]. It was these fluctuations and their effect on the nucleus for homogeneous nucleation that drew the attention of ten Wolde and Frenkel [102]. Their interesting findings have inspired work on nucleation both near the critical point, and below it where the fluctuations are absent but there can be a sudden jump in protein concentration when the protein-rich phase appears. This jump in concentration will cause the environment of a nucleus to change discontinuously and so the barrier to nucleation will also jump. In this section we will first consider nucleation near a critical point, before considering nucleation below the critical point.

The idea of ten Wolde and Frenkel was that near a critical point there are large fluctuations, and as the nucleus is itself a large-amplitude fluctuation, the nucleus could 'piggy-back' on the critical fluctuations, resulting in a lowering of the free energy of the nucleus. Their computer simulations showed that indeed the free-energy barrier is lower than CNT would predict. They also found that the excess number of molecules in the nucleus, n^*, was unusually large and that many of these molecules were not in a crystalline environment. Nuclei near a critical point were further studied by the author, who used first a simple mean-field density-functional theory [103], and then scaling arguments [104]. These calculations showed that the crystalline core of the nucleus could be viewed as a perturbation to the critical fluctuations. Essentially,

the nucleus consists of a crystalline core surrounded by a 'corona' where the density was higher than in the bulk but where there was no crystalline ordering. The number of molecules in this corona was shown to scale as the compressibility [103, 104]. The compressibility diverges at the critical point and therefore so does n^*. Remarkably, as the critical point is approached the number of molecules in the nucleus tends to infinity. According to the nucleation theorem, see section 2.2, the rate of change of the log of the rate should be approximately equal to n^*. Thus as the critical point is approached along a critical isotherm and n^* becomes very large, the barrier height will decrease faster and faster with increasing chemical potential. Density-functional calculations show this clearly [105–107].

So, we have seen that there is interesting new physics to be found in the homogeneous nucleation of one new phase near a critical point associated with another phase transition. However, it is not clear that this has any relevance to nucleation as observed in lysozyme, or any other protein. The fluctuations do not seem to reduce the barrier dramatically, i.e., they will not reduce a barrier that would otherwise be hundreds of kT or more to a level that is low enough to achieve nucleation. Density-functional calculations are much quicker than simulations and so they allow the calculation of the barrier over wide parameter ranges. Shiryayev and Gunton [105] performed such calculations and found that at constant supersaturation the barrier was a few tens of per cent lower near the critical temperature than far above it. They used the density-functional model Talanquer and Oxtoby [107] had developed and studied. These reductions are very modest in comparison to the reductions in the barrier height that can occur near surfaces; see figure 3. Thus the barrier for heterogeneous nucleation may be far below that for homogeneous nucleation even near a critical point.

Before we consider nucleation below the critical point we should mention the work of Pini and co-workers [108]. They considered the effect of fluctuations near a critical point when quite long-range repulsive interactions are introduced into the pair potential. These strongly promote density fluctuations with non-zero wavevector; indeed, if they are strong enough they cause the fluid–fluid transition to transform into microphase separation. Thus it is possible that such fluctuations may lower the nucleation barrier by much larger amounts than do the fluctuations near a conventional critical point. However, there is no evidence for such repulsions near any critical point of a protein solution.

Having considered nucleation near the critical point, we will consider nucleation below the critical temperature. If we start off at low density, outside the coexistence curve of figure 7, then only the crystalline phase can nucleate. However, as the density of the solution is increased into the coexistence region, then either a crystal or a protein-rich solution can nucleate from the protein solution. If the protein-rich solution appears first then the crystal can nucleate from within it. Thus the crystal can form via two successive nucleation steps: nucleation of the protein-rich phase and then nucleation of the crystal in the protein-rich phase.

At coexistence the protein-poor and protein-rich phases have the same chemical potential; therefore, the $\Delta\mu$ in the CNT equation for the barrier, (3), will be the same. However, the interfacial tensions of the nucleus will be different in the two phases. Thus γ in (3) will differ and the nucleation rates will be different in the two phases. We expect it to be lower in the protein-rich phase. This is because the difference in protein density between the crystal and the protein-rich phase will be much smaller than that between the crystal and the protein-poor phase. This will reduce the density gradient across the interface and gradients typically cost free energy [11, 107]. As the free-energy barrier is different in the coexisting phases then as the fluid–fluid transition is traversed the rate of homogeneous nucleation will jump. This has been studied [109]. It was found that the jump in rate is upwards and can be quite large [109]. It could easily result in the rate increasing from being too slow to being observable as a rapid rate.

It should be borne in mind that on increasing supersaturation the phase transition from the protein-poor phase to the protein-rich phase will not be instantaneous. As it is a first-order phase transition it will proceed via nucleation and growth. If the rate of nucleation of the protein-rich phase in the protein-poor phase is comparable to the rate of nucleation of the crystal in the protein-rich phase, then the system may contain a mixture of all three phases. The dynamics of the phase transformation may then be complex. Finally, Auer and Frenkel have also considered the effect of a solid–solid, not fluid–fluid, transition within the fluid–crystal coexistence region. They find unusual critical nuclei and as a result large deviations from the CNT prediction [110, 111].

5.2. Theory of heterogeneous nucleation near a second first-order phase transition

Heterogeneous nucleation of a new phase will also be affected by another phase transition; indeed, because heterogeneous nucleation occurs at a surface, it will be affected by both bulk and surface phase transitions. Thus if we see a protein solution go through its bulk phase separation into protein-rich and protein-poor phases we know that the rates of both homogeneous and heterogeneous nucleation will change. However, if on a surface of a microscopic impurity in the solution there is a surface phase transition this will also change the rate of heterogeneous nucleation. This surface phase transition may be hard to detect in experiment as we expect it to have few consequences other than dramatically changing the rate of heterogeneous nucleation.

As a first-order surface phase transition is crossed the rate of heterogeneous nucleation jumps [112], and is discontinuous, just as the rate of homogeneous nucleation jumps as a bulk first-order transition is crossed. The situation for heterogeneous nucleation near a *bulk* first-order phase transition is more complex [113]. Close to but below the critical point any surface that attracts the protein molecules will be wet by the protein-rich phase [11, 114]. This means that at coexistence between the protein-poor and protein-rich phases a surface in the bulk protein-poor phase will *not* in fact be in direct contact with this phase; there will be a thick layer of the protein-rich phase in contact with the surface and lying between this surface and the protein-poor phase. This is called a wetting layer, and at coexistence it has a thickness that is typically limited by gravity. Wetting layers have been extensively studied; see the reviews [11, 114]. We are interested in surfaces that attract protein molecules because it is there that we expect the nucleation barrier to be low[3].

If surfaces in the protein-poor phase are wet by the protein-rich phase then surfaces in both bulk phases are in contact with the protein-rich phase and hence the rate of heterogeneous nucleation is the same in the two coexisting bulk phases. If the nucleation is occurring on a impurity it does not matter which bulk phase the impurity is in. This implies that there is *no* jump in the nucleation rate as the first-order bulk transition is crossed [113]. If the surface is wet then as the bulk phase transition is approached from the low protein concentration side then a layer of the protein-rich phase forms in contact with the surface. Its thickness then grows continuously as the fluid–fluid transition is approached until at coexistence its thickness is limited only by gravity. It is the continuous growth of the layer of the protein-rich phase that smoothes over the jump in the environment of the nucleus and hence the jump in the rate. See [113] for details. If there is no wetting, which may be the case far below the critical temperature, then in the absence of the wetting layer the environments of the nucleus in the two coexisting phases are different and there is a jump in the rate of heterogeneous nucleation as the bulk phase transition is crossed [113].

[3] It is important to distinguish between surfaces which the protein-rich phase wets, and surfaces which the crystal itself wets. In the second case there is no nucleation barrier, see section 2.1, whereas in the first case the protein-rich phase will form and although this will typically reduce the barrier it will not in general reduce it to zero.

6. Nucleation in colloidal suspensions

Colloidal suspensions have two advantages as systems: the particles are large and slow moving, and the effective interparticle potentials can be controlled. For example, they can be made to interact with a potential that is close to the simple hard-sphere potential. Larger colloidal particles have diameters greater than the wavelength of light and so individual particles of this size can be followed by light microscopy. This was exploited by Gasser *et al* [57], who for the first time actually followed a nucleus growing and surmounting the barrier. They could only do so because the particles were large enough to be visible in a light microscope. Their colloidal particles were spherical and charged. The technique of confocal microscopy has enormous potential as colloidal particles with more complex potentials could be studied and the effect of these potentials on the nucleus studied in detail. Heterogeneous nucleation could also be studied. The ability to control the interaction potential in colloidal systems also makes comparison with simulation and theory easier, which in turn makes it easier to understand the nucleation. For example, if particles are known to interact via a potential that is close to the simple hard-sphere potential then theoretical calculations and simulations may be performed for this potential and the results compared to experiment. Then any difference between the findings in experiment and those in theory must be due to factors such as approximations made in the theory. Also, exact simulations should always agree with the experimental results. This is in contrast to, for example, water, where differences between experiment and simulation may be due to inadequacies in the model used in the simulations. The advantages we have discussed apply not only to nucleation but to other aspects of the phase-transition kinetics. See the recent review of Anderson and Lekkerkerker [115] for an excellent introduction to phase transition kinetics in colloidal suspensions.

It should, however, be borne in mind that synthetic colloidal particles do come with their own problem: polydispersity. Unlike molecules which (excluding isotope effects) are all identical, no two colloidal particles are identical. In particular for spherical colloids there will inevitably be some spread in the diameters. The effect of this has been studied [116]; see also [62, 63].

6.1. Hard spheres

The crystallization of hard-sphere colloids has been extensively studied; see Palberg's review [117]. These experiments have been compared with similarly extensive simulations of mono- and polydisperse hard spheres [62, 63, 118, 119]. Also, Punnathanam and Monson [120] have studied the homogeneous nucleation of binary mixtures of hard spheres. Interestingly, they find that the barrier is higher when the two components fractionate on crystallization, i.e., form crystalline phases with a different composition from that of the fluid. Nucleation in binary mixtures of charged colloidal spheres was studied by Wette *et al* [121].

Auer and Frenkel [118, 119] obtained accurate estimates for the absolute nucleation rate in hard spheres and compared them with experiment [117, 122]. The experimental rate was approximately two orders of magnitude larger than that found in simulation. Although by the standards of no-fitting comparisons between experiment and simulation this is exceptionally good agreement, there is a still a discrepancy. In addition to comparing simulation and experiment, Auer and Frenkel [118, 119] compared simulation data with CNT predictions. They found that CNT underpredicted the simulation barrier height by a few tens of per cent. CNT assumes that the nucleus is spherical, and that the cost of the fluid–crystal interface does not vary with supersaturation. Neither assumption fails dramatically, but also neither is quantitatively accurate. The CNT prediction for the prefactor was also compared with

that from simulation [119]. The CNT prediction is of the correct order of magnitude. The simulation/CNT prediction for the prefactor is between one and three orders of magnitude larger than the values obtained by fitting an expression of the CNT form to the experimental data for hard-sphere colloids [117–119].

The heterogeneous nucleation of hard spheres has also been studied by computer simulation [123]. It was studied at a smooth hard wall. Dijkstra [124] has studied a fluid of hard spheres at such a wall, near and at coexistence with the crystal phase. She finds that the crystal (specifically the (111) plane) wets the hard-wall–fluid interface. Thus, once the wetting layer forms we would not expect a nucleation barrier for crystallization of hard spheres near a hard wall. Auer and Frenkel [123] studied the nucleation of hard spheres near a hard wall without a wetting layer, and found a small nucleation barrier. They worked at a pressure just above the fluid–crystal coexistence pressure. This barrier may be for nucleation of the bulk crystalline phase, of a wetting layer at the wall, or it may be that it is not possible to distinguish between the nucleation of a new bulk phase and of a wetting layer, when we are at a pressure such that the crystal is the equilibrium phase. In any case the nucleation barrier was much smaller than that for homogeneous nucleation at the same supersaturation. Therefore we would expect a fluid of hard spheres in a container with hard walls to freeze starting from these walls.

6.2. Charged spheres

For technical reasons charged spherical particles can be easier to study in experiment than hard-sphere-like colloids. For example, the pioneering work of Gasser *et al* [57] on using microscopy to observe the critical nucleus was done with charged particles. They did not measure the nucleation rate, but Palberg and co-workers have [125, 126]. Charged spheres repel each other via a soft repulsion. The nucleation of spheres which interact via a soft repulsion was studied via simulation by Auer and Frenkel [127]. They found smaller nucleation barriers for soft spheres than for hard spheres. This is consistent with the lower interfacial tensions of soft spheres [127]. It is easy to show that adding a soft repulsion tends to lower interfacial tensions because it actually favours inhomogeneities in the density, such as interfaces [128]. Finally, heterogeneous nucleation has also been studied for charged spheres at a wall [129].

6.3. Non-spherical particles

Relatively little work has been devoted to the study of the nucleation of non-spherical colloids. However, Dogic and Fraden [130, 131] used microscopy to observe some very interesting phenomena in suspensions of long rod-like viral particles that contain added soluble polymer (dextran). These phenomena occur at the transition from a dilute isotropic solution to smectic and crystalline phases. The rod-like viral particles are approximately 10 nm in diameter and 1 μm long. Without the dextran the particles interact with each other via a purely repulsive interaction. The addition of the polymer induces an attraction which is comparable in range to the diameter of the particles and so has a range that is very small in comparison the length of the viral particles. This short range means that there is a deep well in the potential for pairs of rods that lie side-by-side and parallel. The two rods are then within the range of the attraction for their whole length. This deep well favours the formation of layers of parallel rods; hence the smectic ordering. The nucleation dynamics observed by Dogic and Fraden are much richer than those of hard spheres. Schilling and Frenkel [132, 133] have performed theoretical studies of the nucleation of systems of rod-like particles and they have observed multiple nucleation barriers. This proliferation of nucleation barriers when the potential is not spherically symmetric may explain the complex phenomena observed by Dogic and Fraden.

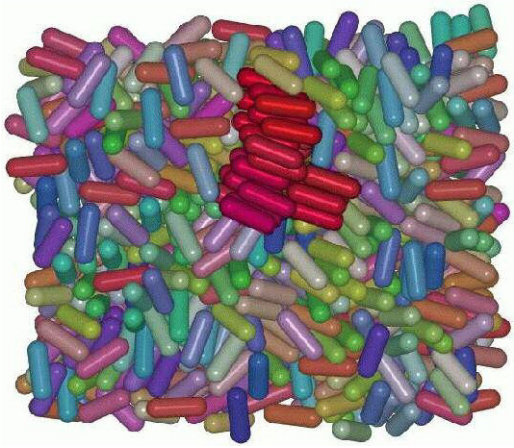

Figure 9. A metastable liquid configuration containing a (smaller than critical) crystalline cluster (the red cylinders). It is taken from work by Schilling and Frenkel [133]. Different shading of the rods indicates different orientations. The particles are hard spherocylinders whose cylindrical part is twice as long as their diameter. The pressure is such that the equilibrium phase is a crystal.

In a study of nucleation in systems of hard spherocylinders, Schilling and Frenkel [133] also found that not only was the nucleus far from spherical but that it was layered; see figure 9. It is likely that the nuclei of particles with highly anisotropic potentials are often not only highly non-spherical but have significant internal structure, such as layers. The internal structure in particular will need to be understood before we are able to understand nucleation in these more complex systems, as it will tend to create additional barriers that will affect the dynamics.

This work on rods is a rare example of a detailed study of non-spherical particles; most of the detailed experimental, simulation and theoretical studies of nucleation have used spherical models. For these simple spherical particles homogeneous nucleation at least is quite simple: there is only one barrier. The work of Dogic and Fraden [130, 131] shows us that when the interaction potential is more complex the phase transition kinetics can be much more complex, and that colloidal systems are ideal systems in which to study this complexity. Relevant timescales are experimentally easily accessible and the particles are large enough for light microscopy to be a powerful tool. It is interesting to speculate that in other systems where the potentials are highly anisotropic, e.g., proteins, nucleation phenomena similar to that observed by Dogic and Fraden for viruses may be occurring but have not been observed due to the lengthscales involved being too small.

7. Conclusion

Having reviewed work on nucleation, the obvious question is: Do we understand nucleation? In a single very simple system, that of monodisperse hard spheres, the answer is a qualified yes. Computer simulation results for homogeneous nucleation agree with experiment to within a couple of orders of magnitude [117–119], which, given the extreme sensitivity of the rate, can be considered quite good agreement. Also, heterogeneous nucleation at a hard wall has been studied [123] and shown to be orders of magnitude faster than homogeneous nucleation. Thus monodisperse hard spheres in a container with hard walls should freeze from the outside in. In protein solutions, a much more complex system, the answer is clearly no. Even for lysozyme, by far the most studied protein, there is still debate about whether the nucleation observed is

homogeneous or heterogeneous [28, 78–80]. Thus the role of impurities is obscure, although lysozyme dimers have been shown to affect the nucleation rate [82], and disordered porous media have been shown to induce nucleation in solutions where the nucleation barrier would otherwise be so large that no nucleation occurs [27, 91]. Also, the non-monotonic variation of the nucleation rate shown in figure 8 is fascinating but we lack a truly microscopic theory of the underlying nucleation behaviour. This is also true of much of the behaviour observed by Dogic and Fraden [130, 131] in their suspensions of long viral particles with added soluble polymer.

Future progress will require quantitative experimental studies of nucleation rates, experimental studies in which nuclei are imaged, as Gasser *et al* [57] did, and detailed simulation studies. In the experimental studies, care should be taken to determine whether the nucleation is homogeneous or heterogeneous. For example, if adding impurities or removing them via purification alters the nucleation rate the nucleation must be heterogeneous. Unfortunately, clearly demonstrating that the nucleation is homogeneous is harder.

Acknowledgments

First of all it is a pleasure to thank D Frenkel for many inspiring discussions on nucleation, and R Piazza for initially suggesting that I write a review on this topic. I am also grateful to P R ten Wolde, P G Vekilov and T Schilling for permission to reproduce figures. I would also like to acknowledge the EPSRC and The Wellcome Trust for funding.

References

[1] Määttänen A, Vehkamäki H, Lauri A, Merikallio S, Kauhanen J, Savijärvi H and Kulmala M 2005 *J. Geophys. Res.* **110** 02002
[2] Dusek U, Frank G P, Hildebrandt L, Curtius J, Schneider J, Walter S, Chand D, Drewnick F, Hings S, Jung D, Borrmann S and Andreae M O 2006 *Science* **312** 1375
[3] Debenedetti P G 1996 *Metastable Liquids* (Princeton, NJ: Princeton University Press)
[4] Oxtoby D W 1992 *J. Phys.: Condens. Matter* **4** 7627
[5] Oxtoby D W 1998 *Acc. Chem. Res.* **31** 91
[6] Kelton K F, Greer A L, Herlach D M and Holland-Moritz D 2004 *MRS Bull.* **29** 940
[7] García-Ruiz J M 2003 *J. Struct. Biol.* **142** 22
[8] Gorti S, Forsythe E L and Pusey M L 2005 *Cryst. Growth Des.* **5** 473
[9] Turnbull D and Vonnegut B 1952 *Ind. Eng. Chem.* **44** 1292
[10] Turnbull D 1950 *J. Chem. Phys.* **18** 198
[11] Bonn D and Ross D 2001 *Rep. Prog. Phys.* **64** 1085
[12] Binks B P and Clint J H 2002 *Langmuir* **18** 1270
[13] Sholl C A and Fletcher N H 1970 *Acta Metall.* **18** 1083
[14] Parry A O, Wood A J and Rascon C 2001 *J. Phys.: Condens. Matter* **13** 4591
[15] Page A J and Sear R P 2006 *Phys. Rev. Lett.* **97** 065701
[16] Valencia A and Lipowsky R 2004 *Langmuir* **20** 1986
[17] Gránásy L, Pusztai T, Börsönyi T, Tóth G, Tegze G, Warren J A and Douglas J F 2006 *J. Mater. Res.* **21** 309
[18] Atkins P W 1986 *Physical Chemistry* 3rd edn (Oxford: Oxford University Press)
[19] Fletcher N H 1958 *J. Chem. Phys.* **29** 572
[20] Padilla P and Talanquer V 2001 *J. Chem. Phys.* **114** 1319
[21] Cacciuto A, Auer S and Frenkel D 2004 *Nature* **428** 404
[22] Oxtoby D W and Kashchiev D 1994 *J. Chem. Phys.* **100** 7665
[23] Viisanen Y, Strey R and Reiss H 1993 *J. Chem. Phys.* **99** 4680
[24] Galkin O and Vekilov P G 1999 *J. Phys. Chem.* B **103** 10965
[25] Bhamidi V, Varanasi S and Schall C A 2005 *Langmuir* **21** 9044
[26] Sear R P 2004 *Phys. Rev.* E **70** 021605
[27] Chayen N E, Saridakis E and Sear R P 2006 *Proc. Natl Acad. Sci.* **103** 597
[28] Sear R P 2006 *J. Phys. Chem.* **110** 21944

[29] Bonn D, Kellay H and Meunier J 1994 *Phys. Rev. Lett.* **73** 3560
[30] Bonn D, Bertrand E, Meunier J and Blossey R 2000 *Phys. Rev. Lett.* **84** 4661
[31] Stange T G, Evans D F and Hendrickson W A 1998 *Langmuir* **13** 4459
[32] Thiele U, Mertig M and Pompe W 1998 *Phys. Rev. Lett.* **80** 2869
[33] Thiele U 2003 *Eur. Phys. J.* E **12** 409
[34] Schick M and Taborek P 1992 *Phys. Rev.* B **46** 7312
[35] Blokhuis E M 1995 *Phys. Rev.* E **51** 4642
[36] Gelb L D, Gubbins K E, Radhakrishnan R and Sliwinska-Bartkowiak M 1999 *Rep. Prog. Phys.* **62** 1573
[37] Lu G Q and Zhao X S 2005 *Nanoporous Materials, Science and Engineering* (London: Imperial College Press)
[38] Woo H-J, Porcheron F and Monson P A 2004 *Langmuir* **20** 4743
[39] Libby B and Monson P A 2004 *Langmuir* **20** 4289
[40] Valiullin R, Naumov S, Galvosas P, Kärger J, Woo H-J, Porcheron F and Monson P A 2006 *Nature* **443** 965
[41] Vishnyakov A and Neimark A V 2003 *Langmuir* **19** 3240
[42] Neimark A V and Vishnyakov A 2006 *J. Phys. Chem.* B **110** 9403
[43] Restagno F, Bocquet L and Biben T 2000 *Phys. Rev. Lett.* **84** 2433
[44] Talanquer V and Oxtoby D W 2001 *J. Chem. Phys.* **114** 2793
[45] Husowitz B and Talanquer V 2004 *J. Chem. Phys.* **121** 8021
[46] Saugey A, Bocquet L and Barrat J L 2005 *J. Phys. Chem.* B **109** 6520
[47] Paul R and Rieger H 2005 *J. Chem. Phys.* **123** 024708
[48] Leung K and Luzar A 2000 *J. Chem. Phys.* **113** 5845
[49] Vishnyakov A and Neimark A V 2003 *J. Chem. Phys.* **119** 9755
[50] Köhler H 1936 *Trans. Faraday Soc.* **32** 1152
[51] Gorbunov B and Hamilton R 1997 *J. Aerosol Sci.* **28** 239
[52] Talanquer V and Oxtoby D W 2003 *J. Chem. Phys.* **119** 9121
[53] Mirabel P, Reiss H and Bowles R K 2000 *J. Chem. Phys.* **113** 8200
[54] Mirabel P, Reiss H and Bowles R K 2000 *J. Chem. Phys.* **113** 8194
[55] Djikaev Y S 2002 *J. Chem. Phys.* **116** 9865
[56] Barrat J L and Hansen J P 2003 *Basic Concepts for Simple and Complex Liquids* (Cambridge: Cambridge University Press)
[57] Gasser U, Weeks E R, Schofield A, Pusey P N and Weitz D A 2001 *Science* **292** 258
[58] Yau S T and Vekilov P G 2000 *Nature* **406** 494
[59] ten Wolde P R, Oxtoby D W and Frenkel D 1999 *Phys. Rev. Lett.* **81** 3695
[60] ten Wolde P R, Oxtoby D W and Frenkel D 1999 *J. Chem. Phys.* **111** 4762
[61] Frenkel D and Smit B 2001 *Understanding Molecular Simulation: From Algorithms to Applications* (New York: Academic)
[62] Auer S and Frenkel D 2004 *Annu. Rev. Phys. Chem.* **55** 333
[63] Auer S and Frenkel D 2005 *Adv. Polym. Sci.* **173** 149
[64] Moroni D, ten Wolde P R and Bolhuis P G 2005 *Phys. Rev. Lett.* **94** 235703
[65] Bolhuis P G, Chandler D, Dellago C and Geissler P G 2002 *Annu. Rev. Phys. Chem.* **53** 291
[66] Pan A C and Chandler D 2004 *J. Phys. Chem.* B **108** 19681
[67] van Erp T S, Moroni D and Bolhuis P G 2003 *J. Chem. Phys.* **118** 7762
[68] Moroni D, van Erp T S and Bolhuis P G 2004 *Physica* A **340** 395
[69] Allen R J, Warren P B and ten Wolde P R 2005 *Phys. Rev. Lett.* **94** 018104
[70] Allen R J, Frenkel D and ten Wolde P R 2006 *J. Chem. Phys.* **124** 024102
[71] Allen R J, Frenkel D and ten Wolde P R 2006 *J. Chem. Phys.* **124** 194111
[72] Valeriani C, Sanz E and Frenkel D 2005 *J. Chem. Phys.* **122** 194501
[73] Chayen N E 2004 *Curr. Opin. Struct. Biol.* **14** 577
[74] Filobelo L F, Galkin O and Vekilov P G 2005 *J. Chem. Phys.* **123** 014904
[75] Galkin O and Vekilov P G 2000 *J. Am. Chem. Soc.* **122** 156
[76] Kulkarni A M and Zukoski C F 2001 *J. Cryst. Growth* **232** 156
[77] Kulkarni A M and Zukoski C F 2002 *Langmuir* **18** 3090
[78] Dixit N M, Kulkarni A M and Zukoski C F 2001 *Colloids Surf.* A **190** 47
[79] Vekilov P G and Galkin O 2003 *Colloids Surf.* A **215** 125
[80] Zukoski C F, Kulkarni A M and Dixit N M 2003 *Colloids Surf.* A **215** 137
[81] Chayen N E and Saridakis E 2001 *J. Cryst. Growth* **232** 262
[82] Burke M W, Leardi R, Judge R A and Pusey M L 2001 *Cryst. Growth Des.* **1** 333
[83] Sear R P 2006 *J. Phys. Chem.* B **110** 4985
[84] Castro M, Sánchez A and Domínguez-Adame F 2000 *Phys. Rev.* B **61** 6579

[85] Castro M 2003 *Phys. Rev.* B **67** 035412
[86] Tsekova D, Popova S and Nanev C 2002 *Acta Crystallogr.* D **58** 1588
[87] Vekilov P G, Feeling-Taylor A R, Petsev D N, Galkin O, Nagel R L and Elison Hirsch R 2002 *Biophys. J.* **83** 1147
[88] Kuznetsov Y G, Malkin A J and McPherson A 2001 *J. Cryst. Growth* **232** 30
[89] McPherson A 1999 *Crystallization of Biological Macromolecules* 1st edn (New York: Cold Spring Harbor Laboratory Press)
[90] McPherson A and Shlichta P 1988 *Science* **239** 385
[91] Chayen N E, Saridakis E, El-Bahar R and Nemirovsky Y 2001 *J. Mol. Biol.* **312** 591
[92] Frenkel D 2006 *Nature* **443** 641
[93] Sear R P 2004 *Europhys. Lett.* **68** 589
[94] Muschol M and Rosenberger F 1997 *J. Chem. Phys.* **107** 1953
[95] Sear R P 1999 *J. Chem. Phys.* **111** 4800
[96] Galkin O and Vekilov P G 2000 *Proc. Natl Acad. Sci.* **97** 6277
[97] Vekilov P G 2004 *Cryst. Growth Des.* **4** 671
[98] Pan W, Kolomeisky A B and Vekilov P G 2005 *J. Chem. Phys.* **122** 174905
[99] Kashchiev D, Vekilov P G and Kolomeisky A B 2005 *J. Chem. Phys.* **122** 244706
[100] Chandler D 1987 *Introduction to Modern Statistical Mechanics* (New York: Oxford University Press)
[101] Kadanoff L P 2000 *Statistical Physics* (Singapore: World Scientific)
[102] ten Wolde P R and Frenkel D 1997 *Science* **277** 1975
[103] Sear R P 2001 *J. Chem. Phys.* **114** 31703
[104] Sear R P 2002 *J. Chem. Phys.* **116** 2922
[105] Shiryayev A and Gunton J D 2004 *J. Chem. Phys.* **120** 8318
[106] Talanquer V 2005 *J. Chem. Phys.* **122** 084704
[107] Talanquer V and Oxtoby D W 1998 *J. Chem. Phys.* **109** 223
[108] Pini D, Ge J L, Parola A and Reatto L 2000 *Chem. Phys. Lett.* **327** 209
[109] Tavassoli Z and Sear R P 2002 *J. Chem. Phys.* **116** 5066
[110] Cacciuto A, Auer S and Frenkel D 2004 *Phys. Rev. Lett.* **93** 166105
[111] Cacciuto A and Frenkel D 2005 *J. Phys. Chem.* B **109** 6587
[112] Sear R P 2002 *Langmuir* **18** 7571
[113] Sear R P 2002 *J. Phys.: Condens. Matter* **14** 3693
[114] de Gennes P G 1985 *Rev. Mod. Phys.* **57** 827
[115] Anderson V J and Lekkerkerker H N W 2002 *Nature* **416** 811
[116] Auer S and Frenkel D 2001 *Nature* **413** 711
[117] Palberg T 1999 *J. Phys.: Condens. Matter* **11** R323
[118] Auer S and Frenkel D 2001 *Nature* **409** 1020
[119] Auer S and Frenkel D 2004 *J. Chem. Phys.* **120** 3015
[120] Punnathanam S and Monson P A 2006 *J. Chem. Phys.* **125** 024508
[121] Wette P, Schöpe H J and Palberg T 2005 *J. Chem. Phys.* **122** 144901
[122] Harland J L and van Megen W 1997 *Phys. Rev.* E **55** 3054
[123] Auer S and Frenkel D 2003 *Phys. Rev. Lett.* **91** 015703
[124] Dijkstra M 2004 *Phys. Rev. Lett.* **93** 108303
[125] Schöpe H J and Palberg T 2002 *J. Phys.: Condens. Matter* **14** 11573
[126] Wette P, Schöpe H J and Palberg T 2005 *J. Chem. Phys.* **123** 174902
[127] Auer S and Frenkel D 2002 *J. Phys.: Condens. Matter* **14** 7667
[128] Sear R P and Frenkel D 2003 *Phys. Rev. Lett.* **90** 195701
[129] Stipp A, Biehl R, Preis T, Liu J, Fontecha A B, Schöpe H J and Palberg P 2004 *J. Phys.: Condens. Matter* **16** S3885
[130] Dogic Z and Fraden S 2001 *Phil. Trans. R. Soc.* A **359** 997
[131] Dogic Z 2003 *Phys. Rev. Lett.* **91** 165701
[132] Frenkel D and Schilling T 2002 *Phys. Rev.* E **66** 041606
[133] Schilling T and Frenkel D 2004 *Phys. Rev. Lett.* **92** 085505