

SUPERCOOLED LIQUIDS AND SUPERCOMPUTERS

BY IVAN SAIKA-VOIVOD AND PETER H. POOLE

The liquid state provides condensed matter physics with some of its longest standing and most perplexing questions^[1,2]. By comparison, the fundamental nature of the other conventional phases, gas and crystal, are much better understood. Our understanding of how the properties of gases and crystals arise from interactions at the molecular level was (and continues to be) facilitated by the availability of idealized but exactly-solvable limiting models as tractable starting points for theory, and by systematic techniques for extending these ideal models to recover the properties of realistic systems. For example, for gases, we can start with the ideal gas, and then make progress toward a real gas by adding terms to a virial expansion. For crystals, the Einstein crystal provides an idealized starting point, and successive improvements can be made by progressing through e.g. the Debye model, to better reveal the thermodynamics of crystals; or Bloch's theorem, to provide a starting point for understanding electronic properties.

For liquids, the situation is different. Liquids lack the discrete symmetry and long-range molecular order of crystals, and so the simplifications exploited in much of solid state physics simply do not apply. Superficially, the structure of a liquid seems to have more in common with that of the gas phase, at least from the standpoint of symmetry. However, the vast difference in density between a gas (well above the critical temperature) and a liquid (near the freezing temperature) precludes the use of the ideal gas as a starting point for any practical approach to studying the liquid. The densities of most liquids near freezing are within ten percent of the corresponding crystal density. Hence, many-body molecular interactions play a dominant role in determining the properties of the liquid state, whereas in a typical real gas these can usually be treated as perturbations to the ideal gas.


In addition, many of the most interesting questions about liquids have to do with their evolution in time, in particular their dynamical behaviour in equilibrium (e.g. diffusion), and with how they behave when they are out of

equilibrium (e.g. during the transformation of a liquid to a solid). Liquids consequently present us with a theoretical "perfect storm": we face all the complexity of a dense, disordered, strongly interacting, many-body system; we are denied simplifications based on symmetry, as in crystals; and we must not only treat the physics of a disordered structure, but also how the structure changes with time. Of course, liquids are not the only physical system to present such barriers to understanding. Indeed, much of modern statistical physics is focussed on systems where time-varying disorder is a central feature (e.g. granular matter, frustrated magnetic systems). Liquids are simply a commonly encountered, and historically important, case.

In this article, we will focus our attention on one regime where these challenges come strongly to the fore: in the supercooled liquid state^[1]. By this, we mean the liquid state that can be observed if a liquid is cooled to a temperature T below the crystal melting temperature T_m . While the supercooled state is a metastable one, in the sense that the crystal has a lower free energy than the liquid, almost all liquids can be studied for some range of T below T_m on a time scale long enough for a metastable liquid-state equilibrium to be established. Interest in this regime derives in large measure from the fact that solid matter, whether crystalline or amorphous, can be formed from the supercooled liquid, in the former case by nucleation, and in the latter via the glass transition. These two solidification mechanisms, while starkly different, depend sensitively on the nature of the supercooled liquid in which they begin, and high performance computing has proven to be an indispensable tool for progress in this area. We will illustrate this by discussing results from our own research and also highlight related studies, especially those carried out in Canada.

COMPUTERS AND LIQUID STATE PHYSICS

Computers have had a singular impact on the development of liquid state physics in general^[2,3]. This is due to the fact that, although the analytical challenges of the liquid state are severe, the physical ingredients of a classical liquid at the molecular level are easy to specify, making an algorithmic approach attractive^[3]. The system's potential energy is usually well-approximated as a sum over a specified two-body interaction function, and Newton's laws suffice to determine the trajectories of the molecules in space, starting from some given initial condition. This specification is straightforward to implement as a computer algorithm, and is known as molecular dynamics (MD).



I. Saika-Voivod
<saika@mun.ca>, Department of Physics and Physical Oceanography, Memorial University of Newfoundland, St. John's NF, A1B 3X7, Canada, and P.H. Poole, Department of Physics, St. Francis Xavier University, Antigonish, Nova Scotia B2G 2W5, Canada

SUMMARY

High performance computing is now central to efforts to resolve long-standing questions on the nature of cold, dense liquids, and on how they transform to crystalline and amorphous solids. We review several current examples.



If one is interested only in structural or thermodynamic properties, the time evolution of the system can be replaced by a stochastic exploration of the system's configuration space, using Monte Carlo (MC) methods.

A computational approach to the liquid state thereby opens up the opportunity to model the behaviour of the system at the molecular level, and study how bulk properties arise from microscopic interactions. Complete knowledge of atomic positions as a function of time are generated in MD simulations, making it possible to evaluate any property accessible in experiments, and many that are not accessible, or at least not yet.

The principal challenges of a molecular-level approach to modelling the liquid state are the limitations on the system sizes and time scales that are accessible using a given generation of computing hardware. In MD simulations of classical liquids, the fundamental time step by which the molecular trajectories are advanced forward in time is usually on the scale of 1 fs. For simulations on a modern single processor of a system having 10^3 molecules, a time scale on the order of 100 ns can be reached if one is willing to wait several weeks for the results. So long as the internal equilibrium relaxation time of the liquid is less than this, reliable results can be obtained. However, there are many processes in supercooled liquids (e.g. crystal nucleation) that may occur only on much longer time scales. Since the liquid-state relaxation slows down as T decreases, this upper time scale also sets a bound on the lowest T that can be reached in equilibrium.

Constraints on the system size depend in part on the time scale that a given study must access: for long time scales, the smallest system sizes are chosen. For short time scale phenomena, system sizes of as large as 10^9 atoms have been realized. While computationally impressive, such systems are a long way from the macroscopic regime. In all simulations of bulk liquids, periodic boundary conditions are used to minimize surface effects. In spite of this, finite-size effects remain a serious challenge for even the most modern simulations. An instructive example is the recent work of Sokolovskii, Thachuk and Patey at UBC [4]. This study examines the influence of system size on the evaluation of tracer diffusion in a hard sphere liquid. The estimation of the diffusion constant of an infinite sized system from finite sized simulations is a long-standing problem, and this work shows how state-of-the-art computing power and careful analysis have finally evolved to the point where a reliable answer can be obtained.

The twin constraints of size and time scale are affected differently by computer hardware and software developments. For large system sizes, parallel computing algorithms are desirable, dividing the work of simulating a single large system over many processors. At the same time, the largest accessible time scale of any simulation, whether serial or parallel, depends on processor speed. The advent in the last decade of very large clusters of fast, inexpensive processors, interconnected by high-bandwidth, low-latency networking, has thus benefitted the simulation of both large systems and long time scales. In

addition, these clusters facilitate studies in which a large number of independent single-processor liquid state simulations are run under different conditions, e.g. of temperature and density in order to evaluate a liquid's equation of state. Such "embarrassingly parallel" parameter-space exploration studies have flourished in the last decade, and have stimulated the development of new computational methods, such as parallel tempering [5].

The above discussion has avoided the question of where one gets the molecular interaction potential required as input to an MD or MC liquid simulation. This is a large and complex topic by itself, and space precludes a full discussion here. For classical simulations (to which we restrict ourselves), these potentials are developed either by fitting the parameters of a model function to a selection of experimentally-known properties (e.g. an atomic radial distribution function, or a melting temperature), or by fitting a model function to a potential energy surface determined quantum mechanically, usually for a small molecular cluster.

The alternative is to implement a fully quantum mechanical approach, i.e. a quantum molecular dynamics (QMD) simulation. In QMD, electronic degrees of freedom are modeled explicitly, and so the molecular interactions are evaluated within the algorithm from first principles. In this way, QMD realizes a qualitatively higher degree of realism and, in addition, allows for the evaluation of the electronic properties of liquids, which are not available from classical simulations. For certain problems, such as the notable work of Stanimir Bonev and co-workers at Dalhousie University on the high pressure properties of liquids such as hydrogen and nitrogen [6,7], QMD is the only reliable way to proceed. However, the computational demands of QMD are much higher than for classical MD. Currently, a large QMD liquid simulation would be of a system of a few hundred molecules over a time scale of tens of ps. At present, this makes QMD unsuitable for most problems related to large length/time scale phenomena in supercooled liquids, e.g. crystal nucleation. However, there is no question that, as computational power continues to progress, QMD studies will systematically displace classical simulations of the liquid state in the years to come.

THERMODYNAMICS AND PHASE DIAGRAMS

As stated in the Introduction, crystal and glass formation depend sensitively on the interplay of both thermodynamic and dynamical properties of the supercooled liquid state. In this section, we focus on some of the thermodynamic aspects. One of the most fundamental thermodynamic descriptors of a liquid, an equation of state (e.g. the pressure P as a function of temperature T and density ρ) is readily evaluated from simulations. However, it is often crucial to determine the thermodynamic relationship of the liquid phase to the crystal. For example, the nucleation rate is a strong function of degree of supercooling, which can only be stated if the coexistence temperature for the liquid and crystal phases is determined.

A suite of methods has therefore been developed to evaluate the free energy of both liquid and crystal phases from simulations, to use this information to locate phase coexistence conditions and, ultimately, to build complete phase diagrams for model substances. Daan Frenkel and coworkers have played a leading role in the development of these methods, and the recent text by Frenkel and Smit is an invaluable resource for any researcher in this area [5].

To evaluate the free energy of any phase, one general approach is to identify a state of the system for which the free energy is known exactly, and then numerically carry out a “thermodynamic integration” (TI) to the particular state of interest in order to find the free energy difference between the exactly known and the desired state. For gases, the low density, ideal gas limit is the natural starting point, followed by a TI along a path that (say) first changes the density to the desired value, and then the temperature. For liquids, the difference between the free energy of the liquid and the ideal gas can be determined by integrating the excess pressure along an isotherm from low densities, where the simulated system’s pressure is well described by a low order virial expansion. As long as no discontinuous phase transitions occur along the chosen path, the absolute free energy can be computed to arbitrary precision.

In computer simulations, the path along which the integration occurs need not be a function of macroscopic variables. For example, it could occur along a path where a parameter in the system Hamiltonian is changing. This approach is exploited to find the free energy of a crystal. For a crystal, a natural starting point is the Einstein crystal of the desired structure; i.e. a crystal in which the molecules do not explicitly interact with each other, but are held near their ideal positions by harmonic springs. The free energy of this system can be computed exactly. A parameter in the system Hamiltonian is then varied so as to “morph” the system from this ideal potential to the real one (usually, at fixed T and ρ), while computing the free energy change along the way. Free energy changes from this state to other T - ρ points can be computed by conventional TI.

Such methods, combined with the common availability of computing clusters with hundreds of processors, have made possible the evaluation of complete equations of state and free energy surfaces for a number of important model systems. These data can then be used to construct extensive phase diagrams. Fig. 1 shows the result of our own work to determine the phase diagram of a commonly studied model of silica, the so-called “BKS” model [8]. The simulated phase diagram, compared to that known from experiment (also shown [9]), both reveals the inadequacies of the model (and thus provides clues for how to improve the model), and clearly identifies regimes of interest, e.g. at what range of T and P the liquid is supercooled, so that crystal nucleation can be studied. As we discuss below, our knowledge of the phase diagram for BKS silica facilitated our subsequent study of the nucleation of the stishovite crystal from the supercooled melt.

Phase diagrams have been determined for a wide range of model materials using these techniques, including those for

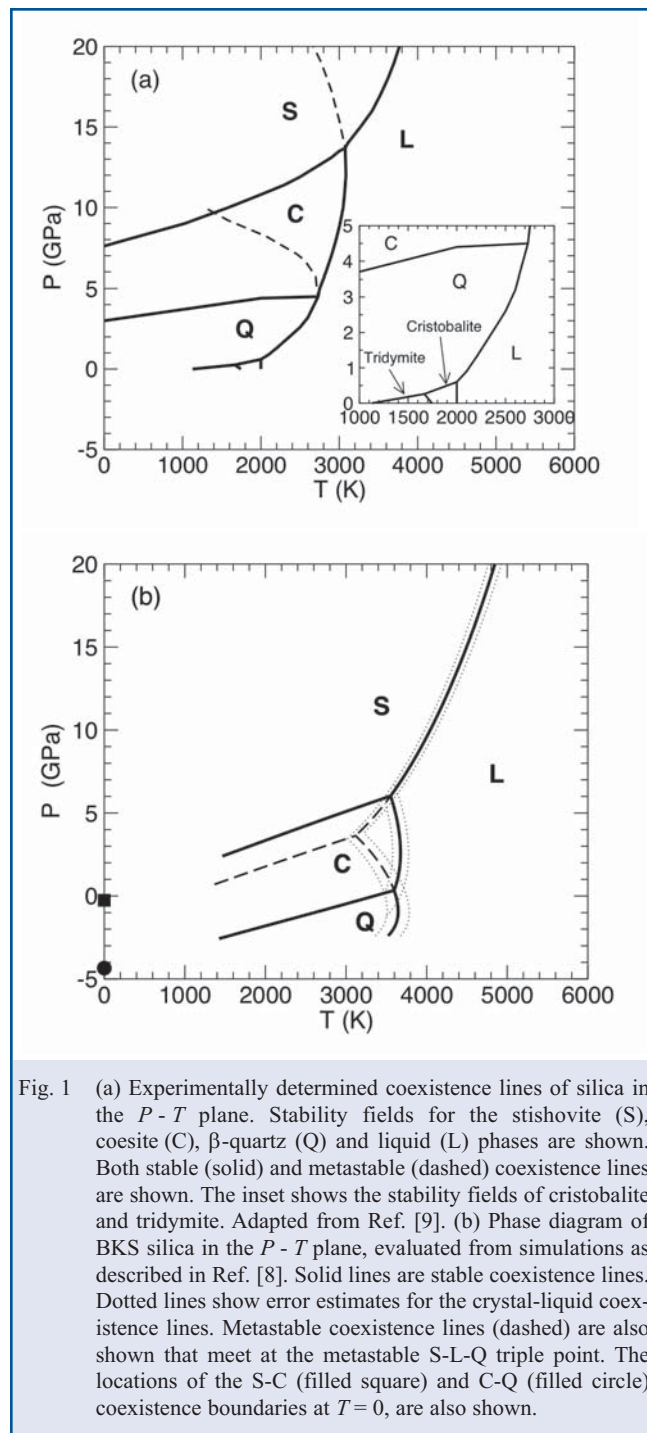


Fig. 1 (a) Experimentally determined coexistence lines of silica in the P - T plane. Stability fields for the stishovite (S), coesite (C), β -quartz (Q) and liquid (L) phases are shown. Both stable (solid) and metastable (dashed) coexistence lines are shown. The inset shows the stability fields of cristobalite and tridymite. Adapted from Ref. [9]. (b) Phase diagram of BKS silica in the P - T plane, evaluated from simulations as described in Ref. [8]. Solid lines are stable coexistence lines. Dotted lines show error estimates for the crystal-liquid coexistence lines. Metastable coexistence lines (dashed) are also shown that meet at the metastable S-L-Q triple point. The locations of the S-C (filled square) and C-Q (filled circle) coexistence boundaries at $T = 0$, are also shown.

several models of water [10]. The thermodynamic properties of supercooled water have been a sustained source of interest for several decades, and simulations have made significant contributions by providing information on states where experiments are challenging: e.g. in the deeply supercooled limit, where fast crystal nucleation pre-empts observation of liquid-state behaviour, and in the regime of negative pressure, where only a few experimental studies have successfully ventured. For example,

simulation results were the basis of the proposal that a first order liquid-liquid phase transition occurs in supercooled water^[11], the influence of which on surrounding states can explain many of water's unusual properties. Evidence now exists for analogous liquid-liquid transitions in a range of substances (e.g. liquid silicon) and simulations have been central to efforts to elucidate this phenomenon^[12]. Similarly, for several years, simulation studies of supercooled water have pointed to the possibility that a minimum of the density (in contrast to the density maximum that occurs at 4°C) occurs in the supercooled regime^[13,14]. Guided in part by these simulation results, experimental evidence for the occurrence in supercooled water of this extremely rare phenomenon has recently been reported^[15]; see Fig. 2.




Fig. 2 Comparison of density vs. temperature curves at ambient pressure for bulk liquid D₂O (open triangles), confined liquid D₂O (filled circles) from Ref. [15], D₂O ice Ih (filled squares), and MD simulations of liquid TIP5P-E water (open diamonds) from Ref. [13]. The density values for the TIP5P-E model (which is a model of H₂O) have been multiplied by 1.1 to facilitate comparison with the behaviour of D₂O. Both a maximum and a minimum of the density occur in simulations and experiment.

DYNAMICS NEAR THE GLASS TRANSITION

With the exception of quantum liquids (i.e. liquid He), there are only two possible fates for a supercooled liquid as T decreases: it will either undergo a first-order phase transition to a crystalline solid, or it will form an amorphous solid, or glass, at the glass transition temperature, T_g ^[16]. Superficially, the glass transition seems to be a purely dynamical transition, unrelated to any thermodynamic process. The viscosity of a liquid increases rapidly as T decreases, and in the absence of crystallization, the time scale for liquid-like structural relaxation eventually exceeds typical observation times. The value of T_g is (somewhat arbitrarily) taken as the T at which the viscosity exceeds 10^{13} poise. Below T_g , the system retains a disordered liquid-like structure, but the mechanical properties become solid-like.

This simple picture of the glass transition was notably critiqued in a 1948 paper by Walter Kauzmann^[17]. The paper describes what has since become known as the “Kauzmann paradox”. Kauzmann pointed out that the heat capacity of a liquid is generally higher than that of the crystal to which it freezes and, as a consequence, the entropy decreases more rapidly in the liquid than in the crystal as T decreases into the supercooled regime. For a wide range of liquids, this results in the thermodynamic behaviour shown schematically in Fig. 3. The liquid entropy, extrapolated to arbitrarily low T , would not only meet the crystal entropy, but even threatens to become zero at finite T . In practice, this “entropy catastrophe” is avoided because the glass transition seems to always intervene, knocking the liquid out of equilibrium, and putting a halt on the further decrease of entropy. The paradox is this: If the glass transition is a purely dynamical phenomenon, how can it be invoked to resolve a purely thermodynamic problem (the entropy catastrophe)? Kauzmann's paradox suggests that thermodynamics must play a role, along with dynamical behaviour, in the physics that underlies the glass transition. This conceptual tension, between the dynamical and thermodynamic underpinnings of glass formation, persists to the present day.




Fig. 3 Schematic behaviour of the entropy for a typical liquid and crystal of the same substance. T_m is the crystal melting temperature. At T_g the liquid falls out of equilibrium and becomes a glass.

In 1995, P.W. Anderson wrote: “The deepest and most interesting unsolved problem in solid state theory is probably the theory of the nature of glass and the glass transition”^[18]. While theories of the glass transition abound, it continues to be true that none is commonly accepted to have “solved” the problem, in the sense of accounting for the complex range of observed behaviour in a unified way. In this context, computer simulations have played a central role in testing theories, and in providing clues for the development of new theories. A prominent example in the 1990's was the simulation work of Kob and Andersen^[19], who used extensive MD simulations of a binary Lennard-Jones liquid to confirm many of the predictions of the mode-coupling theory (MCT) of the glass transition that had

been developed previously by Goetze [20]. This stimulated a great deal of work, both through simulations and experiments, to test the range of applicability of MCT.

The interrelationship of dynamics and thermodynamics in glass-forming liquids has also been explored with much success using simulations [16,21]. For example, a number of theories connecting entropy and diffusion have been proposed, stimulated by the ideas of Adam and Gibbs (AG) in 1965 [22]. Simulations provide a helpful testing ground for such theories because both thermodynamic and transport properties can be evaluated from a single set of runs; in experiments, widely different apparatus are required to access these observables, making systematic studies challenging. Several studies have demonstrated the validity of the AG theory in simulations. In our own work with F. Sciortino on liquid silica, we were also able to draw specific connections to the Kauzmann paradox [23]. We showed that the AG relation is obeyed in liquid silica, and at the same time the T dependence of the configurational entropy exhibits an inflection point that provides the mechanism for this system to avoid Kauzmann's entropy catastrophe. This result from simulation awaits experimental confirmation, since the relevant behaviour occurs in an extremely challenging experimental regime between 3000 and 4000 K.

There has also been ongoing interest in the possibility of finding molecular-level structural features in the liquid associated with the approach to the glass transition [24]. Supercooled liquids are notably homogeneous in a structural sense as they approach T_g . For example, they typically lack any growth of density fluctuations as T decreases, precluding the possibility of thinking of the approach to the glass transition as the approach to a conventional critical point. However, numerous experiments and simulations provide evidence that significant spatial heterogeneities of dynamical properties arise and grow in liquids as $T \rightarrow T_g$. These results indicate that the dynamics in a supercooled liquid does not slow down uniformly in space. Rather, correlated groups of relatively mobile and immobile molecules emerge and grow in size as T decreases. These are of course transient mobility fluctuations, which appear and disappear on the time scale of structural relaxation in the liquid. While in most cases experiments only provide indirect evidence of such "dynamical heterogeneity" (DH), simulations are able to image this phenomenon directly. Simulations carried out by S.C. Glotzer and coworkers have provided particularly clear views of DH [25,26]. In these studies, careful analysis of very long equilibrium MD runs of the binary Lennard Jones liquid showed that a molecule that is significantly more mobile than the average has a higher probability of occurring close to another similarly mobile molecule. These mobile molecules tend to form quasi-one-dimensional "strings" in which molecules move one after another, like dancers in a conga line. These results were subsequently confirmed in experiments on colloids, in which the trajectories of individual colloid particles were recorded and analyzed via confocal microscopy [27]. Much recent work has focussed on how the emergence of the mobility correlations of DH can be incorporated into a broad theory of glass formation.

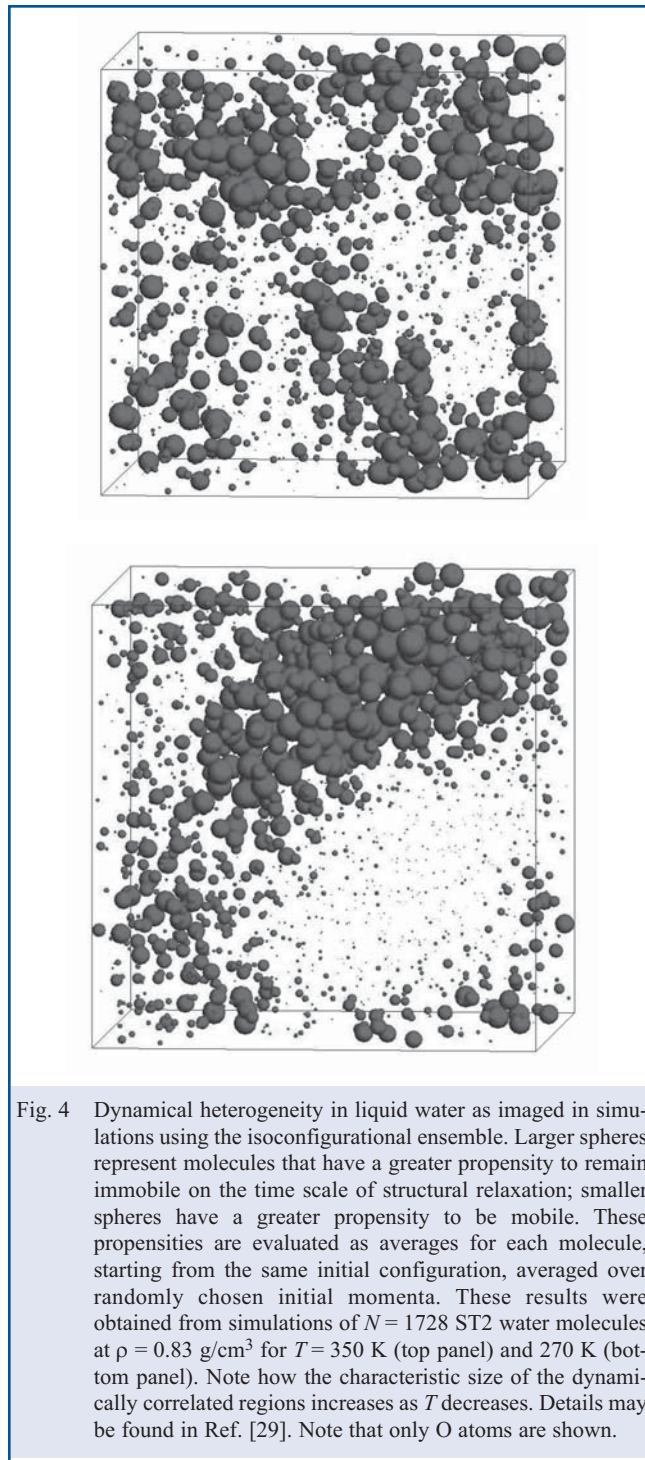


Fig. 4 Dynamical heterogeneity in liquid water as imaged in simulations using the isoconfigurational ensemble. Larger spheres represent molecules that have a greater propensity to remain immobile on the time scale of structural relaxation; smaller spheres have a greater propensity to be mobile. These propensities are evaluated as averages for each molecule, starting from the same initial configuration, averaged over randomly chosen initial momenta. These results were obtained from simulations of $N = 1728$ ST2 water molecules at $\rho = 0.83$ g/cm³ for $T = 350$ K (top panel) and 270 K (bottom panel). Note how the characteristic size of the dynamically correlated regions increases as T decreases. Details may be found in Ref. [29]. Note that only O atoms are shown.

More recently, Harrowell and coworkers developed a novel simulation approach that showed that, despite the absence of an obvious and growing structural heterogeneity in glass-forming liquids, the origins of DH can be ascribed, at least in part, to configurational properties of the liquid state [28]. They define an "isoconfigurational ensemble" of MD simulation trajectories, each starting from an identical equilibrium liquid config-

uration, but in which the molecular velocities are assigned randomly from the Maxwell-Boltzmann distribution. By averaging the displacement of a particular molecule at a given time over all the MD trajectories, the configurationally-induced “propensity” for molecular mobility as a functional of spatial position in the starting configuration can be assessed. In simulations of both 2D soft spheres and 3D liquid water, the resulting spatial maps of “dynamic propensity” affirm the picture of glass-forming liquids becoming progressively more heterogeneous as $T \rightarrow T_g$ [29]; see Fig. 4. The appearance of DH even after this kind of isoconfigurational averaging also suggests that a comprehensive theory of glass formation must be based on both dynamical and configurational ingredients. From the standpoint of computing, the practicality of using an approach such as isoconfigurational averaging is made possible only by the existence of large computing clusters.

CRYSTAL NUCLEATION

In order for a liquid to freeze, nucleation of the new crystalline phase must occur first. We will restrict our discussion to homogeneous nucleation, which takes place within the bulk of the supercooled liquid [1]. Fluctuations in local structure give rise to portions of the liquid that have a high degree of crystalline order. These ordered pockets can be thought of as embryos or nuclei from which the new phase arises. Perhaps surprisingly at first glance, these small embryos tend to shrink and vanish, for although the bulk crystal has a lower free energy than that of the liquid, the interface created between the crystalline embryo and the surrounding liquid makes embryo growth unfavourable from a free energy standpoint.

The idea of the competition between bulk and surface contributions to the free energy of embryo formation (i.e. the work required to form an embryo) is a main ingredient of Classical Nucleation Theory (CNT), a phenomenological theory in which an embryo has a well defined interface with the surrounding liquid. There is no generally accepted microscopic theory of nucleation and so, despite dating back to the 1920’s, CNT forms the theoretical basis for quantitatively understanding nucleation.

In CNT, the work required to assemble an embryo composed of n particles is given by

$$\Delta G(n) = -|\Delta\mu|n + a\gamma n^{2/3} = -k_B T \ln \frac{N_n}{N}, \quad (1)$$

where $\Delta\mu$ is the difference in the chemical potential between the bulk liquid and the bulk crystal, γ is the surface tension, a is a factor that depends on the shape and density of the embryos, N_n is the equilibrium number of embryos of size n present in the liquid, and N is the total number of liquid particles. The generic shape of $\Delta G(n)$ is shown in Fig. 5, where we see a maximum at n^* , the critical embryo size. Embryos must overcome a free energy barrier of height $\Delta G(n^*)$ before it is thermodynamically favourable for them to grow.

The rate of nucleation, or the rate at which embryos cross the

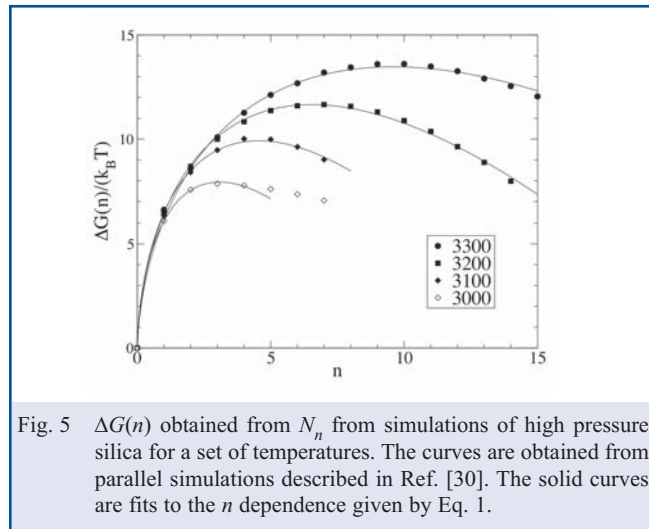


Fig. 5 $\Delta G(n)$ obtained from N_n from simulations of high pressure silica for a set of temperatures. The curves are obtained from parallel simulations described in Ref. [30]. The solid curves are fits to the n dependence given by Eq. 1.

barrier per unit volume, is given by,

$$J = R \exp \frac{-\Delta G(n^*)}{k_B T}, \quad (2)$$

where R is a kinetic prefactor that depends on the dynamics of the supercooled liquid.

The study of nucleation seems ideally suited to computer simulations. One would think that the microscopic level of detail in a MC or MD simulation should enable researchers to systematically peel apart the nucleation process. This is true, but there are challenges nonetheless. Nucleation of a post-critical embryo typically occurs in a metastable supercooled liquid as a rare event, particularly if $\Delta G(n^*)$ is large. For small to moderate supercooling, it may not be feasible to witness even a single nucleation event in even the longest simulations.

Another basic difficulty lies in discerning the embryo from the surrounding liquid particles. In a simple supercooled liquid, the neighbours of a given particle form a fairly ordered environment and, as mentioned earlier, the density mismatch between the liquid and crystal phases is not large. Determining which particles are crystal-like and which are not becomes a subtle task. Nonetheless, satisfactory criteria have been worked out to define local crystalline order, with the help of spherical harmonics. By looking at how this order is correlated between neighbouring particles, it becomes possible to identify embryos and the number of particles they contain. Fig. 6 shows snapshots of embryos taken from our simulations of silica [30].

With the embryos identified, the next step is to be able to drive the system to nucleate. This can be accomplished through biased sampling MC. In this technique, an order parameter, like the size of the largest cluster in the system, is identified. Then, a potential energy term that is a function of the order parameter is added to the model Hamiltonian, and is often taken to be a parabola centred upon a particular cluster size, n_0 . The new addition to the Hamiltonian biases the system to be in a state




Fig. 6 Crystal-like Si atoms in liquid silica. Top left: Sample critical nucleus at 3300 K containing 10 Si atoms. Top right: A snapshot of the growing crystal embryo from a dynamic crystallization simulation at 3000 K when it contains 23 Si atoms. Bottom: Sample end configuration of a crystallization simulation.

containing a cluster of size n_0 . Thus, through biased sampling, we can study at leisure the system when it is in an otherwise improbable state, in our case, states in which large and/or critical embryos are present. By simulating the system at values of n_0 ranging from small to critical to post-critical, we can calculate N_n , which through Eq. 1 determines $\Delta G(n)$

To determine $\Delta G(n)$ for several temperatures, we would run many simulations, each with a different T and n_0 . Perhaps these would all be running at the same time if a computing cluster is available. However, the equilibration of these systems can be greatly sped up by using a technique dubbed parallel tempering. In this scheme, simulations running in parallel are allowed to exchange configurations. The probability with which two processors commit to an exchange is precisely determined by the Boltzmann distribution. Qualitatively, the increase in computational efficiency comes from allowing slow states at low T to benefit from occasional visits to high temperatures, where kinetic barriers are more easily overcome.

Liquid configurations with embryos of critical size can be selected from the biased/tempered MC simulations in order to study their dynamic properties with good statistical sampling. In particular, the rate at which particles attach themselves to a critical embryo is used to calculate the dynamic prefactor in Eq. 2, thus completing the calculation of quantities required by CNT to predict the rate of nucleation.

The development and application of these techniques to nucleation is mostly attributable to Daan Frenkel and coworkers^[31].

They, and now others, have used simulation to drive our understanding of nucleation in numerous systems (argon, hard sphere colloids, NaCl, silica, and carbon, among others), and under various influences, e.g. in the presence of metastable critical points, near interfaces, or under extreme pressure. For example, contrary to recent suggestions, Frenkel was able to show that diamonds are not likely to nucleate in the carbon-rich middle layers of Uranus^[32].

In our work, we have used biased/tempered MC to study nucleation in a model of silica for which we worked out the phase diagram earlier. We focussed on a high pressure regime where nucleation occurs fairly easily, i.e. on a reasonable time scale for simulations. We showed that the form given by CNT for $\Delta G(n)$ holds reasonably well even when the barrier becomes fairly small at large supercooling (see Fig. 5). Additionally, we reached a point where the picture of nucleation begins to change qualitatively and the idea of a limit to liquid metastability may be required to make sense of some of the free energy profiles we calculated.

A number of other efforts are ongoing in Canada to simulate nucleation and crystal growth. For example, Peter Kusalik (formerly at Dalhousie, now at U. Calgary) and coworkers were the first to simulate ice nucleation in the presence of a strong electric field^[33], and more recently have done notable work simulating the interface of a crystal surface as it progresses into the liquid phase^[34]. Nucleation is also being studied in liquid nanoclusters in the group of R.K. Bowles^[35] at the University of Saskatchewan. Freezing of clusters differs from that of bulk liquids in that there is an inherent inhomogeneity in the system, i.e. a significant portion of the particles are on the surface, as well as the fact that there are typically several different structures to which the cluster may freeze at a given T . The frozen cluster structures are not bound to be true periodic crystals, and may have (for example) icosahedral or decahedral structure.

It is intriguing to think about possible connections between nucleation and the glass transition. What impact do dynamical heterogeneities have on the nucleation process? Are the heterogeneities themselves a result of subtle ordering connected with embryo formation? Is the liquid trying to order locally to a structure that cannot fill space? We are engaged in exploring some of these questions, and are encouraged by some hints on the subject now appearing in the literature pertaining to glass formers^[36].

OUTLOOK AND CONCLUSION

Computer simulation has and will continue to play a valuable role in developing our understanding of the supercooled liquid state, the glass transition and crystal nucleation. Simulations have been instrumental in testing microscopic theories of the glass transition, in linking thermodynamics to the glass transition, in calculating material properties from microscopic interactions, and in testing nucleation theories. Despite this progress, fundamental questions remain. For example, there is no generally accepted theory that can tell us why one liquid should form a glass, while another crystallizes easily. It is

inevitable that simulations will play a role in clarifying such questions.

In addition, those contemplating research in this area will benefit from paying attention not only to scientific trends, but also to the technological trends of computing hardware and software. For example, the explosive growth in single-processor speed over the last several decades presently allows us to study classical liquids over nearly 8 orders of magnitude in time, up to almost the μs time scale. More recent improvements in parallel architectures and algorithms have now also allowed the sizes of systems studied to grow dramatically. The focus of current development in processor technology is now shifting to multi-core processors, offering more potential for parallelism (and lower power consumption), but with the speed of single cores not increasing as dramatically as in the past. On its own, this would shift the advantage to the simulation of larger systems (via parallelism), but would slow the increase of the maximum accessible time scale. In another direction, the advent of accelerator cards (e.g. "GPGPUs" based on graphics co-processors) offer the potential for tremendous speed increases with some algorithms.

Finally, we note that our discussion in this article has focussed on simple and network-forming liquids such as water, but the basic ideas behind nucleation and glassy dynamics provide the underpinning for understanding more complex phenomena such as phase transitions and gelation in colloids, macromolecular assembly, protein folding and crystallization, and nanoparticle self-assembly. These are currently studied via simulation and experiments in research groups around the world, and the proliferation of high performance computing facilities will continue to advance simulation as an effective means to transfer our knowledge of basic phenomena to these more complex systems.

ACKNOWLEDGEMENTS

We thank our many collaborators who have provided us over the years so much insight on the liquid state. These include C.A. Angell, R. Bowles, S.C. Glotzer, P. Kusalik, G. Matharoo, M.S.G. Razul, S. Sastry, F. Sciortino, H.E. Stanley and F. Starr. PHP and ISV are both supported by NSERC and ACEnet. PHP also acknowledges the support of the CRC program.

REFERENCES

1. P.G. Debenedetti, *Metastable Liquids: Concepts and Principles* (Princeton University Press, Princeton, 1997).
2. J.-P. Hansen and I.R. McDonald, *Theory of Simple Liquids* (Academic Press, London, 2006).
3. M. Allen and D. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, Oxford, 1989).
4. R.O. Sokolovskii, M. Thachuk and G.N. Patey, *J. Chem. Phys.* **125**, 204502 (2006).
5. D. Frenkel and B. Smit, *Understanding Molecular Simulation* (Academic Press, London, 2002).
6. J.-Y. Raty, E. Schwegler and S.A. Bonev, *Nature* **449**, 448 (2007).
7. S.A. Bonev, E. Schwegler T. Ogitsu and G. Galli, *Nature* **431**, 669 (2004).
8. I. Saika-Voivod, F. Sciortino, T. Grande and P.H. Poole, *Phys. Rev. E* **70**, 061507 (2004).
9. *Silica: Physical Behavior, Geochemistry and Materials Applications*, edited by P. Heaney, C. Prewitt, and G. Gibbs, *Reviews in Mineralogy* Vol. 29 (Mineralogical Society of America, Washington, D. C., 1994).
10. E. Sanz, C. Vega, J.L.F. Abascal, L.G. MacDowell, *Phys. Rev. Lett.* **92**, 255701 (2004).
11. P.H. Poole, F. Sciortino, U. Essmann and H.E. Stanley, *Nature* **360**, 324 (1992).
12. P.F. McMillan, M. Wilson, M.C. Wilding, D. Daisenberger, M. Mezouar and G..N. Greaves, *J. Phys.: Condens. Matter* **19**, 415501 (2007).
13. D. Paschek, *Phys. Rev. Lett.* **94**, 217802 (2004).
14. P.H. Poole, I. Saika-Voivod and F. Sciortino, *J. Phys.: Condens. Matter* **17**, L431 (2005).
15. D. Liu, Y. Zhang, C.-C. Chen, C.-Y. Mou, P.H. Poole, S.-H. Chen, *Proc. Natl. Acad. Sci. USA* **104**, 9570 (2007).
16. P.G. Debenedetti and F.H. Stillinger, *Nature* **410**, 259 (2001).
17. W. Kauzmann, *Chem. Rev.* **43**, 219 (1948).
18. P.W. Anderson, *Science* **267**, 1615 (1995).
19. W. Kob and H.C. Andersen, *Phys. Rev. E* **51**, 4626 (1995); W. Kob and H.C. Andersen, *Phys. Rev. E* **52**, 4134 (1995).
20. W. Götze, in *Liquids, Freezing and Glass Transition*, edited by J.P. Hansen, D. Levesque, and J. Zinn-Justin (North Holland, Amsterdam, 1991), p. 287.
21. D. Wales, *Energy Landscapes* (Cambridge University Press, New York, 2003).
22. G. Adam and J.H. Gibbs, *J. Chem. Phys.* **43**, 139 (1965).
23. I. Saika-Voivod, P.H. Poole and F. Sciortino, *Nature* **412**, 514 (2001).
24. M.D. Ediger, *Annu. Rev. Phys. Chem.* **51**, 99 (2000).
25. C. Donati, J.F. Douglas, W. Kob, S.J. Plimpton, P.H. Poole and S.C. Glotzer, *Phys. Rev. Lett.* **80**, 2338 (1998).
26. S.C. Glotzer, *Journal of Non-Crystalline Solids* **274**, 342 (2000).
27. E.R. Weeks, J.C. Crocker, A.C. Levitt, A. Schofield and D.A. Weitz, *Science* **287**, 627 (2000).
28. A. Widmer-Cooper, P. Harrowell and H. Fynewever, *Phys. Rev. Lett.* **93**, 135701 (2004).
29. G.S. Matharoo, M.S.G. Razul and P.H. Poole, *Phys. Rev. E* **74**, 050502(R) (2006).
30. I. Saika-Voivod, P.H. Poole and R.K. Bowles, *J. Chem. Phys.* **124**, 224709 (2006).
31. P.R. ten Wolde, M.J. Ruiz Montero, and D. Frenkel, *J. Chem. Phys.* **104**, 9932 (1996); S. Auer and D. Frenkel, *J. Chem. Phys.* **120**, 3015 (2004); C. Valeriani, E. Sanz, and D. Frenkel, *J. Chem. Phys.* **122**, 194501 (2005).
32. L.M. Ghiringhelli, C. Valeriani, E.J. Meijer, and D. Frenkel, *Phys. Rev. Lett.* **73**, 055702 (2007).
33. I.M. Svishchev and P.G. Kusalik, *Phys. Rev. Lett.* **99**, 975 (1994).
34. M.S. Gulam Razul, J.G. Hendry and P.G. Kusalik, *J. Chem. Phys.* **123**, 204722 (2005).
35. E. Mendez-Villuendas and R.K. Bowles, *Phys. Rev. Lett.* **98**, 185503 (2007).
36. T. Kawasaki, T. Araki, and H. Tanaka, *Phys. Rev. Lett.* **99**, 215701 (2007).