In and out of equilibrium

J. Kurchan

PMMH Ecole Supérieure de Physique et Chimie Industrielles, 10 rue Vauquelin, 75231 Paris CEDEX 05, France (e-mail: jorge@pmmh.espci.fr)

Albert Einstein's work on brownian motion showed how thermal equilibrium could be brought about by work exchanged through thermal fluctuations and viscous dissipation. Glasses are out-of-equilibrium systems in which this exchange happens at widely different timescales simultaneously. Theory then suggests the fascinating possibility that such behaviour may lead to a more general form of thermalization, in which the effective temperature shared by all components differs at each timescale.

Thermodynamics is a 'gift' of mechanics. Both newtonian and quantum dynamics in many-body systems lead to extremely complicated dynamical evolution. But both also have a feature that is exceptional amongst dynamical systems: in most cases, we can assume that all system configurations of a given energy will be visited with equal probability. When this assumption is valid, we say that the system is in equilibrium. We can then forget the evolution of the system in time, and describe its properties statistically — a procedure leading to thermodynamics. Disregarding the time-dependence involves a certain degree of technical simplification; but much more importantly, the statistical description brings about new, important concepts such as temperature and entropy. The price one pays is that certain quantities (such as diffusivity and viscosity) that do depend on the dynamical details cannot be studied with thermodynamics alone.

Einstein's papers on brownian motion showed a surprising fact for the first time: thermodynamics still contains dynamic information. A particle in suspension in a liquid is rocked by the irregularity of molecular collisions, but the viscous forces produced by the same molecules tend to dampen the motion. Energy is gained through the fluctuations (the irregular collisions) and lost through the dissipation (the viscous drag). The requirement that the particle be in equilibrium with the liquid then implies a balance between the diffusivity of the particle and the viscosity of the liquid — two purely dynamic quantities.

Contrast this situation with the case of glasses. Glasses are systems out of equilibrium, and a thermodynamic description of them is in principle not possible: strictly speaking, we are not even allowed to refer to 'the' temperature of a glass bottle. This is a major failure of thermodynamics.

But as I will discuss below, by taking into account heat exchange through fluctuations and dissipations between the constituents of a glass (with each process operating at different timescales), we are making progress towards a generalization of thermodynamics that could apply to this difficult problem.

Losing equilibrium

If cooled fast enough, liquids can be taken below their freezing temperature without crystallizing. When the temperature of a supercooled liquid is further lowered, the viscosity increases dramatically; but in contrast to the case of crystallization, this change hardly manifests itself in the microscopic structure. Rather, it becomes evident as a sharp slowing down of particle diffusion. The resulting sluggishness of molecular rearrangement makes equilibration of the system difficult. And as temperature is further lowered, there sooner or later comes a point when the liquid cannot keep pace with the changes in the thermal bath, and falls out of equilibrium: a glass is born.

If at any given point in time we keep the temperature fixed and wait, a glass will work its way slowly back to equilibrium — a process known as physical ageing. During ageing, a glass's viscosity gradually increases until it eventually reaches the equilibrium value at the new fixed temperature, but the time needed for this to happen can be astronomical. We are then inevitably led to ask whether there is a temperature below which no amount of waiting will lead to equilibration in a liquid state, and ageing essentially goes on forever — or until crystallization intervenes. There seems to be no answer yet in sight for this hard, inevitably fascinating (although in practice rather irrelevant), question.

Another way of perturbing a supercooled liquid is to force it to flow. When this happens, velocity gradients develop in the liquid, which then responds by lowering its viscosity — a phenomenon present (and welcome) when paint is spread and when blood flows in capillaries. What is important here is that this 'shear-thinning' phenomenon entails a strong deviation from equilibrium. A rough way to understand this is the following: because ageing involves a viscosity increase, a forced decrease in viscosity can be thought of as having a rejuvenating effect. For a glass, this means being forced further away from equilibrium. In a sense that will be made more precise below, a deeply super-cooled liquid 'realises' that it is almost a glass as soon as we try to make it flow.

In practice, we know that a system is not in equilibrium if some quantity (typically energy or density) is not constant in time, or, more generally, if thermal currents inside the system have not yet died away. But as we do not expect any observable of an old piece of glass to be wildly evolving, or that strong thermal currents are still flowing within its bulk, we may wonder why glasses are usually described as systems that are far from equilibrium. A clarifying approach to this question, which takes us back to Einstein's brownian motion papers (and to some recent developments in glass theory), is to think of equilibrium as a situation in which every conceivable type of thermometer coupled to a part of the system reads the same temperature. We declare a system far from equilibrium if two different thermometers indicate substantially different temperatures in the same region of space. This is exactly what happens in glasses.

Because physicists are much more at home when they are able to apply the well-developed equilibrium techniques, the more traditional theoretical approach to glasses has been to consider these systems in an ideal limiting situation, in which they have aged at each temperature for sufficiently long to equilibrate — essentially giving up the description of a main aspect of glassiness. But it should be borne in mind that, the lower the temperature, the more unrealistically long this waiting time is; and that, even if equilibration is achieved, it will be destroyed by shear-thinning as soon as the liquid is made to flow.

Two timescales

If we follow the motion of a particle in a regular crystalline solid, we see it performing fast vibrations around its assigned position in the lattice, and only very rarely swapping sites with a neighbour. In contrast, a particle in a normal liquid wanders away from any given position at a fairly uniform rate. As we take the liquid deeper into the supercooled region, the motion starts to change in several ways. Most strikingly, it splits into a fast vibrational ('cage' or 'rattling') motion of amplitude comparable with the particle size, superimposed on a wandering motion ('structural relaxation') that is on average much slower. As the temperature is further decreased, the rattling continues in much the same way, but the time taken for a particle to wander increases dramatically: the timescales for 'cage' and 'structural' motion become more distinct. (The structural rearrangements become also more spatially correlated: there are small epidemics of mobility continuously spreading and dying out in the sample.) This behaviour is reflected in all observables: local quantities — such as density, magnetization or electric polarization — also fluctuate in two well-separated timescales (Fig. 1).

After a sudden reduction in temperature, the lengthening of the 'slow' timescale is not immediate: the system ages gradually into a situation of slower structural rearrangements, and this is accompanied by the aforementioned increase in viscosity. Conversely, when a supercooled liquid undergoes shear-thinning, the decrease of the viscosity is the result of a shortening of the structural rearrangement timescale, while the fast motion is practically unaffected. All this strongly suggests that if we are trying to understand the elements that are specific to glassiness, especially its out-of-equilibrium nature, it is on the slow part of the motion that we have to concentrate, the fast motion being almost the same as in an equilibrated system. But we also see that the very same actors — the particles — are performing both fast and slow motion simultaneously: is it possible that the same particle is both in and out of equilibrium?

Our way to probe equilibration is by means of thermometers, so to make the discussion more quantitative, we have to think about how temperature is measured within any system.

Thermometers and the fluctuation–dissipation theorem

Consider the simplest kind of thermometer: a harmonic oscillator of frequency Ω, such as a mass attached to a spring (or brownian particles in the springless limit $\Omega = 0$). To measure the temperature (T) of a system, we couple the oscillator to any of its components. In equilibrium, equipartition of energy tells us that, on average, the oscillator's kinetic energy must be $1/2 k_B T$ (where k_B is Boltzmann's constant), so it suffices to measure its mean value over a sufficiently long time to deduce T. Equilibrium thermodynamics guarantees that T will be independent of the oscillator's frequency Ω and also of the details of its coupling to the system.

There are situations in which we wish to couple the thermometer to a microscopic physical quantity X (for example, a local magnetization, electric polarization or density fluctuation) and ensure that the thermometer disturbs the evolution of this quantity as little as possible. At the level of thought experiment, one way to do this is to imagine that we have M copies of the system A_1, A_2, ..., A_M that have gone through the same experimental procedure, and the thermometer couples to the observable X corresponding to each copy of the system simultaneously (Fig. 2): if M is large, then the coupling to each individual copy will be very small. We may assume that X fluctuates around zero, and has on average correlation $<X(t)X(t')> = C_X(t - t')$, or $C_X(\omega)$ in Fourier space. The fluctuations of X act as a noise that pumps energy into the thermometer, just like the irregularities in molecular collisions push a brownian particle.

However, this is not the end of the story. Each system A_i feels the presence of the thermometer as a small perturbation, and this on average modifies the Xs. How much exactly X responds on average at time t to a small perturbation at time t' is measured by the so-called response function $R_X(t - t')$ (or $R_X(\omega)$ in Fourier space). This modification in
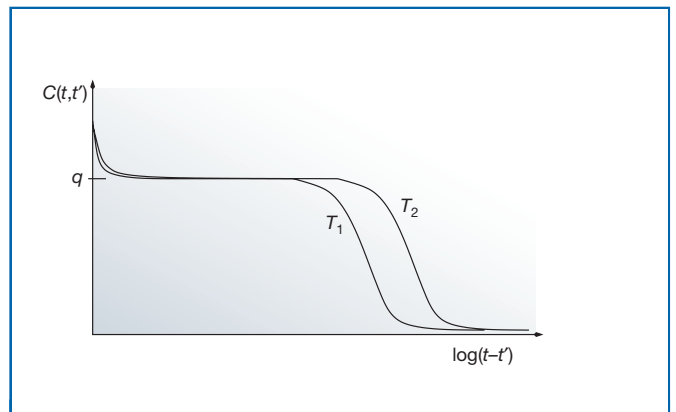

Figure 1 The two-time (t,t') correlation $C(t,t')$ of a typical physical quantity in terms of $\log(t - t')$ for two temperatures $T_1 > T_2$. The initial step ($C > q$) corresponds to the fast vibrational motion, and the second ($C < q$) reflects the slower structural relaxation (see text). The timescale for the latter is not only very sensitive to temperature, but also to shear rate — or, if the system is ageing, to the time elapsed since it fell out of equilibrium..
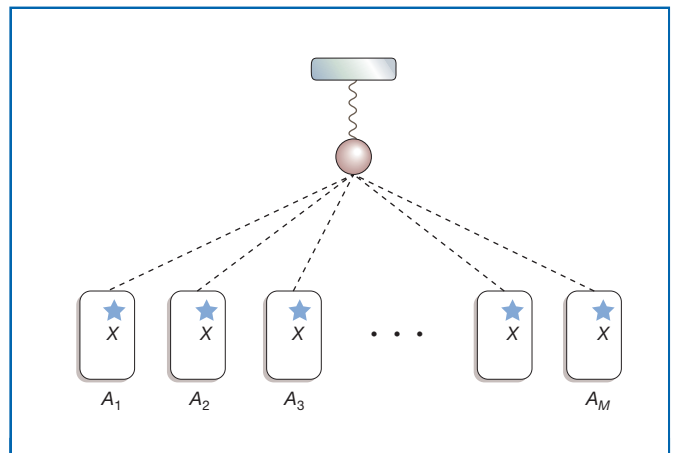

Figure 2 M copies of a system (A_1, A_2,..., A_M). The thermometer (an oscillator) is coupled to the physical quantity X in each copy of the system.

the Xs feeds back into the thermometer, and tends to decrease its energy: this is the 'dissipation', which in the case of the agitated brownian particle represents the role of the liquid's viscosity. In summary, the thermometer feels the 'bare' noise of the observable — the 'fluctuations' (measured by $C_X(\omega)$) — and the echo of its own action on the system (proportional to $R_X(\omega)$). These two contrary effects must be such that they give the right energy as predicted by equipartition, for every conceivable thermometer and observable. A short calculation shows that this can only happen if correlations and responses associated with any observable are proportional: $T = [\omega C_X(\omega)]/[2\mathrm{Im}R_X(\omega)]$. This is the fluctuation–dissipation theorem[1], a generalization of the relation found by Einstein between the diffusivity of brownian particles and the viscosity of the liquid.

Multi-thermalization

As we have just seen, a thermometer that selectively responds to a frequency Ω when coupled to a physical quantity X will indicate a temperature corresponding to the fluctuation–dissipation ratio: $T(\Omega,X) = [\Omega C_X(\Omega)]/[2\mathrm{Im}R_X(\Omega)]$. It is then natural to refer to the quantity $T(\Omega,X)$ as the 'effective temperature' of X at timescale $1/\Omega$. In a system in equilibrium, all the $T(\Omega,X)$ are equal to the ambient temperature T.

In a situation of well-separated timescales, a first result — which can indeed be proved quite generally — is that a thermometer that
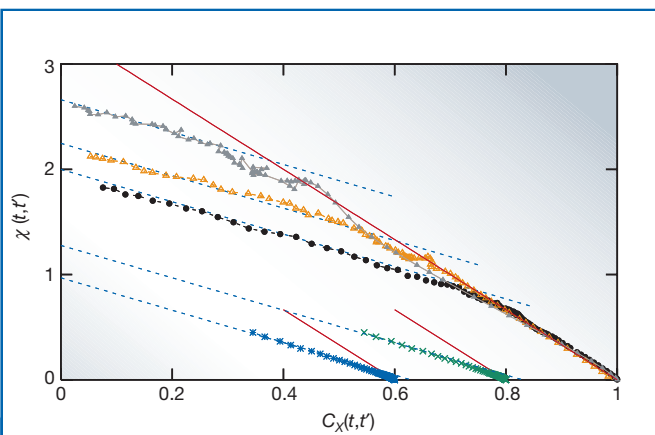
Figure 3 Fluctuation–dissipation plot for a sheared supercooled liquid, taken from ref. 5. The results are from a simulation of a mixture of two types of particle (see text); $\chi(t,t')$ is the integrated response, and $C_X(t,t')$ is the correlation function. Shown in black (filled circles), orange (open triangles) and grey (filled triangles) are the correlation of Fourier components of the structure factor for wave vector $k = 7.47$, $k = 11.22$ and $k = 14.48$, respectively. The crosses in green and blue are plots of the negative of the mobility versus the self-diffusion for each type of particle (the origin is shifted for ease of comparison). The solid lines correspond to the equilibrium value $-1/T$, while the dashed lines are all parallel with the same gradient $\equiv -1/T_{eff}$. The two-temperature model seems to fit the data well. (The fast timescale is not visible in the self-diffusion, as it implies distances too small to be visible on this plot.)

takes heat from the system only through the higher-frequency fluctuations (the ones that are least affected by departures from equilibrium) will indicate the ambient temperature T when coupled to any observable. The heat sensors in our skin are an example, and this explains why glasses feel no hotter to the touch than substances in equilibrium.

The out-of-equilibrium dynamics of glassy systems has been solved within a set of approximations that we shall describe in some detail below[2–4]. A striking outcome is that fluctuations at long timescales comparable to that of structural relaxation are at an effective temperature T_{eff} that is higher than the bath temperature T — the same one for all physical quantities. In an ageing system, T_{eff} is of the order of (but not equal to) the temperature that the system had at the point it fell out of equilibrium, and gradually decreases with time to T if and when the system re-equilibrates. In a flowing supercooled liquid, the effective temperature becomes different as soon as shear-thinning appears. This two-temperature (T, T_{eff}) description is only meaningful to the extent that the timescales of fast 'rattling' motion and slow 'structural rearrangements' remain well separated. (One can of course conceive of cases with more than two well-separated timescales, each with its own effective temperature.) We are still very far from being able to validate this description exactly in a realistic model of a glass, but numerical simulation results for realistic glass models, and to some extent the results of experiments, are encouraging.

A way to interpret the fluctuation–dissipation data at a glance is the following[3]: for a physical quantity X, we measure the correlation function $C_X(t,t')$ and the integrated response:

$$\chi(t,t') = \int_{t'}^{t} d\tau\, R_X(t,\tau)$$

In an equilibrium system, the fluctuation–dissipation theorem states that a parametric plot of $\chi(t,t')$ against $C_X(t,t')$ yields a straight line with slope $-1/T$. If an out-of-equilibrium system has a two-step correlation, as in Fig. 1, and temperatures T and T_{eff} associated with the fast and slow timescales, respectively, the $\chi(t,t')$ versus $C_X(t,t')$ plot will be composed of two straight segments of slopes $-1/T$ and $-1/T_{eff}$ corresponding to short and long times. Clearly, if there is a

single T_{eff} for all physical quantities, the segments associated with the long-time decay of all quantities must be parallel. Figure 3 shows the results of a simulation of a mixture of two types of particle interacting via a Lennard–Jones potential, subjected to shear[5]. Similar results have been obtained for granular matter[6]. The results are quite reassuring.

Experimental results are also encouraging but currently contain much less detail[7–9], mainly owing to difficulties in accessing the response and correlation functions of many observables. Such behaviour is sometimes further obscured by the presence of a strong, intermittent, 'crackling' noise, an interesting and as yet not well-understood phenomenon, which seems to be operating at an entirely different length scale — and is probably beyond the scope of microscopic glass theories[8].

At the phenomenological level, the notion that glasses can be described by two temperatures — the ambient temperature and a higher 'fictive' temperature — has been around for a long time[10]. It seems natural to think that the theoretical developments correspond by and large to that basic idea. It is difficult, however, to make a direct comparison between a theoretical framework that imposes relations and constraints, and a phenomenological approach that has been optimized to fit the experimental data. For example, to be meaningful as temperatures, the effective temperatures have to coincide at each timescale for all observables, a restriction that has been largely abandoned in the 'fictive temperature' phenomenological approaches.

'Mean-bath' theory

It is possible to understand the approximations that yield the effective temperature without going into too much technical detail. The argument also suggests why the same ideas could apply more generally.

In Fig. 2, the systems $A_1, A_2, ..., A_M$ are assumed to be equilibrated at temperature, T, so all the correlations and their associated response functions obey fluctuation–dissipation relations. I have argued above that these relations imply that thermalization is inherited by the thermometer, which in turn implies that all observables belonging to the thermometer itself should also satisfy fluctuation–dissipation relations. We could now couple the thermometer (or a set of such thermometers) to yet another system, and the inheritance of thermalization at temperature T will be carried one generation further. All this is very natural, but we can now ask a less obvious question: what would happen if the systems $A_1, A_2, ..., A_M$ had two timescales with two different effective temperatures? Would this property be inherited in the same way? The answer turns out to be 'yes', provided that the two timescales are widely separated.

Consider now an approximation in which we break the system into components, each of which is assumed to be thermally coupled to the fluctuations of the rest (like the thermometer in Fig. 2), so that a particular component 'feels' the others as a thermal bath. The argument becomes self-consistent when one assumes that all components behave on average in the same way[2]. What we are therefore proposing is just like the familiar mean-field approximation, except that here the subsystems only interact via their fluctuations. What is transmitted is not the average value of the observables, but rather the time-dependent fluctuations around these averages: there is a 'mean thermal bath' rather than a 'mean-field'.

It is now clear that equilibrium is a possibility. Placed in contact with a mean thermal bath that is in equilibrium, a particular component will equilibrate, which in turn justifies the assumption for all the other components. However, if the system is out of equilibrium, either because it is being driven or because it is ageing, and has widely separated timescales, a self-consistent solution with two temperatures is also possible as, by virtue of the 'hereditary' property described above, a component in contact with a bath having two temperatures will also develop two temperatures.

The theoretical situation with this 'multi-thermalization' idea is quite similar to that of ferromagnetism a century ago, when a mean-field approximation was developed that showed how spontaneous magnetization is possible. In that case, each spin is assumed

to interact with all the rest: if the 'other' spins are assumed to have positive magnetization, they force positive magnetization on the selected spin, in turn justifying the initial choice of magnetization for the rest. The lesson that emerged was that a symmetry of the system can be spontaneously broken — but it took many years to show that this could also happen in more realistic lattice models of the spin systems. Similarly, we now know that multi-thermalization schemes are realizable within the mean thermal bath approximations described above, but we have no analytic proof that it happens in an exact solution of a realistic model. The numerical checks are encouraging, but will never tell us whether the multi-thermalization holds strictly, or only as an approximation.

Conclusions

The idea that glasses can be described by two temperatures has a long history. Rethinking the whole question of heat exchange, in a manner close to Einstein's in his work on brownian motion, has made possible a theoretical basis for these old phenomenological ideas. The fluctuation–dissipation results suggest that, in certain cases, there can be an underlying generalization of statistical mechanics with two or more temperatures. Indeed such an approach was proposed for granular matter some time ago[6,11].

In general, what thermodynamics offers is not so much the technical advantage of being able to disregard time — after all, this only reduces the dimensionality of our world from four to three — but rather a framework with powerful notions such as temperature, entropy and equilibration, together with inequalities and arrows of time. Physicists feel so lost without them, that inevitably they talk about 'hot', 'cold', 'entropy' and 'thermalization' even in situations that are far from equilibrium, despite being aware that these terms are not really defined in those cases. Any generalization of thermodynamics that applies to glasses (which represent approximately half the solids that surround us) will be most welcome. ☐

1. Kubo, R., Toda, M. & Hashitsume, N. *Statistical Physics II: Nonequilibrium Statistical Mechanics* (Springer, New York, 1991).
2. Sompolinsky, H. & Zippelius, A. Relaxational dynamics of the Edwards-Anderson model and the mean-field theory of spin-glasses. *Phys. Rev. B* **25,** 6860–6875 (1982).
3. Cugliandolo, L. F. & Kurchan, J. On the out-of equilibrium relaxation of the Sherrington-Kirkpatrick model. *J. Phys. A* **27,** 5749–5772 (1994).
4. Cugliandolo, L. F., Kurchan, J. & Peliti, L. Energy flow, partial equilibration, and effective temperatures in systems with slow dynamics. *Phys. Rev. E* **55,** 3898–3914 (1997).
5. Berthier, L. & Barrat, J.-L. Nonequilibrium dynamics and fluctuation-dissipation in a sheared fluid. *J. Chem. Phys.* **116,** 6228–6242 (2002).
6. Makse, H. A. & Kurchan, J. Testing the thermodynamic approach to granular matter with a numerical model of a decisive experiment. *Nature* **415,** 614–617 (2002).
7. Israeloff, N. E. & Grigera, T. S. Low-frequency dielectric fluctuations near the glass transition. *Europhys. Lett.* **43,** 308–313 (1998).
8. Buisson, L., Ciliberto, S. & Garcimartin, A. Intermittent origin of the large violations of the fluctuation-dissipation relations in an aging polymer glass. *Europhys. Lett.* **63,** 603–609 (2003).
9. Herisson, D. & Ocio, M. Fluctuation-dissipation ratio of a spin glass in the aging regime. *Phys. Rev. Lett.* **88,** 257202 (2002).
10. Tool, A. Q. Relation between inelastic deformability and thermal expansion of glass in its annealing range. *J. Am. Ceram. Soc.* **29,** 240–253 (1946).
11. Edwards, S. F. in *Granular Matter: An Interdisciplinary Approach* (ed. Mehta, A.) 121–140 (Springer, New York, 1994).