# Statistical Physics of Athermal Systems
# Lecture 5

Kabir Ramola
Tata Institute of Fundamental Research, Hyderabad

August 28, 2026

## 1 Vibrational modes at the jamming transition

For a mechanically stable packing, let $\mathbf{R}_0$ denote a local minimum of the potential energy,

$$\nabla_{\mathbf{R}} U|_{\mathbf{R}_0} = 0. \tag{1}$$

Writing a nearby configuration as

$$\mathbf{R} = \mathbf{R}_0 + \mathbf{u}, \tag{2}$$

we expand the potential energy in powers of the displacement $\mathbf{u}$:

$$U(\mathbf{R}_0 + \mathbf{u}) = U_0 + \left.\frac{\partial U}{\partial R_a}\right|_{\mathbf{R}_0} u_a + \frac{1}{2}\left.\frac{\partial^2 U}{\partial R_a \partial R_b}\right|_{\mathbf{R}_0} u_a u_b + O(u^3), \tag{3}$$

where repeated Cartesian indices $a, b$ are summed. Since $\mathbf{R}_0$ is a mechanical equilibrium, the linear term vanishes. We therefore have

$$U(\mathbf{R}_0 + \mathbf{u}) = U_0 + \frac{1}{2} u_a H_{ab} u_b + O(u^3), \tag{4}$$

where the Hessian is defined by

$$H_{ab} \equiv \left.\frac{\partial^2 U}{\partial R_a \partial R_b}\right|_{\mathbf{R}_0}. \tag{5}$$

For $N$ particles in $d$ dimensions, $\mathbf{R}$ has $Nd$ components and $H$ is therefore an $Nd \times Nd$ matrix. The Hessian measures the curvature of the potential-energy landscape about the mechanically stable configuration. Positive eigenvalues correspond to directions in configuration space in which the energy increases quadratically under a small displacement.

The normal modes are obtained by solving

$$H\mathbf{e}_\alpha = \lambda_\alpha \mathbf{e}_\alpha, \qquad \mathbf{e}_\alpha^{\mathsf{T}} \mathbf{e}_\beta = \delta_{\alpha\beta}. \tag{6}$$

For identical particles of mass $m$, the harmonic equations of motion are

$$m\ddot{\mathbf{u}} = -H\mathbf{u}. \tag{7}$$

Expanding the displacement in normal modes,

$$\mathbf{u}(t) = \sum_\alpha q_\alpha(t) \mathbf{e}_\alpha, \tag{8}$$

gives

$$m\ddot{q}_\alpha = -\lambda_\alpha q_\alpha, \tag{9}$$

and hence

$$\omega_\alpha^2 = \frac{\lambda_\alpha}{m}. \tag{10}$$

In a crystal, translation symmetry allows the eigenvectors to be classified by wavevector. The low-frequency modes are acoustic plane waves and the long-wavelength theory is ordinary elasticity. The corresponding density of states is the Debye result,

$$D_{\mathrm{Debye}}(\omega) \sim \omega^{d-1}. \tag{11}$$

In an amorphous solid the Hessian is a disordered matrix. There is no exact crystal momentum, and there is no reason for the eigenvectors to be plane waves on microscopic scales. Nevertheless, translational invariance is recovered after coarse graining, and ordinary elasticity can emerge at sufficiently large length scales.

The interesting question is what happens when the elastic solid is brought to the point where it is only just rigid. This is the jamming transition. At that point the number of constraints is just sufficient to eliminate the extensive floppy subspace, the lowest-frequency part of the spectrum becomes anomalous, and the elastic response becomes strongly non-affine.

Thus a small Hessian eigenvalue corresponds to a soft mode, and a soft mode has a low characteristic frequency. At a mechanically stable minimum all nontrivial eigenvalues must be non-negative. For a system with translational invariance, the uniform translation of all particles does not change the energy, so the Hessian necessarily has $d$ zero modes associated with global translations. These trivial zero modes are excluded when discussing the internal vibrational spectrum.

Far from jamming, the low-frequency modes are acoustic and the density of states has the Debye form. Near jamming, an additional population of anomalous modes develops above a characteristic frequency $\omega^*$. The cleanest way to understand the scale $\omega^*$ is the variational cutting argument.

### 1.1 Variational argument

The following variational construction is due to Wyart (Annales de Physique 30, 1–96 (2005)) and provides a direct connection between isostaticity and the low-frequency spectrum. Consider a region of linear size $\ell$ inside an isostatic packing. Cutting all contacts that cross the boundary removes a number of constraints proportional to the boundary area,

$$N_{\mathrm{cut}} \sim \ell^{d-1}. \tag{12}$$

At isostaticity the cut produces a comparable number of floppy modes. Take one such mode and multiply it by a smooth envelope which vanishes at the boundary,

$$\mathbf{u}_\beta(\mathbf{r}) = C_\beta \sin\left(\frac{\pi x}{\ell}\right) \mathbf{e}_\beta(\mathbf{r}). \tag{13}$$

The envelope introduces gradients of order $1/\ell$. For harmonic stiffness this gives a trial-mode energy of order

$$\delta E \sim \frac{1}{\ell^2}, \tag{14}$$

for a normalized mode. Therefore

$$\omega_\ell \sim \frac{1}{\ell} \tag{15}$$

when the microscopic stiffness is used as the unit of energy.

There are $O(\ell^{d-1})$ such modes in a region of volume $O(\ell^d)$. Thus the fraction of modes below $\omega_\ell$ scales as

$$\int_0^{\omega_\ell} D(\omega)\,\mathrm{d}\omega \sim \frac{\ell^{d-1}}{\ell^d} \sim \frac{1}{\ell} \sim \omega_\ell. \tag{16}$$

If $D(\omega) \sim \omega^q$, then

$$\omega_\ell^{q+1} \sim \omega_\ell, \tag{17}$$

which gives

$$q = 0. \tag{18}$$

Thus the anomalous contribution to the density of states is approximately flat at low frequency, within the regime where the isostatic network description and harmonic expansion apply.

For general contact stiffness $k$, the frequency carries an overall factor $\sqrt{k}$, so the more useful force-law-independent statement is that the crossover frequency satisfies

$$\frac{\omega^*}{\sqrt{k}} \sim \Delta z. \tag{19}$$

Here the mass has been set to unity. Restoring the mass gives the dimensionally explicit form

$$\omega^* \sim \sqrt{\frac{k}{m}}\,\Delta z, \qquad \frac{\omega^*}{\sqrt{k/m}} \sim \Delta z. \tag{20}$$

For harmonic interactions $k$ is constant and therefore

$$\omega^* \sim \Delta z \sim \Delta\phi^{1/2}. \tag{21}$$

For Hertzian interactions, $k \sim \Delta\phi^{1/2}$, giving

$$\omega^* \sim \Delta\phi^{3/4}. \tag{22}$$

The same crossover can be expressed directly in terms of the overdamped dynamics discussed later. Since

$$\lambda^* = m(\omega^*)^2 \sim k\Delta z^2, \tag{23}$$

the corresponding relaxation time is

$$\tau^* \sim \frac{\gamma}{\lambda^*} \sim \frac{\gamma}{k\Delta z^2}. \tag{24}$$

Thus the softening of the Hessian has a dynamical consequence as well as a static one: the same eigenvalues that determine the vibrational frequencies in an inertial system determine the relaxation times in an overdamped system.

## 2 Elastic moduli near jamming

Let $k$ denote the characteristic contact stiffness. The jamming scaling is

$$B \sim k, \qquad G \sim k\Delta z, \qquad \frac{G}{B} \sim \Delta z. \tag{25}$$

For harmonic contacts, $k$ is independent of compression at leading order,

$$B \sim 1, \qquad G \sim \Delta z \sim \Delta\phi^{1/2}. \tag{26}$$

For Hertzian contacts,

$$k \sim \Delta\phi^{1/2}, \tag{27}$$

and therefore

$$B \sim \Delta\phi^{1/2}, \qquad G \sim \Delta\phi. \tag{28}$$

Since $p \sim \Delta\phi^{3/2}$ for Hertzian contacts,

$$B \sim p^{1/3}, \qquad G \sim p^{2/3}. \tag{29}$$

Thus $G/B \to 0$ on approaching jamming. The interaction law sets $k$, while the near-isostatic geometry supplies the factor $\Delta z$ in the shear response.

Continuum elasticity applies on scales large compared with the particle size and the isostatic length:

$$\ell \gg \ell^*. \tag{30}$$

Since $\ell^* \to \infty$ as $\phi \to \phi_J^+$, this regime is pushed to larger scales.

# 3 Lengthscales near jamming

## 3.1 The isostatic length

It is useful to derive the isostatic length explicitly, because the origin of $\Delta z$ is easily obscured by the compact scaling notation. We define

$$\Delta z = z - z_{\text{iso}}, \tag{31}$$

where $z$ is the mean coordination of the load-bearing contact network and $z_{\text{iso}}$ is its isostatic value. For frictionless particles in the thermodynamic limit,

$$z_{\text{iso}} = 2d. \tag{32}$$

Finite systems have corrections associated with boundaries, rattlers, and states of self stress; the scaling argument below concerns the bulk quantity $\Delta z$.

Now consider a compact region of linear size $\ell$. Its volume is $O(\ell^d)$, so the excess connectivity inside the region corresponds to an excess number of constraints of order

$$N_{\text{excess}} \sim \Delta z\, \ell^d. \tag{33}$$

The important point is that this is a bulk contribution: every unit volume contributes an amount proportional to $\Delta z$.

Next imagine cutting the region away from the rest of the packing. This removes only contacts that cross the boundary. The number of such contacts scales with the boundary measure,

$$N_{\text{cut}} \sim \ell^{d-1}. \tag{34}$$

At exact isostaticity the boundary cut creates floppy degrees of freedom because the isolated region has lost precisely the constraints that were attached through the boundary. Above isostaticity, the excess constraints in the interior oppose these floppy motions.

A cut mode can therefore survive as an extended low-energy trial mode only while the number of constraints removed at the boundary is comparable to or larger than the number of excess constraints in the bulk,

$$N_{\mathrm{cut}} \gtrsim N_{\mathrm{excess}}. \tag{35}$$

Using the two scaling estimates,

$$\ell^{d-1} \gtrsim \Delta z\, \ell^{d}. \tag{36}$$

Dividing by $\ell^{d-1}$ gives

$$1 \gtrsim \Delta z\, \ell. \tag{37}$$

The largest scale over which the boundary can still compensate the excess connectivity is therefore

$$\ell^{*} \sim \Delta z^{-1}. \tag{38}$$

The interpretation is worth stating carefully. For $\ell \ll \ell^{*}$, the boundary term is larger than the excess bulk term, so the region can support floppy-like trial motions. For $\ell \gg \ell^{*}$, the excess constraints dominate and the region behaves as an overconstrained elastic body. The isostatic length is therefore a crossover length obtained from constraint counting; it is not the size of an individual defect.

It is also useful to distinguish a floppy mode from a soft normal mode. In the unstressed network described by the compatibility matrix, a floppy displacement $\mathbf{u}$ satisfies

$$\mathsf{C}\mathbf{u} = 0, \tag{39}$$

so that every remaining contact has zero linearized extension, $\delta\ell_{\alpha} = 0$. The stiffness contribution to the quadratic energy then vanishes,

$$\mathbf{u}^{\mathsf{T}} H_{\mathrm{stiffness}} \mathbf{u} = \mathbf{u}^{\mathsf{T}} \mathsf{C}^{\mathsf{T}} \mathsf{K} \mathsf{C} \mathbf{u} = 0. \tag{40}$$

In a real jammed packing the contacts are prestressed, so the full Hessian also contains the geometric prestress term. The modes produced by the cutting argument are therefore best regarded as low-energy trial modes once prestress is restored, rather than as exact zero modes of the full Hessian.

This counting argument can be summarized as follows. The two terms scale differently with $\ell$: the number of removed boundary constraints grows as $\ell^{d-1}$, whereas the number of excess bulk constraints grows as $\Delta z \ell^{d}$. As $\Delta z \to 0^{+}$, a larger and larger region is required before the bulk excess connectivity becomes dominant.

If we move away from isostaticity so that $\Delta z > 0$. A region of size $\ell$ contains an excess number of constraints of order

$$N_{\mathrm{excess}} \sim \Delta z\, \ell^{d}. \tag{41}$$

The boundary cut removes $O(\ell^{d-1})$ constraints. The anomalous construction can survive only while the boundary can compensate the excess constraints,

$$\ell^{d-1} \gtrsim \Delta z\, \ell^{d}. \tag{42}$$

The crossover is therefore

$$\ell^{*} \sim \Delta z^{-1}. \tag{43}$$

Using $\Delta z \sim \Delta\phi^{1/2}$,

$$\ell^{*} \sim \Delta\phi^{-1/2}. \tag{44}$$

This is the isostatic or cutting length. It is a constraint-counting length and should not be confused with the acoustic wavelengths introduced below.

## 4 Pre-stress and marginal stability

Maxwell counting gives a necessary condition for rigidity but says nothing about the sign of the eigenvalues of the prestressed Hessian. Near jamming the positive stiffness contribution and the negative prestress contribution compete.

For an anomalous mode the elastic part scales as $\Delta z^2$ in units of the microscopic stiffness, while the destabilizing prestress contribution scales with the typical compressive strain $e$. The mode frequency can therefore be written schematically as

$$\omega^2_{\mathrm{anom}} \sim \frac{k}{m}\left(A\Delta z^2 - Be\right), \tag{45}$$

with positive constants $A$ and $B$ of order unity in the scaling theory. Stability requires

$$\Delta z^2 \gtrsim e. \tag{46}$$

Near jamming $e \sim \Delta\phi$, giving

$$\Delta z \gtrsim \Delta\phi^{1/2}. \tag{47}$$

The bound is a stability criterion. The relation actually observed for the standard quasistatic ensembles is that the packings remain close to marginal stability:

$$\Delta z \sim \Delta\phi^{1/2}. \tag{48}$$

This is the sense in which the jammed state is marginal: the contact network contains only just enough excess connectivity to balance the destabilizing effect of the forces already stored in the contacts. The inequality is a stability bound. The stronger statement $\Delta z \sim \Delta\phi^{1/2}$ describes the marginally stable states selected by the standard quasistatic preparations near jamming. The stability argument alone does not require every preparation protocol to saturate the bound.

## 5 The Hessian of a central-force network

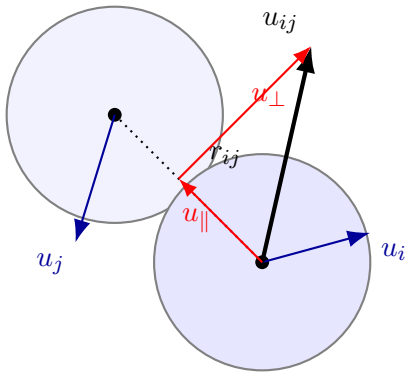


Figure 1: Definition of relative displacement $u_{ij}$, its parallel component $u_\parallel$ along the interparticle axis $r_{ij}$, and its perpendicular component $u_\perp$.

Let two particles interact through a central pair potential $V(r)$. Write

$$r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|, \qquad \hat{\mathbf{n}}_{ij} = \frac{\mathbf{r}_i - \mathbf{r}_j}{r_{ij}}, \tag{49}$$

and let

$$\delta \mathbf{r}_{ij} = \mathbf{u}_i - \mathbf{u}_j \tag{50}$$

be the relative displacement.

Decompose it into longitudinal and transverse parts,

$$\delta r_{ij}^{\parallel} = \hat{\mathbf{n}}_{ij} \cdot \delta \mathbf{r}_{ij}, \tag{51}$$

$$\delta \mathbf{r}_{ij}^{\perp} = \delta \mathbf{r}_{ij} - \hat{\mathbf{n}}_{ij} \delta r_{ij}^{\parallel}. \tag{52}$$

Expanding the pair potential to second order gives

$$\delta^2 V_{ij} = \frac{1}{2} \left[ V''(r_{ij})(\delta r_{ij}^{\parallel})^2 + \frac{V'(r_{ij})}{r_{ij}} |\delta \mathbf{r}_{ij}^{\perp}|^2 \right]. \tag{53}$$

The first term is the ordinary longitudinal stiffness of the contact. There is also a useful matrix form of this decomposition. For a contact $\alpha = (ij)$, define its linearized extension

$$e_\alpha = \hat{\mathbf{n}}_\alpha \cdot (\mathbf{u}_i - \mathbf{u}_j). \tag{54}$$

Collecting all contacts gives

$$\mathbf{e} = \mathsf{C}\,\mathbf{u}, \tag{55}$$

There are two different objects that are denoted by a similar symbol in the following discussion. For a contact $\alpha = (ij)$, $e_\alpha$ is a scalar contact extension. Later, a normal mode is written as a vector $\mathbf{e}_\mu$ satisfying the Hessian eigenvalue equation. To keep the distinction clear, we shall refer to the contact extension as $\delta\ell_\alpha$ whenever it is useful to do so:

$$\delta\ell_\alpha = \hat{\mathbf{n}}_\alpha \cdot (\mathbf{u}_i - \mathbf{u}_j), \qquad H\mathbf{e}_\mu = \lambda_\mu \mathbf{e}_\mu. \tag{56}$$

The first quantity lives in contact space; the second is a displacement pattern in the full particle configuration space. This distinction is important in the counting argument below. where $\mathsf{C}$ is the compatibility matrix. If the contact stiffnesses are assembled in a diagonal matrix $\mathsf{K}$, the unstressed quadratic energy is

$$\delta U = \frac{1}{2}\mathbf{e}^{\mathsf{T}}\mathsf{K}\mathbf{e} = \frac{1}{2}\mathbf{u}^{\mathsf{T}}\mathsf{C}^{\mathsf{T}}\mathsf{K}\mathsf{C}\,\mathbf{u}, \tag{57}$$

so that

$$H_{\text{stiffness}} = \mathsf{C}^{\mathsf{T}}\mathsf{K}\mathsf{C}. \tag{58}$$

The additional geometric term generated by pre-existing contact forces is precisely the prestress contribution discussed above. This makes the connection between Maxwell counting, the compatibility constraints, and the Hessian explicit. The second term comes from the fact that the bond is already carrying a force before the displacement is applied; it is the pre-stress contribution. For a compressed repulsive contact, the force is directed along the line joining the particles and $V'(r_{ij}) < 0$. Hence

$$\frac{V'(r_{ij})}{r_{ij}} < 0. \tag{59}$$

A transverse displacement therefore lowers the quadratic energy. This is why compression destabilizes transverse fluctuations.

The Hessian can be written schematically as

$$H = H_{\text{stiffness}} + H_{\text{prestress}}. \tag{60}$$

Near jamming the two terms become comparable. The first is controlled by the contact stiffness; the second is controlled by the force carried by the contacts. This competition is the origin of marginal stability.

## 6 Affine and non-affine response

Apply a small macroscopic strain $\epsilon$. The particle displacement may be written as

$$\mathbf{u} = \mathbf{u}_{\mathrm{aff}} + \mathbf{u}_{\mathrm{na}}, \tag{61}$$

where $\mathbf{u}_{\mathrm{aff}}$ is the displacement imposed by the macroscopic strain and $\mathbf{u}_{\mathrm{na}}$ is the additional relaxation required to restore force balance. The affine field is the displacement dictated by the imposed boundary deformation. The non-affine part is the particle-scale correction required because the local contact geometry is disordered. The figure is schematic: even compression can contain non-affine motion near jamming, but shear is especially sensitive to these relaxations.

For a homogeneous deformation with deformation gradient $F$, the affine displacement of particle $i$, initially at $\mathbf{R}_i^0$, is

$$\mathbf{u}_i^{\mathrm{aff}} = (F - I)\mathbf{R}_i^0. \tag{62}$$

The actual relaxed displacement is $\mathbf{u}_i$, and the non-affine displacement is defined precisely as

$$\mathbf{u}_i^{\mathrm{na}} = \mathbf{u}_i - \mathbf{u}_i^{\mathrm{aff}}. \tag{63}$$

For simple shear in the $xy$ plane,

$$u_{i,x}^{\mathrm{aff}} = \gamma_{\mathrm{sh}} y_i^0, \qquad u_{i,y}^{\mathrm{aff}} = 0, \tag{64}$$

so the non-affine part is the additional particle motion required after the imposed shear has been applied. It is therefore a relaxation variable, not another externally imposed strain.

A standard local measure of non-affinity is $D^2_{\mathrm{min}}$. For a particle $i$ and its neighbourhood $\partial i$, one determines the local affine transformation $F_i$ that best maps the initial neighbour vectors to the final ones and defines

$$D^2_{\mathrm{min},i} = \min_{F_i} \sum_{j \in \partial i} w_{ij} \left| \mathbf{r}'_{ij} - F_i \mathbf{r}_{ij} \right|^2. \tag{65}$$

If $D^2_{\mathrm{min},i} = 0$, the local deformation is exactly affine; a large value signals a local rearrangement not described by any single affine deformation. This distinction will be useful when we describe plastic events below.

Expanding the energy in the non-affine displacement gives

$$\delta U = \frac{V}{2} C^{\mathrm{Born}} \epsilon^2 + \boldsymbol{\Xi}^{\mathsf{T}} \mathbf{u}_{\mathrm{na}}\, \epsilon + \frac{1}{2} \mathbf{u}_{\mathrm{na}}^{\mathsf{T}} H \mathbf{u}_{\mathrm{na}}. \tag{66}$$

Here $\boldsymbol{\Xi}$ is the force field induced by the imposed affine deformation. Mechanical equilibrium after the strain is imposed requires

$$\frac{\partial \delta U}{\partial \mathbf{u}_{\mathrm{na}}} = 0, \tag{67}$$

which gives

$$H \mathbf{u}_{\mathrm{na}} = -\boldsymbol{\Xi} \epsilon. \tag{68}$$

Removing the translational zero modes, the solution is

$$\mathbf{u}_{\mathrm{na}} = -H^{+} \boldsymbol{\Xi} \epsilon. \tag{69}$$

Substitution into the energy gives

$$C = C^{\mathrm{Born}} - \frac{1}{V} \boldsymbol{\Xi}^{\mathsf{T}} H^{+} \boldsymbol{\Xi}. \tag{70}$$

The second term is negative. Non-affine relaxation always lowers the affine elastic energy. Using the eigenmode expansion

$$H^+ = \sum_{\alpha \in \text{stable}} \frac{|\mathbf{e}_\alpha\rangle\langle\mathbf{e}_\alpha|}{\lambda_\alpha}, \tag{71}$$

we obtain

$$\boldsymbol{\Xi}^{\mathsf{T}} H^+ \boldsymbol{\Xi} = \sum_{\alpha \in \text{stable}} \frac{|\mathbf{e}_\alpha \cdot \boldsymbol{\Xi}|^2}{\lambda_\alpha}. \tag{72}$$

Low-frequency modes therefore have a large effect on the elastic moduli.

This equation is the direct link between the Hessian discussion of Lecture 3 and the jamming transition. The approach to isostaticity produces increasingly soft modes, and those same modes produce increasingly large non-affine corrections to the macroscopic response.

# 7 Relaxation around a jammed state

The Hessian also determines the relaxation of a perturbed packing, providing the dynamical link between the soft modes discussed above and glassy relaxation.

For the overdamped dynamics considered in Lecture 2, after the external driving has been removed we have

$$\gamma \dot{\mathbf{R}} = -\nabla_{\mathbf{R}} U. \tag{73}$$

Let $\mathbf{R}_0$ be a mechanically stable configuration and write a nearby configuration as

$$\mathbf{R} = \mathbf{R}_0 + \mathbf{u}. \tag{74}$$

Since $\nabla U(\mathbf{R}_0) = 0$, the expansion of the force begins with the Hessian,

$$\nabla U(\mathbf{R}_0 + \mathbf{u}) = H\mathbf{u} + O(u^2). \tag{75}$$

The linearized overdamped equation is therefore

$$\gamma \dot{\mathbf{u}} = -H\mathbf{u}. \tag{76}$$

Expand the displacement in Hessian eigenmodes,

$$\mathbf{u}(t) = \sum_\alpha u_\alpha(t)\mathbf{e}_\alpha, \qquad H\mathbf{e}_\alpha = \lambda_\alpha \mathbf{e}_\alpha. \tag{77}$$

Substitution gives one independent equation for every normal mode,

$$\gamma \dot{u}_\alpha = -\lambda_\alpha u_\alpha. \tag{78}$$

Hence

$$u_\alpha(t) = u_\alpha(0) \exp\left(-\frac{\lambda_\alpha}{\gamma} t\right), \tag{79}$$

and the relaxation time of mode $\alpha$ is

$$\tau_\alpha = \frac{\gamma}{\lambda_\alpha}. \tag{80}$$

Thus the same eigenvalue gives $\omega_\alpha^2 = \lambda_\alpha/m$ in the inertial problem and $\tau_\alpha = \gamma/\lambda_\alpha$ in overdamped dynamics.

At $T = 0$ there is no noise-driven barrier crossing. A stable basin is left only when the control parameters or an external drive remove its stability.

For a finite-dimensional system with thermal noise, the Langevin equation is instead

$$\gamma \dot{\mathbf{R}} = -\nabla U + \boldsymbol{\eta}(t), \tag{81}$$

with

$$\langle \eta_a(t)\eta_b(t')\rangle = 2\gamma k_{\mathrm{B}} T\, \delta_{ab}\delta(t-t'). \tag{82}$$

The noise term is what allows activated transitions between basins. The limit $T \to 0$ removes this source of fluctuations, leaving only deterministic relaxation. It is this limit, together with dissipation and preparation by compression or driving, that led us to the athermal jamming problem.

# 8 Glassy dynamics: from fluctuations to relaxation

The glass problem enters naturally once we restore the thermal noise and ask how a dense disordered system relaxes. We will only discuss some of the aspects of this problem.

Consider a system initially prepared in a liquid state and then cooled while keeping the density fixed. The equilibrium state is still a fluid at sufficiently high temperature, but the time required to decorrelate grows rapidly upon cooling. For a fragile glass former this growth is much faster than a simple Arrhenius law,

$$\tau_\alpha(T) \not\sim \exp\left(\frac{E_a}{k_{\mathrm{B}}T}\right) \tag{83}$$

in the regime where the super-Arrhenius increase becomes pronounced.

The glass transition is observed primarily as a dynamical arrest on accessible time scales rather than as an established equilibrium thermodynamic singularity.

A useful object is a two-time correlation function. For a density mode, for example, one may define the intermediate scattering function

$$F(\mathbf{q},t) = \frac{1}{N}\left\langle \rho_{\mathbf{q}}(t)\rho_{-\mathbf{q}}(0)\right\rangle, \qquad \rho_{\mathbf{q}}(t) = \sum_j e^{i\mathbf{q}\cdot\mathbf{r}_j(t)}. \tag{84}$$

In equilibrium, time-translation invariance implies that the correlation depends only on the time difference. We can therefore write $F(\mathbf{q},t)$ rather than $F(\mathbf{q};t,t_w)$.

At high temperature the decay is comparatively simple. On cooling a dense liquid, one often observes a two-step relaxation. Particles first rattle inside transient cages formed by their neighbours. This produces a plateau in the correlation function. At longer times cages rearrange and the system undergoes the final structural, or $\alpha$, relaxation.

A convenient phenomenological representation of the long-time decay is a stretched exponential,

$$F(\mathbf{q},t) \sim \exp\left[-\left(\frac{t}{\tau_\alpha}\right)^{\beta_{\mathrm{KWW}}}\right], \qquad 0 < \beta_{\mathrm{KWW}} < 1, \tag{85}$$

although this form is a fit rather than a microscopic theory.

The growth of $\tau_\alpha$ is accompanied by dynamic heterogeneity. It is useful to define the overlap variable explicitly. Choose a microscopic distance $a$ comparable to the cage size and define the single-particle persistence

$$c_i(t;t_w) = \Theta\left(a - |\mathbf{r}_i(t_w+t) - \mathbf{r}_i(t_w)|\right). \tag{86}$$

The corresponding global overlap is

$$q(t;t_w) = \frac{1}{N}\sum_{i=1}^{N} c_i(t;t_w). \tag{87}$$

Thus $q$ is the fraction of particles that have remained within a distance $a$ of their positions at the beginning of the measurement. The precise choice of $a$ is conventional, but it should lie in the regime that distinguishes caged motion from structural relaxation.

A second common observable is the self-intermediate scattering function,

$$F_s(q,t;t_w) = \frac{1}{N}\sum_{i=1}^{N} \left\langle \exp\left[i\mathbf{q}\cdot(\mathbf{r}_i(t_w+t)-\mathbf{r}_i(t_w))\right]\right\rangle . \tag{88}$$

The two quantities measure related but distinct aspects of structural memory. The overlap counts particles that remain close to their initial positions, whereas $F_s$ probes the phase coherence of their displacements at a specified wavevector.

The growth of $\tau_\alpha$ is accompanied by dynamic heterogeneity: different regions relax at different rates. A standard measure is

$$\chi_4(t;t_w) = N\left[\left\langle q(t;t_w)^2\right\rangle - \langle q(t;t_w)\rangle^2\right]. \tag{89}$$

A large peak in $\chi_4$ signals strong fluctuations in the relaxation dynamics.

The equilibrium counterpart is the fluctuation–response relation. For example, the shear viscosity of an equilibrium fluid may be written as the Green–Kubo integral

$$\eta = \frac{V}{k_{\mathrm{B}}T}\int_0^\infty \mathrm{d}t\ \langle \delta\sigma_{xy}(t)\delta\sigma_{xy}(0)\rangle_{\mathrm{eq}}\,, \tag{90}$$

under the usual conditions for linear response. In an aging glass or an athermal driven material, there is no corresponding equilibrium identity that turns an arbitrary stress correlation into a viscosity or susceptibility. The correlation and the response remain separately well-defined observables, but their relation depends on the dynamical regime.

The dynamical length associated with heterogeneous relaxation need not be the same as a static correlation length. A glass lacks crystalline translation symmetry, but collective elastic response can still emerge on long scales.

## 9 Energy landscapes and inherent structures

The same configuration-space picture applies to thermal glasses: the potential-energy landscape contains many minima separated by barriers, and cooling increases the time spent within restricted regions of configuration space.

Given a thermal configuration $\mathbf{R}$, imagine minimizing the potential energy from that point by the same gradient-descent dynamics used in the athermal problem. The minimum reached is called the corresponding inherent structure. The instantaneous configuration contains vibrational motion inside a basin, while the sequence of inherent structures records the slower structural rearrangements between basins.

This gives a useful interpretation of the basin construction introduced in the previous lecture. The basin $\mathcal{B}_\alpha$ consists of the initial configurations that relax to the same minimum under deterministic minimization. For a thermal system the instantaneous probability of being in a basin is not, in

general, proportional simply to its geometric volume: the Boltzmann factor weights different energies within the basin. If the basin partition function is

$$Z_\alpha^{\mathrm{basin}}(T) = \int_{\mathcal{B}_\alpha} \mathrm{d}\mathbf{R}\, \mathrm{e}^{-\beta U(\mathbf{R})}, \tag{91}$$

then the equilibrium probability of the basin is

$$P_\alpha^{\mathrm{eq}} = \frac{Z_\alpha^{\mathrm{basin}}(T)}{\sum_\beta Z_\beta^{\mathrm{basin}}(T)}. \tag{92}$$

The athermal basin-volume construction is different. There the motion within a basin dissipates deterministically and the final state is selected by the initial distribution and the protocol. There is no thermal integration over fluctuations inside the basin. This is one of the cleanest places to see what is lost when the thermal reservoir is removed.

In finite-dimensional liquids, $T_{\mathrm{MCT}}$ is best regarded as a dynamical crossover: activated processes restore relaxation below the mean-field scale. Preparation and observation times then matter increasingly in the deeply supercooled regime.

### 9.1 Aging and the loss of time-translation invariance

The equilibrium correlation functions used in Lectures 1 and 2 rely on stationarity: correlations depend only on the time difference. After a quench into a glassy regime, the system can continue to evolve slowly, so this property is lost. Let $t_w$ denote the waiting time after the quench. A two-time correlation is then

$$C(t, t_w) = \langle A(t_w + t) A(t_w) \rangle \,. \tag{93}$$

In equilibrium, $C(t, t_w)$ reduces to a function of $t$ alone. In an aging system,

$$C(t, t_w) \neq C(t). \tag{94}$$

The relaxation therefore depends on the age of the system: there is no single time-independent relaxation time describing the subsequent dynamics.

This also gives a direct connection to the fluctuation–dissipation discussion of Lecture 2. In equilibrium, detailed balance relates spontaneous fluctuations to the response to a small perturbation. During aging this equilibrium relation is violated. A fluctuation–dissipation ratio may be defined through

$$R(t, t_w) = -\frac{X(t, t_w)}{k_{\mathrm{B}} T} \frac{\partial C(t, t_w)}{\partial t}, \tag{95}$$

where $X = 1$ in equilibrium and, in general, $X \neq 1$ in an aging regime. Thus the violation of FDT is another consequence of the loss of the equilibrium stationary measure and detailed balance.

For an athermal system at $T = 0$, the thermal definition of $X$ is not applicable. The underlying mechanism is nevertheless the same: the system evolves through a sequence of metastable configurations, and the relaxation time can increase as the system explores progressively more stable regions of configuration space. The statistics of this evolution are determined by the preparation and dynamics rather than by a thermal stationary measure.

## 9.2 Glasses, jamming, and driven athermal dynamics

The glass and jamming transitions describe two related but distinct forms of arrest. In a conventional glass, relaxation becomes increasingly slow as temperature is lowered or density is increased while thermal fluctuations remain present. In an athermal soft-particle system, jamming is most clearly isolated at $T = 0$ and in the quasistatic limit, where increasing the packing fraction produces a mechanically stable contact network at $\phi_J$.

The distinction is therefore dynamical. In a thermal system,

$$\gamma \dot{\mathbf{R}} = -\nabla U + \boldsymbol{\eta}, \tag{96}$$

with

$$\langle \eta_a(t)\eta_b(t')\rangle = 2\gamma k_{\mathrm{B}} T\, \delta_{ab}\delta(t-t'). \tag{97}$$

Thermal noise allows the system to escape from one metastable basin and explore others. At strictly zero temperature,

$$\boldsymbol{\eta} = 0, \tag{98}$$

and the dynamics becomes

$$\gamma \dot{\mathbf{R}} = -\nabla U. \tag{99}$$

Once a local minimum has been reached, the system remains there unless an external perturbation changes its stability.

This is the useful connection between glasses and jamming. A thermal glass is a system with slow motion through a complex energy landscape; an athermal jammed solid is the limit in which thermal barrier crossing has been removed and mechanical stability of individual minima becomes the central issue. The two problems therefore share disorder, metastability, and collective rearrangements, but the mechanism for moving between metastable states is different.

Under an external drive, the overdamped equation becomes

$$\gamma \dot{\mathbf{R}} = -\nabla U + \mathbf{f}^{\mathrm{ext}} + \boldsymbol{\eta}. \tag{100}$$

The corresponding energy balance is

$$\frac{\mathrm{d}U}{\mathrm{d}t} = -\gamma \sum_i |\dot{\mathbf{r}}_i|^2 + \sum_i \mathbf{f}_i^{\mathrm{ext}} \cdot \dot{\mathbf{r}}_i + \sum_i \boldsymbol{\eta}_i \cdot \dot{\mathbf{r}}_i. \tag{101}$$

In a driven steady state, work is continuously supplied by the external drive and removed by dissipation. Detailed balance and equilibrium FDT therefore do not apply.

For athermal quasistatic driving, the noise is absent and the control parameter is changed in small increments,

$$\lambda_n \longrightarrow \lambda_n + \delta\lambda \longrightarrow \mathbf{R}_{n+1}, \tag{102}$$

where $\mathbf{R}_{n+1}$ is obtained after complete mechanical relaxation. The system therefore follows a sequence of local minima until the current minimum loses stability. These dynamics (rheology) are related to the soft modes and Hessian derived earlier: the same eigenvalues that determine the local mechanical stability also determine the relaxation times between successive mechanically stable configurations.

## 10 Mesoscopic rheology and elastoplasticity

Rheology begins when the system repeatedly leaves one mechanically stable basin through plastic rearrangements.

For simple shear, let $\gamma$ be the imposed macroscopic strain and $\sigma$ the corresponding shear stress. In the linear elastic regime,

$$\sigma = G\gamma + O(\gamma^2), \tag{103}$$

where $G$ is the shear modulus of the current configuration. If plastic deformation is present, it is useful to decompose the strain schematically as

$$\gamma = \gamma_{\mathrm{el}} + \gamma_{\mathrm{pl}}, \tag{104}$$

so that the local elastic stress is related to the elastic part of the strain. In the simplest scalar description,

$$\sigma = G\left(\gamma - \gamma_{\mathrm{pl}}\right). \tag{105}$$

For an amorphous solid this relation must be supplemented by the fact that plastic deformation at one location produces an elastic stress change elsewhere.

To make this idea explicit, divide the material into mesoscopic blocks. A block is large enough to contain many microscopic particles, but small compared with the macroscopic sample. Let $\sigma_i$ be its local shear stress and $\sigma_{y,i}$ its local yield threshold. Define the local stability variable

$$x_i = \sigma_{y,i} - |\sigma_i|. \tag{106}$$

The block is stable for $x_i > 0$ and reaches its local yield threshold when $x_i = 0$. The threshold is not a universal microscopic constant: it summarizes the local structure of the glass, and its distribution depends on the preparation history.

Once a block yields, its local plastic strain changes. A simple example of a local viscoplastic rule is

$$\dot{\gamma}_i^{\mathrm{pl}} = \frac{1}{\tau_{\mathrm{pl}}} \operatorname{sgn}(\sigma_i)\left(|\sigma_i| - \sigma_{y,i}\right)\Theta\left(|\sigma_i| - \sigma_{y,i}\right). \tag{107}$$

The local flow rule is model dependent; elastoplastic models share a local stability criterion and a post-yield relaxation rule.

The essential additional ingredient is the long-ranged elastic coupling between blocks. A localized plastic rearrangement produces a change in the surrounding stress field of the form

$$\delta\sigma(\mathbf{r}) \sim G\,\mathcal{G}_E(\mathbf{r} - \mathbf{r}_0)\,\Delta\gamma_{\mathrm{pl}}, \tag{108}$$

where $\mathbf{r}_0$ is the position of the rearrangement. In two dimensions the far-field part of the elastic kernel has the quadrupolar form

$$\mathcal{G}_E(r,\theta) \sim \frac{\cos 4\theta}{r^2}. \tag{109}$$

The precise sign and normalization depend on the convention for shear stress and plastic strain. The angular structure and algebraic decay are fixed by the elastic Green function. This is the same mechanical origin as the long-ranged stress correlations derived in Lecture 4.

An avalanche begins when a block reaches $x_i = 0$. Its stress redistribution can trigger other unstable blocks. At particle scale this is a sequence of minima and rearrangements; at mesoscopic scale it is a cascade of coupled local yield events.

In the quasistatic limit, $\delta\gamma \to 0$ and the system fully relaxes after each increment, giving elastic branches separated by plastic events. At finite shear rate, loading and relaxation occur simultaneously.

For comparison, in an equilibrium liquid the zero-frequency viscosity is constrained by a Green–Kubo relation involving the equilibrium stress autocorrelation. In an athermal system, the steady stress instead results from the balance between loading and dissipation and from the statistics of plastic rearrangements. A flow curve such as

$$\sigma_{\mathrm{ss}}(\dot{\gamma}) = \sigma_y + A\dot{\gamma}^n \tag{110}$$

is often used as a phenomenological parametrization, but the exponent $n$ and even the applicability of this form are model- and material-dependent. It should therefore be regarded as a constitutive form, not a universal consequence of the microscopic equations.