Landscape analysis of constraint satisfaction problems

Florent Krzakala[1] and Jorge Kurchan[2]
[1]*PCT, CNRS UMR Gulliver 7083, ESPCI, 10 rue Vauquelin, 75005 Paris, France*
[2]*PMMH, CNRS UMR 7636, ESPCI, 10 rue Vauquelin, 75005 Paris, France*
(Received 3 April 2007; published 27 August 2007)

We discuss an analysis of constraint satisfaction problems, such as sphere packing, K-SAT, and graph coloring, in terms of an effective energy landscape. Several intriguing geometrical properties of the solution space become in this light familiar in terms of the well-studied ones of rugged (glassy) energy landscapes. A benchmark algorithm naturally suggested by this construction finds solutions in polynomial time up to a point beyond the clustering and in some cases even the thermodynamic transitions. This point has a simple geometric meaning and can be in principle determined with standard statistical mechanical methods, thus pushing the analytic bound up to which problems are guaranteed to be easy. We illustrate this for the graph 3- and 4-coloring problem. For packing problems the present discussion allows to better characterize the *J*-point, proposed as a systematic definition of random close packing, and to place it in the context of other theories of glasses.

DOI: 10.1103/PhysRevE.76.021122 PACS number(s): 02.50.−r, 75.10.Nr, 64.70.Pf, 81.05.Rm

I. INTRODUCTION

Constraint optimization and satisfaction problems are a particular, yet widespread, class. The prototype is the packing problem, in which we are given a fixed volume V and asked to pack N objects of typical size r_o (a scale factor) without overlap, making N or r_o as large as possible. When r_o is increased (with N,V fixed) the problem becomes harder as the number of configurations that solve it decreases, until a point is reached r_o^{pack} beyond which there are no more solutions.

Another celebrated example is the satisfiability, or SAT, problem [1]: we have a set of N Boolean variables $\{x_i=0,1\}_{i=1,\ldots,N}$ and a number $M=\alpha N$ of clauses—when all clauses have exactly K literals, the problem is referred to as K-SAT—which in the present paper we shall assume are the first M of a longer list that is generated at random and stipulated once and for all. We are asked to find logical assignments satisfying

$$F=\bigwedge_{\ell=1}^{M} C_\ell=\bigwedge_{\ell=1}^{M}(z_{i_1}^{(\ell)}\vee z_{i_2}^{(\ell)}\vee\cdots\vee z_{i_k}^{(\ell)}), \qquad (1)$$

where $\wedge$ and $\vee$ stand for the logical "and" and "or" operations, respectively, ℓ labels a set $i_i,\ldots,i_K$ and $z_i^{(\ell)}$ is, depending on ℓ either x_i or its negation (this, and the x_i that participate in each factor is what defines the clause). An assignment of the $\{x_i\}$'s satisfying all clauses is a solution of the K-SAT problem. When the number of clauses $\alpha N=M$ increased, the number of sets $\{x_i\}$'s satisfying all clauses decreases, until a point α^{unsat} is reached beyond which there are no solutions.

The third problem we shall consider is the graph q-coloring problem [1]. We are given a graph with N vertices and are asked to color them with one of q colors so that no two vertices sharing a link have the same color. We shall assume as before that we have a list of links, and we count the number of possible colorings when the graph has the first $M=2\alpha N$ of the list (where the average connectivity of the graph is $c=2\alpha$). As M is made larger, a point α_{uncol} is reached in which no more colorings are possible.

Coloring problems are a particular class of packing problems. Consider the following angle-packing construction: in every node i of a graph we assign an angular variable θ_i. There is a repulsive potential $V(\theta_i-\theta_j)$ between the angles of linked sites as in Fig. 1, it is nonzero where the angle difference $(\theta_i-\theta_j)$ lies between $-2\pi/q$ and $+2\pi/q$ (modulo 2π). For every angle configuration with zero total energy, one can obtain a solution of the q-coloring problem just by assigning to the site i the color numbered by the integer $R=\text{Int}\left(\frac{q\theta_i}{2\pi}\right)$ $(R=0,1,\ldots,q-1)$. Conversely, each coloring solution has at least one zero energy configuration counterpart.

Statistical physics methods are particularly efficient in treating large random problems, especially when mean-field techniques are exact, and indeed several important results were obtained for both satisfiability [2,3] and coloring [4] in this case. The connection with statistical mechanics is usually introduced as follows (see, for example, [3,5,6]): one first defines an energy E such that it is zero when the problem is satisfied and larger otherwise. In the examples above, this could be the number of overlapping spheres, the number of unsatisfied clauses, and the number of links joining vertices with the same color, respectively. Next, one studies for finite temperature $T=1/\beta$ the partition sum over configurations

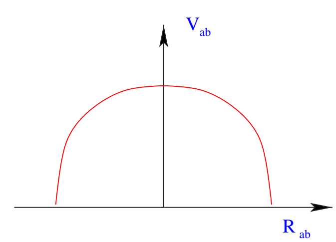

FIG. 1. (Color online) Potential of soft particles with finite-range interactions. It allows to construct a packing version of the coloring problem (see text).

 ©2007 The American Physical Society

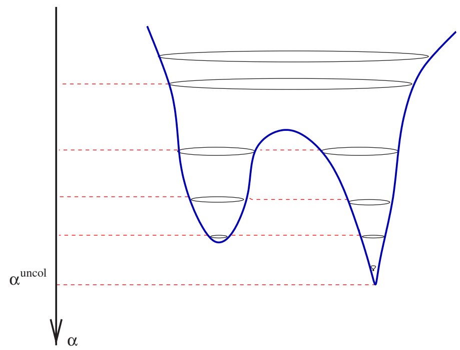

FIG. 2. (Color online) Constructing a pseudoenergy from a constraint satisfaction problem: the regions shown are the satisfied sets at each level of difficulty. The height function is such that at each level the intersection corresponds to a constraint satisfaction problem.

$$Z = \sum_{\mathrm{conf}} e^{-\beta E(\mathrm{conf})}. \tag{2}$$

The results are encoded in the energy and entropy in the zero-temperature limit. One can apply all the analytical techniques of statistical mechanics, and has in addition available all the standard practical methods of annealing energies (in metallurgy as well as in computational physics) that serve at least as a first, general purpose numerical tool.

In this paper we shall follow a different strategy. We first introduce in Sec. II a pseudoenergy defined directly in terms of the constraint satisfaction problem. For the packing problem, consider the set of configurations that satisfies the constraints at a given value of r_o. If we make the problem harder by increasing the common scale of all particles r_o, the set of configurations still satisfying the constraints is strictly a subset of the previous one. Thus, we can construct a single-valued landscape function $\mathcal{E}$ as in Fig. 2. For the SAT and coloring problems we proceed similarly, with now α playing the role of r_o, increased by adding clauses (respectively, links) one by one following the predefined list. *The variable conjugate to $\mathcal{E}$ is in fact a pressure*.

Next, we explore in Sec. III the consequences of making the following landscape hypothesis: *The pseudoenergy $\mathcal{E}$ landscape defined above has in the thermodynamic limit the same qualitative properties as the usual rugged energy landscapes*. This means that pressure plays a similar role as that of an inverse temperature. We shall find that several intriguing results [7] in the geometry of random constraint satisfaction problems (CSP) such as the satisfiability of random formulas or the coloring of random graphs can be readily understood in this language.

In Sec. IV, we describe a family of algorithms that immediately suggests itself: it is simply a quench in the pseudoenergy landscape, and hence by construction polynomial in the system size. At each step one increases the difficulty of the problem, and moves to a configuration in a finite neighborhood so as to restore satisfaction, if there is no such configuration the program stops. We shall show that its performance—where it converges and in what times—can be in principle determined analytically for the very same problems for which analytical statistical mechanical solutions are available. The method can be generalized to slower annealings, or to finite pressure.

Next, we explore in detail the performance of this algorithm for the random q-coloring problem. Although here we are more interested in constructing a benchmark than a competitive numerical strategy, we find that the performance is surprisingly quite good.

In Sec. V we concentrate on hard spheres and recognize that this algorithm yields the procedure of O'Hern *et al.* [8,9] (and more generally the Lubachevsky-Stillinger [10] procedure if a slower annealing is made) to define the so-called "*J*-point," which they propose to identify as random close packing. Comparing with the random satisfaction problems of the preceding sections, we are thus able to give a mean-field realization of the *J*-point scenario [67], and to put it in the context of other special points in glass theory, such as the (putative) ideal thermodynamic glass state, and the (finite-dimensional relic of) the mode-coupling transition.

II. CONSTRAINT SATISFACTION PROBLEMS AND A LANDSCAPE HYPOTHESIS

Let us describe the pseudoenergy in more detail. Consider a fixed set of particles of any shape whose positions and angles are given by the coordinates $\mathbf{x}_1, \ldots, \mathbf{x}_N$, with a common scale r_o, such that, for example, multiplying r_o by 2 doubles the linear size of all particles while preserving their shape. According to Fig. 1, the landscape at a point $E(\mathbf{x}_1, \ldots, \mathbf{x}_N)$ is defined as (minus) the first value of r_o at which there is a particle overlap for these coordinates. (The minus sign is conventional.)

In particular, for spheres of equal size r_o, the pseudoenergy is simply (minus) the smallest interparticle distance in that configuration divided by 2. As such, it is very close to the succession of potentials discussed by Stillinger and Weber [11],

$$E_{\mathrm{SW}} = -\sum_{a \neq b} \frac{1}{|\mathbf{x}_a - \mathbf{x}_b|^n} \quad \text{for } n \to \infty. \tag{3}$$

Here we are considering

$$\mathcal{E} = -N \lim_{n\to\infty} \left\{ \sum_{a \neq b} \frac{1}{|\mathbf{x}_a - \mathbf{x}_b|^n} \right\}^{-1/n}. \tag{4}$$

(The factor N is needed to assure that the quantity is extensive if we make the limit $n \to \infty$ *before* the thermodynamic limit.) *Note that a gradient descent in (3) follows the same path as one in (4).*

For the SAT and the coloring problem the space is discrete, and the difficulty of the problem may be increased by adding clauses and links, respectively. The pseudoenergies are then, by analogy, minus the number of clauses, and minus the numbers of links, respectively. As mentioned above, we are considering in fact a succession of problems, obtained by choosing the first M of a list of clauses or links which is fixed once and for all. Note that if we could reshuffle the list, the net result at finite M would depend on the length of the list itself: for an infinite list we would optimize the choice of

M constraints to make them easily satisfiable—this is called the "annealed" problem in the disordered system literature (not to be confused with the "annealed approximation" of the probabilistic literature, which is something else altogether [12]).

The landscape for the coloring problem is very rugged indeed. We have a graph with M links that is well colored. Let us consider how the pseudoenergy changes when we flip the color of one given vertex. The new pseudoenergy is (minus) the largest value of M' that will make the new graph colorable, so we must take away one by one the links *in reverse order, starting from the last introduced* from M to M', but most of the links deleted will not be even close to the configuration we have flipped, and it will typically take $M-M' \sim O(N)$. We shall see that this poses no problem if one proposes changes only in nearby sites along the tree: thus, jumps will be always to configurations that are close and have lower pseudoenergy, avoiding other ones in the neighborhood that are much higher.

From constraint satisfaction to energy landscapes and back. As mentioned above, in this paper we shall make the landscape hypothesis that the pseudoenergy defined above has in the thermodynamic limit the same qualitative features as expected from a generic rugged energy landscape. We do not know if this is strictly so, but it is clear that the converse is true: *every rugged energy landscape yields, in the microcanonical ensemble, a constraint satisfaction problem* [13]. This is easy to see just reading Fig. 2 in the reverse way: we imagine that the landscape $H(\mathbf{x})$ is given, and by cutting (microcanonical) slices of fixed energy at different depths $\mathcal{E}$, we obtain the CSP of finding $\mathbf{x}$ such that $H(\mathbf{x}) < \mathcal{E}$.

What we shall do in the following section is to briefly review some known facts from complex energy landscapes, and then proceed by cutting microcanonical slices to infer the behavior it implies for random constraint satisfaction problems.

III. ENERGY LANDSCAPE: MICROCANONICAL SLICES

Let us review what we have learned, mostly from mean-field analytic calculations, about rugged energy landscapes. In order to match the constraint satisfaction situation, we shall adopt the rather unusual point of view of making energy slices, rather than constant temperature ones. The states we encounter shall then be regions of the energy shell that are dynamically isolated. Figures 4, 5, and 6 show variations of these constant energy slices: the temperature T is then a variable that classifies the states at that energy: in the landscape of $\mathcal{E}$ it amounts to classifying states according to their pressure (the inverse of the temperature conjugate to $\mathcal{E}$). This is not the usual practice in optimization problems, which is to classify the states according to their internal entropy σ (the logarithm of the number of configurations, $\frac{1}{T}=\frac{\partial\sigma}{\partial E}$).

A. Energy vs free energy barriers

In what follows, we shall frequently allude to states, defined as regions of phase space that are dynamically isolated, in the sense that the system starting a random walk within them will take a long time to exit. "Long" might mean "infinite," or simply diverging with the system size N faster than any power law. Consider, for example, a ferromagnet with continuous spins on a lattice of dimension $d>1$ and linear size L, with nearest-neighbor interaction

$$E = - \sum_{\text{next neighbor}} s_i s_j - A \sum_i (s_i^2 - 1)^2. \tag{5}$$

In terms of the temperature, the phase diagram appears as in Fig. 3: there is a critical temperature below which there are two states of magnetization $M=\Sigma s_i = \pm m L^d$, and above one state of zero magnetization. As is well known, a system at $T_i < T_c$ performing Monte Carlo dynamics starting with magnetization density m will take a time to jump to the state of magnetization $-m$ that diverges in the thermodynamic limit. We can also make a random walk *at constant energy* E_i, and the time to pass from the $+m$ to the $-m$ state is still divergent. Does this mean that the constant energy surface $E(\mathbf{s})=E_i$ is disconnected in two regions? A moment's thought shows that this cannot be the case. The actual saddle point in energy separating the two energy minima is just the energy to create an interface between one-half lattice at magnetization $+m$ and one-half at $-m$, $E_{\text{saddle}} - E_o \sim L^{d-1}$, while $E_i - E_o = O(L^d) \gg E_{\text{saddle}} - E_o$ (see Fig. 3). We conclude that the energy surface is actually guitar shaped, and the reason why the system takes a divergent time to pass from one state to the other is that (for large size L) the neck is exponentially thin.

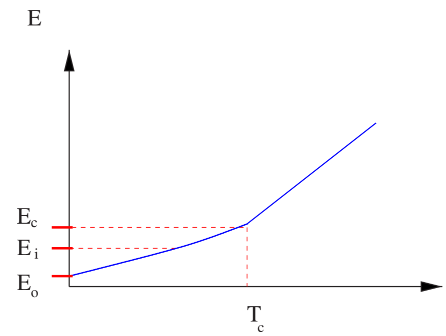

FIG. 3. (Color online) Energy vs temperature of a ferromagnet.

B. Flat bottom states

Consider a system of particles with a finite-range potential as in Fig. 1, confined to a fixed volume. To the extent that the packing without overlaps is possible, it is clear that the total energy landscape will have minima of zero energy which are *flat*. Surprisingly (and confusingly) also the pseudoenergy landscape can have flat bottom states, both in the particle and in the K-SAT and coloring problems, as we shall see in the next section. They turn out to be very important.

C. Mean-field glass landscapes

We shall now give three examples of mean-field landscapes with many states. Mean-field problems in physics correspond to large random problems in optimization and satisfaction: the coloring of large random graphs or the

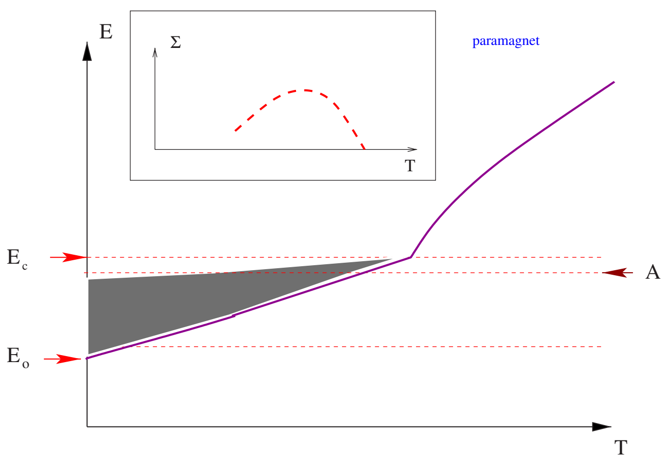

FIG. 4. (Color online) A sketch of the metastable-state space of the Sherrington-Kirkpatrick model: the gray area in the graph corresponds to a region with an exponential number of metastable states. Inset: complexity along a constant-energy surface line A in terms of temperature T (the same graph would be obtained by plotting in terms of state entropy σ).

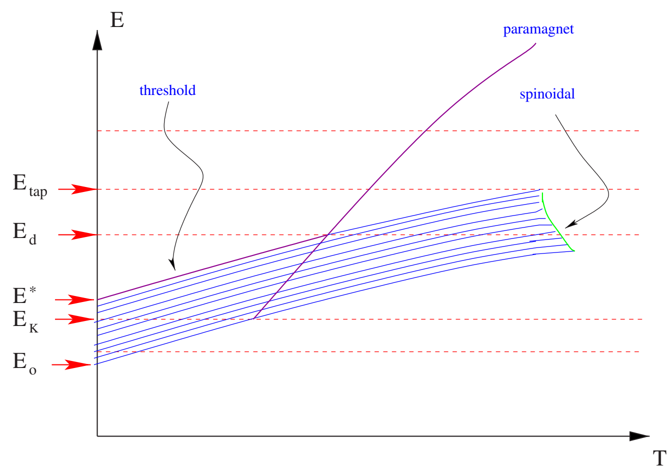

FIG. 5. (Color online) The structure of metastable states of the spherical p-spin model.

satisfiability of random formulas are instances of this class. Indeed these are the models for which we expect a correspondence between landscapes to hold.

The first example (the Sherrington-Kirkpatrick model) has the property that its states having free-energy density higher than the ground state one are no obstacle for the dynamics. This class also includes all models with smooth short range interactions in finite dimensions, like the Edwards-Anderson model.

The second example, the p-spin model is the precise opposite: a naive dynamics is trapped in a threshold level where the first, highest minima are encountered—this happens at energy density *above* the lowest.

The third example, the Ising p-spin model, is the most typical of random optimization problems: it has a combination of transparent high states that are avoided by the dynamics, but it also has a threshold level at finite energy density above the lowest, below which the dynamics only goes in times that diverge with the size.

1. The Sherrington-Kirkpatrick class

The Sherrington-Kirkpatrick model [6] (SK) is the most studied spin glass, it consists of ±1 spins coupled by random, fully connected interactions

$$E = \sum_{ij} J_{ij} s_i s_j. \tag{6}$$

It was soon realized [14,15] that the free-energy landscape had many metastable states (see [6]), a situation depicted in Fig. 4. The dynamics in the thermodynamic limit has been also solved [16], and surprisingly one finds that at each constant temperature the energy *density* tends for large times to the equilibrium one. In other words: *all but the very lowest metastable states are transparent to the dynamics*, if one wishes to compute the energy *per spin* with any percentage accuracy, this can be done in polynomial time (although the true ground state might be harder to find). It took over a decade to clarify completely this from a pure landscape point of view [17].

Because the states are not strong dynamic traps, a perturbation with small random forces that do not derive from a potential—a weak stirring—immediately sets the system into motion: the whole state structure is washed away.

2. Spherical and Ising p-spin glass

The spherical p-spin model $(p>2)$ has a Hamiltonian

$$E = \sum_{i_1,\ldots,i_p} J_{i_1,\ldots,i_p} s_{i_1} \cdots s_{i_p}, \tag{7}$$

with the s_i satisfying a spherical constraint. It has a well-studied [18,19] structure of states shown in Fig. 5. Each line in the figure symbolizes a state (a free-energy minimum). These do not merge and keep their order. There is a *threshold* level above which all the stationary points are unstable. States under the threshold are stable, the more so the deeper, and those that are just below are *marginal*, the spin-glass susceptibility diverges within them [20]. A system sitting on the threshold states will be set into motion by small nonconservative stirring forces (another manifestation of their marginality), but this is not the case with deeper states which are rigid.

The zero-temperature intercept of the threshold energy E^* will be of particular importance for us. Just above this point are the overwhelming majority of energy barriers [20], and this means, according to the discussion above, that for all energy slices higher than E^* the states are in fact connected by narrow bridges.

The out of equilibrium dynamics of this model is well studied [21], a quench starting from a random configuration never goes below the threshold level. At zero temperature, a quench ends in the energy E^*. In fact, no (instance independent) external field protocol is known that leads below the threshold: temperature cycling, magnetic field changes, even quantum annealing [22] have been tried, but the system always ends up on the threshold level at the end of the cycle.

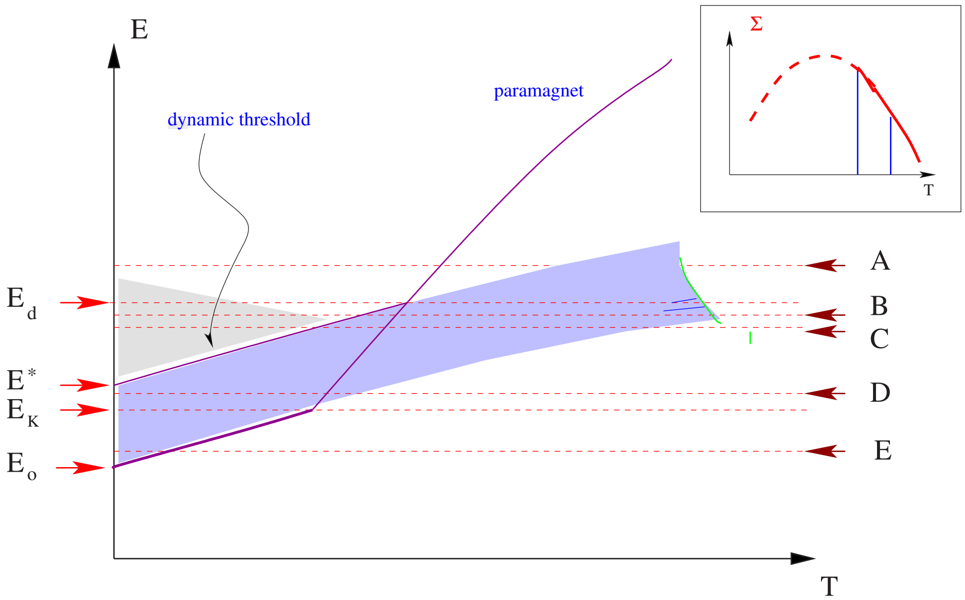

FIG. 6. (Color online) A sketch of the organization of states in the Ising p-spin model. Inset: complexity in slice C; the two vertical lines indicate the dynamic threshold (left-hand side) and the typical configurations (right-hand side), the dashed line symbolizes states that are transparent to the dynamics. The situation in slices $A-E$ have been all found in constraint satisfaction problems for different levels of constraints (see text).

The thick line of Fig. 5 labeled paramagnet corresponds to the equilibrium high-temperature phase. Above E_d this is a big state that dominates the dynamics and contains the overwhelming majority of configurations. Between E_K and E_d the paramagnet is fractured in many states sharing each a negligible volume, the probability of two configurations chosen at random of being in the same state is zero. Below E_K the lowest states dominate, there are still many states having a non-negligible volume, but the sharing of volume is so inegalitarian that two configurations chosen at random have a finite probability of being in the same state.

Microcanonical slices at different energy levels give us a first glimpse of what happens in a constraint satisfaction problem. The number of states, as calculated in [19], grows with decreasing temperature (at constant energy), up to $T=0$—or the intersection with the threshold line, if this happens first. A system prepared at any energy $E^* < E < E_d$ performing energy-conserving dynamics [23] will end in the threshold level. If, on the contrary, $E > E_d$, the dynamics will evolve in the paramagnetic state.

Note the presence at $E_d < E < E_{TAP}$, in a range in which the paramagnet dominates, of states that are completely irrelevant from the dynamic (or static) point of view: these were found in [7,24,25] for the coloring and SAT problems in a given range of parameters.

The spherical p-spin model is somewhat special, and in fact the generic situation is richer. For example, the Ising version of the p-spin model has a structure sketched in Fig. 6. In addition to the states that resemble those of the spherical case, there is a set of high-lying states that resemble those of the SK model [26] and are transparent to the dynamics. A small stirring perturbation wipes out the high-lying states, while the deep nonmarginal ones are stable. There is still a dynamically defined threshold level above which states are marginal [27,28], but it is still not clear if there is a purely local characterization of it (but at any rate an exact dynamic calculation is always possible [16]). We shall in the last section also propose a static—though nonlocal—calculation to achieve the same goal.

Making microcanonical sections as in Fig. 6, we see some of the possibilities for the organization of states and their respective dynamical and equilibrium properties. The inset of Fig. 6 shows a sketch of section along the line C: the states with smallest internal entropy correspond to the gray region of dynamically transparent SK-like states. This feature has been found in K-SAT [29] and in the coloring problem [30]. The typical configurations (vertical line to the right-hand side in the inset of Fig. 6) are those that maximize $\sigma+\Sigma$ and do not coincide in general with the threshold states [31].

In any case, the level achieved with a zero-temperature annealing E^* may (and as we shall see will) be lower than E_d and even E_K. For example, in the p-spin model these points can be calculated analytically [18,21,32] and the results are plotted in Fig. 7.

The important conclusion of this figure is that *if the task is to find a configuration with energy as small as possible, generically neither E_d nor E_K pose a serious limit in themselves,* because E^* may be smaller. *The same is true, for the same reason, for constraint satisfaction problems*. This becomes more so for the smaller the values of p—and eventu-

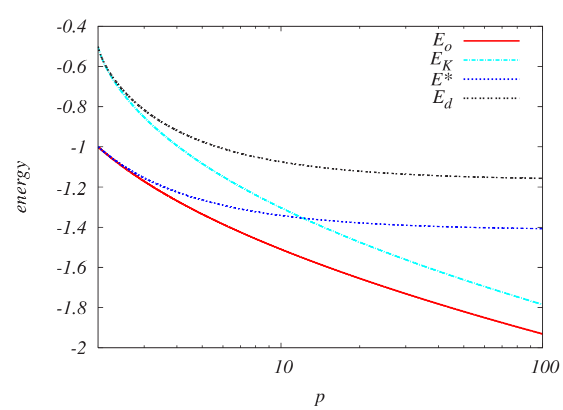

FIG. 7. (Color online) The energies E_d, E_K, E^*, and E_o for the p-spin model as a function of p (from [32]). A simple gradient descent program goes to E^* which is beyond E_d, E_K for $p<13$.

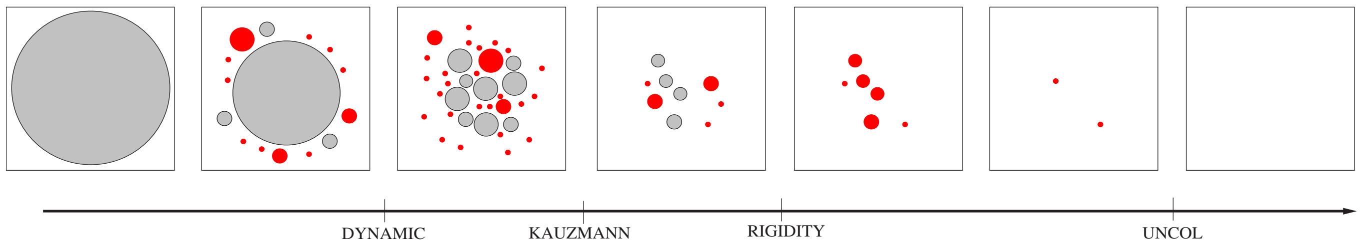

FIG. 8. (Color online) Sketch of the set of solutions in random CSP when the connectivity is increased (adapted from [7,24]). From left-hand side to right-hand sides, in order of increasing difficulty: (1) easy problem, the whole solution space is connected; (2) a negligible fraction of solutions become disconnected; (3) dynamic transition α_d, the whole space breaks into dynamically isolated regions, two configurations at random belong to different regions; (4) replica symmetry breaking, Kauzmann point α_K, the regions become so rare that the probability of two configurations being in the same region is now nonzero; (5) equilibrium rigidity point α^{hard}. Typical configurations have frozen variables; (6) best packing, unSat, unCol ($\alpha_{\text{pack}}, \alpha_{\text{unsol}}, \alpha_{\text{uncol}}$), the last configuration satisfying the constraints disappears. Red (dark) and grey regions symbolizes clusters with and without frozen variables, respectively.

ally K and q in SAT and coloring. We shall see this in the coloring problem in the next section.

3. Random constraint satisfaction problems

The random satisfiability and the random graph coloring problem have been studied using the cavity method [2,4] and their landscapes when varying the connectivity was characterized in [7,24,25]. It is interesting to observe that they follow our landscape hypothesis: they behave indeed like microcanonical slices of the typical energy landscapes we have been discussing. This is best seen in Fig. 8.

For the coloring problem, the col-uncol transition appears at $\alpha_{\text{uncol}}(3)=2.355$ and $\alpha_{\text{uncol}}(4)=4.45$ [4]. The other critical connectivities in the properties of the landscape were found and computed recently [24]. Using three colors, one has $\alpha_d(3)=2$ (a result already obtained in [33]), $\alpha_K(3)=2$ and $\alpha^{\text{hard}}(3)=2.33$. For the 4-coloring problem, the critical values are $\alpha_d(4)=4.175$, $\alpha_K(4)=4.23$, and $\alpha^{\text{hard}}(4)=4.42$.

IV. ALGORITHMS: ARKLESS STRATEGY FOR FLOOD VICTIMS

Consider a rugged landscape that is being slowly flooded. We adopt the following strategy: we stay still until the shore reaches our feet. We then move in small steps away from the water, just enough each time to stay dry. At a certain point, the patch where we are standing becomes an island, all land bridges have been flooded: this is the dynamic transition level. We keep on moving uphill, just avoiding the waterline. Our island may further divide into smaller ones, but our choice of island is dictated by the evolution of the shoreline. Our fate is sealed when our island disappears altogether: this may happen all of a sudden if its top is flat, or gradually if the top is rounded. Note that nothing guarantees that we have ended in the highest summit, so our survival level is just locally—but not globally—optimal. Clearly, the dry land is the set of configurations satisfying the constraints, and the level of the water is the number of clauses, links or particle size, which we increase gradually, while at the same time moving in small steps to remain satisfied. The point reached with this algorithm is image (5) of Fig. 8.

This is the fast procedure, a slightly more sophisticated one is at each level of the water, not just to stay at the shoreline, but to explore randomly all the available land, always without crossing water. This means that on occasion we shall take advantage of a land bridge to go (randomly) from a promontory to another.

A composite strategy is to explore all the available land at one level of water, and then for all subsequent levels just follow the direct lazy strategy of moving just what is necessary.

Now, from the point of view of pseudoenergy landscape (which is reversed with respect to the flood analogy, with summits becoming valleys), in continuous cases it is clear that the fast algorithm is *just a gradient descent*, and in general a zero temperature quench, while the procedure in which we take time to diffuse at constant energy, is a slow annealing [34], in this case we must specify how slow. The last, composite strategy, is again a rapid quench, but this time starting from an equilibrium configuration at some level.

A. Hard sphere procedures

For the packing of hard spheres, the two procedures mentioned above have been used by Lubachevsky and Stillinger [10] (the slow annealing), and by O'Hern *et al.* [8] (the rapid quenched). In both cases one inflates all spheres simultaneously by a very small amount, and then uses some repulsion to eliminate any overlap (wet toes, in the analogy above) that might have been generated [9]. In the annealed version, one also takes time to diffuse without changing the radii or allowing overlaps.

Another possibility is to adapt the composite procedure used by Sastry *et al.* [35] to analyze (true) energy landscapes: one starts from a fully equilibrated configuration at a given (low) level of packing, and from there onwards performs a fast quench, without diffusing at each step more than is necessary to eliminate overlaps.

The reason why O'Hern *et al.* used the rapid quench rather than an annealing is that their purpose was to define a point in parameter space (the J-point) in an unambiguous way. Had they used an annealing procedure, a different point would have been obtained for each annealing time—not to speak about problems of crystallization in a monodisperse

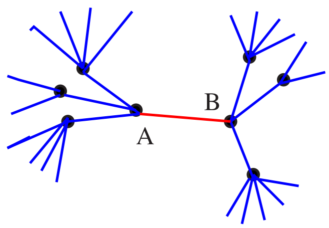

FIG. 9. (Color online) Adding a link to a tree.

sphere case (and, indeed, the infinitely slow annealing limit would lead to a packing fraction essentially corresponding to equilibrium).

As mentioned above, one of the aims of this paper is to put this J-point in the context of the rest of glass theory. We shall discuss this in further detail in Sec. V. Let us note here that we can apply exactly the procedure of O'Hern *et al.* (or Lubachevsky-Stillinger) for the angle-packing problem described in the introduction for a given, fixed graph, starting from $q=\infty$ and decreasing q. The smallest integer value of q attained gives then a realization of the coloring problem. We shall not pursue this line here, but rather attack the coloring problem from a different angle.

B. Introducing temperature in pseudoenergy landscape

The pseudoenergy landscape suggests how to construct a Monte Carlo and/or parallel tempering program for hard spheres.

Consider monodisperse hard spheres—the generalization to polydisperse is straightforward—at inverse temperature (i.e., pressure) β. Given a configuration, its pseudoenergy is proportional to (minus) the smallest interparticle distance $2r_o$. We choose a particle at random and displace it by a random amount. The minimal interparticle distance $2r_o^{\text{old}}$ may have changed: if it is so it must be due to the distance between the particle just moved and one of its neighbors becoming some value $r_o^{\text{new}} < r_o^{\text{old}}$ We thus accept the motion if $r_o^{\text{new}} \geq r_o^{\text{old}}$, otherwise we accept it with probability $e^{-2}\beta(r_o^{\text{old}} - r_o^{\text{new}})$. This is the Monte Carlo procedure. One can now run several copies in parallel at different β and implement the usual parallel tempering procedure.

C. Discrete lattice procedures

We have chosen to exemplify the algorithm in the random q-coloring problem. We are given a long list of links, and add them one by one. Suppose that we have a graph with $M-1$ links, and a configuration of colors for the vertices satisfying the constraints. We now add the Mth link from the list: with probability 1 it will locally appear as in Fig. 9. As we add the link, it may be that vertices A and B are of different colors, in which case we proceed. If, on the contrary, the colors of A and B coincide, we must modify at least one of the two. This will typically create problems in the tree that has roots in them, so we shall need to modify colors up the branches. If this can be done with a finite number of trials, we do it, compute the total number of steps it took, and proceed with the next link addition. In practice, for the readjustment of the trees with root A or B, we use a Walk-COL algorithm introduced in [24], which is the coloring version of the celebrated Walk-SAT program [36,37] This is perhaps not ideal, but it definitely provides an upper bound of the number of flips needed. Moreover, this strategy is straightforward to implement numerically. An important property of this algorithm is that it is *focused*: it changes only frustrated variables and it thus acts locally starting from the initial frustrated link, and does not perturb the solution far from this link.

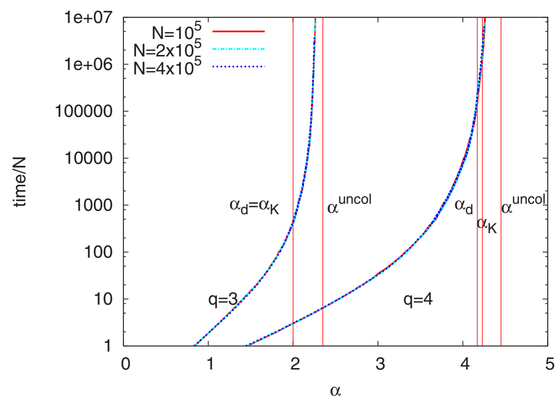

FIG. 10. (Color online) Performance of the recursive algorithm for the $q=3,4$ random coloring problem. The time needed to find a proper coloring diverges at connectivity $\alpha_* < \alpha_{\text{uncol}}$ in both cases, but the algorithm is able to find colorings in linear time beyond the clustering α_d and Kauzmann α_K transitions.

In this form, the algorithm is a *recursive version* of the method called Incremental-SAT [39] in the computer science community, in which one starts from a SAT formula and its known solution (or a graph and its proper q-coloring) and then adds one formula (or a new link for coloring) and must find a solution to the new problem. This is usually done using Walk-SAT, as we do, although this procedure has not been proposed as far as we know as a general way to obtain solutions starting from scratch. The point reached by our algorithm can thus be seen as a easy-hard transition for the recursive, incremental COL problem.

D. Coloring random graphs

We now discuss our results and show our data for the 3-coloring and the 4-coloring in Fig 10. It turns out that the number of steps needed (the number of recolorings close to A or B to get back to satisfaction) grows on average and diverges at a given connectivity α^*, where our program stops.

The fact that the rearrangement needed grows with the connectivity is not surprising in view of the results of Montanari and Semerjian [40,41], who showed that the minimal number of rearrangements needed to satisfy constraints following a change of color of a random vertex diverges like a power law at a certain connectivity. The distance along the graph of the rearrangements is a power law as well. This critical connectivity corresponds in fact to the rigidity transition α^{hard} considered in [24] where frozen variables (or hard fields in the cavity language) appear that fix the value of a finite fraction of variables for all configurations within a state. After a situation with frozen variables is reached, an additional link will have a finite probability of connecting two variables frozen to the same color, thus rendering unsatisfied all configurations within the state. States with hard fields have a high mortality rate: after an extensive number of link additions, the survival probability of such a state is exponentially small in N.

In fact, here the situation is somewhat different to the one in [40,41], in that we are not considering an equilibrium configuration, but one obtained by a succession of link additions and rearrangements. The actual value α^* that our program can reach need not coincide exactly with $\alpha^{\rm hard}$, because hard fields appear at somewhat lower connectivities in our out-of-equilibrium procedure.

It can be shown that for large connectivity, all clusters have frozen variables beyond a certain connectivity (see, for instance [24,25], or [42] for rigorous proof in SAT): this puts a strict limit to the algorithm. For the 3- and 4-coloring problem, it has been shown (see again [24]) that all states have frozen variables for $\alpha \geqslant \alpha^{\rm hard}$.

Because we can scale the curves with system size, and we consider only the regions where the curves superpose properly, we can in fact infer the behavior for $N=\infty$. We do that in Fig. 11. In the 3-coloring, the algorithm reaches in time linear in N a value $\alpha^* \approx 2.275$ which is beyond the clustering and Kauzmann transition $\alpha_d=\alpha_K=2$. It is however, as expected, below the appearance of frozen variables in the thermodynamic states that take place at the rigidity transition $\alpha^{\rm hard}=2.33$ (and of course below $\alpha_{\rm uncol}=2.345$).

In the 4-coloring, the algorithm reaches in time linear in N a value $\alpha^* \approx 4.31$ which is again well above $\alpha_d=4.175$ and $\alpha_K=4.23$ (but still systematically below $\alpha^{\rm hard}=4.42$ and $\alpha_{\rm uncol}=4.45$).

The integrated number of steps up to connectivity α is $(\alpha^*-\alpha)^{-\mu}$ with $\mu \sim 0.28$. In Fig. 11 we show for reference the case $q=3$, which is somewhat more dubious because it is not clear yet to which class—from the replica theory point of view—it belongs. In any case, the value $\alpha^* \sim 2.275$, far beyond $\alpha_d=\alpha_K=2$, is quite close (but still smaller) to the one ($\alpha_{\rm SP}=2.3$) obtained in Ref. [4] with the best algorithm—survey propagation (SP) [2]—for smaller sizes.

Before concluding this section, we wish to compare for this case the present algorithm with the belief propagation (BP) algorithm [7,24]. We start with a configuration obtained with BP at a given α (which we shall try to make as large as possible) [43], and then improve the procedure by adding links one by one. The result is shown on the right-hand side of Fig. 11: at the beginning there is a very fast progress, but the asymptotic α^* is not significantly different—at least within our precision. We conclude that starting from a BP solution does not seem to improve the scaling with the system size.

To explain this, we can think in the flood analogy of BP as starting us somewhere in the middle of an island. As the level of water keeps on going up, at the beginning we have little to do to stay dry, and the program proceeds rapidly. Only when the many-dimensional shore reaches us do we begin to move in a complex manner, and the system slows down.

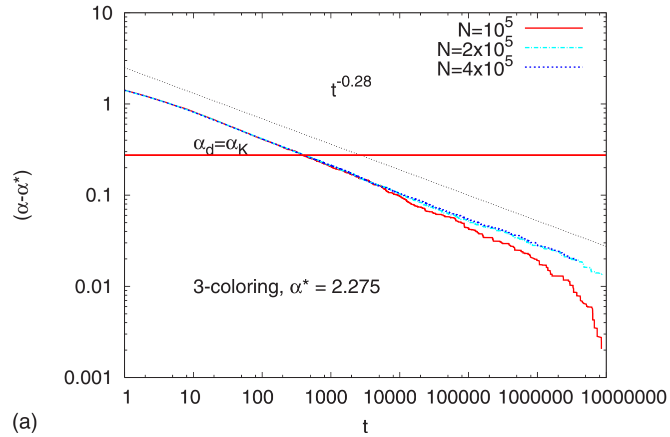

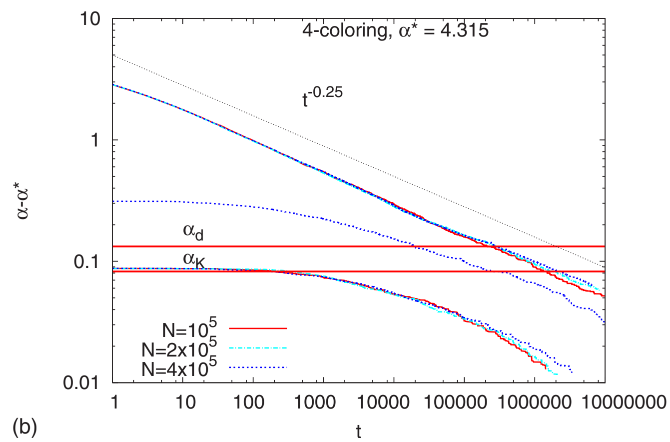

FIG. 11. (Color online) Estimate for the asymptotic α^* of the algorithm. Left-hand side: $q=3$ random coloring problem. Right-hand side: $q=4$ random coloring problem. The two lower families of curves correspond to procedures starting from configurations obtained using the belief propagation algorithm of Ref. [24]: the initial advantage seems eventually lost.

Scalings with system size. By performing at each step Walk-COL starting from a site at the end of a new link, we have used on average a certain function $G(\alpha^*-\alpha)$ color flips each time we add one link, a quantity of order one. The program time $t(\alpha)$ (the total number of steps up to connectivity α) thus increases by $G(\alpha^*-\alpha)$ each time that α increases by $1/N$,

$$\frac{dt}{d\alpha} = NG(\alpha^* - \alpha) \tag{8}$$

and $t(\alpha)$ is N times the integral of G. Thus, the curves $t(\alpha)/N$ for different N should collapse, which we check in Fig. 10. The fact that the total time scales with N holds to the extent that the number of rearrangements is, for large graphs, independent of the graph size.

Note that nothing guarantees that $G(\alpha^*-\alpha)$, defined on the basis of Walk-COL, is the minimum average number of flips needed to make the configuration satisfied after adding a link, but just an upper bound. On the other hand, in a locally treelike structure, the maximal number of moves with any program is to try all the flips of sites at distance along the graph smaller or equal than $G(\alpha^*-\alpha)$, that is $\sim\alpha^{qG(\alpha^*-\alpha)}$, still a quantity independent of N. Thus, the fact that the local structure of the graph is treelike has helped us to obtain a program linear in N.

If we did not know that we had a treelike structure—for example, if some loops are present—we would have to explore *all* configurations in phase space with a given number of flips, trying all the phase space points at progressively larger phase-space distances until satisfaction is restored. In this way, $G(\alpha^*-\alpha)$ flips would be obtained after at most $\sim N^{qG(\alpha^*-\alpha)}$ trials, and the total integrated time would be

$$t \sim AN\int_0^\alpha d\alpha' N^{qG(\alpha^*-\alpha')} \sim \left(\frac{AN}{\ln N}\right)\frac{N^{qG(\alpha^*-\alpha)}}{qG'(\alpha^*-\alpha)}. \quad (9)$$

This is the absolute worst the program will do for the random graphs discussed here, not using either Walk-COL or the knowledge of the local treelike structure. Below α^*, the performance is still a power law in N.

E. Marginality, hard fields and rattlers

As we have seen in the preceding sections, when we add links to the coloring problem (or clauses to the SAT problem) each time the number of changes needed to keep the system satisfied on average increases. This happens because as we approach the point α^* larger and larger clusters of vertices must be moved together, changes that keep the system satisfied must be done in a correlated manner in each cluster. The same happens with the sphere-packing problem: the clusters of particles that are in contact, and must be displaced together at each step, increases with packing and at the critical value of r_o the whole structure percolates. By the time α^* (or the J-point, in a packing problem) is reached, the correlation diverges, and the algorithm virtually stops. Some important remarks are as follows.

(i) The density of links at which this happens *depends on the procedure*, the point α^{hard} obtained at equilibrium need not coincide with the one obtained through the sequence of additions of links we have been doing here. (This is why in Fig. 11 we have tried to compare the divergence obtained by addition of links starting from low link density, with the divergence obtained starting from a very connected graph colored with a BP program. Admittedly, the difference, if any, is not large.) For a particle system, the packing fraction at which this percolation happens also *depends on the procedure*.

(ii) For the overwhelming majority of runs of the algorithm, when the level α^* is reached, the addition of an extensive number of links requires a divergent number of recolorings. However, it will generally happen that a finite fraction of vertices can still be recolored in many ways, without destroying the satisfaction. For example, the intermediate vertices in a short loop made of a succession of vertices of connectivity 2 can be recolored in many ways. Similarly, when a particle system reaches r_o, there still are particles, or groups of particles (the rattlers) which can move inside cages made by the others.

To summarize: the SAT, coloring, and packing problems all have critical situations that are process dependent in which a subset of the elements form a backbone subjected to hard fields from their neighbors and rattlers that still have freedom. This point corresponds to the easy/hard transition for the recursive incremental SAT (or Col) algorithm described here. If we look at the pseudoenergy of a state, with large probability the number of configurations at the critical level of difficulty is still large (due to the freedom of rattlers), but it disappears completely at a level of difficulty just higher. In other words: seen from the point of view of the pseudoenergy, the minima have a flat bottom consisting of the rattler configurations [44].

V. THE *J*-POINT

In a series of papers [8,45] the highest density configuration of hard spheres reached by inflating as in Sec. IV A have been studied [46]. The point where this happens has been called "*J*-point." As we noted above, they did not consider a slow annealing, as in the Lubachevsky-Stillinger procedure, because this would make the system evolve in an annealing-time-dependent manner, ultimately allowing it to crystallize. For spheres in three dimensions the J-point packing fraction turns out to be very close to the one quoted as random close packing [9], and whether one can identify J-point and random close packing is a debated question [47]. The interest of the J-point lies at any rate in the fact that it is *critical*, it has associated with it a diverging length and soft modes [8,48]. There is also good evidence that for spherical particles it is *isostatic*, i.e., that the number of contacts is just the number it takes to immobilize the system, without redundant static conditions.

A question that immediately arises is whether the J-point packing corresponds to the ideal glass phase, that is, the best packed amorphous state *without crystallization*. Putting the question this way, we immediately run into difficulties because by considering different degrees of crystallinity, we can obtain a continuum of denser packings [49]. The random models studied here allow us to avoid—at least to fix ideas—the conceptual complication of crystallization, and to concentrate on the glassy aspects [50,51]. We shall hence first discuss the SAT and coloring, which by construction have no ordered states. On the latter the procedure of Ref. [8] can be repeated in the angle packing version without modifications.

In SAT and coloring we have already pointed out that the recursive incremental procedure of adding difficulty in a constraint satisfaction problem yields a critical point with a diverging length, given by the range of rearrangements necessary each time to satisfy the constraints (see Refs. [40,41]). This point, which we identify as the *mean-field version of the J-point*, also corresponds to the appearance of frozen variables in the state [52].

In the framework of the pseudoenergy landscape the procedure leading to the J-point amounts to a deep (zero temperature) quench in this complex landscape, starting from a random configuration. In this language, it is an infinite temperature inherent structure. Clearly, there is no reason for this procedure to converge to the (glassy) ground state (otherwise it would be the ultimate solution to the coloring and SAT problem): minima with the largest basin of attraction are generically not the deepest.

Running the program several times starting from different initial conditions leads to different end densities [8], distributed around a typical value, the closer the larger the system: The pseudoenergy view allows us to recognize this as the standard situation with the inherent structure energy distribution [53].

Another thing to notice is that configuration reached in the J-point is *not* typical of the given packing fraction, just as an inherent structure reached by a quench in an energy landscape is not the typical one of the corresponding energy. One way to emphasize this is to envisage the analog of the protocol of Ref. [35]: starting from *equilibrium* configurations at different initial packing fractions (or values of r_o), one applies the inflating procedure of O'Hern *et al.* If the landscape hypothesis advocated here holds, the final packing fraction (or r_o) reached should depend on the initial one in much the same way as in Ref. [35] the inherent structure energies depend upon initial equilibrium energies (in finite dimensions crystallization effects must be taken care of [50]). It would be interesting to study other amorphous configurations with the same packing fraction as the J-point, but reached after different annealing protocols, and check whether they are not critical. This is what happens in the mean-field models discussed above: their *typical* configurations at a packing fraction equal to that of the J-point are *not* critical, even in this situation in which there is no possibility of crystalline order.

VI. ANALYTICAL COMPUTATIONS OF PERFORMANCE: DYNAMICS AND CAUSAL REPLICA CHAINS

The statistical mechanical description of complexity, in particular based on the replica trick, concerns the average geometric structures in a given subset of phase space. This information is local, in the sense that it involves the structure of typical states and not their whole basin of attraction. Sometimes, this local information is enough to infer where a dynamic process starting from a random configuration will go. This is the case, for example, of dynamics of certain mean-field models, where marginal states are chosen. However, it is not clear in general up to what point local information of states is sufficient. One alternative that has been widely discussed in the literature, is to solve the averaged dynamics (for a review of this approach, see Ref. [54]). Here we shall discuss a different but related approach, which is a generalization of several works in the replica literature of the 1980s and 1990s [55–57] and in optimization problems [31].

Let us first consider an energy landscape of discrete variables. We start by choosing a random configuration $\mathbf{s_1}$ (here we use boldface to denote N-dimensional vectors). Next, keeping this one fixed, we search for another configuration $\mathbf{s_2}$ at distance Δ_1 from the first. Keeping these two fixed, we choose a third configuration $\mathbf{s_3}$ at distance Δ_2 from the second, and so on. In this way, we construct a causal chain, since the subsequent links do not affect the preceding ones. We must demand, for example, that each new replica be at an energy as small as possible (subject to its constraints): this is done by thermalizing each link at a very low temperature. Next, we must make the distance between links tend to zero, and the number of links to infinity. From the point of view of replica theory, this is just a generalization of the effective potential method [57], and is also closely related to the discussion in Refs. [55,56]. We write a partition function

$$Z = \lim_{n_1\to 0} \ddots \lim_{n_R\to 0} \int d\mathbf{s_1}^{\alpha_1} d\mathbf{s_2}^{\alpha_2} \ddots d\mathbf{s_R}^{\alpha_\mathbf{R}} \delta(\text{chain}) \times \exp\left\{ -\beta_1 \sum_{\alpha_1=1}^{n_1} E(\mathbf{s_1}^{\alpha_1}) - \beta_2 \sum_{\alpha_2=1}^{n_2} E(\mathbf{s_2}^{\alpha_2}) - \ddots - \beta_R \sum_{\alpha_R=1}^{n_R} E(\mathbf{s_R}^{\alpha_\mathbf{R}}) \right\}, \tag{10}$$

where $\mathbf{s_1}^{\alpha_1}, \mathbf{s_2}^{\alpha_2} \ddots \mathbf{s_R}^{\alpha_\mathbf{R}}$ are all N-dimensional vectors of spins. $\mathbf{s_1}^{\alpha_1}$ is replicated n_1 times, $\mathbf{s_2}^{\alpha_2}$ n_2 times, etc. Spins are organized in a hierarchical manner: one of the n_1 replicas $\mathbf{s_1}^{\alpha_1}$ has an offspring of n_2 of the $\mathbf{s_2}^{\alpha_2}$. One of the n_2 of these has in turn an offspring of n_3 of the $\mathbf{s_3}^{\alpha_3}$, and so on. The term δ(chain) is a product of δ's imposing that the rth offspring is at a fixed distance Δ_r from its parent.

For a constraint satisfaction problem, the problem can be recast as follows. Taking the example of the coloring problem, we suppose we have a graph with connectivity per node $\alpha_{\max}$. The links are given by an $N \times N$ symmetric matrix J_{ij} with elements one and zero. An ordering is introduced by defining the symmetric matrix C_{ij}, with elements taken at random with distribution uniform in $(0, \alpha_{\max})$. The number of miscolorings at connectivity α is then

$$M(\mathbf{s},\alpha) = \sum_{ij} J_{ij} \Theta(\alpha - C_{ij}) \delta_{s_i s_j}, \tag{11}$$

where $\Theta(x)$ is 1 if x is positive, and zero otherwise. One repeats now the argument as before, for the case of a chain of replicas with progressively larger values of α, keeping the total number of errors equal to zero.

In practice, one can make the computation for chains of a small number of links, and extrapolate the result to the limit of continuous chain. For systems on the Bethe lattice, the replica calculation should be translated into the cavity language, which is more flexible.

VII. CONCLUSIONS

We have adopted a unified point of view of hard-particle glasses and the constraint satisfaction problems of computer science and taken both to a (pseudo) energy landscape setting. On the one hand, the glass literature has long been a source of inspiration and methods to attack satisfaction problems. On the other hand, the statistical mechanics of optimization problems provides nontrivial yet solvable models for which questions on the jamming transition can be answered exactly.

In both cases, there are questions—such as a rigidity (hard field) threshold, the presence of rattler and a backbone of fixed particles (or Boolean variables, or colors)—peculiar to the fact that the task is posed as a zero temperature problem of optimizing parameters while respecting a constraint. The pseudoenergy language helps to bring back many of these results to a problem of relaxation in a complex landscape.

A. Satisfaction

The statistical mechanic point of view of satisfaction has been mainly concentrated on computation of static snapshot properties of phase space. This approach has been very fruitful and has suggested powerful algorithms [2]. Here, instead, we take the approach of defining an algorithm whose properties can in principle be computed analytically, at least in the same cases for which a static solution is possible. This change of perspective reflects a similar passage from a static to a dynamic approach in the glass literature [54], and has also been pursued in the statistical mechanics of optimization [58]. In this paper we have only stated but not completed the analytical calculation, but we have shown numerically that problems are easy beyond the clustering transition, and with the coloring problem that in some cases even beyond the Kauzmann one. Indeed, one of the conclusions is that a rather excessive importance has been attributed to the clustering (dynamic) transition point: this is the point up to which simple programs can easily sample *all* configurations [59], but not the last point at which it can easily obtain at least *one* configuration—this happens at α^*. It is then, upon reflection, not so surprising that relatively straightforward algorithms such as Walk-SAT [38] find easily solutions almost up to the satisfiability threshold.

B. Jamming transition

When the J-point was introduced [8] and shown to correspond to a packing fraction close to the empirical value of random close packing, the question arose of its meaning in terms of the glass state. What the present approach underlines (but was already implicit in [10]), is that the nature of this point is the same as that of a zero-temperature quench in a rugged energy landscape [60]. Deeper, more compact, points can be reached with other annealing procedures—*even in models without a crystalline state*, as the examples we have considered show. This suggests a picture of the transition in the temperature–shear-rate–density space [61], rather than as a single surface, as a multilayered onion, each level being reached after a different protocol.

There is density at which a hard particle system has an experimentally unaccessible relaxation time scale: this happens at a density distinctly smaller than the J-point one. Indeed, the present analysis shows that the J-point is necessarily denser than the mode-coupling transition density, in models in which such a transition exists and is sharp. On the other hand, the question of which point is denser, the J-point or the Kauzmann transition point depends on the system and the dimensionality—again based on what can be inferred from models for which the latter transition can be shown to exist [62]. Note that, if the J-point is denser than the glass transition point, this is at the expense of being out of equilibrium.

An interesting aspect of the J-point is that it is critical. Here we have found that the same is also true for mean-field models (as had already been found in the q-core problem [63]), but this criticality does not correspond to the thermodynamic glass transition itself (which in mean-field exists), but is that of the zero-temperature threshold states [21] already familiar from the mean-field glassy dynamics. In fact, in any of the problems discussed here, if the procedure of approach to the J-point is stopped when the correlation length has reached a certain large value, and a thermalization at fixed r_o (or α) is applied subsequently, as the system tends towards equilibrium the pressure drops and the correlation length *diminishes*, rather than growing.

The study of problems with disorder raises doubts on the fact that isostaticity of a frictionless particle system is a necessary (apart from sufficient [45,64], although see [65]) condition for criticality. This may be checked just by testing the J-point procedure with a packing of nonspherical (e.g., ellipsoidal) particles, which are generically hypostatic [66]: if the J-point for ellipsoids turns out to be critical, then it will be clear that isostaticity is not a necessary condition for criticality.

Finally, the analogy between particle and satisfaction problems suggests that one may perhaps build a mean-field, random first-order theory of rigidity (and to a certain extent plasticity) by considering systems which develop hard fields, and subjecting them to driving, nonconservative forces. The coloring Hamiltonian, or its angle packing version, is a good candidate for such a strategy.

ACKNOWLEDGMENTS

The authors wish to thank B. Chakraborty, S. Franz, A. Liu, M. Mézard, R. Monasson, A. Montanari, F. Ricci-Tersinghi, S. Sastry, G. Semerjian, L. Zdeborová, and R. Zecchina for discussions about these issues.

[1] M. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory and NP-Completeness* (Freeman, San Francisco, 1979); C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, New York, 1994).

[2] M. Mézard and R. Zecchina, Phys. Rev. E **66**, 056126 (2002); M. Mézard, G. Parisi, and R. Zecchina, Science **297**, 812 (2002).

[3] R. Monasson and R. Zecchina, Phys. Rev. E **56**, 1357 (1997).

[4] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, Phys. Rev. Lett. **89**, 268701 (2002).

[5] Y. Fu and P. W. Anderson, J. Phys. A **19**, 1605 (1986).

[6] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).

[7] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová, Proc. Natl. Acad. Sci. U.S.A. **104**, 10318 (2007).

[8] C. S. O'Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, Phys. Rev. E **68**, 011306 (2003).

[9] For compression procedure closely related to the experimental one, and yielding densitites close to random close packing, see

Hernán A. Makse, David L. Johnson, and Lawrence M. Schwartz, Phys. Rev. Lett. **84**, 4160 (2000).
[10] B. D. Lubachevsky and F. H. Stillinger, J. Stat. Phys. **60**, 561 (1990).
[11] F. H. Stillinger and T. A. Weber, J. Chem. Phys. **83**, 4767 (1985).
[12] A.-S. Sznitman, *Brownian Motion, Obstacles and Random Media* (Springer, Berlin, 1998).
[13] This is perhaps related to the formalism in P. M. Duxbury, D. J. Jacobs, M. F. Thorpe, and C. Moukarzel, Phys. Rev. E **59**, 2084 (1999). A potential is also defined in C. Brito and M. Wyart, Europhys. Lett. **76**, 986 (2006).
[14] C. De Dominicis, M. Gabay, T. Garel, and H. Orland, J. Phys. (France) **41**, 922 (1980).
[15] A. J. Bray and M. A. Moore, J. Phys. C **13**, L469 (1980).
[16] L. F. Cugliandolo and J. Kurchan, J. Phys. A **27**, 5749 (1994).
[17] The literature on the metastable states of the Sherrington-Kirkpatrick model covers a quarter of a century. For recent works, see A. Crisanti, L. Leuzzi, G. Parisi, and T. Rizzo, Phys. Rev. B **70**, 064423 (2004); M. Muller, L. Leuzzi, and A. Crisanti, *ibid.* **74**, 134431 (2006); A. Cavagna, I. Giardina, G. Parisi, Phys. Rev. Lett. **92**, 120603 (2004). The barrier to escape a state of free energy per spin f scales as $B(f-f_0)^{-1/3}$, where f_0 is the free energy density of the lowest state; a finite number of flips will unstabilize states with $f>f_0$, see T. Aspelmeier, R. A. Blythe, A. J. Bray, and M. A. Moore, Phys. Rev. B **74**, 184411 (2006). Thus, these calculations are in agreement with the conclusions from the dynamics.
[18] J. Kurchan, G. Parisi, and M. A. Virasoro, J. Phys. I **3**, 1819 (1993).
[19] A. Crisanti and H.-J. Sommers, J. Phys. I **5**, 805 (1995).
[20] J. Kurchan and L. Laloux, J. Phys. A **29**, 1929 (1996).
[21] L. F. Cugliandolo and J. Kurchan, Phys. Rev. Lett. **71**, 173 (1993).
[22] L. F. Cugliandolo, D. Grempel, G. Lozano, and H. Lozza (unpublished).
[23] This can always be done by doing Monte Carlo dynamics with a temperature adjusted to keep energy constant.
[24] L. Zdeborová and F. Krzakala, e-print arXiv:0704.1269v1.
[25] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian (unpublished).
[26] A. Montanari and F. Ricci-Tersenghi, Eur. Phys. J. B **33**, 339 (2003).
[27] In the thermodynamic limit, it is not at present clear whether the point reached after an annealing from high temperatures to a given low temperature $T<T_d$, followed by a long wait at T depends on the speed of the annealing. If this were so, then the threshold needs a process-dependent definition in this model [28].
[28] A. Montanari and F. Ricci-Tersenghi, Phys. Rev. B **70**, 134406 (2004); B. Capone, T. Castellani, I. Giardina, and F. Ricci-Tersenghi, *ibid.* **74**, 144301 (2006); A. Barrat, S. Franz, and G. Parisi, J. Phys. A **30**, 5593 (1997).
[29] Andrea Montanari, Giorgio Parisi, and Federico Ricci-Tersenghi, J. Phys. A **37**, 2073 (2004).
[30] Florent Krzakala, Andrea Pagnani, and M. Weigt Phys. Rev. E **70**, 046705 (2004).
[31] R. Monasson and D. O'Kane, Europhys. Lett. **27**, 85 (1994); G. Biroli, R. Monasson, and M. Weigt, Eur. Phys. J. B **14**, 551 (2000); M. Mézard, M. Palassini, and O. Rivoire, Phys. Rev. Lett. **95**, 200202 (2005); also see Ref. [7].
[32] A. Barrat, e-print arXiv:cond-mat/9701031.
[33] J. van Mourik and D. Saad, Phys. Rev. E **66**, 056120 (2002).
[34] Here "annealing" refers to letting a system adapt slowly to each new energy, and is not to be confused with the "annealed approximation," which consists of optimizing both configuration *and* problem simultaneously.
[35] S. Sastry, P. G. Debenedetti, and F. H. Stillinger, Nature (London) **393**, 554 (1998).
[36] B. Selman, H. A. Kautz, and B. Cohen, Proceedings of the AAAI-94, Seattle, 1994.
[37] The algorithm is the adaptation of an efficient strategy introduced for satisfiability in [38]: We choose randomly one of the spins that has the same colors as one, or more, of its neighbors and change its color randomly: the move is accepted with a probability 1 if this improves the coloring—if more spins are satisfied—and with a probability α otherwise. The parameter α is tuned for more efficiency, see [24] for details.
[38] John Ardelius and Erik Aurell, Phys. Rev. E **74**, 037702 (2006).
[39] J. Whittemore, J. Kim, and K. Sakallah "SATIRE: a new incremental satisfiability engine," in Proceedings of the 38th conference on design automation, 2001, pp. 542–545; N. En and N. Srensson, "Temporal induction by incremental SAT solving," First international workshop on bounded model checking, ENTCS, Vol. 89(4).
[40] A. Montanari and G. Semerjian, Phys. Rev. Lett. **94**, 247201 (2005); J. Stat. Phys. **124**, 1572 (2006).
[41] In fact the model considered in [40] is somewhat special. More general and detailed results are obtained for the coloring and the satisfiability problem in G. Semerjian, e-print arXiv:0705.2147.
[42] D. Achlioptas and F. Ricci-Tersenghi, Report No. CC/0611052.
[43] BP alone does not give a solution to the problem, but only estimates the probability that a given variable takes a given color over the whole set of solutions. Finding the solution thus requires a decimation of these variables (see [24] for details). To find a proper coloring of a graph of N variables, we applied the following procedure recursively N times: (i) run BP and (ii) fix the most biased variable to its most likely value. We have also tried to fix the variables according to their BP probability estimate—and not systematically to its most biased value—but, interestingly, the iterative application of Walk-COL is less effective when starting from these solutions. This is a sign that the first procedure gives solutions that belong to larger clusters.
[44] The actual density of rattlers and its dependence on the annealing procedure for hard spheres is still a matter of debate.
[45] For an overview, see M. Wyart, Ann. Phys. (Paris) **30**, 1 (2005).
[46] They obtain what seems to be an equivalent situation with a different deflating process, but we shall not discuss that one here.
[47] Reference [8] was criticized in A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly, Phys. Rev. E **70**, 043301 (2004). See the reply: C. S. O'Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, *ibid.* **70**, 043302 (2004).
[48] M. Wyart, S. R. Nagel, and T. A. Witten, Europhys. Lett. **72**, 486 (2005).
[49] S. Torquato, T. M. Truskett, and P. G. Debenedetti, Phys. Rev.

Lett. **84**, 2064 (2000).
[50] In finite dimensional systems one may minimize the conceptual complications related to crystallization is several ways: (i) considering high enough dimensions where it is conjectured that the basin of attraction of amorphous packings dominates [51], (ii) working with dynamics that artificially imposes constant values of parameters of crystallinity (denoted Q in [49]), (iii) using a mixture of particles that either does not crystallize or if it does its order is so complicated as to be dynamically irrelevant.
[51] M. Skoge, A. Donev, F. H. Stillinger, and S. Torquato, Phys. Rev. E **74**, 041127 (2006).
[52] The transition where this happens along an equilibrium path was computed in [24]. Here the idea is the same but along a nonequilibrium path.
[53] W. Kob, F. Sciortino, and P. Tartaglia, Europhys. Lett. **49**, 590 (2000).
[54] L. F. Cugliandolo, e-print arXiv:cond-mat/0210312.
[55] G. Parisi and M. A. Viarsoro, J. Phys. (Paris) **50**, 3317 (1989).
[56] R. Monasson, Phys. Rev. Lett. **75**, 2847 (1995).
[57] S. Franz and G. Parisi, Physica A **261**, 317 (1998).
[58] S. Cocco and R. Monasson, Theor. Comput. Sci. **320**, 345 (2004).
[59] The clustering point is the celebrated mode-coupling (dynamic) transition point of glass theory. This has been recently extended to dilute systems in Andrea Montanari and Guilhem Semerjian, J. Stat. Phys. **125**, 23 (2006).
[60] One may then ask if an infinite temperature inherent structure of, say, a Lennard-Jones system, should also be critical. The studies of vibration modes of inherent structures of temperature T seem to show an excess of very low frequency modes at higher starting temperatures T (see Fig. 3 of Ref. [53]), but this requires a more detailed study.
[61] A. J. Liu, and S. R. Nagel, Nature (London) **396**, 6706 (1998).
[62] See the discussion for hard spheres in F. Zamponi, Phys. Rev. E **75**, 043101 (2007).
[63] J. M. Schwarz, A. J. Liu, and L. Q. Chayes, e-print arXiv:cond-mat/0410595.
[64] C. F. Moukarzel, Phys. Rev. Lett. **81**, 1634 (1998).
[65] An additional puzzle is brought about by the replica calculations, where states are isostatic but not critical, see G. Parisi and F. Zamponi, J. Stat. Mech.: Theory Exp.(2006) 3017; J. Chem. Phys. **123**, 144501 (2005).
[66] P. M. Chaikin, A. Donev, W. N. Man, F. H. Stillinger, and S. Torquato, Ind. Eng. Chem. Res. **45**, 6960 (2006); A. Donev, R. Connely, F. H. Stillinger, and S. Torquato (unpublished).
[67] A true mean-field calculation for the J-point of hard spheres should also be possible [65].