Determining Energy Barriers by Iterated Optimization: The Two-Dimensional Ising Spin Glass

C. Amoruso and A. K. Hartmann
Institut für Theoretische Physik, Universität Göttingen, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany

M. A. Moore
School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, U.K.

Energy barriers determine the dynamics in many physical systems like structural glasses, disordered spin systems or proteins. Here we present an approach, based on subdividing the configuration space in a hierarchical manner, which leads to upper and lower bounds for the energy barrier separating two configurations. As an application, we consider Ising spin glasses, where the energy barriers which need to be surmounted in order to flip a compact region of spins of linear dimension L are expected to scale as L^{ψ}. The fundamental operation needed is to perform a constrained energy optimization. For the the two-dimensional Ising spin glass we use an efficient combinatorial matching algorithm, resulting in the non-trivial numerical bounds $0.25 < \psi < 0.54$.

PACS numbers: 75.10.Nr, 75.40.Mg, 02.10.Jf

I. INTRODUCTION

A unifying concept in the physics of disordered systems is the notion of energy landscapes. The dynamics of these systems is determined by (free) energy barriers. Prominent examples include spin glasses[1], structural glasses[2] and folding of proteins[3]. In the language of theoretical computer science, the barrier problem belongs[4] to the fundamental class of non-deterministic polynomial (NP) hard problems[5]. As a consequence all known algorithms determining lowest barriers run with a time which grows exponentially with system size.

In this paper, we will present a hierarchical approach to calculating minimum barriers. The basic operational tools needed are combinatorial optimization algorithms to perform constrained energy optimizations. The algorithm presented below will result in the number of constraints increasing with the level in the hierarchy[6]. As an illustrative application, we study the prototypical two-dimensional Ising spin glass, where non-trivial bounds for the barriers have not previously been obtained.

In the droplet[7] or scaling theory of Ising spin glasses the low-energy excitations are compact droplets. The creation of a droplet (a region of reversed spins of linear extension L) results in the formation of a domain wall around the reversed spins, its energy scaling as L^{θ}. The system orders at low temperatures only if $\theta > 0$. The domain wall is fractal and has area L^{d_S}, with $d-1 \leq d_S \leq d$. The dynamics of the system are controlled by the height of the barriers which have to be crossed to create such droplets. It is generally assumed that the barrier to be surmounted to create a droplet of linear extent L has an energy which scales as L^{ψ} where $\theta \leq \psi \leq d-1$ for dimension d. The argument for these inequalities goes as follow: the barrier must be at least as large as the energy required to create the domain, hence $\theta \leq \psi$; the upper limit is due to the fact that the barrier must be lower than the energy of a compact droplet with a non-fractal surface containing the same number of reversed spins[7]. While for a directed polymer in a random system it was shown that $\psi = \theta$[8], for Ising spin glasses ψ appears to be an independent exponent.

While many numerical studies have been done to calculate the exponent θ, we are aware of only one direct numerical estimate of ψ and that was for two dimensions[9]. It was equal to the upper bound i.e., $\psi = d-1$, but only small systems were studied ($L \leq 6$). As mentioned, the barrier problem is NP hard and it has been difficult even to find good approximate algorithms for barriers. This is presumably the reason for the paucity of studies of this exponent. Recently Drossel and Moore instead of attempting to calculate the barrier exactly, placed bounds on its energy[10]. They showed that for the hierarchical lattice (that is constructed iteratively by replacing each bond by 2^d bonds) there exists a lower bound on ψ which coincides with the upper bound $d-1$, hence $\psi = d-1$. In the same spirit as their work, we will also obtain upper and lower bounds on the energy of the barriers but for the more physically relevant case of the square lattice. Rather than studying droplets of order of the system size L, we will consider the computationally simpler task of determining upper and lower bounds for the barrier separating the two ground states (GSs) related by a flip of all spins. This represents the highest barrier in the system and thus determines the longest timescale of the dynamics. It is equivalent to sweeping a domain wall through the system. Both types of barriers should be equivalent for two-dimensional spin glasses since domain walls and droplets seem to share the same exponents[11].

In the next section we explain the general idea of our algorithm. We describe the general case, but for illustration we give examples in the chemical language, where a molecule changes its state, e.g. the transition of a protein between two conformations. This is probably more illustrative for the general reader. Then we introduce the spin-glass model which we study here and show how the algorithm is implemented for this case. In the main section, we present the results for the barriers of two-dimensional Ising spin glasses. We conclude with a summary and a discussion.

II. ALGORITHM

A. General outline

The algorithm works in the following way. Given are an energy function $H(c)$ of configurations c, e.g. the (free) energy of a molecule depending on the positions c of all atoms, and two initial configurations $c_1^{(0)}, c_2^{(0)}$. For simplicity, we consider in the following just energies instead of free energies. We assume that we have a good minimization algorithm *opt* available which finds ground states $\min_c H(c)$. We also assume that the algorithm allows us to introduce constraints C', e.g. restricting some bond angles or inter-atomic distances to fixed values or within an interval, i.e. the result $opt(C')$ depends on C'. For the case of molecules, there is in fact no efficient method to generate exact ground states, but we can assume that we can get reasonable good approximations using steepest descent methods like conjugate gradient, possibly combined with other heuristics. It is important to say that for the two-dimensional spin-glass problem we are finally considering here, there is in fact a fast polynomial algorithm, which we use to calculate exact ground states, see below.

The general goal is to find a (reaction) path in configuration space between $c_1^{(0)}, c_2^{(0)}$ whose energy $E_{\max}$ at the highest point along the path is as low as possible, i.e. a minimum path in that sense. The energy barrier is then $\Delta E_B = E_{\max} - H(c_1^{(0)})$. The standard way to describe the transition between the two configurations is by using a reaction coordinate τ, that changes along the path. This can be a suitably chosen bond angle or inter-atomic distances in the case of chemical molecules. We always assume that reaction coordinates are available that change strictly monotonously during the transition, e.g. if a side group of the molecule rotates by 240°, a bond angle is better suited than some inter-atomic distances. Also for the case of two-dimensional spin glasses there are such natural coordinates, see below. If one knows the reaction path fully, $\tau \in [\tau_1, \tau_2]$ just parametrizes the configurations $c(\tau)$ along the reaction path, i.e. the full configurations are parametrized by τ and $E_{\max} = \max_\tau\{H(c(\tau))\}$.

In our case we do *not* know the reaction path in advance, we only have $c(\tau_1) = c_1^{(0)}$, $c(\tau_2) = c_2^{(0)}$. Our approach here is to find a good approximation of the minimum reaction path by first selecting one "physical" reaction coordinate τ which characterizes the transition very well, e.g. a bond angle, which changes most during the transition. One can now get a rough idea of the transition, if one varies now τ between τ_1 and τ_2 and generates a finite set of values $\tau \in \{\tau_1 = \tau^0, \tau^1, \ldots, \tau^{k-1}, \tau^k = \tau_2\}$. For each of these values, using *opt* we then obtain minimum energy configurations c^i $(i = 1, \ldots, k)$ with a minimization under the constraint $\tau = \tau^i$. Hence, we get a set of energies $H(c^i)$, which is a coarse-grained picture of the energy landscape encountered during the reaction path, see upper figure in Fig. 1. This coarse-grained path is determined by the set of constraints C^i which fix the path at $\tau = \tau^i$. Since the description of the path via constraints is more general, and in particular more suitable for our spin-glass problem, we will use this description. The reaction path so obtained is a reasonable good approximation for the true minimum reaction path, because a ground state was calculated (under the given constraints) for each intermediate point. This is the *first basic idea* of our approach. This first coarse-grained path we denote as the path on *level one*. Nevertheless, the quality of the result depends on a good choice of τ, which might be difficult for molecules, but is straightforward for the case of two-dimensional spin glasses.

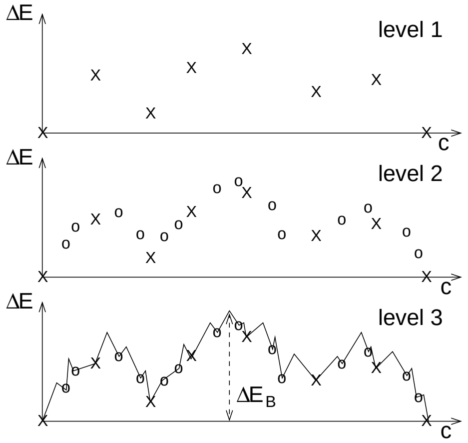

FIG. 1: Excitation energies ΔE for sequences of configurations c generated by the algorithm.

Now, it might be the case for two adjacent configurations c^i, c^{i+1}, that these configurations are again separated by a high energy barrier, e.g. another sub-group of the molecule might first need to be rotated aside. For continuous degrees of freedom this becomes obvious, if the interpolation between the two configurations leads to ridiculous high energies (such as would occur if one attempted to move some atoms through other atoms). For discrete degrees of freedom, when assuming a dynamics consisting of changes of single variables, an energy barrier might occur, if c^i, c^{i+1} differ by more than one variable. In general for such a case, one can apply the same algorithm recursively for all adjacent pairs c^i, c^{i+1} where it is necessary. This is the *second basic idea* of our approach. Hence, one obtains the reaction path with higher level of details, i.e. a reaction path on *level two*, and also a refined picture of the energy landscape, see middle figure in Fig. 1. Note that within the optimization the constraints which have appeared on the upper level will be also suitably included, i.e. $\tau \in [\tau^i, \tau^{i+1}]$ when looking at the case $c^i \to c^{i+1}$.

The whole approach is iterated on lower levels, leading to finer and finer descriptions of the reaction path and of the energy landscape along the path, until one is sure that between two adjacent configurations there is no high energy barrier, see lower figure in Fig. 1. For the case of systems with discrete degrees of freedom, as it is the case for the two-dimensional spin glasses, this is achieved when one arrives at the level of single-variable changes, i.e. single-spin flips here.

To formally summarize, the algorithm is based on two fundamental operations. Note that for the following description we again describe the restriction of the accessible configuration space via sets of constraints. Hence sequences of configurations along a reaction path are given in terms of sequences of constraints, each constraint being equivalent to the configuration obtained after a minimization under the constraints:

- *seq*(c_1, c_2, C) finds sequences $\{C^1, C^2, \ldots, C^k\}$ of constraints interpolating between the two given configurations c_1, c_2 under an already given set C of constraints (empty when called on level one). In general, the number k of elements of the sequence, i.e. of interpolation steps along the reaction path, may also be determined by *seq*, which is possible for problems having discrete degrees of freedom, as for Ising spin glasses. For continuous degrees of freedom, one would have to limit or preselect the number of elements.

  Note that the choice of *seq* determines mainly the quality of the reaction path. The better *seq*, the better the final result.

- *opt*(C') finds the optimum energy of the system for the given set C' of constraints. In case the optimization problem is hard, e.g. for proteins and other molecules, this could be a heuristic as mentioned above.

The algorithm hierarchically constructs a path in configuration space. It first calls *seq* for $c_1^{(0)}, c_2^{(0)}$ without any constraints. For each set C^l of constraints returned by *seq*, *opt*$(C \cup C^l)$ is called (C are the constraints from the previous level, modified slightly to allow for variations between certain limits as discussed above), resulting each time in optimum configurations under the given constraints. This sequence of configurations is the first level of the hierarchy, see again Fig. 1. The reaction coordinate chosen on level one *must* change from τ_1 to τ_2 when moving from $c_1^{(0)}$ to $c_2^{(0)}$ also for *all* other possible reaction paths. Since there are no other constraints present on level one, and because the configurations are obtained from a (presumably exact) optimization, *any* reaction path must overcome an energy barrier at least as high as the energy barrier encountered on level one. Hence one has obtained already a lower bound for the true minimum barrier energy.

To construct the path in finer detail, for each consecutive pairs of constraints obtained on the first level the same procedure is applied again. This results in a more detailed path in configuration space (level 2). The process is continued hierarchically until one arrives on the level of changes of single degrees (for discrete systems) or when one has the desired level of detail (continuous systems), i.e. when linear interpolation is feasible. The highest energy encountered on the final level is an upper bound for the true minimum barrier. First, it is an upper bound, because *any* reaction path provides an upper bound. Second, this bound is not tight, because the algorithm cannot guarantee that e.g. the choices made on level one do not lead to higher energy barriers on lower levels. It could be that if one chooses a path with comparatively higher energy on level one, that this leads to a path where the final energy barrier is comparatively lower. The true minimum barrier can presumably only be found by considering basically all possible paths.

B. Two-dimensional spin glasses

To show how the algorithm can be implemented explicitly we apply it to two-dimensional Ising spin glasses. The model consists of $N = L^2$ spins $S_i = \pm 1$ on a square lattice with free boundary conditions in both directions. The Hamiltonian is

$$H = -\sum_{\langle ij \rangle} J_{ij} S_i S_j \,, \tag{1}$$

where the sum runs over all pairs of nearest neighbors $\langle ij \rangle$ and the J_{ij} are the quenched random spin couplings. We will consider a Gaussian distribution of couplings with zero mean and unit width, which guarantees that the system just has two GSs $\{\sigma^0\}$ and $\{-\sigma^0\}$. We define the barrier to be the energy required to invert all spins in the GS within the context of single spin flip dynamics, as in[9], i.e. the reaction path goes from $\{\sigma^0\}$ to $\{-\sigma^0\}$. Each trajectory of reversal is characterized by the highest maximum in the trajectory having energy ΔE_B, in excess of the GS energy, and the minimum barrier height $E_B = \min(\Delta E_B)$, where the minimum is over all the possible trajectories.

All spin configurations can be described by the set of spins having $\{\sigma^0\}$ orientation, while the remaining spins have the $\{-\sigma^0\}$ orientation. Hence, each configuration is equivalent to a set of domain walls (DWs) separating the two GS orientations. Since these so defined DWs have always positive energies, our search for low-energy paths can be restricted to configurations which are separated by just *one* DW. Therefore, each configuration is described by one DW w here. The path from $\{\sigma^0\}$ to $\{-\sigma^0\}$ will be equivalent to a DW sweeping – say – from left to right through the system via single spin flips. For our implementation, we will allow only that each spin flips exactly once along the reaction path, i.e. the DW cannot fluctuate back and forth. This is another major approximation and as a consequence we can finally only obtain an upper bound for the minimum barrier height ΔE_B instead of the true one.

The basic idea of the implementation of the approach for two-dimensional spin glasses is to use the position of one part of the DW wall as reaction coordinate. For level one of the energy-landscape hierarchy, we fix ("pin") the DW - say - at different positions at the top boundary of the system and move this position from left to right through the system. The remaining part of the DW is allowed to adjust freely such that its energy becomes minimal. One higher levels of the iteration, a DW is moved between two given DWs which are adjacent on the preceeding level. The process is repeated hierarchically until a sequence of N single spin flips from $\{\sigma^0\}$ to $\{-\sigma^0\}$ is obtained.

Now we describe the *seq* and *opt* operations. We start with the *opt* operation, because the algorithm we use imposes some restriction on the kind of constraints we can use to describe the reaction path. Hence, the nature of these constraints then determines the way we implement the *seq* operation.

The *opt* operation consists here of the calculation of the GS $\{\sigma^0\}$. This can be done in a time, which grows polynomially in the number N of spins, by mapping the GS problem to a graph-theoretical problem, the so-called *minimum-weight perfect matching* problem. For the details of this well-established method, see Refs. 6,12. The very first step for the treatment for each realization of the disorder is to calculate its true GS $\{\sigma^0\}$ without any constraints. Then $\{\sigma^0\}$ and $\{-\sigma^0\}$ are the start and end points of the reaction path to be constructed.

The most natural way to realize constraints is to fix the orientations of some spins. Unfortunately, this is not possible within the matching approach, because it does not work in the presence of local or global magnetic fields. The only thing one can do is to change some of the interaction bonds J_{ij}. Hence, constraints are included here by introducing *hard bonds*, which are bonds of very high magnitude, e.g. $|J_{ij}| = 10^9$ *replacing* the original bonds between neighboring spins i, j. The high magnitude of the bonds ensures that they are satisfied in all following calculations of configurations. Note that hard bonds fix only the *relative orientation* of two adjacent spins, not the absolute orientation of single spins. Here, as a general rule, the signs of *most* hard bonds are chosen such that they enforce the same relative orientation of the two spins as in the original GS. Furthermore, to induce DWs with respect to the GS, a few of the hard bonds will be *inverted*, i.e. they ensure a relative orientation of adjacent spins, which is opposite to the GS orientation. As we will see below, the remaining part of the algorithm (i.e. the operation *seq*) consists mainly of choosing which hard bonds to introduce, and which hard bonds to invert.

To understand the way the hard bonds are introduced, we look at the following example. Assume that we have calculated the GS of a given realization. Now we introduce a horizontal line of hard bonds connecting all spins in the top row. If we calculated with this changed realization a new GS, we would get exactly the same GS again. To get something new, we furthermore invert exactly one of the hard bonds, say the bond J_{l_0,l_0+1} somewhere in the middle of the top row, see Fig. 2. Now we compute the GS of this changed realization. The inverted hard bond guarantees that spins l_0 and l_0+1 have opposite relative orientations to the original GS. Hence, a DW originates between spins l_0 and l_0+1. The other (non-inverted) hard bonds guarantee that all other spins at the top row have the same relative orientations as in the GS. This means, the DW has to leave the system at one of the other boundaries, e.g. at the bottom row as shown in the example. Since the DW is obtained from another GS calculation, its energy is minimized under the constraint that its starting location is pinned between spin l_0 and l_0+1 and that it does not terminate at the top row. Everything else will be adjusted such that the energy is minimal. In the resulting configuration, all spins left of the DW will have one GS orientation and all spins right of the DW the opposite GS orientation.

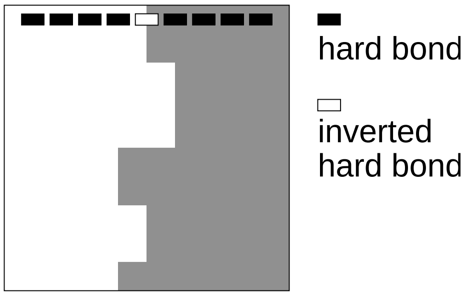

FIG. 2: Example of the generation of a DW via introducing hard bonds (small black rectangles) and one inverted hard bond (small white rectangle), after the GS of the system has been obtained. The following GS calculation including the hard bonds generates two domains, one where the spins have the GS orientation (white area), one where the spin have the opposite GS orientation (dark area). The DW separates these domains and is forced to run through the inverted hard bonds, and it must not run through the other hard bonds.

The basic idea to construct the reaction path on level one is to consider all possible positions of the inverted hard bond along the horizontal line, running from left to right, such that the DW is swept from left to right as well. This generates a sequence of DWs, see Fig. 3, describing the path on level one. The refinement of the path and the energy landscape is obtained by recursively finding sequences of DWs on level $n+1$ which interpolate between adjacent DWs on level n again using straight lines along which a inverted hard bond is moved. The only additional generalization we apply is that we do not only interpolate between DW configurations along horizontal, but also along vertical lines. Furthermore, we also remark that the energies obtained on level one are not necessarily true lower bounds on the barrier height, because we choose to sweep the DW from left to right, but it might be better to sweep it from bottom to top, or even diagonally. Nevertheless, from symmetry arguments, the results we denote as "lower bound" should be very close to true lower bounds.

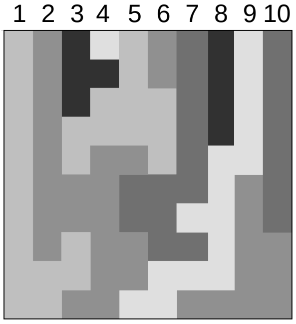

FIG. 3: A sample sequence of DWs generated by considering all different positions of the inverted hard bond along the horizontal line, corresponding to level one of calculation of the reaction path. The different shades regions show the different parts of the domains flipped when the ith bond is inverted. Hence, for the inverted bond at the leftmost position (connecting spins 1 and 2), the resulting GS exhibits a domain consisting of part 1 with one GS orientation, and another domain consisting of parts 2-10 with the other GS orientation. For the inverted bond connecting spins 2 and 3, the two domains consist of parts 1,2 and 3-10, respectively, etc.. The sequence of DWs generated in this way describes on a coarse-grained level the order in which one moves from one GS to the other GS.

Note that also the initial situation on the first level of the above described hierarchy, i.e. the GSs $c_1^{(0)} = \{\sigma^0\}$ and $c_2^{(0)} = -\{\sigma^0\}$ can be described in the DW picture via two straight-line DWs being located at ("behind") the left and right borders of the system, respectively. One can imagine that these initial configurations correspond to the situations where the DW has not yet entered the system, and where the DW has finally left the system, respectively. Hence, the task on all level of the hierarchy is to interpolate between two DWs.

This means in general, two domain walls w_1, w_2 are passed to *seq*. Furthermore, we pass constraints C such that the bonds left of w_1 and right of w_2 are hard bonds (i.e. no hard bonds at all on level one). These hard bonds cause the spins left of w_1 and right of w_2 to be frozen independently into relative GS orientations. Thus, all subsequent DWs calculated on higher levels will run between w_1 and w_2. This guarantees in the end that the domain wall does not fluctuate back and forth. Technically, it also accelerates the subsequent GS calculations considerably, because, hard bonds do not have to be considered during the minimization. The basic idea is now, to generate a (partial) reaction path interpolating between w_1 and w_2, i.e. to sweep a DW from w_1 to w_2 from left to right, by pinning it at one location through an inverted hard bond and moving the inverted

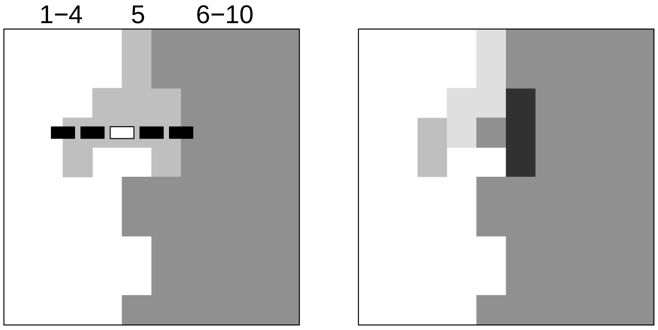

FIG. 4: Example of the operation of the algorithm. Left: two adjacent domain walls (generated on level one) shown together with the line of hard bonds (horizontal black line) generated by *seq*, with the inverted hard bond (white part of line) (i_k, i_{k+1}), corresponding to C^k. The hard bonds left of the left DW and right of the right DW are not shown for clarity. After a new GS calculation with additional constraints C^k, a DW will result, which runs between the two shown DWs and exactly through the inverted hard bond. The resulting sequence of "interpolating" domain walls obtained after the calls to *opt* could look like that on the right, where the different domains are shown in different shades for better visibility.

hard bond such that it moves the DW from w_1 to w_2. This is achieved here by *seq* by determining the longest straight (horizontal or vertical) separation between w_1 and w_2 which lies completely inside the region without hard bonds. Let $k+1$ be the length of this separation. To construct sequences of constraints leading to this behavior, we add (inverted) hard bonds, such that the DW is effectively moved. Formally, we do this by denoting with $i_1, \ldots, i_k, i_{k+1}$ the spins along this line, and by i_0 and i_{k+2} the spins which are reached when extending the separation in both directions by one spin. Then, *seq* finally returns the sequence $\{C^1, \ldots, C^k\}$ of constraints, which are defined as follows: C^l $(l = 1, \ldots, k)$ consists of the hard bonds $\{(i_0, i_1), (i_1, i_2), \ldots, (i_k, i_{k+1}), (i_{k+1}, i_{k+2})\}$[13] where exactly the bond (i_l, i_{l+1}) is inverted, the others are not inverted, see Fig. 4.

Thus, for the subsequent calls to *opt* with the constraints $C^l \cup C$, each time a domain wall is generated, which lies in the region between w_1 and w_2 and is forced by the inverted hard bond to run exactly between spins i_l and i_{l+1}, the DW is pinned between i_l and i_{l+1}. Going with l from 1 to k, each time calculating the GS, results in a sequence of domain walls interpolating between w_1 and w_2. For all pairs of adjacent domain walls generated in this way, the same procedure is continued recursively at the next level.

The hierarchical procedure stops for all pairs w_1, w_2 which are separated only by a line of spins of width 1. For this case one can just flip in the region between w_1 and w_2 one spin after the other starting at one end. Hence, the full algorithm finally outputs a sequence of single spin flips corresponding to the movement of a domain wall from one end of the system to the other.

For each flipped spin the energy of the configuration is calculated, and at the end the maximum $\Delta E_{\max}$ among the N energies is considered. An example of the corresponding energy landscape sampled by this sequence is shown in Fig. 5.

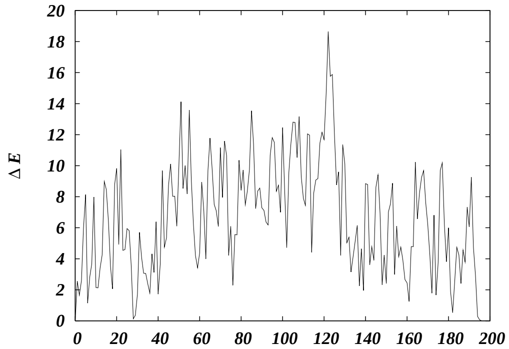

FIG. 5: An example of the energy landscape explored by our algorithm for $L = 14$; the exponent ψ is associated with the scaling of the maximum energy found in the sequence.

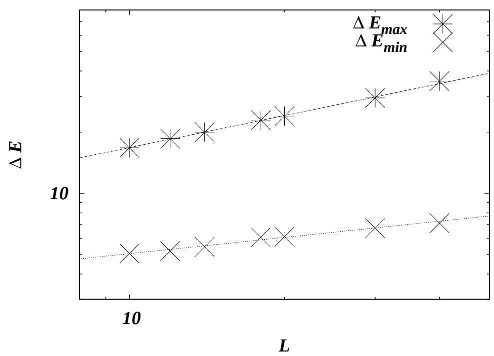

FIG. 6: Upper line: average barrier energy as obtained by the algorithm, being an upper bound for the true lowest barrier height. Lower line: lower bound for the barrier energy as obtained from the first level of the hierarchy.

III. RESULTS

We applied our method for system sizes in the range $8 \leq L \leq 40$, and for each L used 1000 independent realizations of the bonds J_{ij}. The two averaged quantities $\langle \Delta E_{min} \rangle$ (the lower bound, obtained from the first level) and $\langle \Delta E_{max} \rangle$ (the upper bound) as a function of system size are shown in Fig. 6. Both quantities can be fitted by an algebraic function $\sim L^{\psi}$, where $\psi_{lw} = 0.25 \pm 0.01$ and $\psi_{up} = 0.54 \pm 0.01$. There could of course be a procedure to obtain a better path in configuration space, yielding a barrier which grows more slowly than $L^{0.54}$. Anyway, our approach is to our knowledge the only numerical estimate of the energy barriers exponent giving an upper bound different from the trivial one $d-1$. Note that the exact values for small sizes, which were also restricted to each spin flipping once[9], lie basically between our upper and lower bound, so it is not clear which bound is the better approximation.

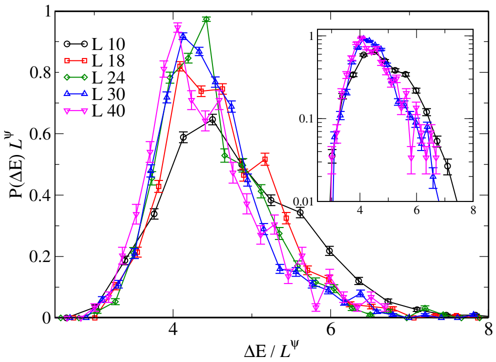

FIG. 7: (Color online) Rescaled probability distribution of the upper bound for the barriers heights corresponding to $\psi = 0.54$. The inset shows the same data (but not for all system sizes, for clarity) with logarithmic y-scale.

There are experimental estimates of ψ for two-dimensional Ising spin glass systems in the literature. Dekker et al.[14] reported experimental verification of activated dynamics in a $d = 2$ system obtaining $\psi = 0.9$. Schins et al.[15] find $\psi = 1.0 \pm 0.1$ by studying aging via the low-frequency ac susceptibility giving credence to the claim that ψ is equal to its upper bound. The numerical study of Gawron et al.[9] gave $\psi = 0.9 \pm 0.1$ in 2 dimension, but that of Berthier and Bouchaud[16] had $\psi \sim 1.0$ and $\psi \sim 2.3$ in 3 and 4 dimensions respectively. We remark that our approach has the advantage that is does not suffer from equilibration problems and it allows the study of large systems at least in two dimensions. Note that the algorithm finds energically low-lying paths because, by iteratively applying the matching algorithm, we are producing large scale changes to the domain wall, resulting in a value $\psi < d - 1$. We have also studied the barrier associated with a straight domain wall and found that the expected $B \sim L$ behavior is approached by $L \geq 40$, being sublinear for $L < 40$.

We also studied the distributions of the barriers heights over different random samples. In Fig. 7 the probability distributions rescaled by L^{ψ} for the upper bounds is shown for different sizes L, while Fig. 8 shows the same data for the lower-bound results. In both cases the result is compatible with an approach to a fixed shape in the thermodynamic limit. In particular this means that the barrier energy is not self-averaging because the width of the rescaled distributions does not decrease with growing system size. Note that both these distributions show the existence of some very large barriers in some samples, see insets of Figs. 7 and 8. Nevertheless, the distributions decrease faster than exponentially in the tails.

Finally, we also studied geometrical properties of the domain walls at the energy barriers. The domain walls

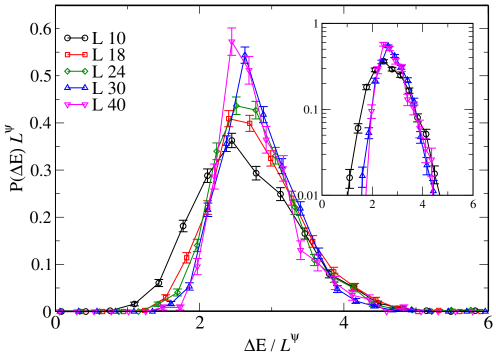

FIG. 8: (Color online) Rescaled probability distribution of the lower bound for the barriers heights corresponding to $\psi = 0.25$. The inset shows the same data (but not for all system sizes, for clarity) with logarithmic y-scale.

are fractal with a an average length of order $P \sim L^{d_S}$, with $d-1 \leq d_S \leq d$. For the small sizes accessible within this framework, we obtained $d_S = 1.273 \pm 0.007$ for the upper bound domain walls, see Fig. 9. This is very close to the values found in previous work for domain walls of minimum energy[17–20]. For the domain wall representing the lower bounds, we obtained in the same way $d_S = 1.278 \pm 0.004$.

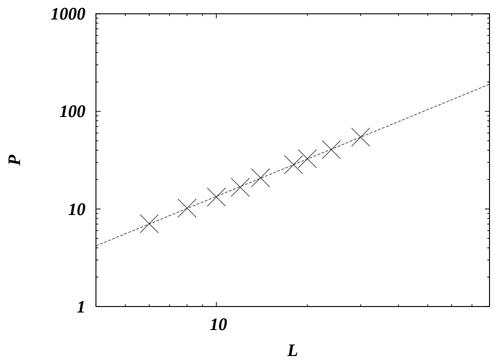

FIG. 9: The perimeter (length) of the surface of the domain wall as a function of system size. corresponding to the upper bound (maximum) energy. Error bars are smaller than symbol sizes. Data are fitted by a power law $P \sim L^{d_S}$ with $d_S = 1.273 \pm 0.07$.

IV. CONCLUSIONS

To conclude we have introduced a hierarchical algorithm to compute upper and lower bounds on energy barriers. Applying the algorithm to two-dimensional Ising spin glasses, we find that the minimum barrier energy is bounded above by the scale $L^{0.54}$ and below by $L^{0.25}$. The power of our approach becomes visible by noting we can study much larger system sizes than before and by the fact that our numerical upper bound is significantly less than the rigorous upper bound (and the value for the hierarchical lattice), $d-1$. Hence, we suspect that ψ is an entirely non-trivial exponent. At the present time there seems to be no theoretical approach available which might explain a non-trivial value of ψ.

Acknowledgments

We thank A. Bray for critically reading the manuscript. We have obtained support from the *VolkswagenStiftung* (Germany) within the program "Nachwuchsgruppen an Universitäten", from the Paderborn Center for Parallel Computing in Germany, and from the European Community via the DYGLAGEMEM, High-Level Scientific Conferences (HLSC) and the "Exystence" programs.

[1] Reviews on spin glasses can be found in: K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986); M. Mezard, G. Parisi, M. A. Virasoro, *Spin glass theory and beyond*, (World Scientific, Singapore 1987); K. H. Fischer and J. A. Hertz, *Spin Glasses*, (Cambridge University Press, Cambridge 1991); A. P. Young (ed.), *Spin glasses and random fields*, (World Scientific, Singapore 1998).

[2] Reviews on glasses can be found in C. A. Angell, Science **267**, 1924 (1995); P.G. DeBenedetti, *Metastable liquids*, (Princeton University Press, 1996); S. Franz, S. C. Glotzer, and S. Sastry (eds.), *Unifying Concepts in Glass Physics*, J. Phys. Cond. Matt. **12**, 6295-6682 (2000)

[3] T. E. Creighton (ed.), *Protein folding*, (Freeman, New York 1992).

[4] A. A. Middleton, Phys. Rev. E **59**, 2571 (1999).

[5] M. R. Garey and D. S. Johnson, *Computers and intractability* (Freeman, San Francisco, 1979).

[6] A. K. Hartmann and H. Rieger, *Optimization Algorithms in Physics*, (Wiley-VCH, Berlin 2001).

[7] D. S. Fisher and D. A. Huse, Phys. Rev. B **38**, 373 and 386 (1988).

[8] L. V. Mikheev, B. Drossel, and M. Kardar, Phys. Rev. Lett. **75**, 1170 (1995); B. Drossel, J. Stat. Phys. **82**, 431 (1996).

[9] T. R. Gawron, M. Cieplak and J. R. Banavar J. Phys. A:Math. Gen. **24** L127 (1991)

[10] B. Drossel and M. A. Moore, Phys. Rev. B **70**, 064412 (2004).

[11] A. K. Hartmann and M. A. Moore, Phys. Rev. Lett. **90**, 127201 (2003).

[12] I. Bieche, R. Maynard, R. Rammal, and J. P. Uhry, J. Phys. A **13**, 2553 (1980); F. Barahona, R. Maynard, R. Rammal, and J.P. Uhry, J. Phys. A **15**, 673 (1982); U. Derigs and A. Metz, Math. Prog. **50**, 113 (1991).

[13] If one of the "extra" spins i_0 or i_{k+2} does not exists, the corresponding bonds do not occur. This happens always on the first level, where the separation spans the entire system. i.e. in this case the line of hard bonds starts at i_1 and ends at i_{k+1}.

[14] C. Dekker, A. F. M. Arts, H. W. de Wijn, A. J. van Duyneveldt, and J. A. Mydosh, Phys. Rev. Lett. **61**, 1780 (1988).

[15] A. G. Schins, E. M. Dons, A.F. M.Arts, H. W. de Wijn, E. Vincent, L. Leylekian and J. Hammann, Phys. Rev. B **48**, 16524 (1993).

[16] L. Berthier and J-P. Bouchaud Phys. Rev. B **66**, 054404 (2002).

[17] A.J. Bray and M.A. Moore, Phys. Rev. Lett. **58**, 57 (1987).

[18] H. Rieger, L. Santen, U. Blasum, M. Diehl, M. Jünger, and G. Rinaldi, J. Phys. A **29**, 3939 (1996).

[19] M. Palassini and A.P. Young, Phys. Rev. B **60**, R9919 (1999).

[20] A.A. Middleton, Phys. Rev. B **63**, 060202(R) (2001).