

Backtracking Dynamical Cavity Method

Freya Behrens¹,[✉] Barbora Hudcová,^{2,3} and Lenka Zdeborová¹

¹*Statistical Physics of Computation Laboratory,*

École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

²*Algebra Department, Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic*

³*Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University, Prague, Czech Republic*



(Received 31 March 2023; revised 21 June 2023; accepted 24 July 2023; published 21 August 2023)

The cavity method is one of the cornerstones of the statistical physics of disordered systems such as spin glasses and other complex systems. It is able to analytically and asymptotically exactly describe the equilibrium properties of a broad range of models. Exact solutions for dynamical, out-of-equilibrium properties of disordered systems are traditionally much harder to obtain. Even very basic questions such as the limiting energy of a fast quench are so far open. The dynamical cavity method partly fills this gap by considering short trajectories and leveraging the static cavity method. However, being limited to a couple of steps forward from the initialization, it typically does not capture dynamical properties related to attractors of the dynamics. We introduce the backtracking dynamical cavity method that instead of analyzing the trajectory forward from initialization, it analyzes the trajectories that are found by tracking them backward from attractors. We illustrate that this rather elementary twist on the dynamical cavity method leads to new insight into some of the very basic questions about the dynamics of complex disordered systems. This method is as versatile as the cavity method itself, and we hence anticipate that our paper will open many avenues for future research of dynamical, out-of-equilibrium properties in complex systems.

DOI: [10.1103/PhysRevX.13.031021](https://doi.org/10.1103/PhysRevX.13.031021)

Subject Areas: Complex Systems, Statistical Physics

I. INTRODUCTION

The cavity method is one of the main analysis tools to investigate equilibrium properties of disordered and complex systems. It has been introduced in a series of seminal works as an alternative to the replica method for mean-field models of spin glasses [1]. Subsequent key developments on diluted lattices also known as sparse random graphs [2,3] and the link between the cavity method and message-passing algorithms [4,5] have led to an explosion of applications of the method in systems on sparse random structures, such as error-correcting codes, random constraints satisfaction problems, random graphs coloring, or community detection to mention just a few of many; see, e.g., the textbook [6]. Results obtained using the cavity method are in many cases exact in the thermodynamic limit which is particularly appealing for theoretical studies in computer science and mathematics.

Many questions about complex systems of current interest are, however, not concerned with equilibrium properties but with dynamical, out-of-equilibrium ones.

An exact analysis of dynamical properties is much more challenging compared to the equilibrium ones. Let us give two concrete examples of some very basic questions about dynamics that are so far open and that the method proposed in this paper resolves.

Example 1: Consider the antiferromagnetic Ising model or a spin glass with random ± 1 interactions on a random d -regular graph of n nodes. Consider then the dynamics where at each time step every spin aligns with their magnetic field or remains in case the field is zero. We initialize each spin randomly. To which value of energy does such a dynamics converge at large times when $n \rightarrow \infty$?

Example 2: Consider now the ferromagnetic Ising model on a random d -regular graph, the same dynamical process but initialized at magnetization $-1 < m < 1$. For what values of $m \geq 0$ does the dynamics go to the homogeneous all $+1$ configuration, and for what values of $m \geq 0$ does it go elsewhere when $n \rightarrow \infty$? What other attractor types does the dynamics converge to for other values of m ?

While these are very basic questions that could be studied numerically in an undergraduate class on statistical mechanics, the asymptotically exact answer is so far not known even for random graphs for which many static properties are known exactly in the thermodynamic limit via the cavity method [1–3].

The main contribution of this paper is to present a method to answer dynamical questions such as the above

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

by quantifying the basin of attraction of different types of attractors for deterministic dynamics. We call it the backtracking dynamical cavity method (BDCM). This method provides a solution in the sense that for models on random graphs in the limit $n \rightarrow \infty$ it gives a closed-form analytical prescription of how to compute the desired values. This leads, for instance, to the value of the limiting energy from Example 1 for random regular graphs and to exhibiting different types of attractors and dynamical phase transitions between them in Example 2.

The main idea behind the backtracking dynamical cavity method is simple. We start with the established idea of the dynamical cavity method (DCM) [7–11] that considers the trajectory of a spin for a finite number of time steps T . It considers this trajectory as an augmented T -dimensional spin variable and applies the traditional static cavity method to this trajectory variable. The dynamical cavity method provides an exact description of the dynamics as long as the system is large $n \rightarrow \infty$ and the time $T = O(1)$ finite. Evaluating the corresponding equations is in general exponentially costly in T and thus limits the choice of T . Consequently, properties that require not-so-small values of T cannot be analyzed using this method unless one resolves to approximations.

The key twist in the backtracking dynamical cavity method is that instead of taking T steps from the beginning of the trajectory, we take T steps from the attractor (thus, tracking the dynamics back). This way, we can access properties of the attractors and their basins of attraction T steps back in an asymptotically exact manner without further approximations. We show that by exploring the last $T = O(1)$ steps of the dynamics, the backtracking dynamical cavity method is able to provide answers to the two examples posed above. What comes as a surprise to us, is that looking at T step backward covers a basin of attraction of entropy (logarithm of the number of configurations in that basin) very close to the total entropy of all initial conditions already for very moderate values of T .

While the existing DCM is able to access properties that happen in the first few steps of the dynamics [7–11], and with approximations it is also able to describe qualitatively correctly large time behavior even for local observables [12–17], it does not provide asymptotically exact results about the attractors of the dynamics or their basin of attraction. BDCM does exactly that, describing the last steps of the dynamics. Moreover, as we see in examples below, only a few steps back into the basin of attraction may already exhibit qualitative properties of the complete basin of attraction. We illustrate this, in particular, on the majority rule where the types of attractors found from initial configurations with different magnetizations already show when we step into the basin of attraction by only one step.

Finally, we want to emphasize that the idea of looking at the last T steps of the dynamics rather than the first T steps

is very generic, and open questions about the properties of attractors are abundant in the study of dynamics of complex systems. We thus anticipate that the BDCM will become one of the key analytical methods in the field. Possible applications include training dynamics of artificial neural networks where we would want to study the basin of attraction of a region with good generalization properties, social dynamics on networks where we may want to know what type of Nash equilibria will be reached, gene regulatory networks where attractors correspond to cell types, or various types of far-from-equilibrium physical systems where different attractors may correspond to different phases. The backtracking idea can be applied not only in conjunction with the dynamical cavity method, but also, for instance, within the dynamical mean-field theory [18] that has been influential in the study of strongly correlated electron systems or neural networks.

II. SETTING AND NOTATION

By an undirected graph of size n , we understand the tuple $G = (V, E)$ where $V = \{1, \dots, n\}$ is the set of nodes and $E = \{(i, j) | i, j \in V\}$ is the set of edges. For each node $i \in V$, we define the neighborhood of i to be the set $\partial_i = \{j | (i, j) \in E\} \subseteq V$ with the degree of i as $d(i) = |\partial_i|$. We say a graph is d -regular if each node has degree d . Each node i of the graph G can be assigned in one of the discrete states in a set S , $x_i \in S$. Such an assignment then represents a configuration $\mathbf{x} = x_1, \dots, x_n \in S^n$. By \mathbf{x}_{∂_i} , we mean the subset of the configuration that includes all neighbors of node i .

We consider time-discrete dynamical processes operating on configurations of a graph G with n nodes. The state of each node gets updated synchronously, and the update depends on the node's own state and the state of its neighbors. The dynamical rule is specified for each node individually using the local update function $f_i: S^{1+|\partial_i|} \rightarrow S$. This gives rise to a global mapping $F: S^n \rightarrow S^n$ governing the dynamics of the system. For a configuration $\mathbf{x} \in S^n$, the i th node with neighborhood $\partial_i = (i_1, \dots, i_{d(i)})$ gets updated according to

$$[F(\mathbf{x})]_i = f_i(x_i, x_{i_1}, \dots, x_{i_{d(i)}}).$$

To describe the global dynamics, the symbol $\underline{\mathbf{x}}$ denotes a sequence of configurations from S^n ; i.e., $\underline{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^t)$ for some $t \in \mathbb{N}$. We define the configuration graph as an oriented graph whose nodes are the configurations from S^n with edges of the form $(\mathbf{x}, F(\mathbf{x}))$, $\mathbf{x} \in S^n$. If $\underline{\mathbf{x}}$ satisfies that $\mathbf{x}^{i+1} = F(\mathbf{x}^i)$ for each i , we call it the trajectory of length t starting from the initial configuration \mathbf{x}^1 . Since the configuration space is finite, each long enough trajectory becomes eventually periodic. We call the preperiod of the sequence the transient and its periodic part the attractor

or limit cycle. For an attractor, the set of all configurations converging to it is called its basin of attraction.

In this paper, we consider the majority dynamics in models with Ising variables and random ± 1 edge weights (covering the Ising ferromagnet, antiferromagnetic, and a spin glass). Such a dynamics has attractors of length $c \in \{1, 2\}$, which is due to an elegant argument on decreasing energy functions by [19,20]. The number of attractors and short limit cycles for closely related models have been studied, e.g., in Refs. [21–23]. Their basin of attraction has, as far as we know, not been studied analytically, and we use this as an example of applications of the backtracking dynamical cavity method developed in this paper.

III. BACKTRACKING DYNAMICAL CAVITY

A. General idea

The key idea of the BDCM is the fact that it acts on static objects that track the dynamics backward from the attractors instead of forward from arbitrary initial states. To formalize this, we define a (p/c) backtracking attractor to be a trajectory of length p that leads into a limit cycle of length c on the configuration graph. As we increase the length of the incoming trajectory p , such an analysis incorporates a growing fraction of an attractor's basin and will illuminate important dynamical questions.

B. The distribution of backtracking attractors

For a given global update rule F , path length p , and cycle size c , our goal is to analyze the properties of the attractors and their transients. To do this, we introduce a probability distribution over all sequences of configurations $\underline{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^p, \mathbf{x}^{p+1}, \dots, \mathbf{x}^{p+c}) \in (S^n)^{p+c}$ as follows:

$$P(\underline{\mathbf{x}}) = \frac{1}{Z} \mathbb{1}[F(\mathbf{x}^{p+c}) = \mathbf{x}^{p+1}] \prod_{t=1}^{p+c-1} \mathbb{1}[F(\mathbf{x}^t) = \mathbf{x}^{t+1}]. \quad (1)$$

Here, $\mathbb{1}(\cdot)$ is the indicator function which is 1 if the Boolean statement is true and 0 otherwise; Z is the normalization constant of the probability distribution.

A sequence $\underline{\mathbf{x}}$ has only nonzero measure if it is consistent with the time evolution of the global update rule due to the term $\prod_{t=1}^{p+c-1} \mathbb{1}[F(\mathbf{x}^t) = \mathbf{x}^{t+1}]$. The boundary condition $F(\mathbf{x}^{p+c}) = \mathbf{x}^{p+1}$ ensures that this trajectory of configurations ends up in a limit cycle of length c . Consequently, only (p/c) backtracking attractors can have a nonzero measure in the distribution (1).

Analogous to the classical cavity method for statical analysis, the goal is then to compute the free entropy density $\Phi = (1/n) \log(Z)$, i.e., the logarithm of the number of sequences that are valid backtracking attractors. Then, Φ can be viewed as a proxy for the size of an attractor's basin.

C. Adding observables

A key virtue of the BDCM is that we can obtain this entropy Φ conditioned on backtracking attractors with specific properties, e.g., fixed energy or magnetization or magnetization in the attractor. This conditioning can be achieved by flexibly weighting the sequences $\underline{\mathbf{x}}$ in the probability distribution according to the relevant observable. Concretely, one adds the factor

$$e^{-\sum_k \lambda_k \Xi_k(\underline{\mathbf{x}})} \quad (2)$$

on the right-hand side of Eq. (1) and adjusts the normalization Z accordingly; exactly K observables $\Xi_k(\underline{\mathbf{x}})$ are added as summary statistics of the backtracking attractors. Each observable has an associated parameter $\lambda_k \in \mathbb{R}$ which acts as a temperature from a physics perspective, or as a Lagrangian multiplier viewing the observables as optimization constraints. We use the notation uppercase $\Xi_k(\underline{\mathbf{x}})$ for the function acting on a trajectory $\underline{\mathbf{x}}$. The notation lowercase ξ_k is used when the value of $\Xi_k(\underline{\mathbf{x}})/n$ is fixed to ξ_k , usually as a constraint and intensive quantity. Formally, we define the number of valid backtracking attractors conditioned on fixed observables as $\mathcal{N}(\xi_1, \dots, \xi_K) = e^{ns(\xi_1, \dots, \xi_K)}$ so that s is their entropy. Then, the following relation between the entropy s and the normalization constant Z including the extra factor (2) holds

$$Z = e^{n\Phi(\lambda_1, \dots, \lambda_K)} = \int_{\underline{\mathbf{x}} \in B} e^{-\sum_k \lambda_k \Xi_k(\underline{\mathbf{x}})} \quad (3)$$

$$= \int \left[\prod_k d\xi_k \right] e^{n[s(\xi_1, \dots, \xi_K) - \sum_k \lambda_k \xi_k]}, \quad (4)$$

where the set B is the set of all valid (p/c) backtracking attractors. In the large system limit, when $n \rightarrow \infty$, applying the saddle-point method on the right-hand side gives an explicit form of the entropy

$$s(\hat{\xi}_1, \dots, \hat{\xi}_K) = \Phi(\lambda_1, \dots, \lambda_K) + \sum_k \lambda_k \hat{\xi}_k \quad (5)$$

under the condition that for all $k = 1, \dots, K$,

$$\left. \frac{\partial s(\xi_1, \dots, \xi_K)}{\partial \xi_k} \right|_{\xi_k = \hat{\xi}_k} = \lambda_k, \quad (6)$$

$$\frac{\partial \Phi(\lambda_1, \dots, \lambda_K)}{\partial \lambda_k} = -\hat{\xi}_k = -\frac{1}{n} \langle \Xi_k \rangle. \quad (7)$$

The $\langle \cdot \rangle$ brackets define an average over the probability measure induced by Eq. (1). As it is infeasible to directly compute Z and Φ due to the high-dimensional integral over $\underline{\mathbf{x}} \in \mathbb{R}^{(p+c) \times n}$ when $n \rightarrow \infty$, we compute the leading order

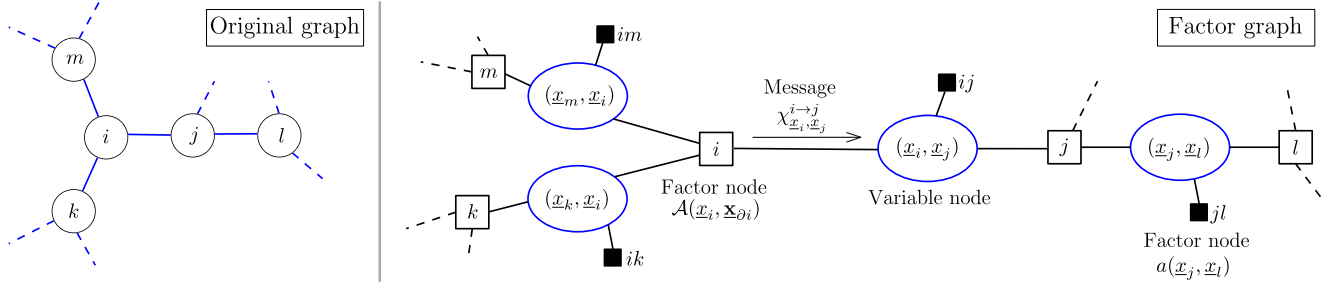


FIG. 1. A subgraph of the factor graph for the BDCM on a six-regular graph. Left: original graph. Right: factor graph in the edge dual space. The tuples in the round variable nodes can take on values of all possible trajectories \underline{x} . The factor nodes on tuples can enforce the constraints on the variables on their own. The factor nodes between d tuples correspond to the consistency constraint of the local update rule f_i . Messages χ are sent back and forth between the nodes.

(in n) of the free entropy using the replica-symmetric cavity method or equivalently belief propagation (BP) [6].

D. Factorization over the graph

For the cavity method to be exact, one requires a probability distribution with a treelike graphical model. To create such a graphical model for our distribution (1), we need two properties to factorize: the global rule F and the observables Ξ .

First, the constraint on the global rule F factorizes on the local node neighborhoods as

$$\mathbb{1}[F(\mathbf{x}) = \mathbf{x}'] = \prod_{i=1}^n \mathbb{1}[f_i(x_i, \mathbf{x}_{\partial i}) = x'_i], \quad (8)$$

which holds since we define F in terms of the local rules f_i .

We assume that the observables can be factorized similarly, i.e., that we can decompose them as a sum over functions on a single node or edge sequences. When $\underline{x}_i = (x_i^1, \dots, x_i^{p+c})$ is the sequence of states of a single node i in \underline{x} we define the node-localized or edge-localized factorization of an observable as

$$\Xi(\underline{x}) = \sum_{i \in V} \tilde{\Xi}(\underline{x}_i), \quad \Xi(\underline{x}) = \sum_{(ij) \in E} \bar{\Xi}(\underline{x}_i, \underline{x}_j). \quad (9)$$

where $\tilde{\Xi}_k: S^{p+c} \rightarrow \mathbb{R}$ and $\bar{\Xi}_k: S^{p+c} \times S^{p+c} \rightarrow \mathbb{R}$. The application examples in our work require four different

observables: the magnetization of the initial configuration m_{init} , the average magnetization in the attractor m_{attr} , the energy of the configuration after t time steps e^t , and for $c \geq 2$ the fraction of changing nodes (rattlers) in the attractor ρ :

$$m_{\text{init}}(\underline{x}) = \frac{1}{n} \sum_{i \in V} x_i^1, \quad (10)$$

$$m_{\text{attr}}(\underline{x}) = \frac{1}{n} \sum_{i \in V} \frac{1}{c} \sum_{t=p+1}^{p+c} x_i^t, \quad (11)$$

$$e^t(\underline{x}) = \frac{1}{m} \sum_{(ij) \in E} x_i^t x_j^t, \quad (12)$$

$$\rho(\underline{x}) = \frac{1}{n} \sum_{i \in V} \mathbb{1} \left[1 \leq \sum_{t=p+1}^{p+c-1} \mathbb{1} [x_i^t \neq x_i^{t+1}] \right]. \quad (13)$$

Each property naturally factorizes either on the nodes or edges. While we do not consider observables that factorize on local neighborhoods $(\underline{x}_i, \underline{x}_{\partial i})$, they can be easily integrated into the framework.

Using these factorizations of F and Ξ_k , the distribution over sequences \underline{x} from Eq. (1) can be factorized over the graph to read

$$P(\underline{x}) = \frac{1}{Z} \prod_{i \in V} \underbrace{\left[e^{-\sum_{\tilde{k}} \lambda_{\tilde{k}} \tilde{\Xi}_{\tilde{k}}(\underline{x}_i)} \mathbb{1} [f_i(x_i^{(p+c)}; \mathbf{x}_{\partial i}^{(p+c)}) = x_i^{p+1}] \right]}_{\mathcal{A}_i(\underline{x}_i, \underline{x}_{\partial i})} \prod_{t=1}^{p+c-1} \mathbb{1} [f_i(x_i^t; \mathbf{x}_{\partial i}^t) = x_i^{t+1}] \prod_{\{ij\} \in E} \underbrace{\left[e^{-\sum_{\bar{k}} \lambda_{\bar{k}} \bar{\Xi}_{\bar{k}}(\underline{x}_i, \underline{x}_j)} \right]}_{a(\underline{x}_i, \underline{x}_j)},$$

where \tilde{k} and \bar{k} are meant to sum only over the observables that are node and edge localized, respectively. This distribution defines a probabilistic model that can be represented as a factor graph where variables are the local

sequences \underline{x}_i . The factors $\mathcal{A}(\underline{x}_i, \underline{x}_{\partial i})$ and $a(\underline{x}_i, \underline{x}_j)$ ensure that only (p/c) attractors have a nonzero probability and are biasing toward a given observable if λ_k is nonzero. However, the implicit factor graph is not locally treelike: If

node i and j are connected by an edge, they appear together in the two factors $\mathcal{A}_i, \mathcal{A}_j$. Hence, for every edge (i, j) , there is a loop of length 4 connecting $x_i \leftrightarrow \mathcal{A}_i \leftrightarrow x_j \leftrightarrow \mathcal{A}_j \leftrightarrow x_i$ (see Fig. 5 in Appendix A). Then, the factor graph is incompatible with an asymptotically exact application of BP. Nonetheless, by moving to the edge dual representation of the graph, these small loops can be eliminated (the resulting factor graph shows in Fig. 1; for examples of a similar dual construction, see, e.g., Refs. [11,22]). In the dual space, the variables of the factor graph are tuples of node trajectories $(\underline{x}_i, \underline{x}_j)$ for all i and j that neighbor on the original graph.

E. BP equations

As a consequence, the factor graph has the same structure as the original graph. This leads to BP fixed-point equations with messages of the form

$$\chi_{\underline{x}_i, \underline{x}_j}^{i \rightarrow j} = \frac{1}{Z^{i \rightarrow j}} a(\underline{x}_i, \underline{x}_j) \sum_{\underline{\mathbf{x}}_{\partial i}} \mathcal{A}_i(\underline{x}_i, \underline{\mathbf{x}}_{\partial i}) \prod_{k \in \partial i \setminus j} \chi_{\underline{x}_i, \underline{x}_k}^{k \rightarrow i},$$

which may be iterated on a given graph until convergence. At convergence, the BP result for the free entropy follows as

$$n\Phi_{\text{BP}} = \sum_{i \in \mathcal{V}} \log(Z^i) - \sum_{(ij) \in E} \log(Z^{ij}), \quad (14)$$

$$Z^i = \sum_{\underline{x}_i, \underline{\mathbf{x}}_{\partial i}} \mathcal{A}_i(\underline{x}_i, \underline{\mathbf{x}}_{\partial i}) \prod_{j \in \partial i} \chi_{\underline{x}_i, \underline{x}_j}^{j \rightarrow i}, \quad (15)$$

$$Z^{ij} = \sum_{\underline{x}_i, \underline{x}_j} a(\underline{x}_i, \underline{x}_j) \chi_{\underline{x}_i, \underline{x}_j}^{i \rightarrow j} \chi_{\underline{x}_j, \underline{x}_i}^{j \rightarrow i}. \quad (16)$$

We can compute the entropy $s(\xi_1, \dots, \xi_K)$ of the number of valid configurations according to Eq. (5), as the constraints in Eq. (7) are fulfilled by the fact that we require the BP messages to have converged; they are satisfied at the fixed point. Note that both the length of the trajectory p and the size of the limit cycle c need to be constant in n , as otherwise, the limit $n \rightarrow \infty$ becomes problematic.

F. Simplification for random regular graphs

The previous equations simplify considerably when we consider regular graphs where all local degrees are d . Furthermore, from hereon we assume that the same local update rule $f: S^{d+1} \rightarrow S$ is used for every node, and we consider only rules that are independent of the neighbors ordering. By this permutation symmetry, all BP messages become the same locally as

$$\chi_{\underline{x}, \underline{y}}^{\rightarrow} = \chi_{\underline{x}_i, \underline{x}_j}^{i \rightarrow j} \quad \forall i, j = 1, \dots, n,$$

and the BP messages are updated according to

$$\chi_{\underline{x}, \underline{y}}^{\rightarrow} = \frac{1}{Z^{\rightarrow}} a(\underline{x}, \underline{y}) \sum_{\underline{\mathbf{y}}_{[d-1]}} \mathcal{A}(\underline{x}, \underline{\mathbf{y}}_{[d-1]}) \prod_{\underline{\mathbf{y}} \in \underline{\mathbf{y}}_{[d-1]}} \chi_{\underline{y}, \underline{x}}^{\rightarrow}, \quad (17)$$

where $\underline{\mathbf{y}}_{[d-1]}$ are the trajectories $(\underline{y}_1, \dots, \underline{y}_{d-1})$ of the $d-1$ neighbors that are relevant for the local update f . The free entropy density can be computed as

$$\Phi_{\text{BP}} = \log(Z^{\text{fac}}) - \frac{d}{2} \log(Z^{\text{var}}), \quad (18)$$

$$Z^{\text{fac}} = \sum_{\underline{\mathbf{y}}_{[d]}} \mathcal{A}(\underline{x}, \underline{\mathbf{y}}_{[d]}) \prod_{\underline{\mathbf{y}} \in \underline{\mathbf{y}}_{[d]}} \chi_{\underline{x}, \underline{y}}^{\rightarrow}, \quad (19)$$

$$Z^{\text{var}} = \sum_{\underline{x}, \underline{y}} a(\underline{x}, \underline{y}) \chi_{\underline{y}, \underline{x}}^{\rightarrow} \chi_{\underline{x}, \underline{y}}^{\rightarrow}. \quad (20)$$

Eventually, this moves iterating $O(n)$ messages on a full graph to iterating only $O(1)$ messages until convergence. In addition, for random regular graphs, there are typically no short loops of length $O(\log n)$, which permits the application of BP in the first place.

G. Solving the equations

For general graphs, the complexity of solving the BDCM equations grows exponentially in $dT = d(p+c)$. Similar to the dynamical cavity method, it is thus prohibitive to analyze exactly long paths p or large cycles c unless one makes approximations [12–15] or restricts oneself to oriented graphs [8], graphs with asymmetrically weighted edges [9], or unidirectional dynamics with absorbing states [11,24]. In this paper, we do not make any such assumptions or approximations. The problems we address in the next section can be solved using the BDCM directly, thanks to the following properties: First, for the considered examples, the cycle size c is in $\{1, 2\}$ [19,20]. Second, we empirically observe that the dynamics converge in logarithmic time of the system size n , so short path lengths p are sufficient to observe interesting properties (see the transient lengths in Fig. 2). Finally, the local update rules are independent of the order of the neighborhood which removes the exponential dependence on d via dynamical programming [16]. Overall, we then obtain a time complexity of $O(d2^{(p+c)})$ per iteration of Eq. (17). Depending on the problem, this allows us to obtain exact results for up to $p \leq 8$ readily.

IV. APPLICATIONS

A. Limiting energy of a quench

As a first application of BDCM, we consider the question posed in the Introduction as Example 1, i.e., the limiting energy of the considered dynamics in the

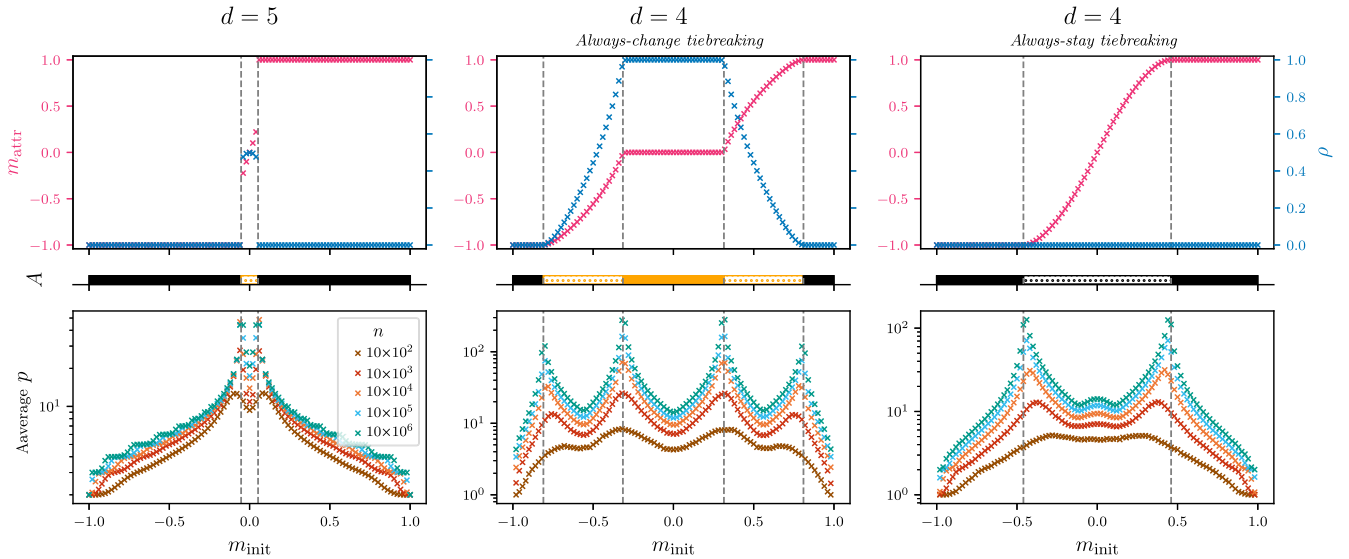


FIG. 2. Empirics of dynamical phase transitions for majority rules. We sample trajectories and attractors starting from random initializations with varying magnetization m_{init} on instances of random regular graphs of varying sizes n . On finite systems, every deterministic dynamics reaches a limit cycle in finite time. Top row: properties of the sampled attractors for $n = 10^5$; the magnetization of the attractor m_{attr} and the fraction of rattlers ρ as a function of m_{init} . Middle row: combined, the values m_{attr} and ρ are linked to the four attractor types as defined in the main text. Bottom row: the average transient length p . The positions of the dynamical phase transitions are marked with dashed lines estimated by the divergences in the transient lengths (see Figs. 7 and 8 in Appendix C).

antiferromagnetic Ising model. We note that due to the universality properties discussed, e.g., in Refs. [22,25], the limiting energy will be the same in the Ising spin glass, i.e., the model with random ± 1 interactions with the zero-temperature dynamics. For simplicity of the exposition of our results, in what follows, let us describe the case of the antiferromagnet.

The Ising antiferromagnet on a random regular graph has energy $e(\mathbf{x}) = (2/dn) \sum_{(ij) \in E} x_i x_j$. This is the energy that an antiferromagnet at zero temperature is naturally minimizing. A quench is a zero-temperature dynamics where at every time step every spin turns in the direction of its local magnetic field or remains unchanged if the local magnetic field is zero. Initializing spins at random, this dynamics starts at zero energy and decreases the energy to a value that we aim to compute. For the antiferromagnet, this process corresponds to a specific instance of a local synchronous update rule where each spin (or node) takes the minority state in its neighborhood.

We now use the BDCM and compute the size of the basin of attraction after a path of length p of all single-point attractors ($c = 1$) that have a given energy $e_p^* := e_p^{p+1}$. In Table I, we report for every path length $p < 8$ the energy e_p^* that maximizes the size of the basin of attraction, i.e., the entropy $e_p^* = \max_e s_p(e)$. We report also the associated maximal entropies $s_p = s(e_p^*)$. Concretely, the results are obtained as follows: We numerically find a solution of Eq. (17) via fixed-point iteration, which in turn gives us the value of Φ_{BP} in Eq. (18). The value of the energy e is then obtained from Eq. (7) using this fixed point. Since the

Lagrangian parameter ξ during the fixed-point iteration is set to zero, the resulting fixed point will be a local maximum of the entropy, and thus the energy is the energy which a typical attractor has. This procedure is equivalent to the maximization of the entropy over the energy but numerically simpler.

To interpret the results, let us first look at the entropies s_p that quantify how many initial configurations end up in a

TABLE I. Normalized entropy $s_p/\log(2)$ of the basins of attraction p steps backward from an attractor obtained from BDCM on the antiferromagnet on a four-regular random graph. Column e_p^* gives the energy of the attractor for which this entropy is reached. Column e'_p gives for comparison the energy of the last configuration of a randomly initialized trajectory after p steps obtained with the DCM. We compare this with the empirically obtained energy of the attractor from graphs of size $n = 10^5$. Table III in the Appendix provides analogous results for $d = 6$.

$d = 4$	BDCM		DCM
p	$s_p/\log(2)$	e_p^*	e'_p
0	0.6026	-0.3616	0.0000
1	0.8679	-0.4759	-0.2812
2	0.9516	-0.5156	-0.4239
3	0.9812	-0.5331	-0.4945
4	0.9925	-0.5411	-0.5261
5	0.9970	-0.5447	-0.5392
6	0.9988	-0.5464	-0.5444
7	0.9995	-0.5471	-0.5463
Empirical \bar{e}_∞			-0.5475(1)

pointlike attractor in p or fewer steps. Remarkably, when stepping away, only $p = 4$ steps backward in time from any point attractor; our analysis of the entropy s_p shows that one can already reach more than 99% of the full entropy of the configuration space. With three additional steps, $p = 7$, the covered fraction is at more than 99.9%. Thus, the size of the basin of attraction for point attractors under this rule quickly encompasses almost all the entropy as p is increased.

We next focus on the value of the energy e_p^* and compare it with the final energy \tilde{e}_∞ obtained numerically on systems of size $n = 10^5$. We note that the considered synchronous dynamics converges to an attractor, and we thus define the stopping time of the simulation as the time when the attractor is reached. We see that e_p^* converges closely to the empirically obtained energy for already very moderate values of p . The value of the energy matches in three digits after only $p = 7$ steps away from the attractor. At this point, the fraction of the basin of attraction covered is up to 99.9% and provides a nice measure of how close to the limiting $p \rightarrow \infty$ result the value e_p^* is.

In Table I, we also compare to the results of the standard forward DCM [7–11] for increasing lengths of trajectories p . We note that as far as we know, this energy has not been evaluated before using the DCM and is thus a result of independent interest. We see that the values of the energies also converge to the empirical value very fast, but slightly slower than the BDCM that we propose here. Moreover, the forward DCM does not come up with the natural measure of convergence provided by the value of the entropy s_p . Overall, this example serves to illustrate the main conceptual differences between the BDCM and DCM in a concise manner. In the Appendix, we report an analog of Table I for six-regular random graphs.

Finally, we comment on the fact that we use the replica-symmetric version of the cavity method for the reported results. Since we are stepping back from the attractors, it could be that describing the statistics of the attractors requires replica symmetry breaking. Following the standard literature on spin glassed on sparse random graphs, we analyze the stability toward replica symmetry breaking [1,2] (more details on the stability in the Appendix) to conclude that the reported results are stable where it was possible to check them using the population dynamics method ($p \leq 5$ for $d = 4, 6$ and $p \leq 4$ for $d = 8, 10$). Also, the empirical results still closely match the ones obtained from the replica-symmetric BDCM; see Table II. This is quite interesting, as in Table II we also give the energy e_{stab} below which replica symmetry breaking (RSB) needs to be considered at equilibrium. We obtain $e_{\text{stab}}(d = 4) = -0.5774$ which is below the energy reached by the quench \tilde{e}_∞ (as analyzed above). However, for $d \geq 6$ the energy reached by a fast synchronous quench (analyzed with the RS approach that is stable toward RSB) is lower than the equilibrium energy at which the RSB needs to be taken into account. Since the fast

TABLE II. We compare the energy to which the synchronous dynamics on d -regular graphs with always-stay tiebreaking converges to empirically, \tilde{e}_∞ (see Appendix C 1 b) and the energy predicted by the BDCM for path length $p = 4$ and the associated entropy. This is compared to energies e_{stab} below which equilibrium properties are described with replica symmetry breaking [2], and the corresponding ground-state energy e_{GS} obtained using the one-step replica symmetry-breaking ansatz computed in Refs. [2,25].

d	Energy after a quench			Equilibrium	
	\tilde{e}_∞	e_4^*	$s_4/\log(2)$	e_{stab}	e_{GS}
4	-0.5475	-0.5411	0.992	-0.5774	-0.7365
6	-0.4764	-0.4656	0.981	-0.4472	-0.6097
8	-0.4283	-0.4151	0.969	-0.3780	-0.5317
10	-0.3930	-0.3785	0.958	-0.3333	-0.4775

quench does not follow equilibrium configurations, it thus seems that it goes to out-of-equilibrium parts of the phase space—the set of attractors of the randomly initialized quench—that are replica symmetric, while the majority of configurations at that same energy (the equilibrium) require RSB. This is quite a surprising behavior, perhaps reminiscent of other problems, such as random graph coloring, where simple algorithms were shown to be finding valid colorings even in regions where the equilibrium is described by replica symmetry breaking [26,27]. A closer investigation of these replica-symmetric subspaces of the RSB equilibrium phase is left for future work.

B. Dynamical phase transition for majority rules

As the second illustration of the BDCM, we consider the ferromagnetic Ising model and dynamics corresponding to the majority rule. The questions we investigate here can find applications, e.g., in generalized bootstrap percolation [28], the zero-temperature Glauber dynamics [29,30], models of segregation [31], density classification for cellular automata [32,33], opinion dynamics [10], or local versions of max or min cut [22,25].

We consider three basic types of deterministic majority dynamics depending on the degree of the nodes and the type of tiebreaking:

- (i) Odd-degree, simple majority rule: At each time step, each spin turns in the direction of the majority of its neighbors.
- (ii) Even-degree, always-change tiebreaking type: Each spin turns to the majority among its neighbors. In the case of balance among the neighbors, the spin always changes to the opposite value from the previous time step.
- (iii) Even-degree, always-stay tiebreaking type: Each spin turns to the majority among its neighbors. In the case of balance among the neighbors, the spin always remains at the same value as in the previous time step.

We then investigate the type of attractors to which the dynamics converges when initialized at random but with a magnetization m_{init} fixed between -1 and 1 . We remind that the majority dynamics always converges to attractors of length either 1 or 2 [19,20]. We distinguish among four types of attractors, specifically the following ones:

- (i) ■ **Homogeneous stable:** These are length-1 attractors with almost all spins either $+1$ or -1 ; $m_{\text{attr}} \in \{-1, +1\}$ and $\rho = 0$.
- (ii) ⊠ **Mixed stable:** These are length-1 attractors with a finite fraction of spins in both $+1$ and -1 ; $m_{\text{attr}} \in (-1, +1)$ and $\rho = 0$.
- (iii) ⊠ **Partially rattling:** These are length-2 cycles where a finite fraction of nodes is not changing during the cycle; $\rho \in (0, 1)$.

- (iv) ■ **All rattling:** These are length-2 attractors where almost all nodes are switching during the cycle; $\rho = 1$.

Note that each of the observed attractors falls under exactly one of these four categories. We emphasize that our definition makes the distinction between m_{attr} and ρ only when they correspond to a finite fraction $\Theta(n)$ of the nodes. This disregards a subleading number $o(n)$ of nodes that might be of a different sign in a homogeneous stable attractor, or $o(n)$ nodes that are not rattling in the all-rattling attractor.

To make the connection with Example 2 from the Introduction, the homogeneous stable attractor corresponds to the all-1 configuration. The question is then, what is the least biased value of the initial magnetization

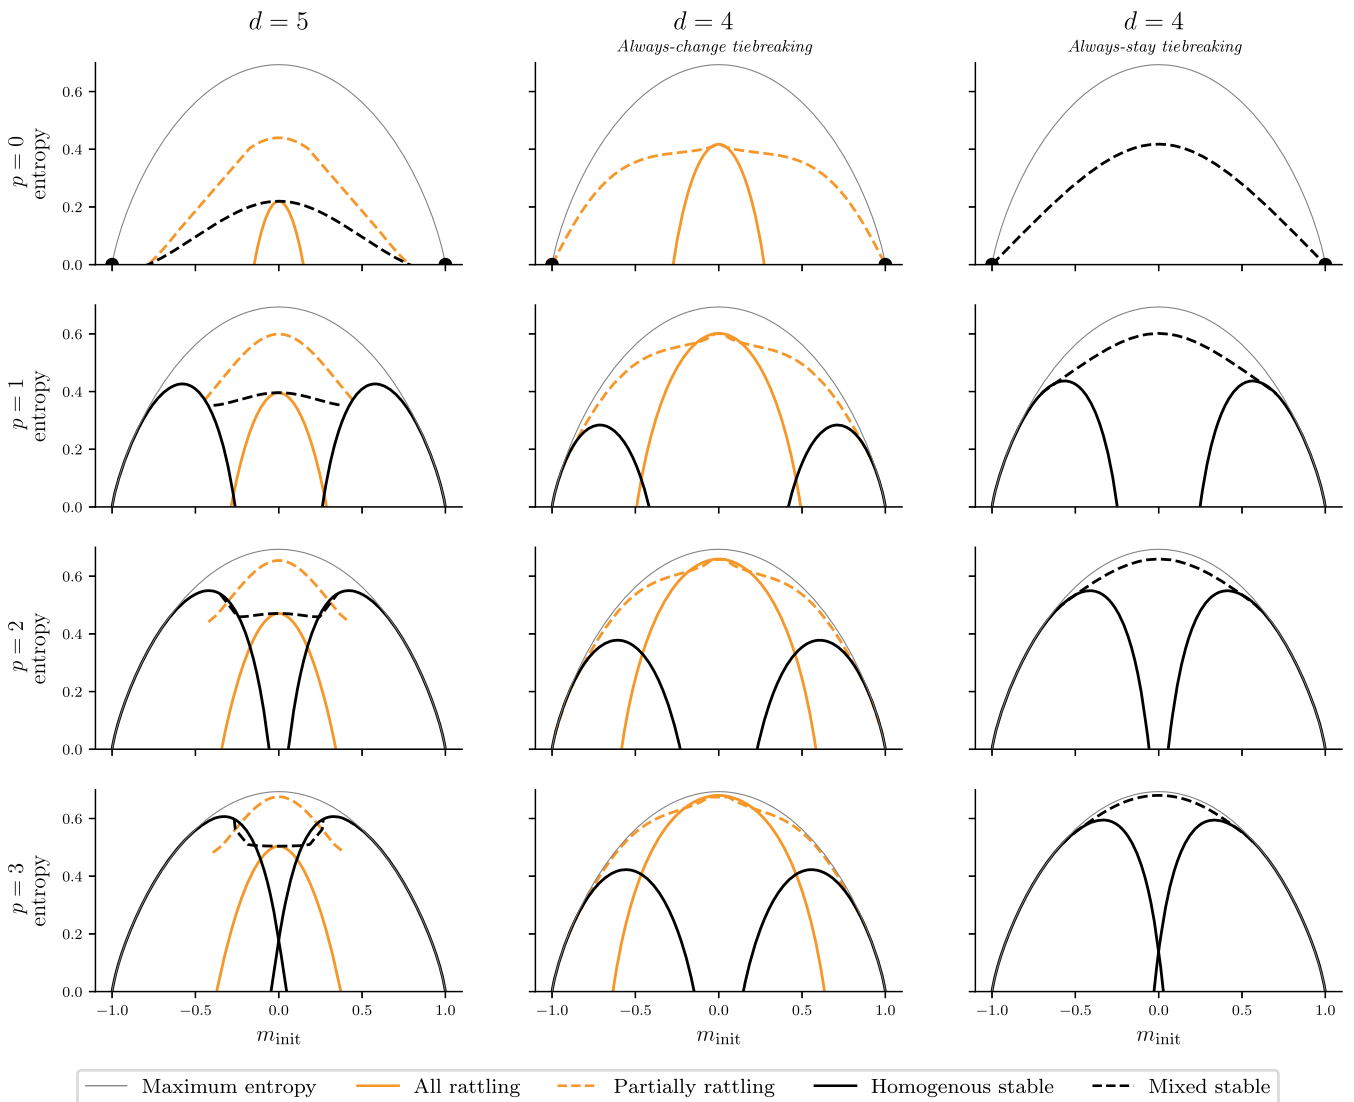


FIG. 3. The BDCM entropy of different attractors and path lengths p for majority rules. Comparison of the BDCM for the majority on regular graphs with degree $d = 5$ and 4 with always-stay and always-change tiebreaking. The points in the first row where $p = 0$ indicate that the entropy of the homogeneous point attractors is exactly zero at $m_{\text{init}} = \pm 1$.

so that the dynamics converge with a high probability to such an attractor?

On d -regular graphs, we can first observe numerically that depending on the initial magnetization m the three dynamical rules converge with high probability to one of the four types of attractors defined above, as shown in Fig. 2 in the upper panel for degree $d = 4$ and 5. For all three considered dynamical rules, we see that for large enough initial magnetization the dynamics converges to the homogeneous stable attractor. For the simple majority dynamics, odd-degree d , initial magnetization close enough to zero converges to the partially rattling cycle. For the always-stay tiebreaking dynamics, an initial magnetization close enough to zero converges to the mixed stable attractor. For the always-change tiebreaking dynamics, initial magnetization close enough to zero converges to the all-rattling cycle, but an intermediate value of magnetization converges to the partially rattling cycle. When we plot the length of the transient to reach the attractor, Fig. 2 lower panel for different graph sizes and degree $d = 4$ and 5, we observe logarithmic divergences of the transient lengths at values of the initial magnetization corresponding to those where the type of attractors changes. In statistical physics, a diverging timescale is usually associated with a phase transition, in this case, a dynamical phase transition. Note that simply counting the attractors of various types, as done, e.g., in Refs. [21–23], does not lead to any sensible explanation of these dynamical phase transitions. We now illustrate how to use the BDCM method to explain and quantify them.

In the BDCM, we compute the size of the basins of attraction of the various types of attractors, or in other words, the entropy of the (p/c) backtracking attractors $\underline{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^{(p+c)})$, conditioned on which type of attractor is expressed in $(\mathbf{x}^{p+1}, \dots, \mathbf{x}^{p+c})$. We are able to separate the different types of attractors by introducing m_{init} , m_{attr} , and ρ as observables in the BDCM. We set c according to the attractor length and threshold the m_{attr} and ρ to analyze each type separately. Since we solve the BP equations numerically, this amounts to thresholding the observables m_{attr} and

ρ with a $\varepsilon = 10^{-8}$. We can isolate the homogeneous and all-rattling attractors by conditioning in the BP update on the homogeneous attractors by forbidding all messages $\chi_{\underline{x} \rightarrow \underline{y}}$ with $x^{p+1} \neq +1$ with $c = 1$ for the all +1 and similar for the all -1 attractor. For the all-rattling attractor, we similarly forbid any $\chi_{\underline{x} \rightarrow \underline{y}}$ with $x^{p+1} = x^{p+2}$ for $c = 2$.

Figure 3 depicts the entropy of the basin of attraction for each of the four types of attractors for path lengths $p < 4$ toward the attractor, in four different line types each for one type of attractor, as a function of the initial magnetization. The values of the entropy correspond to the exponent in the number of configurations of magnetization m_{init} that converge after p steps to an attractor of the corresponding type. When a line for a given type of attractor is not present, it means that this attractor with high probability does not exist for that case.

It is remarkable to note that already with $p = 1$ we observe the qualitatively correct picture where the empirically observed attractors indeed correspond to those of the largest entropy. Also, the value of the largest entropy is already relatively close to the total entropy at the corresponding magnetization. The values of the initial magnetization where the maxima change for $p = 1$ are of course only rough approximations of those at $p \rightarrow \infty$, but the qualitative behavior for the three types of dynamics agrees with the one observed empirically.

For path lengths $p = 1, 2, 3$, we observe that the points where the maximum entropy at that p is reached by a different attractor type are getting closer as p grows to the empirically observed value that would correspond to $p \rightarrow \infty$, as reported in Table IV. The convergence is not as fast as we observed, e.g., for the values of the energy in the previous section, but the fact that the maximum entropy converges rather fast to the total entropy indicates the qualitative correctness of the picture.

Next to the values of the entropies, the BDCM also readily provides the values of the attractor magnetization m_{attr} and the fraction of rattlers in the attractor ρ . These values are plotted in Fig. 4 (full lines) for the attractors




FIG. 4. Comparison between BDCM and empirics for $p = 3$ and majority rules. For a given m_{init} , we compare the magnetization and rattlers predicted from the dominating BDCM fixed point, with the respective empirical value. The empirical results are obtained on graphs of size $n = 100$, and only backtracking attractors with transients of length $p = 3$ are sampled.

that correspond to the largest value of the entropy of the basin of attraction after $p = 3$ steps backward from the attractor. We observe discontinuities in these parameters at the initial magnetization where the type of attractor changes. These data compare qualitatively well with Fig. 2 that gives the numerical values for $p \rightarrow \infty$. On small graphs, we can also sample very many initial conditions that lead to attractors after $p = 3$ steps. Doing so, we compare in Fig. 4 with the empirically obtained values of m_{attr} and ρ observing an excellent quantitative agreement with the theory. The discontinuities are smoothed due to finite-size effects.

V. CONCLUSION

In this paper, we introduce the BDCM on sparse random graphs for models with synchronous discrete-time deterministic dynamics on discrete variables. We illustrate the method on the problem of computing the limiting energy of a quench, finding cases where the quench goes below the energy that marks the onset of replica symmetry breaking at equilibrium, yet the space of attractors the quench converges to is replica symmetric. We also use the method to characterize dynamical phase transitions occurring as the magnetization of the initial configuration changes in majority-driven dynamics.

Here we discuss possible extensions and avenues for future work. Generalization to dense graphs and continuous variables will require constructing a backtracking version of the dynamical mean-field theory. Such a generalization will open the way to studying limiting dynamics of quenches in dense spin glasses or those of gradient descent in the training of neural networks.

The effects of replica symmetry breaking can be incorporated straightforwardly following the lines developed in Refs. [2,3]. Future work will investigate glassy examples where this is relevant.

Another avenue for development is the generalization of the BDCM to stochastically evolving dynamical systems. The dynamical cavity method can be generalized to stochastic dynamics, but more work will be needed to replace the simple counting of states in the basin of attraction with other

free-energy-like notions that will be able to pin which of the trajectories are those from random configurations.

The BDCM shares all the limitations of the usual cavity method [1–3] in terms of the structure of the interactions that is restricted to mean-field-type geometries. A clear limitation of the method is the fact that the time $T = O(1)$ and solving the corresponding equations for large values of T becomes cumbersome, yet more work can be done to investigate more efficient solvers for large values of T .

The code for the solver is available on GitHub [34].

ACKNOWLEDGMENTS

We thank Guilhem Semerjian, Ginestra Bianconi, and Alejandro Lage Castellanos for pointing out relevant references to us, as well as the anonymous referees who provided valuable feedback that improved the quality of the paper. Our work was supported by the Czech project AI&Reasoning, Project No. CZ.02.1.01/0.0/0.0/15_003/0000466 and the European Regional Development Fund, and by Grant No. SVV-2020-260589.

APPENDIX A: FACTOR GRAPH FOR THE BDCM

In Fig. 1, the factor graph with dual variables $(\underline{x}_i, \underline{x}_j)$ is shown, which is subsequently used to derive the BDCM equations. This dual view allows one to prevent short loops in the factor graph; when the naive construction is used where a variable node contains exactly one nodes trajectory \underline{x}_i , both the factor nodes i (representing the update constraint on node i and its neighborhood) and the factor nodes ij (representing the observables on edges) need to be connected to the relevant variables. As shown in Fig. 5, this construction leads to short loops of length 4, which can be prevented by using the dual representation in Fig. 1.

APPENDIX B: STABILITY OF THE BDCM FIXED POINTS TOWARD RSB

Since the RS ansatz that we follow in this paper may not be correct, we check whether the fixed points obtained from the BDCM in our results are stable toward replica




FIG. 5. Naive construction of a factor graph, which contains only a single node trajectory \underline{x}_i in each variable node. This leads to loops of length 4, a problem for applying BP.

TABLE III. Same as Table I for random regular graphs of degree $d = 6$, except that additionally e_p^{start} is provided, the energy at the start of the path for the BDCM.

$d = 6$	BDCM			DCM
p	$s_p / \log(2)$	e_p^*	e_p^{start}	e'_p
0	0.5542	-0.3138	-0.3138	0.0000
1	0.8223	-0.4079	-0.1520	-0.1953
2	0.9205	-0.4407	-0.0865	-0.3100
3	0.9617	-0.4568	-0.0531	-0.3803
4	0.9807	-0.4656	-0.0337	-0.4226
5	0.9901	-0.4705	-0.0216	-0.4477
6	0.9949	-0.4732	-0.0138	-0.4617
Empirical \tilde{e}_∞				-0.4764(1)

symmetry breaking via population dynamics and its convergence analysis.

We initialize the population dynamics with 300 BP messages initialized Gaussian independent and identically distributed. and then normalized to 1. In every iteration, 80% of the messages are updated according to Eq. (17) applied to $d - 1$ randomly selected neighbors. This process is run until convergence. We then check whether the distribution concentrates on a delta function identical to the fixed point found from the RS, or whether it converged to a nondelta distribution over messages. If the first is the case, we say the distribution is stable to RSB, otherwise, it is unstable.

For the experiments from Table I, we validate that the fixed points are indeed stable for $p \leq 5$ and for Table II for $p \leq 4$. Since for the larger p the computational time is exceedingly large, we are not able to verify these results. For the BDCM fixed points used in Fig. 3, we check the stability of the dominating fixed point in each of 80 equally sized intervals we divided $m_{\text{init}} \in \{-1, +1\}$ into, for $p = 0, 1, 2, 3$.

APPENDIX C: ADDITIONAL RESULTS AND SUPPORTING EMPIRICS

1. Limiting energy of a quench

a. Energetic results for $d = 6$

In Table III, we present the BDCM and DCM results for the energy of the antiferromagnet or a spin glass on a six-regular random graph, analogous to Table I from the main text. As before, for large enough p the BDCM and DCM are close to the energy that is observed empirically. Again, the BDCM is slightly more precise given the same number of steps into the attractor's basin p . The size of the basin of attraction as measured by the normalized entropy $s_p / \log(2)$ converges very fast to 1 but slightly slower in p than for $d = 4$. For the interested reader, we also provide the energies measured by the BDCM at the start of the path going into the attractor, i.e., the energy e_p^{start} . The observation that the e_p^{start} quickly grows to zero, the energy of a random configuration, implies that it takes only a logarithmic number of steps between the initial random configuration and the final energy. This can be viewed as an alternative measure to the entropy that allows one to assess the quality of the BDCM prediction after p steps.

b. Empirical results

To determine the limiting energy of the quench empirically, we sample 2048 random regular graphs per graph size n . We initialize them with a random configuration where the number of $+1$ and -1 spins is equal. The synchronous dynamics are run until convergence, and we report the sample average for the energy in the first configuration of the attractor. These empirical results for $p \rightarrow \infty$ are shown in Fig. 6. We extrapolate \tilde{e}_∞ in terms of n ; this is used to compare against the BDCM and DCM results in Tables I and III.

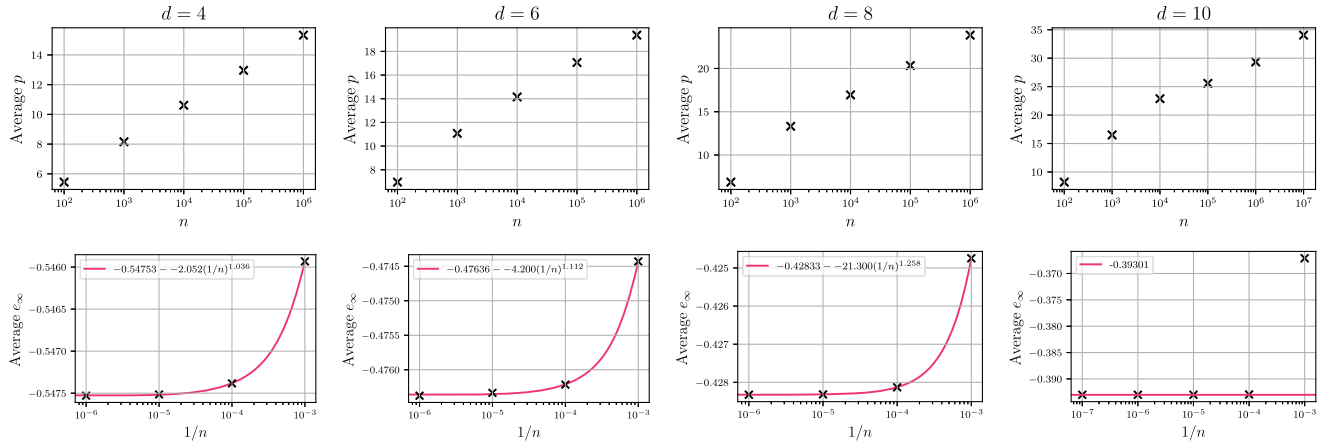


FIG. 6. Empirical dynamics of the quench. Results for the energy reached from a randomly sampled balanced initial configuration on d -regular graphs with n nodes. Top row: average transient length until an attractor is reached. We see a growth logarithmic in n . Bottom row: average energy in the attractor for different n . The red line extrapolates the energy for $n \rightarrow \infty$.




FIG. 7. Transient lengths: the average transient length maximized over the magnetization m_{init} to determine the position of the dynamical phase transitions. We show the four types of dynamical phase transitions from Table IV. The position $\tilde{m}_{\infty}^*(n)$ of the maximal average transient spike is shown and extrapolated to $n \rightarrow \infty$. We use these results as a reference for the empirical phase transitions at $p \rightarrow \infty$.




FIG. 8. Enlargement of the transient lengths for different dynamical phase transitions. According to the four types of transitions defined in Table IV, this plot shows enlargements of the transients obtained in the same manner as in the lower panel of Fig. 2. Every sample is the average over the dynamics run on 2048 graph instances with random initializations.

c. The energetic landscape of the spin glass

In Table II, we compare the energies reached by the quench and obtained from the BDCM with characteristics of the equilibrium energetic landscape. Note that the energy reached by the synchronous dynamics for $d \geq 6$ is lower than the energy below which replica-symmetric breaking is required to describe the equilibrium e_{stab} . However, the BDCM is still stable toward RSB in the investigated cases of $d = 6$ and $p \leq 5$ and $d = 8, 10$ with $p \leq 4$. For larger values, the stability check we perform is numerically too costly, and we anticipate we would also find stable results.

2. Dynamical phase transitions for majority rules

In Figs. 2 and 3, we find four different dynamical phase transitions. By the ± 1 symmetry, we look only at the transitions occurring for $m_{\text{init}} > 0$. These are

- (1) $\otimes \leftrightarrow \blacksquare$ for $d = 5$.
- (2) $\blacksquare \leftrightarrow \otimes$ for $d = 4$, always-change tiebreaking.
- (3) $\otimes \leftrightarrow \blacksquare$ for $d = 4$, always-change tiebreaking.
- (4) $\otimes \leftrightarrow \blacksquare$ for $d = 4$, always-stay tiebreaking.

Note that bounds on the transition (1) as well as its empirical positions were investigated thoroughly for a range of different d in Ref. [10].

TABLE IV. Dynamical phase transition for fixed p via BDCM. The table shows all of the different types of dynamical phase transitions that are observed in Fig. 3 (considering the ± 1 symmetry). Four different dynamical phase transitions occur for the majority rules between different types of attractors: \blacksquare homogenous stable, \otimes mixed stable, \otimes partially rattling, and \blacksquare all rattling. In addition, we show the size of the basin of attraction taken into account proportional to the maximal entropy $H(m)$ for a configuration of a given magnetization m . We compare the results to the empirically observed transition at $p \rightarrow \infty$ (from Fig. 7).

p	$d = 5$		$d = 4$ always change				$d = 4$ always stay	
	(1) $\otimes \leftrightarrow \blacksquare$		(2) $\blacksquare \leftrightarrow \otimes$		(3) $\otimes \leftrightarrow \blacksquare$		(4) $\otimes \leftrightarrow \blacksquare$	
	m_p^*	$s(m_p^*)/H(m_p^*)$	m_p^*	$s(m_p^*)/H(m_p^*)$	m_p^*	$s(m_p^*)/H(m_p^*)$	m_p^*	$s(m_p^*)/H(m_p^*)$
1	0.443	0.627	0.132	0.832	0.902	0.925	0.617	0.880
2	0.302	0.795	0.179	0.912	0.872	0.968	0.496	0.951
3	0.231	0.874	0.200	0.945	0.855	0.983	0.457	0.977
\tilde{m}_{∞}^*	0.0528(1)		0.312(1)		0.81(1)		0.46(1)	

a. Empirical results

We obtain the accurate positions of the dynamical phase transitions above for $p \rightarrow \infty$ by numerical simulations. Figure 8 shows an enlargement for the lower panel of Fig. 2—the averages of the transient lengths p as a function m_{init} for different sizes n . Averages are taken over 4096 samples of random regular graphs and initial configurations. Then, Fig. 7 shows the extrapolation of the position of the maxima for large n .

b. Transitions from the BDCM

We give the locations of dynamical phase transitions shown in Fig. 3 in Table IV. The values of the transitions for small values of p with the BDCM are not very close to the empirically found transitions \tilde{m}_{∞}^* . However, as p grows they become more accurate. This is in line with the observation that for the example of the quench, the fraction of the basin of attraction was much closer to 1 than it is in the examples shown here.

-
- [1] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* (World Scientific Publishing Company, Singapore, 1987), Vol. 9.
- [2] M. Mézard and G. Parisi, *The Bethe Lattice Spin Glass Revisited*, *Eur. Phys. J. B* **20**, 217 (2001).
- [3] M. Mézard and G. Parisi, *The Cavity Method at Zero Temperature*, *J. Stat. Phys.* **111**, 1 (2003).
- [4] Y. Kabashima and D. Saad, *Belief Propagation vs. Tap for Decoding Corrupted Messages*, *Europhys. Lett.* **44**, 668 (1998).
- [5] M. Mézard, G. Parisi, and R. Zecchina, *Analytic and Algorithmic Solution of Random Satisfiability Problems*, *Science* **297**, 812 (2002).
- [6] M. Mézard and A. Montanari, *Information, Physics, and Computation*, Oxford Graduate Texts (Oxford University Press, New York, 2009).
- [7] J. Hatchett, B. Wemmenhove, I. P. Castillo, T. Nikolettopoulos, N. Skantzos, and A. Coolen, *Parallel Dynamics of Disordered Ising Spin Systems on Finitely Connected Random Graphs*, *J. Phys. A* **37**, 6201 (2004).
- [8] I. Neri and D. Bollé, *The Cavity Approach to Parallel Dynamics of Ising Spins on a Graph*, *J. Stat. Mech.* (2009) P08009.
- [9] K. Mimura and A. C. C. Coolen, *Parallel Dynamics of Disordered Ising Spin Systems on Finitely Connected Directed Random Graphs with Arbitrary Degree Distributions*, *J. Phys. A* **42**, 415001 (2009).
- [10] Y. Kanoria and A. Montanari, *Majority Dynamics on Trees and the Dynamic Cavity Method*, *Ann. Appl. Probab.* **21**, 1694 (2011).
- [11] A. Y. Lokhov, M. Mézard, and L. Zdeborová, *Dynamic Message-Passing Equations for Models with Unidirectional Dynamics*, *Phys. Rev. E* **91**, 012811 (2015).
- [12] E. Aurell and H. Mahmoudi, *Dynamic Mean-Field and Cavity Methods for Diluted Ising Systems*, *Phys. Rev. E* **85**, 031119 (2012).
- [13] P. Zhang, *Inference of Kinetic Ising Model on Sparse Graphs*, *J. Stat. Phys.* **148**, 502 (2012).
- [14] G. Del Ferraro and E. Aurell, *Dynamic Message-Passing Approach for Kinetic Spin Models with Reversible Dynamics*, *Phys. Rev. E* **92**, 010102(R) (2015).
- [15] T. Barthel, *Matrix Product Algorithm for Stochastic Dynamics on Networks Applied to Nonequilibrium Glauber Dynamics*, *Phys. Rev. E* **97**, 010104(R) (2018).
- [16] G. Torrisi, R. Kühn, and A. Annibale, *Uncovering the Non-Equilibrium Stationary Properties in Sparse Boolean Networks*, *J. Stat. Mech.* (2022) 053303.
- [17] C. J. Hurry, A. Mozeika, and A. Annibale, *Dynamics of Sparse Boolean Networks with Multi-Node and Self-Interactions*, *J. Phys. A* **55**, 415003 (2022).
- [18] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, *Dynamical Mean-Field Theory of Strongly Correlated Fermion Systems and the Limit of Infinite Dimensions*, *Rev. Mod. Phys.* **68**, 13 (1996).
- [19] E. Goles-Chacc, F. Fogelman-Soulie, and D. Pellegrin, *Decreasing Energy Functions as a Tool for Studying Threshold Networks*, *Discrete Appl. Math.* **12**, 261 (1985).
- [20] B. Derrida, *Dynamical Phase Transitions in Spin Models and Automata* (CEA Centre d'Etudes Nucleaires de Saclay, Paris, 1989).
- [21] A. Bray and M. Moore, *Metastable States, Internal Field Distributions and Magnetic Excitations in Spin Glasses*, *J. Phys. C* **14**, 2629 (1981).
- [22] F. Behrens, G. Arpino, Y. Kivva, and L. Zdeborová, *(Dis)assortative Partitions on Random Regular Graphs*, *J. Phys. A* **55**, 395004 (2022).
- [23] S. Hwang, E. Lanza, G. Parisi, J. Rocchi, G. Ruocco, and F. Zamponi, *On the Number of Limit Cycles in Diluted Neural Networks*, *J. Stat. Phys.* **181**, 2304 (2020).
- [24] F. Altarelli, A. Braunstein, L. Dall'Asta, and R. Zecchina, *Optimizing Spread Dynamics on Graphs by Message Passing*, *J. Stat. Mech.* (2013) P09011.
- [25] L. Zdeborová and S. Boettcher, *A Conjecture on the Maximum Cut and Bisection Width in Random Regular Graphs*, *J. Stat. Mech.* (2010) P02020.
- [26] D. Achlioptas and C. Moore, *Almost All Graphs with Average Degree 4 Are 3-Colorable*, in *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2002), pp. 199–208, 10.1145/509907.509940.
- [27] L. Zdeborová and F. Krzakala, *Phase Transitions in the Coloring of Random Graphs*, *Phys. Rev. E* **76**, 031131 (2007).
- [28] A. N. Zehmakan, *Two Phase Transitions in Two-Way Bootstrap Percolation*, in *Proceedings of the 30th International Symposium on Algorithms and Computation (ISAAC 2019)*, Leibniz International Proceedings in Informatics Vol. 149, edited by P. Lu and G. Zhang (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2019), pp. 5:1–5:21, 10.4230/LIPIcs.ISAAC.2019.5.
- [29] R. Morris, *Zero-Temperature Glauber Dynamics on \mathbb{Z}^d* , *Probab. Theory Relat. Fields* **149**, 417 (2011).

- [30] M. Damron and A. Sen, *Zero-Temperature Glauber Dynamics on the 3-Regular Tree and the Median Process*, *Probab. Theory Relat. Fields* **178**, 25 (2020).
- [31] T. C. Schelling, *Dynamic Models of Segregation*, *J. Math. Sociol.* **1**, 143 (1971).
- [32] R. H. Schonmann, *On the Behavior of Some Cellular Automata Related to Bootstrap Percolation*, *Ann. Probab.* **20**, 174 (1992).
- [33] A. Bušić, N. Fatès, J. Mairesse, and I. Marcovici, *Density Classification on Infinite Lattices and Trees*, in *LATIN 2012: Theoretical Informatics*, Lecture Notes in Computer Science Vol. 7256, edited by D. Fernández-Baca *et al.* (Springer, Berlin, Heidelberg, 2012), 10.1007/978-3-642-29344-3_10.
- [34] <https://github.com/SPOC-group/backtracking-dynamical-cavity>.