Gibbs Measures and Dismantlable Graphs

Graham R. Brightwell *
Dept. of Mathematics
London School of Economics
Houghton St.
London WC2A 2AE England

Peter Winkler†
Bell Labs 2C–365
Lucent Technologies
700 Mountain Ave.
Murray Hill, NJ 07974 U.S.A.

August 6, 1999

Abstract

We model physical systems with "hard constraints" by the space $\mathrm{Hom}(G,H)$ of homomorphisms from a locally finite graph G to a fixed finite constraint graph H. Two homomorphisms are deemed to be adjacent if they differ on a single site of G.

We investigate what appears to be a fundamental dichotomy of constraint graphs, by giving various characterizations of a class of graphs that we call *dismantlable*. For instance, H is dismantlable if and only if, for every G, any two homomorphisms from G to H which differ at only finitely many sites are joined by a path in $\mathrm{Hom}(G,H)$. If H is dismantlable, then, for any G of bounded degree, there is some assignment of activities to the nodes of H for which there is a unique Gibbs measure on $\mathrm{Hom}(G,H)$. On the other hand, if H is not dismantlable (and not too trivial), then there is some r such that, whatever the assignment of activities on H, there are uncountably many Gibbs measures on $\mathrm{Hom}(\mathsf{T}_r,H)$, where T_r is the $(r+1)$-regular tree.

1 Introduction

Continuing a theme begun in our earlier paper [4], we investigate models which exhibit what physicists sometimes call "hard constraints"—forbidden configurations, in which (for example) adjacent particles are not permitted to have certain pairs of spins. In the classical (ferromagnetic) Ising model, adjacent particles are discouraged from having opposing spins, since such opposition increases the energy of a configuration, making it a less likely state; this is a "soft" constraint. In contrast, the hard-core lattice gas model studied e.g. by Dobrushin ([5]) and Van den Berg and Steif ([3]) strictly forbids configurations in which adjacent sites are both occupied. The Widom–Rowlinson model, introduced in [12], has two types of particle, with different types not permitted on adjacent sites. In a combinatorial setting, the condition that a graph be properly n-colored is another example of a hard constraint—as a physical system, this is the n-state antiferromagnetic Potts model at zero temperature.

Suppose we are given a (possibly infinite) locally finite graph G (the *board*) whose nodes represent particles, or "sites". We model hard constraints by means of a finite graph H (the *constraint*

*Author's net address: g.r.brightwell@lse.ac.uk; the research of this author was supported by DIMACS during two visits to Bell Laboratories

†Author's net address: pw@lucent.com

graph), whose nodes may be thought of as different spins, or, as we prefer, different "colors". Adjacent sites of G may receive colors i and j only if i and j are adjacent nodes of H; in particular both sites may have color i only when in H there is a loop at node i. Thus a legal coloring is no more or less than a homomorphism from G to H: that is, a map ϕ from the sites of G to the nodes of H such that if u is adjacent to v in G (written $u \sim v$) then $\phi(u) \sim \phi(v)$ in H.

The constraint graph for the hard-core model is simply the graph with two adjacent nodes, one looped: coloring a site with the unlooped node corresponds to the site being occupied by a particle, and the constraint that adjacent sites may not be occupied is modeled by the absence of a loop on that node. The constraint graph for the Widom–Rowlinson model consists of three looped nodes 1, 0, and -1, with only the pair $(1, -1)$ non-adjacent. For n-coloring, the constraint graph is the unlooped complete graph K_n.

Among legal configurations, relative likelihood is determined by positive reals, called "activities", assigned to the colors. Thus, suppose two legal configurations differ only at site u; if the activity of the color of u in the first configuration is twice that of the color of u in the second, then the first configuration is twice as likely. A *Gibbs measure* is, slightly loosely, a probability measure on the set $\mathrm{Hom}(G, H)$ satisfying such "local" conditions on its conditional probabilities. We shall be more precise shortly.

Dobrushin [5] proved that, provided $\mathrm{Hom}(G, H)$ is non-empty, there is always a Gibbs measure, for any set of activities. Frequently there is more than one Gibbs measure, even for very simple constraint graphs and boards. For instance, the hard-core model on a 3-regular tree admits more than one Gibbs measure, provided the ratio of the activity of the unlooped node to the activity of the looped node is sufficiently high (see e.g. [8] or [4]).

It can be seen (using, e.g. Van den Berg's "paths of disagreement" criterion [2]—see Theorem 7.1 below) that, in both the hard-core and Widom–Rowlinson models on *any* board of bounded degree, there is *some* set of positive activities giving rise to a unique Gibbs measure. However, this is not the case for n-colorings. The purpose of this paper is to identify a distinction between two classes of constraint graphs. For a *dismantlable* constraint graph, we always get uniqueness of Gibbs measures for some choice of the activities, whereas for a non-dismantlable constraint graph we do not.

This distinction between the two classes of constraint graph seems to us to be quite fundamental, and we give several other ways in which their behavior differs. Our main result is Theorem 4.1 below, giving a large number of equivalent characterizations of dismantlable graphs, in terms of homomorphisms, Gibbs measures, and even pursuit games.

The next two sections introduce all the concepts required to state Theorem 4.1, and most of the rest of the paper is devoted to its proof. One more equivalent condition is discussed in the final section.

2 Notation and Preliminaries

We begin with an introduction to, and formal definitions of, the concepts we are studying. This is a slightly shortened version of material from our earlier paper [4].

We will frequently abuse notation by confusing a graph with its set of nodes. There are two roles for graphs in our framework: the board G and the constraint graph H. Each is an undirected graph, possibly with some loops (edges with both ends at the same node), but no multiple edges.

For the most part, we use standard graph theory notation and terminology. In particular, $|G|$ denotes the number of nodes of a (finite) graph G, and, for a set U of nodes of G, $G \setminus U$ is the graph obtained by deleting all nodes of U and incident edges. The set of neighbors of a node x is denoted $N(x)$, and we also write

$$N_m(x) = \{y \in G : d(x,y) \le m\}$$

where $d(x,y)$ is the length of a shortest path between x and y, and m is a non-negative integer. Similarly, for a set U of vertices of G, we set

$$N_m(U) = \{y \in G : d(u,y) \le m \text{ for some } u \in U\}.$$

Note that, if U is finite and G is locally finite, then each set $N_m(U)$ is finite.

We constrain a *board* to be a countable, connected, locally finite, loopless graph with at least two nodes. We will study homomorphisms from boards, which we will always denote G, to constraint graphs, which will always be denoted H. In the most commonly studied settings, G is an infinite, highly symmetric graph like the d-dimensional cubic lattice $\mathbb{Z}^d$ or the regular r-branching tree T_r, uniquely defined by being connected, $(r+1)$-regular and cycle-free. The nodes of the board G will be called "sites" to distinguish them from the nodes of H. We will tend to use letters u, v, w to denote sites of G. A finite set U of sites of G will be called a *patch* and, again, we will deliberately confuse U with the subgraph of G induced by U. The (exterior) *boundary* ∂U of U is the set of sites $G \setminus U$ which are adjacent to at least one site of U.

In contrast to the board, the constraint graph H will be finite and usually small, with nodes $1, 2, \ldots, n$ represented by variables i, j or k. The constraint graph H will often have loops at some or all of its nodes; the loops are important.

We denote by $\mathrm{Hom}(G,H)$ the graph whose nodes are homomorphisms from G to H, with $\alpha \sim \beta$ when α and β differ on at most one site of G. (There is no requirement that the two values at that site be adjacent in H.) We denote the image under α of u by $\alpha(u)$, so that for $\alpha \in \mathrm{Hom}(G,H)$, $u \sim v$ in G implies $\alpha(u) \sim \alpha(v)$ in H.

A *set of activities* for a constraint graph H is a function $\lambda: \ H \to \mathbb{R}^+$ from the nodes of H to the positive reals, two such being regarded as equivalent if they differ by a constant factor. The value λ_i of λ at a node i is called its "activity", and will represent the relative probability of i as an image.

When G is finite and H and λ are given, we define the *multiplicative measure* m_G to be the probability measure on $\mathrm{Hom}(G,H)$ given by

$$\Pr_{m_G}(\{\alpha\}) := \frac{1}{Z} \prod_{u \in G} \lambda_{\alpha(u)}$$

where Z is the necessary normalizing constant,

$$Z := \sum_{\phi \in \mathrm{Hom}(G,H)} \prod_{u \in G} \lambda_{\phi(u)} \ .$$

If all the λ_i's are equal then m_G is the uniform distribution on $\mathrm{Hom}(G,H)$.

The measure m_G enjoys the following property, which we will call the "one-site condition": if Ψ is any event that fixes the colors (that is, images of the homomorphism in H) of all neighbors of a site u, then

$$\Pr_{m_G}(\phi : \phi(u) = i \mid \Psi) = \frac{\lambda_i}{\sum_{j \in J} \lambda_j}$$

where J is the set of colors adjacent to all of the colors assigned by Ψ to the neighbors of u, and $i \in J$. In other words, the conditional probabilities for "eligible" colors of u are proportional to their activities.

Now suppose that we are given an infinite board G, along with a constraint graph H and a set λ of activities as before. If U is a subset of G and $\phi \in \mathrm{Hom}(G,H)$, we denote by $\phi \restriction U$ the restriction of ϕ to U; thus $\phi \restriction U \in \mathrm{Hom}(U,H)$. If A is an event of the form

$$A = \{\phi \in \mathrm{Hom}(G,H) : \ \phi \restriction U \in \Phi\}$$

for some patch U and some $\Phi \subseteq \mathrm{Hom}(U,H)$, then we call A a "patch event". We equip $\mathrm{Hom}(G,H)$ with the σ-field (denoted by $\mathcal{F}$) generated by the patch events, and consider henceforth only measures μ on $\langle \mathrm{Hom}(G,H), \mathcal{F}\rangle$ such that $\mu(\mathrm{Hom}(G,H)) = 1$.

The following definition generalizes the notion of a multiplicative measure to infinite boards, by asserting that the conditional behavior of a measure on each patch is exactly what it should be. We fix but suppress reference to G, H and λ, and define $U^+ := U \cup \partial U$ for any patch U. (So $U^+ = N_1(U)$.)

Definition 2.1. A measure μ on $\mathrm{Hom}(G,H)$ is a *Gibbs measure* (for λ) if for any finite $U \subseteq G$, and for μ-a.e. $\psi \in \mathrm{Hom}(G,H)$,

$$\Pr_{\mu}\Big(\phi : \ \phi = \psi \mid \phi\restriction(G \setminus U) = \psi\restriction(G \setminus U)\Big) = \Pr_{m_{U^+}}\Big(\phi : \ \phi\restriction U = \psi\restriction U \mid \phi\restriction\partial U = \psi\restriction\partial U\Big) \ .$$

In other words, the probability distribution of a random ϕ inside a patch U, conditioned on its values outside U, depends only on its values on the boundary of U. Furthermore, the conditional distribution is the same as for the finite graph U^+.

The one-site condition mentioned earlier is just the special case of the Gibbs condition above where U consists of a single vertex. One problem that we shall explore shortly is that of determining when the one-site condition suffices to ensure that a measure is a Gibbs measure.

Gibbs measures do exist for any G, H and λ, as shown by the following special case of a theorem of Dobrushin [5]. A (simple) proof of this special case appears in [4].

Theorem 2.2. *Let H be a constraint graph with a set of activities λ, and let G be a board for which* $\mathrm{Hom}(G,H)$ *is non-empty. Then there exists at least one Gibbs measure for λ on* $\mathrm{Hom}(G,H)$.

3 A Point Process, a Pursuit Game, and a Structural Property

In this section, we introduce three more concepts which will turn out to be relevant to our discussion.

We define a point process $\mathcal{P} = \mathcal{P}(G,H,\lambda)$ whose state space is $\mathrm{Hom}(G,H)$ as follows: each site u of G "fires" independently and with an exponential waiting time whose mean is the degree $d(u)$ of u. Whenever a site (say, u) fires, a node (say, i) of H is selected at random with probability proportional to its activity. If i is permissible as the image of u, the current homomorphism is altered accordingly; otherwise it is left unchanged.

More formally, suppose the process $\mathcal{P} = \mathcal{P}(G,H,\lambda)$ is in state ϕ at the instant when site u is fired, and that node i is randomly selected, as above. If $N(i) \subseteq \phi(N(u))$ and i is not already equal to $\phi(u)$, ϕ is altered to ϕ' where $\phi'(u) = i$ and $\phi' = \phi$ on $G \setminus \{u\}$. Otherwise $\mathcal{P}$ remains in state ϕ.

The point of having the mean time between firings depend on the degree of the site is that otherwise, in certain bizarre cases where G does not have bounded degree, the state of the system might not be well defined. With the definition above, however, it is easily checked that the states which fire in some short interval of time do not "percolate"—that is, there is no infinite path consisting of fired states. In particular, given the state of the system at time 0, the state of any site u at time $\frac{1}{2}$ depends on only a finite number of firings and is thus uniquely defined.

Of course, if the state ϕ_t of the system is well defined for $t \in [0, \frac{1}{2}]$ then it is well defined for all $t \geq 0$. We may thus speak of a distribution μ on $\mathrm{Hom}(G,H)$ being stationary if, when ϕ_0 is drawn from μ, the distribution of ϕ_t is again μ for all $t > 0$.

Let us suppose first that G is finite; then stationarity reduces to a condition we have already considered.

Lemma 3.1. *If G is finite then a distribution μ on* $\mathrm{Hom}(G,H)$ *is stationary if and only if it satisfies the "one-site condition" of §2.*

Proof. We show first that if μ satisfies the one-site condition ("1SC") then it is stationary; in fact this conclusion does not depend on the mechanism which decides which site fires next, as long as sites fire one at a time and in such a way that "next" is well defined, and firing times are independent of state. We show that even *given* that the site next to fire is u, that if the probability distribution of the state ϕ before firing is μ then the distribution of the state ϕ' after firing is again μ.

This is easy because if $J := \bigcap_{v\sim u} N_H(\phi(v))$ is the set of colors permitted for u by the ϕ-colors of u's neighbors, then

$$\begin{aligned}
\mu'(\phi) &= \sum_{j\in J} \frac{\lambda_{\phi(u)}}{\sum_{j\in J}\lambda_j}\mu(\phi[u\to j]) \\
&= \sum_{j\in J} \frac{\lambda_{\phi(u)}}{\sum_{j\in J}\lambda_j}\frac{\lambda_j}{\lambda_{\phi(u)}}\mu(\phi) \\
&= \mu(\phi)
\end{aligned}$$

where $\phi[u\to j](u) := j$ and $\phi[u\to j](v) := \phi(v)$ for $v \neq u$.

To show that the 1SC is *necessary* for stationarity, we do need that each state has positive probability of firing. Then, if $C_1, \ldots, C_k$ are the connected components of $\mathrm{Hom}(G,H)$ and the initial state of the process is in C_j, the process constitutes an irreducible Markov chain with unique stationary state m_j. Since the multiplicative measure m on $\mathrm{Hom}(G,H)$, restricted to C_j, satisfies 1SC, m_j is precisely that measure.

It follows that any stationary measure for the point process on G is a convex combination $\mu = \sum_{i=1}^k \alpha_i m_i$ of the m_i's, and therefore also satisfies the 1SC. □

Now suppose that G is infinite, but that the set of sites S which fire between times 0 and 1 does not percolate. Suppose μ satisfies the 1SC and let A be a patch event on the patch U. Let $V \supset U$ be a finite patch for which no site in ∂V fires between times 0 and 1. Then $\mathrm{Pr}_1(A) = \mathrm{Pr}_0(A)$ by applying the theorem to the finite graph $V \cup \partial V$. Since such a V exists with probability 1, and the patch events form a basis for $\mathrm{Hom}(G,H)$, we conclude that μ is stationary.

We have not been able to prove the converse (that stationarity implies the 1SC for infinite G) although this may well be true. When the issue of stationarity arises later we will confine ourselves to the case of finite G.

The next concept we introduce is a cop-and-robber game studied by Quilliot [11] and by Nowakowski and Winkler [9]. Two players, a cop $\mathcal{C}$ and a robber $\mathcal{R}$, compete on a fixed, finite, undirected graph H. The cop begins by placing herself at a node of her choice; the robber then does the same. Then the players alternate beginning with $\mathcal{C}$, each moving to an adjacent node. The cop wins if she can "capture" the robber, that is, move onto the node occupied by the robber; $\mathcal{R}$ wins by avoiding capture indefinitely. In doing so $\mathcal{R}$ is free to move (or even place himself initially) onto the same node as the cop, although that would be unwise if the node were looped since then $\mathcal{C}$ could capture him at her next move.

Evidently the robber can win on any loopless graph by placing himself at the same node as the cop and then shadowing her every move; among graphs in which every node is looped, $\mathcal{C}$ clearly wins on paths and loses on cycles of length at least 4. (In the game as defined in [9, 11], there is in effect a loop at every node of H.)

It is convenient to say that the cop wins if the robber cannot move, i.e. if the robber begins on a loopless isolated point. Thus $\mathcal{C}$ wins when H has only one node, looped or not, and on a pair of unconnected nodes, she wins unless they are both looped.

The graph on which the game is played is said to be *cop-win* if $\mathcal{C}$ has a winning strategy, *robber-win* otherwise.

Finally, we give a structural description of a class of graphs. We shall see that this is exactly the class of cop-win graphs, as well as being the class of constraint graphs for which every stationary distribution for the point process $\mathcal{P}$ is always a Gibbs measure, and also the class of constraint graphs for which there is always some set of activities yielding a unique Gibbs measure.

Suppose that i and j are nodes of a graph H such that $N(i) \subseteq N(j)$. Then the map taking i to j, and every other node of H to itself, is a homomorphism from H to $H \setminus \{i\}$. We call this a *fold* of the graph H.

A finite graph H is *dismantlable*[1] if there is a sequence of folds reducing H to a graph with one node (looped or not). The name is chosen (in preference to "cop-win") in order to stress the structure, rather than the game, and in fact appears already in the literature (see e.g. [1, 10]). Relative to our usage, however, previous articles have considered only the case where all nodes of the graph are regarded as having loops.

In fact, the only case where a sequence of folds reduces H to an unlooped node is when H is a set of isolated unlooped nodes—in which case we call H *trivial.* (In fact, we are most interested in connected graphs H with at least two nodes.) If i is a looped node of H such that there is a sequence of folds reducing H to i, then we call i a *persistent* node of H.

A finite graph H is *stiff* if it has more than one node, and no folds, i.e., no pair of nodes (i, j) with $N(i) \subseteq N(j)$. It is of course quite possible for a graph to be non-dismantlable without being stiff: we shall give a little more information on the structure of a general non-dismantlable graph in the next section. Fig. 1 below shows some dismantlable and non-dismantlable graphs, including stiff and non-stiff examples of the latter.

4 An Equivalence Theorem

We can now state our main result.

Theorem 4.1. *The following are equivalent, for finite graphs H.*

1. *H is dismantlable.*

2. *H is cop-win.*

3. *For every finite board G, $\mathrm{Hom}(G, H)$ is connected.*

[1] Spelling authorities indicate that a verb ending in "le" *usually* keeps the "e" when the suffix "able" is appended, as in "handleable". Since retention is not, apparently, mandatory, we have maintained the spelling that appears in the mathematics literature.

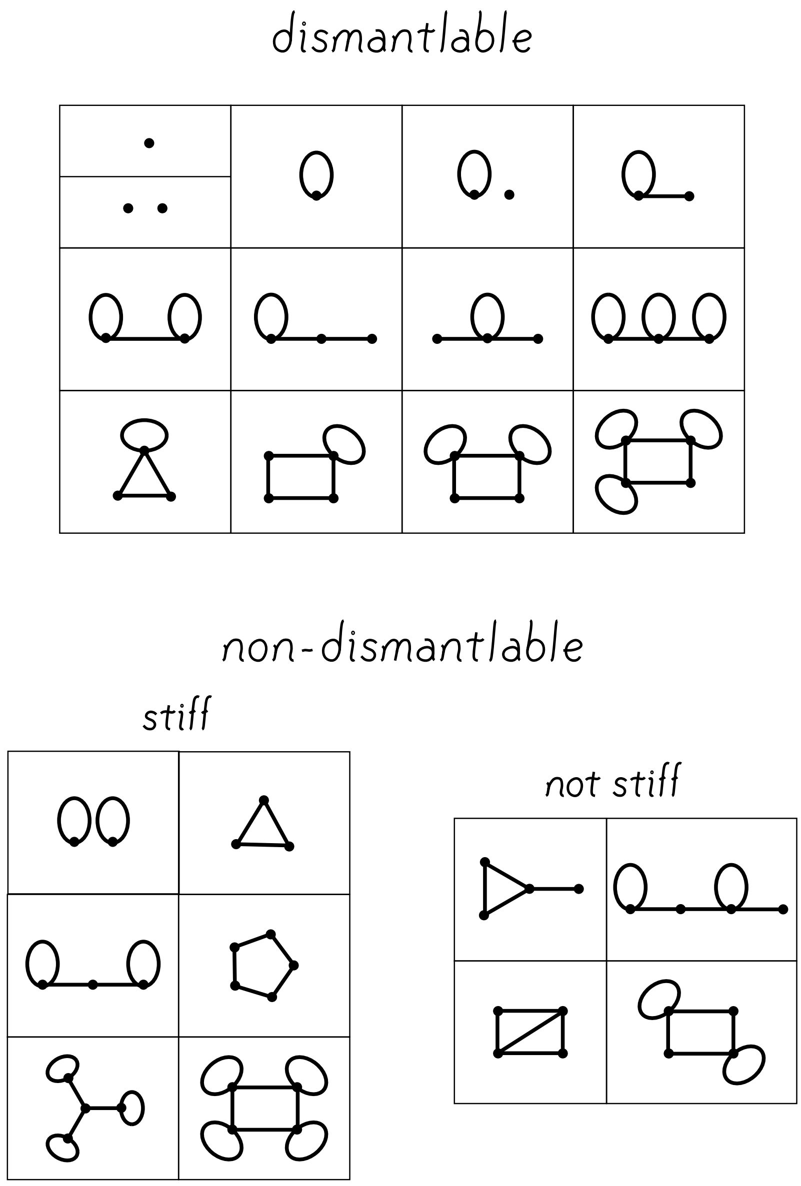

Figure 1: Dismantlable and non-dismantlable graphs

4. *For every board G, and every pair $\phi, \psi \in \mathrm{Hom}(G,H)$ agreeing on all but finitely many sites, there is a path in $\mathrm{Hom}(G,H)$ between ϕ and ψ.*

5. *There is some positive integer m such that, for every board G, every pair of sets U and V in G at distance at least m, and every pair of maps $\phi, \psi \in \mathrm{Hom}(G,H)$, there is a map $\theta \in \mathrm{Hom}(G,H)$ such that θ agrees with ϕ on U and with ψ on V.*

6. *For every positive integer r, and every pair of maps $\phi, \psi \in \mathrm{Hom}(\mathsf{T}_r, H)$, there is a site u in T_r with $\phi(u) \neq \psi(u)$, a patch U containing u, and a map $\theta \in \mathrm{Hom}(\mathsf{T}_r, H)$ which agrees with ψ on $\mathsf{T}_r \setminus U$ and with ϕ on u.*

7. *For every board G and set λ of activities, if μ is a measure on $\mathrm{Hom}(G,H)$ satisfying the one-site condition, then μ is a Gibbs measure.*

8. *For every finite board G and set λ of activities, every stationary distribution for the point process $\mathcal{P}(G,H,\lambda)$ is a Gibbs measure.*

9. *For every board G of bounded degree such that $\mathrm{Hom}(G,H)$ is non-empty, there is a positive set λ of activities such that there is a unique Gibbs measure on $\mathrm{Hom}(G,H)$.*

10. *Either H has no edges or, for every r, there is a positive set λ of activities such that there is a unique Gibbs measure on $\mathrm{Hom}(\mathsf{T}_r, H)$.*

For instance, note that the constraint graphs for the hard-core and Widom–Rowlinson models are dismantlable, whereas the unlooped complete graph is not. Theorem 4.1 thus gives a wide variety of ways in which the first two models will differ from the model of random graph coloring.

One could write down many other equivalent statements. For instance, notice that (6) is a very weak version of a very special case of (5), so that any statement intermediate between the two will also give a characterization of dismantlable graphs. Yet another equivalent statement will be given as Theorem 9.2 later.

We shall prove Theorem 4.1 in stages, stating the various results separately as we go along, sometimes in slightly stronger forms. Note that, if H is a trivial dismantlable graph, then it satisfies all the above conditions, mostly vacuously, so we may always assume in what follows that H has at least one edge or loop. Moreover, it is immediate that the addition or removal of isolated, loopless nodes of H does not affect the truth or falsity of any of the statements, and it is simple to check that all the statements are false if H has more than one non-trivial component (for (4) and (7), consider any finite board). Thus we may assume whenever it is convenient that H is connected.

Some implications among (1)–(10) are instant. We noted above that (5) implies (6), and it is also clear that (4) implies (3), and that (9) implies (10). We have already seen, via Lemma 3.1, that (7) implies (8).

For the remainder of this section, we consider only (1) and (2). The equivalence of these two conditions, providing a structural characterization of cop-win graphs, generalizes the results of [9, 11], where cops and robbers were permitted to remain at a current node; thus there, in effect, all nodes were automatically looped.

Theorem 4.2. *Let H be a finite, connected graph. Then H is cop-win if and only if it is dismantlable.*

Proof. Let us first assume that H is dismantlable, and show that H is cop-win, by induction on the order of H. With our conventions $\mathcal{C}$ does win on one-node graphs. Suppose H has nodes i and j such that $N(i) \subseteq N(j)$ and $H \setminus \{i\}$ is cop-win: we need to prove that H is cop-win. Let S be a strategy for $\mathcal{C}$ that is winning on $H \setminus \{i\}$. We modify S for operation on H as follows: whenever $\mathcal{R}$ is on i, $\mathcal{C}$ plays as if he were on j. Eventually she will either capture the robber or move to j when he is on i, in which case she will win on the next move.

The converse is equally easy. Since the game has only a finite number of states, if H is cop-win then the cop can win in (say) s steps. Thus there is a position (say, robber at i and cop at j) with robber to move from which $\mathcal{C}$ can win in one step. This means that $N(i) \subseteq N(j)$. We claim $H \setminus \{i\}$ is also cop-win; if not, then a winning strategy for $\mathcal{R}$ on $H \setminus \{i\}$ also works on H, as long as the robber regards a cop at i as if she were at j. $\square$

The equivalence of (1) and (2) for disconnected graphs follows immediately: isolated nodes do not affect the play, while if there are two components each containing an edge, then $\mathcal{R}$ wins by moving to a different component from that chosen by $\mathcal{C}$ on the first move.

The proof of Theorem 4.2 shows that to determine whether a given graph H is cop-win, it suffices to identify *any* i, j pair with $N(i) \subseteq N(j)$ and remove i, repeating until only one node remains (in which case H is cop-win) or until no such pair exists, i.e., we reach a stiff graph. Thus the cop-win property, and hence dismantlability, is recognizable in polynomial time.

It may aid understanding later to consider the structure of a general non-dismantlable graph, and it is convenient to do this here. Evidently a stiff graph is non-dismantlable, but more generally there is some sequence of folds that can be made before reaching a stiff graph, and this sequence of folds will typically not be unique. However, we show that the stiff graph eventually reached by a sequence of folds is unique up to isomorphism.

Accordingly, suppose that H is a non-dismantlable graph, and consider some sequence of folds on H reducing it to a stiff graph J. Now define the family of sets $(C_j)_{j\in J}$ recursively as follows.

(i) For every $j \in J$, $j \in C_j$.

(ii) For $h \in H$ and $j \in J$, if h has a neighbor in C_i for every $i \in J$ with $i \sim j$, then $h \in C_j$.

(iii) The C_j are the unique minimal sets subject to (i) and (ii).

Lemma 4.3. *The sets C_j are disjoint. Furthermore, if j and j' are not adjacent in J, then there are no edges between C_j and $C_{j'}$ in H.*

Proof. Consider the composition $\psi : H \to J$ of the folds in the sequence. The map ψ is a retract, i.e., it is a graph homomorphism whose restriction to J is the identity. We claim that $C_j \subseteq \psi^{-1}(j)$ for each j; this clearly implies the result.

Suppose that h has a neighbor g_i in $\psi^{-1}(i)$ for every $i \in J$ with $i \sim j$. Since ψ is a homomorphism, $\psi(h) \sim \psi(g_i) = i$ for each such i. Then $N_J(j) \subseteq N_J(\psi(h))$. Since J has no folds, this implies that $j = \psi(h)$, i.e., $h \in \psi^{-1}(j)$. Thus the sets $(\psi^{-1}(j))_{j\in J}$ satisfy (ii) and (since ψ is a retract) also (i). By (iii), we have $C_j \subseteq \psi^{-1}(j)$ for each j, as desired. $\square$

It is perhaps worth stressing that some vertices of H may be in none of the C_j. For a simple example, consider the case of an unlooped triangle with one pendant edge. Then J consists of the three vertices of the triangle, but the other vertex does not belong to any of the C_j.

Theorem 4.4. *Let H be a non-dismantlable graph. Then, up to isomorphism, there is a unique stiff graph J that can be reached by a sequence of folds from H.*

Proof. Consider some specific sequence of folds from H resulting in a stiff graph J, and let the sets C_j, $j \in J$, be defined as above. We claim that any sequence of folds from H preserves an isomorphic copy $\hat{J}$ of J, with the node $\hat{j}$ of $\hat{J}$ corresponding to $j \in J$ being in C_j, for each j. Indeed, suppose that H can be folded to a graph L containing a suitable copy $\hat{J}$ of J, but that, in L, some node $\hat{j} \in \hat{J}$ can be folded to another node k, i.e., $N_L(\hat{j}) \subseteq N_L(k)$. In particular, k is adjacent to all the neighbors of $\hat{j}$ in $\hat{J}$, i.e., all the nodes $\hat{i}$, for i a neighbor of j in J, and so $k \in C_j$. Now, by Lemma 4.3, k is not adjacent to those nodes of $\hat{J}$ in some $C_{j'}$ with j not adjacent to j'. Thus, replacing $\hat{j}$ by k gives us a suitable copy of J in the fold of L.

We have shown that any sequence of folds from H results in a graph containing a copy of J. But this is true for any stiff graph that can be obtained by a sequence of folds from H, so in fact we always obtain exactly a copy of J, as required. □

The stiff graph J of Theorem 4.4 is reminiscent of the "core" of a graph as defined in [7].

5 Connectedness of the Homomorphism Graph

In this section, we show the equivalence of conditions (1), (3), and (4) of Theorem 4.1, and we also show that (1) implies (5) and (6). As already noted, it is immediate that (4) implies (3), and that (5) implies (6).

Our first result in this section shows that (3) implies (1).

Theorem 5.1. *If H is non-dismantlable, then there is a finite board G for which* $\mathrm{Hom}(G, H)$ *has at least two components.*

Proof. Let H be non-dismantlable, and suppose that nonetheless $\mathrm{Hom}(G, H)$ is connected for all finite boards G; let $n = |H|$ be minimal with respect to these properties.

If there are nodes i and j of H with $N(i) \subseteq N(j)$, then $H \setminus \{i\}$ is also non-dismantlable. In this case, we claim that the connectivity of $\mathrm{Hom}(G, H)$ implies connectivity of $\mathrm{Hom}(G, H \setminus \{i\})$. To see this, define, for $\phi \in \mathrm{Hom}(G, H)$, the map $\phi' \in \mathrm{Hom}(G, H \setminus \{i\})$ by changing all i's to j's in the image. If α and β are two maps in $\mathrm{Hom}(G, H \setminus \{i\})$, then we may connect them by a path $\phi_1, \ldots, \phi_t$ in $\mathrm{Hom}(G, H)$; now we observe that the not-necessarily distinct sequence of maps $\phi'_1, \ldots, \phi'_t$ connects α and β in $\mathrm{Hom}(G, H \setminus \{i\})$. This contradicts the minimality of H, so we may assume from now on that there is no such pair of nodes i and j, i.e., H is stiff.

We now break the proof into two cases, depending on whether or not H has a looped node. Suppose first that the set L of looped nodes in H is non-empty, and define a finite graph G as follows. The nodes of G are those of H together with a second copy i' of each node $i \in L$, and the edges of G are

$$\{\{i, j\} : \ i \sim j \text{ in } H\} \cup \{\{i, i'\} : \ i \in L\} \cup \{\{i', j\} : \ \ i \sim j \text{ in } H \text{ and } i \in L\} \ .$$

Thus each looped node of H is replaced in G by two adjacent, unlooped nodes connected as before. We let $\iota \in \mathrm{Hom}(G, H)$ be the "pseudo-identity" map sending i and i' to i for $i \in L$, and j to j for $j \notin L$.

We claim that ι is an isolated point of $\mathrm{Hom}(G, H)$; otherwise for some $u = i$ or $u = i' \in G$ and $i, j \in H$ the alteration ι' of ι sending u to j instead of i is a homomorphism. But then $N(i) \subseteq N(j)$, a contradiction.

Observe that ι is not the only member of $\mathrm{Hom}(G, H)$, since the constant map $\kappa : \ G \to \{k\}$ is in $\mathrm{Hom}(G, H)$ for any $k \in L$. Thus $\mathrm{Hom}(G, H)$ is disconnected after all.

Now suppose that our graph H has no loops, and let G be the 'weak' square of H, that is, the graph whose nodes are ordered pairs (i_1, i_2) of nodes of H with $(i_1, i_2) \sim (j_1, j_2)$ just when $i_1 \sim j_1$ and $i_2 \sim j_2$. There are two natural homomorphisms from G to H, the projections π_1 and π_2, where $\pi_1(i_1, i_2) = i_1$ and $\pi_2(i_1, i_2) = i_2$; we claim that π_1 is an isolated point of the graph $\mathrm{Hom}(G, H)$.

If not, there is a map π' taking (say) (i_1, i_2) to $k \neq i_1$ and otherwise agreeing with π_1. Let j_2 be a fixed neighbor of i_2 and j_1 any neighbor of i_1. Then $(i_1, i_2) \sim (j_1, j_2)$ hence $k \sim j_1$. We have shown that every neighbor of i_1 is also a neighbor of k, contradicting the assumption that H is stiff.

This completes the proof of the theorem. □

Next, we prove a lemma about dismantlable graphs that will be used to show that (1) implies (4) and (5), and later (7) and (9) as well. The lemma is stated so as to cover all the various uses to which we will put it, at the expense of brevity.

Lemma 5.2. *Let H be a non-trivial dismantlable graph with $|H| = n$, and let j be a persistent node of H. Let G be any board, and U any subset of G. Let ϕ be any map in $\mathrm{Hom}(G, H)$. Then there is a homomorphism $\theta \in \mathrm{Hom}(G, H)$ such that $\theta\restriction(G \setminus N_{n-2}(U)) = \phi\restriction(G \setminus N_{n-2}(U))$, $\theta(v) = j$ for every $v \in U$, and $\phi^{-1}(j) \subseteq \theta^{-1}(j)$.*

Furthermore, if U is finite, then there is a path in $\mathrm{Hom}(G, H)$ connecting θ and ϕ all of whose intermediate elements agree with θ and ϕ on $G \setminus N_{n-2}(U)$. Finally, if in addition $\psi \in \mathrm{Hom}(G, H)$ agrees with ϕ on $G \setminus U$, then ψ is also connected to θ by a path with this property.

Proof. Since j is a persistent node of the dismantlable graph H, we can find a sequence of graphs $H = H_n, H_{n-1}, \dots, H_1 = \{j\}$, with $|H_k| = k$ for each k, and a sequence of pairs of nodes $(i_n, j_n) \in H_n, (i_{n-1}, j_{n-1}) \in H_{n-1}, \dots, (i_2, j_2) \in H_2$, such that $N(i_k) \subseteq N(j_k)$ in H_k and $H_{k-1} = H_k \setminus \{i_k\}$ for each k, $1 < k \leq n$.

We define a sequence of maps $\phi = \phi_n, \phi_{n-1}, \dots, \phi_1 = \theta$ in $\mathrm{Hom}(G, H)$ with the property that $\phi_k(N_{k-1}(U)) \subseteq H_k$ for each k, as follows.

Suppose that ϕ_k is already defined and that $\phi_k(N_{k-1}(U)) \subseteq H_k$. For each $u \in N_{k-2}(U)$ with $\phi_k(u) = i_k$, we put $\phi_{k-1}(u) = j_k$; for all $v \notin N_{k-2}(U)$, and all $v \in G$ with $\phi_k(v) \neq i_k$, we let $\phi_{k-1}(v) = \phi_k(v)$. Every neighbor of a site in $N_{k-2}(U)$ is in $N_{k-1}(U)$ and thus its image under ϕ_k lies in H_k; since $N(i_k) \subseteq N(j_k)$ in H_k, the changes are legitimate and $\phi_{k-1} \in \mathrm{Hom}(G, H)$. Furthermore we have $\phi_{k-1}(N_{k-2}(U)) \subseteq H_{k-1}$ and $\phi_k^{-1}(j) \subseteq \phi_{k-1}^{-1}(j)$. If U is finite, then so are all the $N_k(U)$; in this case the necessary changes can be made one at a time, and there is a path between ϕ_k and ϕ_{k-1} in $\mathrm{Hom}(G, H)$ which never makes changes away from the set $N_{n-2}(U)$.

By the time we reach $\phi_1 = \theta$, we have a map which is identically equal to the single node j of H_1 on $N_0(U) = U$, as desired. Also, the set $\phi_k^{-1}(j)$ is non-decreasing throughout the process.

Finally, note that the value of $\theta(v)$ on a fixed site v in $G \setminus U$ depends only on the value of $\phi(v)$. Thus if ψ is another map agreeing with ϕ on $G \setminus U$, then it too is connected to θ by a path of the same form. □

It is an immediate consequence of Lemma 5.2 that every non-trivial dismantlable graph H has property (4) of Theorem 4.1. It is also easy to deduce that H has property (5) as well, as we now see.

Theorem 5.3. *Let H be a non-trivial dismantlable graph with $|H| = n$. Let G be any board, with subsets V and W such that $d(v, w) \geq 2n - 1$ for all $v \in V$ and $w \in W$. Let ϕ and ψ be any maps in $\mathrm{Hom}(G, H)$. Then there is $\theta \in \mathrm{Hom}(G, H)$ agreeing with ϕ on V and with ψ on W.*

Proof. Let j be a persistent node of U, and apply Lemma 5.2 with $U = G \setminus N_{n-1}(V)$ to get a homomorphism θ_1 that agrees with ϕ on V, and is identically j on $G \setminus N_{n-1}(V)$. Similarly, there is a homomorphism θ_2 that agrees with ψ on W, and is identically j on $G \setminus N_{n-1}(W)$. Note that there are no edges between $N_{n-1}(V)$ and $N_{n-1}(W)$, so the map θ defined by setting $\theta(u) = \theta_1(u)$ if $u \in N_{n-1}(V)$, and $\theta(u) = \theta_2(u)$ otherwise, is a homomorphism with the required properties. □

6 The One-Site Condition

We have already seen that condition (7) of Theorem 4.1 is at least as strong as (8); we now show that (4) is equivalent to both (7) and (8).

Theorem 6.1. *If the finite graph H is dismantlable then for any board G, finite or infinite, and any set λ of activities for H, every measure satisfying the one-site condition is a Gibbs measure. If H is not dismantlable then there is a measure on a finite board which satisfies the one-site condition but is not a Gibbs measure.*

Proof. Let us assume first that H is dismantlable and that G and λ are given, along with a measure μ on $\mathrm{Hom}(G,H)$ satisfying the one-site condition. We wish to show that μ is a Gibbs measure.

Accordingly, take any finite subset $U \subset G$. The definition of the one-site condition allows for a μ-null set of homomorphisms to misbehave at each site. For $u \in U$, let $\mathcal{E}_u$ denote the set of homomorphisms ψ such that, for some homomorphism ψ' agreeing with ψ except on u, we have

$$\frac{\mathrm{Pr}_\mu\left(\phi:\ \phi=\psi \mid \phi\restriction(G\setminus\{u\})=\psi\restriction(G\setminus\{u\})\right)}{\mathrm{Pr}_\mu\left(\phi:\ \phi=\psi' \mid \phi\restriction(G\setminus\{u\})=\psi\restriction(G\setminus\{u\})\right)} \neq \frac{\lambda_{\psi}(u)}{\lambda_{\psi'(u)}}.$$

Then each $\mathcal{E}_u$ is μ-null. Now let $\mathcal{E}$ denote the set of homomorphisms ψ agreeing with some ψ' in some $\mathcal{E}_u$ on $G \setminus N_{n-2}(U)$. Then $\mathcal{E}$ is again a μ-null set.

We claim that, for any $\psi \notin \mathcal{E}$,

$$\Pr_{\mu}\left(\phi:\ \phi=\psi \mid \phi\restriction(G\setminus U)=\psi\restriction(G\setminus U)\right) = \Pr_{m_{U+}}\left(\phi:\ \phi\restriction U=\psi\restriction U \mid \phi\restriction\partial U=\psi\restriction\partial U\right).$$

This is equivalent to saying that, for any $\psi \notin \mathcal{E}$, and any $\psi' \in \mathrm{Hom}(G,H)$ agreeing with ψ on $G \setminus U$, we have

$$\frac{\mathrm{Pr}_\mu\left(\phi:\ \phi=\psi \mid \phi\restriction(G\setminus U)=\psi\restriction(G\setminus U)\right)}{\mathrm{Pr}_\mu\left(\phi:\ \phi=\psi' \mid \phi\restriction(G\setminus U)=\psi\restriction(G\setminus U)\right)} = \prod_{u\in U}\frac{\lambda_{\psi(u)}}{\lambda_{\psi'(u)}}.$$

The left-hand ratio here is equal to

$$\frac{\mathrm{Pr}_\mu\left(\phi:\ \phi=\psi \mid \phi\restriction(G\setminus N_{n-2}(U))=\psi\restriction(G\setminus N_{n-2}(U))\right)}{\mathrm{Pr}_\mu\left(\phi:\ \phi=\psi' \mid \phi\restriction(G\setminus N_{n-2}(U))=\psi\restriction(G\setminus N_{n-2}(U))\right)},$$

since dividing numerator and denominator by $\mathrm{Pr}_\mu\Big(\phi:\ \phi\restriction(G\setminus U)=\psi\restriction(G\setminus U) \mid \phi\restriction(G\setminus N_{n-2}(U)) = \psi\restriction(G\setminus N_{n-2}(U))\Big)$ gives the original ratio.

Similarly, the one-site condition (or, more precisely, the definition of $\mathcal{E}$) tells us that we have

$$\frac{\mathrm{Pr}_\mu\left(\phi:\ \phi=\alpha \mid \phi\restriction(G\setminus N_{n-2}(U))=\alpha\restriction(G\setminus N_{n-2}(U))\right)}{\mathrm{Pr}_\mu\left(\phi:\ \phi=\alpha' \mid \phi\restriction(G\setminus N_{n-2}(U))=\alpha\restriction(G\setminus N_{n-2}(U))\right)} = \frac{\lambda_{\alpha(u)}}{\lambda_{\alpha'(u)}},$$

whenever α and α' differ only on $u \in U$, and agree with ψ on $G \setminus N_{n-2}(U)$.

By Lemma 5.2, ψ is connected by a path to ψ' in $\mathrm{Hom}(G,H)$, with all intermediate steps agreeing off $N_{n-2}(U)$. The result now follows by multiplying up the appropriate ratios along such a path.

It remains to show that if H is non-dismantlable then the one-site condition is strictly weaker, even on finite boards, than the Gibbs condition. This is an easy task in view of Theorem 5.1. Let G be any finite graph for which $\mathrm{Hom}(G,H)$ is disconnected, and let C be a component of $\mathrm{Hom}(G,H)$. Fix a set of activities for H and let m be the corresponding Gibbs measure. Now define a new measure m_C by doubling the relative probability of every map in C, i.e., put

$$\Pr_{m_C}\Big(\phi:\ \phi=\psi\Big)\propto K_\psi\prod_{j\in J}\lambda_{\psi(j)}$$

where $K_\psi=2$ for $\psi\in C$ and 1 otherwise. Then m_C satisfies the one-site condition but is different from the unique Gibbs measure μ. $\square$

7 Uniqueness of Gibbs Measures

Now we show that (1) implies (9), and therefore also (10). The proof uses Lemma 5.2 once more, and also the following very nice result of Van den Berg [2], here stated in our hard-constraint context.

Theorem 7.1. [Van den Berg (1993)] *Let G be any board, H any constraint graph and λ any set of activities for H. Suppose that μ_1 and μ_2 are two different Gibbs measures on* $\mathrm{Hom}(G,H)$ *and let ϕ_i be drawn independently from μ_i, $i=1,2$. Then with probability >0 there is an infinite path P in G such that ϕ_1 and ϕ_2 disagree on every site of P.*

Theorem 7.2. *Let H be any dismantlable graph and G any board of bounded degree. Then there is a set of activities for H for which there is only one Gibbs measure on* $\mathrm{Hom}(G,H)$.

Proof. Choose a persistent node j in H, set $n=|H|$, and set $r=\max\{2,\Delta(G)-1\}$, where $\Delta(G)$ is the maximum degree of G. We define λ by $\lambda_j:=r^4n^{4r^n}$, and $\lambda_i:=1$ for $i\neq j$.

Now suppose there are two different Gibbs measures, μ_1 and μ_2, and let ϕ_1 and ϕ_2 be selected from them as in the statement of Theorem 7.1. We want to show that, in fact, with probability 1 there is no "infinite path of disagreement" in G.

Let P^k be any path in G on $k\geq 2$ sites; we claim that with high μ_1-probability, more than half the sites of P^k are mapped to the persistent node j.

Let ψ be any map in $\mathrm{Hom}(G,H)$ such that $\psi(v)=j$ for at most half of the sites v of P^k. By Lemma 5.2, there is a map $\theta\in\mathrm{Hom}(G,H)$ which agrees with ψ on $G\setminus N_{n-2}(P^k)$, is identically j on P^k, and satisfies $\theta^{-1}(j)\supset\psi^{-1}(j)$. Now the Gibbs property for μ_1 assures us that (except on a null set of homomorphisms ψ)

$$\frac{\Pr_{\mu_1}\Big(\phi:\ \phi=\theta\mid\phi\restriction(G\setminus N_{n-2}(P^k))=\psi\restriction(G\setminus N_{n-2}(P^k))\Big)}{\Pr_{\mu_1}\Big(\phi:\ \phi=\psi\mid\phi\restriction(G\setminus N_{n-2}(P^k))=\psi\restriction(G\setminus N_{n-2}(P^k))\Big)}\geq\lambda_j^{k/2}\ ,$$

so we have that, for a.e. ψ,

$$\Pr_{\mu_1}\Big(\phi:\ \phi=\psi\mid\phi\restriction(G\setminus N_{n-2}(P^k))=\psi\restriction(G\setminus N_{n-2}(P^k))\Big)\leq\lambda_j^{-k/2}\ .$$

Now, it follows on integration that, for any homomorphism α from the finite set $N_{n-2}(P^k)$ to H that colors at most half of the sites of P^k with j,

$$\Pr_{\mu_1}\Big(\phi:\ \phi\restriction N_{n-2}(P^k)=\alpha\Big)\leq\lambda_j^{-k/2}\ .$$

The number of sites in $N_{n-2}(P^k)$ is no more than when G is the $(r{+}1)$-regular tree T_r, so that

$$|N_{n-2}(P^k)| \le k + (kr-k+2) + r(kr-k+2) + r^2(kr-k+2) + \cdots + r^{n-1}(kr-k+2)$$
$$= kr^n + 2(r^n-1)/(r-1) \le 2kr^n.$$

It follows that the total number of maps in $\mathrm{Hom}(N_{n-2}(P^k), H)$ is less than n^{2kr^n}. Therefore the μ_1-probability that at most half of the sites of P^k are colored j is at most $n^{2kr^n}\lambda_j^{-k/2} = r^{-2k}$.

The same reasoning applies to μ_2, thus for any fixed path P_k of length $k \ge 2n$, with probability at least $1 - 2r^{-2k}$ both ϕ_1 and ϕ_2 will map more than half of P_k to node j and will thus agree on some site of P_k.

Since there are at most $(r+1)r^{k-2}$ paths on k sites emanating from any given site u, the probability that some P_k starting at u is a path of disagreement tends to zero as k tends to infinity. Thus with probability 1 there is no *infinite* path of disagreement starting at u. Summing over all u, we deduce that there is no infinite path of disagreement anywhere, and now application of Theorem 7.1 completes the proof. □

It is perhaps worth pointing out that a straightforward adaptation of the proof gives us more concrete information about how to construct a set of activities λ for which there is a unique Gibbs measure on $\mathrm{Hom}(G,H)$—we say that such a λ is *forcing* for $\mathrm{Hom}(G,H)$. Indeed, one can prove the following result, via an almost identical proof.

Theorem 7.3. *Let j be a persistent node in a dismantlable graph H, and fix an integer Δ. Then, for any set of activities λ on H, there is a real value x such that if λ' is a set of activities which agrees with λ on $H \setminus \{j\}$ and satisfies $\lambda'_j > x$ then, for any board G of maximum degree at most Δ, λ' is forcing for* $\mathrm{Hom}(G,H)$.

This covers various known results about the hard-core and Widom–Rowlinson models (see, e.g., [3]), proved in essentially the same manner. In the hard-core model, the looped node is persistent, while all three nodes are persistent in the Widom–Rowlinson constraint graph.

8 Frozen Gibbs Measures

To complete our proof of Theorem 4.1, we need to show that each of (6) and (10) implies (1). These results will follow from the existence of a rather bizarre family of Gibbs measures on $\mathrm{Hom}(\mathsf{T}_r, H)$, which in the case where H is stiff are concentrated on a single homomorphism from the $(r+1)$-regular tree T_r to H.

We define a *frozen* coloring of a board G by a constraint graph H to be a homomorphism $\psi \in \mathrm{Hom}(G,H)$ such that, for any patch $U \subset G$, the only homomorphism $\phi \in \mathrm{Hom}(G,H)$ such that $\phi\restriction(G \setminus U) = \psi\restriction(G \setminus U)$ is $\phi = \psi$ itself. This is equivalent to saying that the partial coloring $\psi\restriction\partial U$ determines $\psi\restriction U^+$.

For $\psi \in \mathrm{Hom}(G,H)$, define the probability measure μ_ψ on $\mathrm{Hom}(G,H)$ by $\mathrm{Pr}_{\mu_\psi}\Big(\phi:\ \phi = \psi\Big) = 1$; i.e., μ_ψ is concentrated on the single homomorphism ψ.

Theorem 8.1. *For ψ a frozen coloring in* $\mathrm{Hom}(G,H)$*, and λ any set of activities on H, the measure μ_ψ is a Gibbs measure on* $\mathrm{Hom}(G,H)$.

Proof. Consider any finite subset U of G. Since ψ is frozen, $\psi\restriction U^+$ is the only coloring of U^+ consistent with $\psi\restriction\partial U$, i.e.,

$$\Pr_{m_{U^+}}\Big(\phi:\ \phi\restriction U^+ = \psi\restriction U^+ \mid \phi\restriction\partial U = \psi\restriction\partial U\Big) = 1\ ,$$

in agreement with μ_ψ. $\square$

Theorem 8.2. *Let H be a stiff graph of maximum degree $\Delta \geq 2$. Let λ be any set of activities on H, and r be any integer $\geq \Delta$. Then there are uncountably many frozen colorings $\psi \in \mathrm{Hom}(\mathsf{T}_r, H)$.*

Proof. We may assume that H is connected, since if not then we can consider homomorphisms from some connected component of maximum degree 2 to T_r. Since H is stiff and connected, with a node of degree at least 2, it does not contain any nodes of degree 1.

We construct a homomorphism ψ from T_r to H as follows. Start with any node x_0 of T_r, and color it with any node i_0 of H. Then we color the $r+1$ neighbors of x_0, in such a way that each neighbor of i_0 in H is used at least once. We continue to construct the coloring ψ, working out from x_0. Having colored a site x of T_r, with $\psi(x) = i$, all the r neighbors of x not on the unique path from x_0 are still uncolored: color them in such a way that each neighbor j of i is used on at least one neighbor of x. This will always be possible, since x has $r \geq \Delta$ such neighbors. Proceeding in this way, we construct some homomorphism $\psi \in \mathrm{Hom}(\mathsf{T}_r, H)$. Obviously there are uncountably many choices for ψ. We claim that ψ is frozen.

It suffices (taking supersets if necessary) to show that for any connected patch U containing x_0, the restriction to U of a homomorphism $\psi \in \mathrm{Hom}(\mathsf{T}_r, H)$ is uniquely determined by $\psi \restriction \partial U$. Suppose this assertion to be false and let U be a minimum-size counterexample. Let x be a site in U of maximum distance from x_0, so that all of x's "children" relative to the root x_0 lie in ∂U.

Let $\psi, \psi' \in \mathrm{Hom}(\mathsf{T}_r, U^+)$ agree on ∂U, but not on U, and let $i = \psi(x)$. By construction, for every neighbor j of i in H, there is a neighbor y of x in ∂U with color j. Thus $\psi'(x)$ has to be adjacent to every neighbor j of i, but since H is stiff, this means that $\psi'(x) = i$. But then $U \setminus x$ is already a violating patch, contradicting the minimality of U. $\square$

If H is stiff with maximum degree 1, then it is a disjoint union of edges. In this case, for any r, the choice of color of any single site of T_r determines the entire homomorphism, so there are $|H| \geq 2$ elements of $\mathrm{Hom}(\mathsf{T}_r, H)$, and all are frozen.

Theorems 8.1 and 8.2 show that stiff graphs do not satisfy property (10) of Theorem 4.1. Moreover, they do not satisfy property (6) either: take any frozen coloring ψ of T_r by a stiff graph H, and any $\phi \neq \psi$. Then for any site u with $\phi(u) \neq \psi(u)$, and any patch U containing u, there is no common extension of $\psi \restriction (\mathsf{T}_r \setminus U)$ and $\phi \restriction \{u\}$.

Now consider a general non-dismantlable graph H. Take a sequence of folds reducing H to a stiff subgraph J. The composition α of folds in the sequence is a retract from H to J, i.e., a homomorphism from H to J such that $\alpha \restriction J$ is the identity. Then every homomorphism ϕ in $\mathrm{Hom}(G, H)$ induces a homomorphism $\alpha\phi \in \mathrm{Hom}(G, J)$.

Now take ψ to be a frozen coloring of some T_r by J, and let U be any patch of T_r. Suppose that $\phi \in \mathrm{Hom}(\mathsf{T}_r, H)$ has $\alpha\phi \restriction \partial U = \psi \restriction \partial U$; then $\alpha\phi$ is a homomorphism from T_r to J agreeing with ψ on ∂U, so it must also agree with ψ on all of U^+.

This already suffices to show that (6) implies (1) in general, in other words that a non-dismantlable graph H never satisfies (6). Proceeding as before, take a frozen coloring ψ of T_r by J, and a different $\phi \in \mathrm{Hom}(T_r, J)$. For any patch U, any extension θ of $\psi \restriction (\mathsf{T}_r \setminus U)$ must have $\alpha\theta = \psi$, so θ cannot agree with ϕ on any site where ψ and ϕ differ.

Now fix any set λ of activities on H. We can now follow a standard technique to construct a Gibbs measure giving probability 1 to the set $\mathcal{S}$ of homomorphisms $\phi \in \mathrm{Hom}(\mathsf{T}_r, H)$ with $\alpha\phi = \psi$. To do this, we construct a sequence of measures (μ^k) from the fixed homomorphism ψ. Fix a "root" x_0 of T_r, let $U = \{x_0\}$, and consider the sets $N_k(U)$, for k a non-negative integer. The measure μ^k gives positive probability only to homomorphisms agreeing with ψ outside $N_k(U)$; for such a homomorphism θ, we set

$$\Pr_{\mu^k}\Big(\phi:\ \phi=\theta\Big) = \Pr_{m_{N_{k+1}(U)}}\Big(\phi:\ \phi=\theta\restriction N_{k+1}(U) \mid \phi\restriction\partial N_k(U)=\theta\restriction\partial N_k(U)\Big)$$

(note that $N_{k+1}(U) = (N_k(U))^+ = N_k(U) \cup \partial N_k(U)$); in other words, we fix $\phi = \psi$ on $\partial N_k(U)$, and then select ϕ according to the multiplicative measure $m_{N_{k+1}(U)}$, conditioned on the boundary values. It is a well-known result (see, e.g., [5] or [6]) that the limit of any convergent subsequence of the μ^k, in the usual compact topology, is a Gibbs measure on $\mathrm{Hom}(G, H)$.

The discrete measures μ^k are positive only on homomorphisms ϕ agreeing with ψ outside a finite subset, all of which are in $\mathcal{S}$; hence the limit μ of any convergent subsequence of the μ^k is a Gibbs measure such that $\mu(\mathcal{S}) = 1$, as claimed.

This shows that, for any non-dismantlable graph H, and any set λ of activities on H, there are multiple Gibbs measures, at least one for each frozen coloring of the derived stiff graph J. This shows that (10) implies (1).

This completes the chain of implications required to prove Theorem 4.1.

9 Branching Random Walks

If the constraint graph H is non-dismantlable, among the consequences is the failure of property (6) of Theorem 4.1, which can be thought of as showing the possibility of "long-range influence". In this case, one can construct a homomorphism ψ from some suitably highly branching T_r to H such that, no matter how far from the root one goes, the values of ψ at that distance convey "hard" information, sufficing to rule out some root color.

We conclude with one more characterization of dismantlable constraint graphs, showing that this phenomenon is not so freakish as one might think, and does not rely on our ability to construct homomorphisms carefully. We shall show that, in the natural setting of branching random walks, we see long-range influence with high probability, provided the branching number is large enough (in fact, the number required is only modestly larger than that required for the existence of frozen homomorphsims).

Given a weight vector $\vec{w} = (w_1, \ldots, w_n)$ assigning positive reals to the nodes of H, a node-weighted *random walk* on H is a Markov chain whose states are the nodes of H and whose transition probabilities are $p_{ij} := w_j/z_i$ when $j \sim i$ and 0 otherwise, where $z_i := \sum_{k\sim i} w_k$.

We may think of a random walk on H as a token which steps randomly from node to node along the edges of H, with probabilities weighted by $\vec{w}$. If the token is replaced by an amoeba, which divides r ways before each step, the result is an r-*branching* random walk on H. Each amoeba-child steps independently of its siblings, and indeed of every other amoeba on H, and many may occupy the same node simultaneously.

If we agree to place the initial amoeba at a node drawn from the stationary distribution, and divide $r+1$ ways at step 1 only and r ways thereafter, then the branching random walk defines a probability measure $\mu_{\vec{w}}$ on $\mathrm{Hom}(\mathsf{T}_r, H)$. To obtain a random ϕ from $\mu_{\vec{w}}$, we fix a root u of T_r and let $\phi(u)$ be the position of the initial amoeba. That amoeba's $r+1$ children determine the values of ϕ on the neighbors of u, and so forth.

In fact, $\mu_{\vec{w}}$ is a particularly nice Gibbs measure for a certain set of activities. A measure on $\mathrm{Hom}(\mathsf{T}_r, H)$ is said to be *simple* if whenever $\phi(v)$ is fixed at some site v, the behavior of ϕ on each of the $r+1$ components of $\mathsf{T}_r \setminus \{v\}$ is independent. It is *invariant* if it is unchanged by any of the (many) automorphisms of T_r.

Theorem 9.1. ([4]) *Fix r and a constraint graph H with set of activities $\lambda = (\lambda_1 \dots \lambda_n)$. Then the simple invariant Gibbs measures on* $\mathrm{Hom}(\mathsf{T}_r, H)$ *are precisely the measures $\mu_{\vec{w}}$ arising from branching random walks on H with weight vector $\vec{w}$ satisfying $\lambda_i = w_i/z_i^r$. Moreover, there is always at least one such weight vector.*

The main objective of [4] is to characterize the graphs H for which there can be more than one simple invariant Gibbs measure for the same set of activities. That will not be the issue here.

Unless H is bipartite an ordinary random walk on H exhibits no long-range memory; in other words, all information about the state of the walk at time 0 is lost as time advances. In a branching random walk memory may persist (and does, with high probability, when r is large). When H is non-dismantlable something even more startling occurs.

Theorem 9.2. *Let H be a non-dismantlable graph with weight vector $\vec{w}$. Then there is an integer r, a real $\delta > 0$ and a node i of H such that for any t, the state of an r-branching random walk on H at time t is, with probability $> \delta$, inconsistent with initial state i.*

Thus we have (with probability bounded away from 0) *hard* information about the initial state of the branching random walk, even after an arbitrarily large amount of time is past. This is much stronger than the negation of conditions (5) or (6) of Theorem 4.1; in particular if we take an r-branching random walk on a dismantlable graph then there is a time t past which *any* initial state is *always* possible.

Proof of Theorem 9.2. Given the non-dismantlable constraint graph H we define J and the C_j ($j \in J$) as in §4. Let T_r^t be the subtree of T_r consisting of the root and all sites at distance at most t from the root. Thus there are $(r+1)r^{t-1}$ leaves in T_r^t, and all other sites have degree $r+1$. Given a coloring $\phi \in \mathrm{Hom}(T_r^t, H)$, consider the following labeling scheme, using the vertices of J as labels.

(i) If a leaf has a color from C_j, it receives label j. Otherwise it is unlabeled.

(ii) For each $j \in J$, if a non-leaf x has successors y_i labeled i, for each $i \in J$ adjacent to j, then x receives label j. Otherwise x is unlabeled.

Lemma 9.3. *If a site x receives label j, then $\phi(x) \in C_j$.*

Proof. We verify this inductively, working back from the leaves. The statement is certainly true if x is a leaf.

If a non-leaf x receives label j, then it has successors y_i labeled i, and thus—by the induction hypothesis—with $\phi(y_i) \in C_i$, for each $i \sim j$. The site x must have a color adjacent to all the colors $\phi(y_i)$, and this color will therefore be in C_j, by definition of C_j. □

We now take an r-branching node-weighted random walk on H (r to be chosen later) in accordance with the scheme for generating a map $\phi \in \mathrm{Hom}(\mathsf{T}_r, H)$ from the distribution $\mu_{\vec{w}}$. For each t, $\phi\restriction\mathsf{T}_r^t$ records the history of the walk to time t, and on it we may carry out the labeling described above. Clearly, if a site receives a label in the labeling of T_r^t, then it also receives a label in the labeling of T_r^s for each $s < t$. We can then label T_r by giving a site a label j if it receives label j in the labeling of *every* T_r^t. Let $\mathcal{L}$ be the event that the root is labeled.

Unless H is dismissably uninteresting we cannot expect $\mathcal{L}$ to have probability 1, since if j is a vertex j of J of degree greater than 1 there is a positive probability that the root is colored j and all its immediate successors have the same color; this prevents the root being labeled.

We claim, however, that for r sufficiently large, $\mathcal{L}$ has positive probability. Our estimates for how large an r is required will naturally depend on the transition probabilities in the random walk.

Accordingly, we choose $\varepsilon > 0$, and subsets A_j of the C_j with the property that, whenever $i \sim j$ in J and $h \in A_i$, we have

$$\frac{1}{z_h} \sum_{g \in A_j, g \sim h} w_g \geq \varepsilon.$$

In other words, whenever the random walk is in some state in A_i, and $j \sim i$, the probability that it will step to a state in A_j is at least $\varepsilon > 0$. (For instance, we could take $A_j = \{j\}$ for every j, or we could take $A_j = C_j$ for every j, but it may well be that some intermediate choice allows a higher value of ε.)

Lemma 9.4. *Let H, J, $\vec{w}$, $\varepsilon > 0$ and the sets A_j, for $j \in J$, be as above, and let Δ be the maximum degree in J.*
(i) Set $\alpha = \alpha_\Delta$ to be the unique root of

$$\alpha e^{-\alpha} = \frac{1}{e\Delta}.$$

Suppose that $r \geq \alpha/\varepsilon$. Then, for each j, the probability that the root is labeled, conditional on it receiving a color from A_j, is at least $1 - 1/\alpha$.
(ii) For $\Delta \geq 3$, and

$$r \geq \frac{1}{\varepsilon} \frac{\log \Delta + \log\log \Delta}{1 - 1/\log \Delta},$$

the probability that the root is labeled, conditional on it receiving a color from some A_j, is at least $1 - 1/\log \Delta$.

In either case, $\Pr(\mathcal{L}) > 0$.

Proof. For convenience, we ignore one branch of the tree leading from the root, so that every site has forward-degree r. This clearly does not increase the probability that the root is labeled.

Now, for $t \geq 0$, we consider the labeling of the tree T_r^t, generated from the colors at its leaves. Let p_t be the minimum, over all colors i in $\bigcup_{j \in J} A_j$, of the probability that the root is labeled in the labeling of T_r^t, conditional on its color being i. The probability that the root is labeled in T_r, conditional on its color being from a given A_j, is then at least $\lim_{t\to\infty} p_t$.

With all probabilities taken according to the branching random walk on T_t^r, we have, for any $i \in J$ and $h \in A_i$,

$$\begin{aligned}
&\Pr(\text{root is labeled} \mid \text{root has color } h) \\
&\quad \geq \; 1 - \sum_{j \sim i} \Pr(\text{no successor is labeled } j \mid \text{root has color } h) \\
&\quad = \; 1 - \sum_{j \sim i} \left(1 - \Pr(\text{given successor is labeled } j \mid \text{root has color } h)\right)^r.
\end{aligned}$$

Now, if the root has a color h from A_i, then a given successor x has a color from A_j with probability at least ε, and the probability that x then actually receives the label j is at least p_{t-1}, since the subtree rooted at x is a copy of T_r^{t-1} with one branch removed. Thus we have

$$\Pr(\text{root is labeled} \mid \text{root has color } h) \geq 1 - |N(i) \cap J| \left(1 - \varepsilon p_{t-1}\right)^r,$$

for every $i \in J$ and $h \in A_i$. Therefore

$$p_t \geq 1 - \Delta \left(1 - \varepsilon p_{t-1}\right)^r \geq 1 - \Delta e^{-\varepsilon p_{t-1} r}.$$

We need to check that this recursion gives the conclusion of the theorem in the two cases. For (i), let us assume that α is as given, that $r\varepsilon \geq \alpha$, and that $p_{t-1} \geq 1 - 1/\alpha$. (This is certainly valid for $t = 1$, since $p_0 = 1$.) Now we have

$$p_t \geq 1 - \Delta e^{-\varepsilon p_{t-1} r} \geq 1 - \Delta e^{-\alpha(1-1/\alpha)} = 1 - e\Delta\left(\frac{1}{e\Delta\alpha}\right) = 1 - \frac{1}{\alpha}.$$

Hence, by induction, $p_t \geq 1 - 1/\alpha$ for every t, as required.

Now suppose that $\Delta \geq 3$ and that r satisfies the bound of (ii). Again, we have $p_0 = 1$. If $p_{t-1} \geq 1 - 1/\log\Delta$, then $\varepsilon p_{t-1} r \geq \log\Delta + \log\log\Delta$, so

$$p_t \geq 1 - \Delta e^{-\log\Delta - \log\log\Delta} = 1 - 1/\log\Delta.$$

Hence, by induction, $p_t \geq 1 - 1/\log\Delta$ for every t. □

It is now easy to complete the proof of Theorem 9.2. If r is chosen large enough, as in Lemma 9.4, then there is positive probability that the root will receive a color from some A_j, and that it will then be labeled. In this case, by Lemma 9.3, colors not in C_j are ruled out as colors of the root by the colors of the sites at any given distance from the root. □

We note for convenience that $\alpha_2 \simeq 2.68$, $\alpha_3 \simeq 3.29$, $\alpha_4 \simeq 3.69$, $\alpha_5 \simeq 3.99$. For higher values of Δ, the bound in part (ii) of Lemma 9.4 indicates a reasonable approximation.

To take a specific example, we return to the graph-coloring case where H is an unlooped complete graph K_n, and all node-weights are equal (forcing all activities to be equal as well). Then $J = H$, and each C_j – hence necessarily each A_j – consists of the single vertex j. We can take $\varepsilon = 1/\Delta = 1/(n-1)$, and hence there is positive probability that the root is labeled provided that $r \geq (n-1)\alpha_{n-1}$. For $n > 3$, it suffices to have

$$r \geq (n-1)\frac{\log(n-1) + \log\log(n-1)}{1 - 1/\log(n-1)} .$$

For this example, it is easy to check that this bound on r is at least asymptotically best possible (as $n \to \infty$).

To be even more specific, the reconstruction strategy succeeds on $H = K_3$, with uniform activities, when $r \geq 6$, with probability at least 0.6. Moreover, one can write down the precise recursion for the p_t in this case, namely

$$p_t = h(p_{t-1}) \equiv 1 - 2\left(1 - \frac{p_{t-1}}{2}\right)^r + (1 - p_{t-1})^r .$$

For $r = 5$, the p_t tend to the largest root of $p = h(p)$, which is at approximately $p = 0.8988$. For $r = 3$ or 4, there is no root of $p = h(p)$, and the p_t tend to 0, so the probability of being able to reconstruct the root with certainty tends to 0 for $r \leq 4$.

References

[1] H.-J. Bandelt and H.M. Mulder, Helly theorems for dismantlable graphs and pseudo-modular graphs, in: *Topics in Combinatorics and Graph Theory (Oberwolfach, 1990)*, Physica, Heidelberg 1990, 65–71.

[2] J. van den Berg, A uniqueness condition for Gibbs measures with application to the 2-dimensional Ising antiferromagnet, *Comm. Math. Phys.* **152** (1993), 161–166.

[3] J. van den Berg and J.E. Steif, Percolation and the hard-core lattice gas model, *Stochastic Proc. and their Appls.* **49** (1994), 179–197.

[4] G. Brightwell and P. Winkler, Graph homomorphisms and phase transitions, *J. Comb. Theory* Series B, to appear.

[5] P.L. Dobrushin, The description of a random field by means of conditional probabilities and conditions of its regularity, *Thy. of Prob. and its Appls.* **13** #2 (1968), 197–224.

[6] H.-O. Georgii, *Gibbs Measures and Phase Transitions*, de Gruyter, Berlin (1988).

[7] P. Hell and J. Nešetřil, The core of a graph, *Disc. Math.* **109** (1992), 117–126.

[8] F.P. Kelly, Stochastic models of computer communication systems, *J. R. Statist. Soc. B* **47** #3 (1985), 379–395.

[9] R. Nowakowski and P. Winkler, Vertex-to-vertex pursuit in a graph, *Discrete Math.* **43** (1983), 235–239.

[10] N. Polat, A Helly theorem for geodesic convexity in strongly dismantlable graphs, *Disc. Math.* **140** (1995), 119–127.

[11] A. Quilliot, Homomorphismes, points fixes, rétractions et jeux de pousuite dans les graphes, les ensembles ordonnés et les espaces métriques, Thése d'Etat, Université de Paris VI, Paris, France (1983).

[12] B. Widom and J.S. Rowlinson, New model for the study of liquid-vapor phase transition, *J. Chem. Phys.* **52** (1970), 1670–1684.