Indecomposable Coupled Map Lattices with Non-unique Phase

R.S.MacKay

Mathematics Institute, University of Warwick, Coventry CV4 7AL, U.K.
mackay@maths.warwick.ac.uk

1 Introduction

A compact topologically mixing uniformly hyperbolic attractor for a diffeomorphism of a finite-dimensional manifold with Hölder continuous derivative carries a unique probability measure describing the long-time statistics of the forward orbits of almost all initial conditions in its basin. This was proved [Si] by converting the orbits of the dynamical system into symbol sequences from a finite alphabet (indexed by time) and considering them as states of a statistical mechanical spin chain with a certain interaction energy which decays exponentially with separation, so leading to a unique phase. The results have been extended to many classes of non-uniformly hyperbolic system by different approaches (e.g. [Vi]).

What about spatially extended dynamical systems? By analogy with statistical mechanics, could increasing from one dimension (time) to more than one (time plus space) permit non-unique probabilistic behaviour (even if the system is indecomposable in an appropriate sense at the topological level)? This question was raised by [Ha, BuS], for example.

Because of the statistical mechanics analogy, I shall refer to the possible probabilistic behaviours of dynamical systems as *phases*. These are probability distributions on the set of all orbits, not just on the state space.

Although infinite systems are probably a fiction, they are a good idealisation to study for such questions, because (i) under conditions to rule out associated spin chains with infinitely many states per site or slowly decaying interaction, finite-dimensional topologically mixing systems have unique phase, and (ii) non-unique phase for an infinite system can be reflected in extraordinarily long transients for large but finite versions (for a proved example, see [T+]).

For many researchers the interest in spatially extended dynamical systems focusses on partial differential equations (PDEs), but analysis of PDEs is highly technical, which makes it difficult to address the essential phenomena. The question of non-unique phase is already interesting for spatially discrete systems, however, such as networks of oscillators [Ku]. I shall address the question in the simplest context: coupled map lattices (CML) [Ka84], where time also is taken to be discrete.

Despite much numerical work suggesting non-unique phase can occur in CML (e.g. [MH, CM] and Just's lectures in this school) and some papers claiming to prove it or make it plausible for some examples ([Sa] comes very close), the paper [GM] appears to be the first to give a proved example. I do not count examples where the non-uniqueness arises purely from a bifurcation at the topological level, or where it refers to escape measures on unstable sets (e.g. [JP] and Just's lectures).

As has been seen in other lectures in this school, piecewise affine systems are much easier to analyse and yet can exhibit a lot of the phenomenology of smooth ones. Following [GM], I shall construct a variety of piecewise affine CML with provably non-unique phase. The examples illustrate not just non-unique invariant measure but also the phenomena of non-trivial collective behaviour, eternal transience and an arrow of time. Although somewhat artificial, it is better to have some concrete examples than nothing at all. An analogy is the 2D Ising model of statistical mechanics, which is a toy model of a ferromagnet but gave the first context in which it was proved that non-unique phase can arise (I don't count mean field approximations).

The examples all have "symbolic dynamics", which again makes them easier to analyse. Systems with symbolic dynamics have a particular class of phase called *Gibbs phases*. I will argue that for systems with symbolic dynamics, Gibbs phases form the physically relevant class, describing the statistics of local observables on the orbits of all "uncertain" initial conditions started a long time in the past.

In more detail, Section 2 begins by giving four examples of CML with provably non-unique Gibbs phase. The key idea is to simulate probabilistic cellular automata (PCA) for which analogous non-uniqueness has already been proved. Then I give an appropriate definition of indecomposability and show that these examples are indecomposable. Next, an example is given with an absorbing subset which looks as if it ought to attract the orbit of almost every initial condition, yet there is a large subset of initial conditions which never discover it, despite the system being "pre-indecomposable". Then I show how to build invertible versions of these examples. Finally, I make a time-reversible indecomposable CML with non-reversible phases.

Section 3 motivates the concept of Gibbs phase for CML with symbolic dynamics, by building up from the simpler cases of Markov chains, ordinary dynamical systems and PCAs. Some readers may prefer to read this section before section 2, but I consider it better to go straight to the examples, which have obviously relevant non-unique phases, before going into details of what class of phases should be considered for general CML with symbolic dynamics. The discussion is by no means complete, but I hope it helps.

The notes conclude with some challenges and an appendix on the concept of reversibility for Markov chains and PCA.

2 A Selection of Examples

2.1 CML with Ferromagnetic phases and variants

Example 1 Let "space" $S = \mathbb{Z}^2$ and "time" $T = \mathbb{Z}$. For each "site" $s = (s_1, s_2) \in S$, define its neighbourhood $N(s) = \{s, s + (1,0), s + (0,1)\}$ (the "North-East-Centre" (NEC) neighbourhood). Take "local state space" $M = [0,1]$. For "lattice state" $x = (x_s)_{s \in S} \in M^S$ define "lattice symbolic state" $\sigma \in \{+,-\}^S$ by (for each $s \in S$):

$$\sigma_s(x) = \begin{cases} + \text{ if } x_s \geq \frac{1}{2} \\ - \text{ if } x_s < \frac{1}{2} \end{cases} .$$

For $\varepsilon \in (0,1]$, apply at each site a piecewise affine map from Fig. 1 according to the symbolic state of the NE neighbours. The slopes are $\frac{2}{\varepsilon}$ and $\frac{2}{2-\varepsilon}$. This defines a CML, even though the influence of neighbours is via only their symbolic state rather than continuously (like the more familiar discrete diffusion). Note that in contrast to common usage, ε represents deviation from deterministic symbolic dynamics rather than coupling strength; the uncoupled case is $\varepsilon = 1$. Given initial condition $x^{-N} \in M^S$ at some large negative time $-N$, denote the state at times $t \in T$ with $t \geq -N$ by x^t.

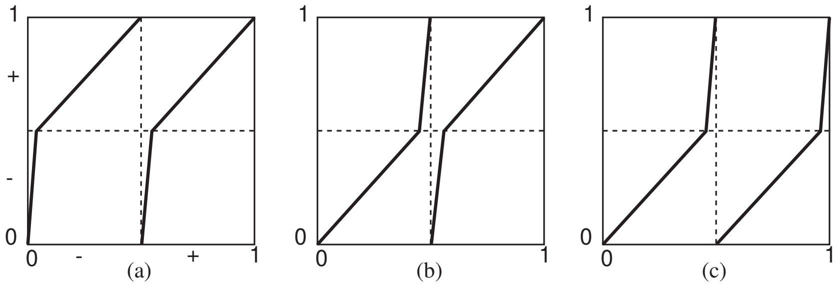

Fig. 1. Local maps for a CML with ferromagnetic phases for $\varepsilon < \varepsilon_c$ (here $\varepsilon = 0.13$), for symbolic state of the NE neighbours: (a) $++$, (b) $+-$ or $-+$, (c) $--$.

For ε near 1, this CML has a unique Gibbs phase G. In particular, by the $x \mapsto 1 - x$ symmetry of the CML, $G(\sigma_s(x^t) = +) = \frac{1}{2}$ for all $(s,t) \in S \times T$. In contrast, for ε small enough, there is more than one Gibbs phase. The set of Gibbs phases for a system is convex, so can be described by the *pure* (or "extremal") Gibbs phases, those which can not be written as a non-trivial convex combination of others. Thus I will make the statement in terms of pure Gibbs phases [GM]:

There exists $\varepsilon_c \in (0,1)$ (numerically about 0.18 [BG]) such that for all $\varepsilon \in [\varepsilon_c, 1]$ there is a unique Gibbs phase, whereas for each $\varepsilon \in (0, \varepsilon_c)$ there are pure Gibbs phases $G^{\pm}(\varepsilon)$ and $c(\varepsilon) > \frac{1}{2}$ such that for all $(s,t) \in S \times T$,

$$G^+\left(\sigma_s(x^t)=+\right)=G^-\left(\sigma_s(x^t)=-\right)=c;$$

furthermore, $G^\pm$ are space-time translation invariant.

By analogy with statistical mechanics, I say the phases $G^\pm$ are "ferromagnetic". To simplify the following discussion of this example I'll assume that $G^\pm$ are the only pure Gibbs phases (though I'm not aware of a proof).

What is the implication of this result for typical initial conditions? As I'll explain in Section 3, for spatially infinite systems I don't know how to interpret "typical", so instead I think in terms of "uncertain" initial conditions. An *uncertain initial condition* is a probability distribution for initial conditions which has absolutely continuous marginal on each finite subset L of S, whose densities ϱ_L satisfy some further conditions. Then the answer is that for any bounded subset Λ of space-time $S\times T$ and initial time sufficiently far in the past, the probability distribution for the components in Λ of the orbits is some convex combination $p_+G^+ + p_-G^-$ with $p_\pm \geq 0$, $p_+ + p_- = 1$, depending on the initial condition, initial time and Λ. In words, you should see samples from the probability distribution G^+ with probability p_+ and from G^- with probability p_-. For most initial conditions drawn from the initial probability distribution, I think this translates into seeing large patches with the statistics of G^+ and large patches with those of G^-, separated by narrow wandering domain walls, though the logical implication is not clear to me.

Example 2 The next example is given by Fig. 2. For $\varepsilon \in (0,\varepsilon_c)$ it has two pure Gibbs states G^{es} (for "even sites") and G^{os} (for "odd sites") with

$$G^{es}\left(\sigma_s(x^t)=(-)^{s_1+s_2}\right)=G^{os}\left(\sigma_s(x^t)=-(-)^{s_1+s_2}\right)=c$$

(the same ε_c and $c(\varepsilon)$ as above). So you see chessboard patterns. The phases are "antiferromagnetic".

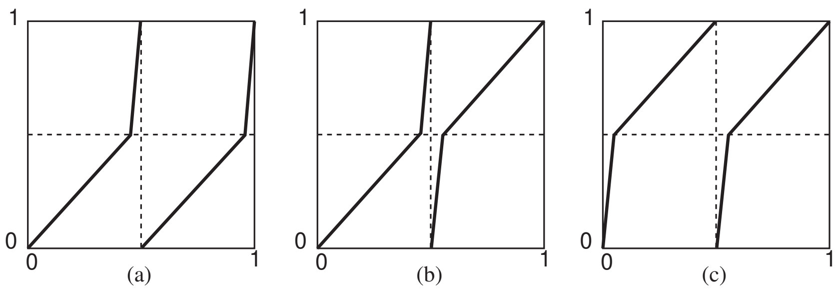

Fig. 2. Local maps for a CML with antiferromagnetic phases for $\varepsilon<\varepsilon_c$, for symbolic state of the NE neighbours: (a) $++$, (b) $+-$ or $-+$, (c) $--$.

Example 3 The third example uses the maps in Fig. 3 (equivalently, interchange ε with $2-\varepsilon$ in the first example). For $\varepsilon \in (0,\varepsilon_c)$ it has two extremal

Gibbs states G^{et} (for "even time") and G^{ot} (for "odd time") with

$$G^{et}\left(\sigma_s(x^t) = (-)^t\right) = G^{ot}\left(\sigma_s(x^t) = -(-)^t\right) = c.$$

So a period-2 behaviour emerges, alternating between a predominance of $+$ and $-$. The phases are "period-2 ferromagnetic". They are examples of "non-trivial collective behaviour" [CM, Ka90] or "asymptotic periodicity" [LM]. This is a type of non-uniqueness that would not be apparent if one restricted attention to invariant measures, and yet is highly significant. Hence I think it important to look at measures on the set of orbits, not just invariant ones (on the state space).

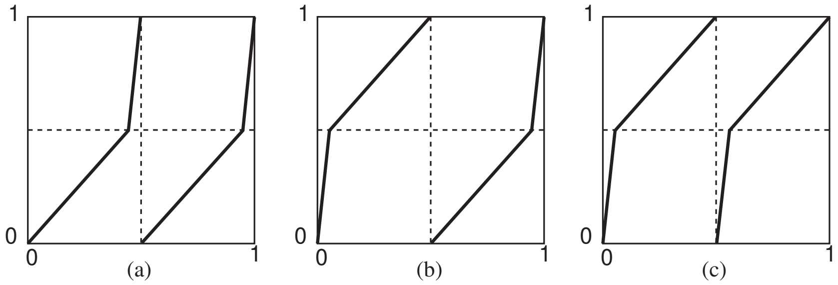

Fig. 3. Local maps for a CML with period-2 ferromagnetic phases for $\varepsilon < \varepsilon_c$, for symbolic state of the NE neighbours: (a) $++$, (b) $+-$ or $-+$, (c) $--$.

Example 4 Finally, use the maps in Fig. 4 (equivalently, interchange ε with $2-\varepsilon$ in the second example). For $\varepsilon \in (0, \varepsilon_c)$ it has two extremal Gibbs states G^e and G^o (for "even" and "odd" space-time) with

$$G^{e}\left(\sigma_s(x^t) = (-)^{s_1+s_2+t}\right) = G^{o}\left(\sigma_s(x^t) = -(-)^{s_1+s_2+t}\right) = c.$$

These phases are "period-2 antiferromagnetic".

2.2 How the constructions work

The idea of these constructions is that (i) there are indecomposable probabilistic cellular automata (PCA) with non-unique Gibbs phase, and (ii) any PCA can be simulated by some CML.

A PCA is a stochastic process on an infinite product of finite state spaces generated by transition probabilities given by independent spatially local rules. A good reference on theory and examples of PCA is [T+].

The examples of the previous subsection are based on a case of Toom's *majority voter PCA*. Its state space is $\{+,-\}^S$ with $S = \mathbb{Z}^2$. Let m be the majority sign in the NEC neighbourhood of s. Then the σ_s^t are updated independently and synchronously by transition probabilities:

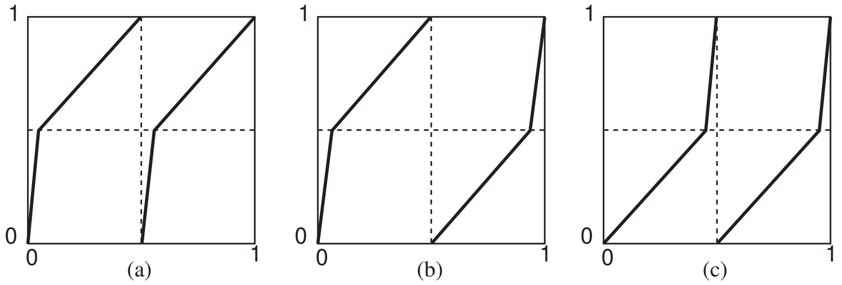

Fig. 4. Local maps for a CML with period-2 antiferromagnetic phases for $\varepsilon < \varepsilon_c$, for symbolic state of the NE neighbours: (a) $++$, (b) $+-$ or $-+$, (c) $--$.

$$P(\sigma_s^{t+1} = m|\sigma^t) = 1 - \frac{\varepsilon}{2}$$
$$P(\sigma_s^{t+1} = -m|\sigma^t) = \frac{\varepsilon}{2}.$$

Toom proved that there is $\varepsilon_c \in (0, 1)$ such that starting from the all $+$ state at $t = 0$, for $\varepsilon \in (0, \varepsilon_c)$ there is $c > \frac{1}{2}$ such that for all $s \in S$ and $t \in Z_+$, $P(\sigma_s^t = +) \geq c$ (and converges to c as $t \to \infty$) [T+, LMS]. I will not go into the proof here, except to remark that it uses a stochastic analogue of monotonicity, a concept that has been seen to be useful in other lectures at this school. Specifically, for all $s \in S$,

$$P(\sigma_s^{t+1} = +|\sigma^t = \sigma') \geq P(\sigma_s^{t+1} = +|\sigma^t = \sigma'')$$

whenever $\sigma' \geq \sigma''$ in the partial order defined by $\sigma'_s \geq \sigma''_s$ for all $s \in S$ in the order $+ > -$. This property is called "attractivity" in some of the literature.

To simulate a PCA by a CML, if the PCA has N states for each site ($N = 2$ here) and transition probabilities $p_{ij}(n)$ from state i to j at site s given the state n of the neighbours (excluding s itself), use for each n a piecewise affine map of $M = [0, 1]$ to itself subdivided into N intervals of equal length to represent the states, with slopes $\frac{1}{p_{ij}(n)}$ from the i^{th} subinterval to the j^{th}. One can allow the lengths of the subintervals to vary, with corresponding modifications of the slopes (e.g. [Sa]), but ours seems the easiest general prescription (Sakaguchi's produces a second order CML, i.e. x_s^{t+1} depends on x_s^{t-1} as well as x^t, and is less flexible in the set of PCAs it can simulate).

In what sense does the CML simulate the PCA? Firstly, if one considers probability distributions on M^S with constant marginal density on each set corresponding to the symbolic state of a finite subset of S, then the induced map on the densities preserves this property and is precisely the action of the PCA on probabilities. There is a deeper sense, however. Under iteration of the induced map on probability measures on M^S, for any finite subset $L \subset S$, the marginal density ρ_L of any uncertain initial condition on $[0, 1]^S$ becomes asymptotically constant on each cylinder set (i.e. set corresponding to a given symbolic state on L) [GM].

Knowledgeable readers might have expected me to use a 2D kinetic Ising model instead of Toom's PCA. This will be done in Section 2.7, but Toom's PCA has the advantage of robustness: one can modify the transition probabilities somewhat and still keep the phenomenon of non-unique phase (e.g. [BG]). This is in contrast to kinetic Ising models where breaking the $+/-$ symmetry reduces them generically to a unique phase.

Toom's PCA is 2D. Are there 1D examples with non-unique phase (and yet indecomposable in a good sense)? Yes. Indeed there is a 1D PCA with infinitely many pure phases [Ga], but it is complicated and not entirely explicit (see [Gr] for an introduction to the paper)! It would be nice to have a simpler 1D example, and two pure phases would suffice.

2.3 Indecomposability

It is clear that non-unique phase can arise in trivial ways if we do not require the system to be indecomposable in some sense. If a dynamical system has more than one attractor there will be at least one phase for each. Even if there is one attractor but it consists of several pieces which are cyclically permuted there will be at least one phase for each of the pieces in which the orbits can start (some people would want to identify them because they give the same time-averages but I think it is useful to distinguish them).

The same goes for Markov chains. There it is usual to require irreducibility and aperiodicity, because any Markov chain can be reduced to pieces with these properties. A Markov chain is *irreducible* if it has no proper forward invariant subset; equivalently, if for all states i, j it is possible to get from i to j in a positive number of steps. An irreducible Markov chain is *aperiodic* if there is N such that for all i, j it is possible to get from i to j in precisely N steps (if not then it is the "composition" of a deterministic periodic cycle with an aperiodic irreducible Markov chain). An aperiodic irreducible finite-state Markov chain has a globally attracting invariant measure, thus a unique phase, whereas a period-M irreducible finite-state Markov chain has a unique invariant measure but M phases.

It is appropriate here to make a warning about terminology. Many probabilists (though not all, e.g. [Ke]) say a system with a globally attracting invariant measure is *ergodic*, whereas dynamicists say it is *mixing* and use the word "ergodic" for an invariant measure that can not be decomposed into two (i.e. an extremal invariant measure).

For dynamical systems, a good notion of indecomposability is *topological mixing*: for any two non-empty open sets A, B in the state space, the image of A intersects B at all large enough times. This could be applied to CML by considering product topology on M^S.

A stronger notion of indecomposability, however, is Bowen's *specification property* (e.g. [KH]): for all $\delta > 0$ there exists $D_0 \in \mathbb{Z}_+$ such that for all finite sequences of orbit segments $x(i), i = 1, \ldots, I$ on consecutive time intervals

$[a_i, b_i]$ with separations $a_{i+1} - b_i \geq D_0$ there exists an orbit y agreeing to within δ with $x(i)$ at each time in $[a_i, b_i]$ for each $i \in \{1, \ldots, I\}$.

For CML, an appropriate version of the specification property should refer to space as well as time. Thus I say a CML on a space S is *indecomposable* if for all $\delta > 0$ there exists $D_0 \in \mathbb{R}_+$ such that for all finite sets of bounded subsets $A_i \subset S \times T$, $i = 1, \ldots, I$, with separations at least D_0, and orbits $x(i)$, there exists an orbit y agreeing to within δ with $x(i)$ on A_i for each i.

The *separation* of two subsets $A, B \subset S \times T$ is defined by $D(A, B) = \inf\{d(a, b) : a \in A, b \in B\}$. For metric d on $S \times T$ one can use $d((s,t),(s',t')) = d_S(s, s') + |t' - t|$ (or the maximum) where d_S is any metric on S such that the neighbourhoods $N(s)$ have bounded diameter.

This is a stronger definition than used in [GM] but still holds for many appropriate CML with symbolic dynamics, including our Examples 1–4. For these examples, let $\lambda = \frac{2}{2-\varepsilon}$ and choose $D_0 > 2\frac{\log 1/\delta}{\log \lambda}$. Let B_i be the $D_0/2$-neighbourhood of A_i. Let σ be the symbolic state of $x(i)$ on each B_i, extended arbitrarily on the rest of $S \times T$. Let y be the orbit corresponding to σ. Each y_s^t is determined by $\Sigma_s^t = (\sigma_r^u)_{r \in N(s), u \geq t}$, because for any $\tau \geq t$,

$$y_s^t = g_t(\ldots(g_\tau(x_s^{\tau+1})\ldots),$$

where g_u denotes the inverse branch of the local map corresponding to the symbolic state on $N(s)$ at time u. The g_u are contractions by a factor of at least λ, so by taking $\tau \to \infty$ we obtain a unique y_s^t for each choice of Σ_s^t, and all those with given σ on $N(s) \times [t, \tau]$ are contained in an interval of length at most $\lambda^{t-\tau}$. In particular, y is determined to within δ on each site of A_i by σ on B_i, so agrees with $x(i)$ to within δ on each site of A_i.

The above definition of indecomposable can also be applied to PCA by replacing "orbit" by "allowed realisation" and dispensing with δ (it is similar to the "strong irreducibility" of [BS]), though probabilists often ask for a stronger property to rule out events with vanishing probability (see the literature on "positive rates", e.g. [Gr]).

2.4 CML with Eternal Transients

An interesting phenomenon discovered numerically by [CK] and again by Livi et al. (see [GLP] for recent work of theirs and references), is that some spatially extended systems with a small attracting subset can take extraordinarily long times to discover the attractor; instead they bounce around for a long time in a chaotic transient. This suggests that for the infinite system there is an additional pure phase to that on the attracting subset. I present here a proved example of this.

Example 5 Let $S = \mathbb{Z}$ and $N(s) = \{s, s+1\}$ for each $s \in S$. For $x_s \in [0, 1]$ define $\sigma_s = \pm$ according as $x_s \geq \frac{1}{2}$ or $< \frac{1}{2}$. Then apply the local maps indicated in Figure 5, where the slopes are $\frac{1}{1-\varepsilon}$ and $\frac{1}{\varepsilon}$, for ε small enough.

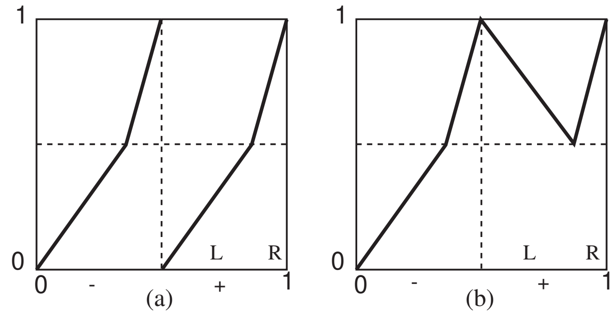

Fig. 5. Maps defining a CML with eternal transient phase for ε small enough (drawn with $\varepsilon = \frac{1}{4}$), to be applied when the right-hand neighbour is (a) $-$, (b) $+$.

The subset $[\frac{1}{2}, 1]^S$ is invariant, so there is a Gibbs phase (in fact pure and corresponding to the Bernoulli measure $B(1-\varepsilon, \varepsilon)$ on it), and it would seem that it ought to absorb the orbits of most initial conditions. This is indeed the case if ε is not too small, but remarkably for ε small enough there is another pure Gibbs phase, supported on the complement.

The construction is based on *Stavskaya's PCA* (a case of the discrete-time asymmetric contact process, also known as directed site percolation) on $\{+,-\}^S$ where

$$P(\sigma_s^{t+1} = +|\sigma^t) = \begin{cases} 1 \text{ if } \sigma_s^t = \sigma_{s+1}^t = +, \\ \varepsilon \text{ otherwise} \end{cases},$$

(the remaining probability $1-\varepsilon$ in the second case going to $\sigma_s^{t+1} = -$), for which corresponding behaviour has been proved if ε is small enough (≤ 0.09, though numerically it is good up to about 0.31 [T+]). To make the construction fit within our general framework of expanding maps with symbolic dynamics, we subdivide the $+$ interval at $1-\frac{\varepsilon}{2}$ into two subintervals $L \cup R$ and use a map that expands each of L and R to $[\frac{1}{2}, 1]$ for the $++$ case, so that we obtain symbolic dynamics with alphabet $\{-, L, R\}$. The same phenomenon would occur, however, if the map of $[\frac{1}{2}, 1]$ in the $+$ case were replaced by any other with uniform density as mixing measure.

It would be interesting to analyse examples where the map of $[\frac{1}{2}, 1]$ in the $+$ case has a globally attracting fixed point, to exhibit the phenomenon of "stable chaos" of [GLP].

Many variants of this example can be constructed, using more general discrete-time contact processes or PCAs associated to directed percolation, though I am not aware of proofs of non-unique phase for most of these (there is a proof for a broad class of continuous-time contact processes, e.g. [Lig]).

2.5 Pre-indecomposability

The eternal transient example is not indecomposable, since it possesses the forward invariant subset $[\frac{1}{2},1]^S$. In [GM] we got round this by allowing D_0 to depend on the diameters of the A_i, which allows us to wait for influence of distant $-$ to propagate in and exploit the discontinuity, but I think that obscured the point that really the CML should be considered only "pre-indecomposable" and it is not appropriate to ask for more.

I'll explain the concept first in the context of Markov chains. I say a Markov chain is *pre-indecomposable* if it has an aperiodic irreducible component which can be reached from anywhere. For a finite-state Markov chain this is a necessary and sufficient condition for unique phase.

Similarly, I'll say a dynamical system is *pre-indecomposable* if it has an attractor A with the specification property and for all $\delta > 0$ there exists $\tau > 0$ such that for every orbit up to time 0 and orbit on A from time τ there exists one within δ of each. Thus we can include the basin of the attractor and furthermore the closure of the basin. For $C^{1+\alpha}$ systems in finite dimensions I think this suffices for unique phase (Bowen's fat horseshoe example shows that the Hölder condition can not be dropped).

Thus I'll say a CML is *pre-indecomposable* if it has an indecomposable invariant subset A and for all $\delta > 0$ there exists τ such that for every orbit x up to time 0 and orbit y on A from time τ onwards there exists an orbit z within δ of each.

This holds for Example 5: take $A = [\frac{1}{2},1]^S$ and $\tau > 2\frac{\log 1/\delta}{\log \lambda}$ where $\lambda = \frac{1}{1-\varepsilon}$, then let σ equal that for x up to time $\tau/2$ and that for y thereafter, and let z be the orbit with symbol table σ, which exists because the only forbidden transition is $++ \to -$.

2.6 Construction of Invertible examples

So far, the examples have been non-invertible. This would not disturb researchers in dissipative PDE (like reaction-diffusion systems), but for other contexts like nonlinear wave equations or oscillator networks, it would be better to have invertible examples.

A simple way to make an invertible CML to simulate a PCA is to replace the expanding 1D maps by "baker's maps": take local state space M to be $Q = [0,1]^2$ with coordinates (x,y), and construct a map which on the x coordinate is as before and on the y coordinate is just (for 2-state PCA)

$$y' = \begin{cases} y/2 \text{ for } x < 1/2, \\ (1+y)/2 \text{ for } x \geq 1/2 \end{cases}.$$

The Gibbs phases for this CML are just those for the expanding map examples in x, supplemented by the distribution in y corresponding to the possible symbol histories of x, converted to binary expansions. For example, for G^+

in Example 1, for any (s,t) the marginal on x_s^t has a step function density with values $c(\varepsilon), 1-c(\varepsilon)$ for $x>\frac{1}{2}$, $x<\frac{1}{2}$, and the marginal on y_s^t is a singular continuous distribution because the symbol histories form a source with entropy (in base-2) $h<1$, thus for large N most of the probability is supported on about 2^{Nh} of the 2^N intervals $[\frac{p-1}{2^N},\frac{p}{2^N}]$ for $p=1,\ldots,2^N$, whose total length $2^{N(h-1)}\to 0$ as $N\to\infty$, yet there are no atoms.

Note that in general for an invertible system one obtains different Gibbs potentials for forward and backward evolution (see Sec. 3.8), and hence in general different sets of Gibbs phases. For the invertible CML just constructed, there is a unique backwards Gibbs phase and it is uniform in y.

An alternative construction, for which the local maps are continuous (though the CML still has discontinuities with respect to coupling), is to take piecewise affine continuous deformations of hyperbolic automorphisms of a 2-torus, with continuously deformed Markov partition, but there are restrictions on the possible transition probability matrices, so I won't pursue it.

2.7 Reversible CML with non-reversible phases

In some contexts (e.g. much of Hamiltonian mechanics), the dynamics is not only invertible but also reversible. A dynamical system is *reversible* if it is invertible and conjugate to its inverse via an involution (a map R on the state space whose square is the identity). So for any motion there is an equally possible reversed motion. The prime example is celestial mechanics: if one were to reverse the momenta of all the planets one would obtain the time-reversed motion.

The basic laws we use for the universe are reversible, yet we see a very definite arrow of time. This is expressed in various ways like the second law of thermodynamics, cause precedes effect, and outgoing radiation conditions. Why are only motions compatible with these conditions apparent and not their time-reverses?

The prevailing view (e.g. [Pe]) attributes the arrow of time to a very special choice of initial condition (big bang), but I consider that unsatisfactory as an explanation, because it is still locked in the "cause precedes effect" mindset. In contrast, I suggest it is a result of non-unique space-time Gibbs phase at some level, probably quantum gravity, selecting one light-cone in preference to the other in large space-time patches.

Here I make a toy model to illustrate the possibility. I use chaotic dynamics instead of quantum physics to provide the stochastic effect (as done in [Be] for other purposes).

The starting point is synchronous 2D kinetic Ising models. These are PCA (of period-2 rather than autonomous, but that does not cause essential trouble) which are reversible (see the Appendix for a discussion of what it means for a PCA to be reversible) and which have precisely two pure phases. I'll simulate this by a reversible CML. The phases are reversible but by making a suitable skew product over this CML, I'll then make a reversible CML with non-reversible phases.

For kinetic Ising model, take state space $\{+,-\}^S$ with $S=\mathbb{Z}^2$ and define the neighbourhood $N(s)$ of $s\in S$ to consist of $s\pm(1,0)$ and $s\pm(0,1)$. Choose $J\in\mathbb{R}_+$. Say a space-time site (s,t) is "even" or "odd" according as s_1+s_2+t is even or odd. For (s,t) even, let $m=\sum_{r\in N(s)}\sigma_r^t$ and "Glauber dynamics"

$$P(\sigma_s^{t+1}=\pm|\sigma^t)=\frac{1}{2}(1\pm\tanh(Jm))=\frac{\exp\pm Jm}{\exp Jm+\exp -Jm},$$

and for (s,t) odd, let $\sigma_s^{t+1}=\sigma_s^t$. This PCA has the nice features that it is reversible and its Gibbs phases are in one-to-one correspondence with those for the equilibrium statistical mechanics of the 2D Ising model at temperature 1, with potential $-\sum_s \frac{J}{2}\sigma_s\sum_{r\in N(s)}\sigma_r$ (the factor $\frac{1}{2}$ is to compensate for counting each bond twice). For J large enough, the 2D equilibrium Ising model has two pure phases, thus so does the kinetic Ising model. There are other transition rates that achieve the same properties (e.g. Metropolis rates). Also the kinetic Ising model can be reformulated as an autonomous PCA by going to a diagonally moving frame (Ex 1.4b of [T+]), so one could produce an autonomous CML to simulate it, but the moving frame complicates the reversibility properties.

The kinetic Ising model can be simulated by an expanding CML using the construction of Section 2.2 again. Indeed, a variant of the method was used by Sakaguchi [Sa] long before our work, but with a variable partition and constant slope on each partition element instead of fixed partition and varying the slopes of the map within each partition element, which to my mind makes the non-uniqueness of phase less marked because both phases have uniform marginal density on all finite subsets. Also, it is not clear whether the paper contains a complete proof, and the construction is less flexible than ours. Finally, Sakaguchi's CML is second-order, meaning that to determine x^{t+1} one needs both x^t and x^{t-1} (though through only the symbolic state of x^{t-1}), so to make a first-order CML out of it, one would have to augment the local state, which becomes messy.

We want to simulate the kinetic Ising PCA by a reversible CML, however. The following variant of our baker's map construction produces one.

Example 6 Take $S=\mathbb{Z}^2$, local state space $M=[-1,0]^2\cup[0,1]^2$ (really, I mean disjoint union, so there are two distinct points $(0,0)$) with coordinates (x,y), for each $s\in S$ let $N(s)=\{s\pm(0,1),s\pm(1,0)\}$, and let $\sigma_s=\pm$ according as $(x_s,y_s)\in[0,1]^2$ or $[-1,0]^2$. Choose $J\in\mathbb{R}_+$. For (s,t) even, evaluate $m(s)=\sum_{r\in N(s)}\sigma_r\in\{-4,-2,0,2,4\}$ and $p_\pm=\frac{1}{2}(1\pm\tanh Jm)$, and apply the following map to (x_s,y_s) (see Fig. 6):

$$\text{For } 0<x<p_- \quad \begin{cases} x'=-\frac{x}{p_-}\\ y'=-p_+y\end{cases}$$

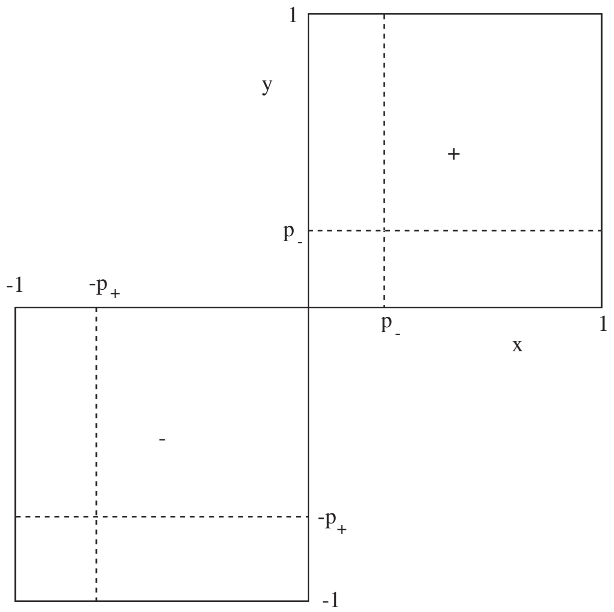

Fig. 6. The domains (vertical rectangles) and ranges (horizontal rectangles) of the affine pieces of the local maps for Example 6.

$$\begin{aligned}
&\text{For } p_- < x < 1 && \begin{cases} x' = 1 - \frac{1-x}{p_+} \\ y' = 1 - p_+(1-y) \end{cases} \\
&\text{For } -p_+ < x < 0 && \begin{cases} x' = -\frac{x}{p_+} \\ y' = -p_- y \end{cases} \\
&\text{For } -1 < x < -p_+ && \begin{cases} x' = -1 + \frac{1+x}{p_-} \\ y' = -1 + p_-(1+y) \end{cases} .
\end{aligned}$$

For (s,t) odd, apply the identity map.

This CML is reversible with respect to $R : x \leftrightarrow y$ and $t \mapsto 1-t$ (as the CML is non-autonomous, the notion of reversibility needs a slight modification, hence the necessity to specify which reflection to take in time). It is also indecomposable (similar proof to Example 1). For $J > J_c = \log(1+\sqrt{2})$, there are $c > \frac{1}{2}$ (in fact, $c = \frac{1}{2}(1+(1-1/\sinh^4 J)^{\frac{1}{8}})$) and precisely two pure phases $G^{\pm}$, and $G^{\pm}(\sigma_s^t = \pm) = c$. Both phases are reversible, however, because they are essentially just the Gibbs phases of the kinetic Ising PCA, which are reversible.

To make a reversible CML with non-reversible Gibbs phases, I first note that Example 6 is also reversible with respect to $R' : (x,y) \mapsto (-y,-x)$, $t \mapsto 1-t$, and $G^{\pm}$ are not reversible with respect to this involution. I am reluctant to claim this as a valid example of reversible CML with non-reversible phases, however, because one could say I just chose the wrong involution. So I seek to modify Example 6 to lose R-reversibility but keep R'-reversibility. This can be done by a skew-product over Example 6, where dynamics in additional variables z is added.

Example 7 In addition to the (x, y) dynamics of Example 6, choose an invertible map g of a manifold Z which is not reversible with respect to any involution (this rules out Z being finite), take local state space $([0,1]^2 \cup [-1,0]^2) \times Z$ with coordinates (x, y, z), at odd (s, t) take $z' = g(z)$ if $\sigma = +$ and $z' = g^{-1}(z)$ if $\sigma = -$, and at even (s, t) take $z' = z$. This is reversible with respect to $R'' : (x, y, z) \mapsto (-y, -x, z)$ but not with respect to any involution which interchanges x and y. If g has unique and distinct forward and backward phases, then the resulting CML is indecomposable (because given $\delta > 0$ one can insert a bounded band of symbols $+$ or $-$ to move z_s from one phase to the other), and has pure phases corresponding to $G^{\pm}$ on (x, y) and the forwards (respectively backwards) phase of g, because with G^{+}-probability 1, for each site s, $\sum_{t=0}^{T} \sigma_s^t \to +\infty$ as $T \to +\infty$, so z_s^t moves to the forward phase of g as t increases, and similarly for G^{-}. These phases are not reversible.

Concretely, one can take $Z = [-1, +1]$ and g to be any diffeomorphism of $[-1, +1]$ with fixed point set $\{-1, +1\}$ and the product of the slopes at the ends not equal to 1. The final condition rules out existence of any smooth reversor, but one might prefer to rule out all continuous reversors. Then one could take something like $Z = \mathbb{B}^3$, the unit ball in three dimensions, and g to be a Bernoulli diffeomorphism on its boundary S^2 (which can be made by smoothing a pseudo-Anosov map of S_4^2, the sphere with four punctures, following Katok) and to attract all the interior to a fixed point at the origin. Alternatively, one could take $Z = \mathbb{T}^2$ and g a non-reversible uniformly hyperbolic map: non-reversibility can be achieved by making the forward and backward SRB measures not mappable onto each other.

Although Example 7 solves the problem posed, it is not ideal. Firstly, it leaves the realm where Gibbs phases are defined, since the skew products are not hyperbolic in any good sense: the z dynamics is a sequence of g and g^{-1}, so any amount of expansion in Z can be undone by a suitable symbol sequence in the future. Secondly, my original aim was to make examples with phases carrying a thermodynamic arrow of time or selecting a radiation condition, but z has no interpretation as entropy and there is no coupling of the z dynamics between sites. Perhaps a thermodynamic arrow of time could be achieved by a space-time analogue of the interface phases of the 3D Ising model, and a radiation condition by skew-product with a suitable deterministic cellular automaton.

An interesting question is whether examples of reversible CML with non-reversible phases can be made that are also symplectic, because that would suggest that non-reversible phase could occur in spatially extended Hamiltonian systems.

3 Motivation for Gibbs phases

3.1 Probabilistic behaviour

For a deterministic dynamical system, each initial condition determines a unique trajectory, but it often happens that many initial conditions lead to similar trajectories as time goes to infinity. This holds in a strong sense for initial conditions in the basin of an attracting fixed point. It holds in a slightly weaker sense for initial conditions in the basins of different points of the same attracting periodic orbit, because their orbits become similar only if one allows to compare suitable time-shifts. For chaotic systems (meaning at this stage that they have sensitive dependence on initial conditions), trajectories of typical pairs of initial conditions are unlikely to converge together, but the trajectories of many initial conditions may still be similar to each other in a statistical sense. The tidiest way to say this is that the probability distribution of the state at time t with respect to any member of a specified class of measures on initial conditions may converge in weak-* as $t \to \infty$ to a common probability measure (a sequence μ_n of measures on a topological space M *converges in weak-** if for all continuous functions $\phi : M \to \mathbb{R}$ the sequence of expectations $\mu_n(\phi)$ converge; the terminology comes from viewing measures as a subset of the dual of the space of continuous functions on M).

For example, for $f : x \mapsto 4x(1-x)$ on $[0,1]$ and any absolutely continuous measure μ on $[0,1]$, i.e. one which gives probability zero to all sets of Lebesgue measure zero, then μ is represented by a density $\varrho \in L^1[0,1]$ with $\int \varrho(x)dx = 1$. Define the "transfer" operator T on densities ϱ by $\int g(x)(T\varrho)(x)dx = \int g(f(x))\varrho(x)dx$ for all $g \in L^\infty[0,1]$; then $T^n\varrho \to \bar{\varrho}$ as $n \to \infty$ where

$$\bar{\varrho}(x) = \frac{1}{\pi\sqrt{x(1-x)}},$$

meaning $\int g(x)(T^n\varrho)(x)dx \to \int g(x)\bar{\varrho}(x)dx$ for all $g \in L^\infty[0,1]$. Note that it follows that for any $\tau > 0$, the probability distribution for $(x^t, x^{t+1}, \ldots, x^{t+\tau})$ converges as $t \to \infty$ (in weak-* topology) to that generated by f from $\bar{\varrho}$.

I think about this is in terms of the orbits of uncertain initial conditions. If one does not have infinite precision over the initial condition then it is more reasonable to consider the initial condition as a probability measure, perhaps tightly peaked around some point but nonetheless absolutely continuous. This point of view allows me to say that the orbits of all uncertain initial conditions have the same statistics in the above example, whereas all I can claim for exact initial conditions is that almost all lead to the same statistics: it is essential to remove a set of Lebesgue measure zero of exceptional orbits, like those which converge to an unstable fixed point or periodic orbit.

One can make examples, however, with a two-band attractor or more than one attractor, for which such convergence does not happen unless one starts in the same basin. Such examples have more than one probabilistic behaviour.

I call probability measures on realisations in time originating from nice enough initial measures far in the past *phases* (not to be confused with phase in the sense of position in a cycle or in a bundle over a circle). They are sometimes called "states" but I reserve this word for the points of the state space of the dynamics (confusingly called "phase space" in the older literature!). Their time-averages are often called natural or physical or SRB measures.

Any convex combination of phases is also a phase: for example start with initial probability distribution supported partly in the basin of one attractor and partly in that of another, but the interest focusses on the *pure phases*, which are the phases which can not be expressed as a convex combination of others (sometimes called *extremal phases*).

3.2 Ordinary Symbolic dynamics

Some dynamical systems have the nice feature that there is a partition of the state space into regions and a directed graph between the regions such that every path in the graph (called a *symbol sequence*) occurs for some orbit and there is precisely one orbit for each path in the graph. Thus to each symbol sequence σ there corresponds an initial condition $x^0(\sigma)$. The regions are allowed to overlap along "negligible" sets, so orbits which pass through such overlap sets may have more than one symbol sequence, but this is a relatively unimportant issue. The system is said to have *symbolic dynamics*.

Symbolic dynamics occurs for every locally maximal compact uniformly hyperbolic set for a C^1 diffeomorphism, and also for suitably constructed maps with jump discontinuities, like expanding maps for which each interval of continuity is mapped onto a union of intervals of continuity.

3.3 Gibbs phases for Markov chains

Before considering the asymptotic probabilistic behaviour of deterministic dynamical systems with symbolic dynamics, let's study the same question in the simpler context of a Markov chain.

It is usual to specify a Markov chain by giving the forward transition probabilities $p_{\sigma\sigma'}$ from states σ to σ', which define probabilities

$$P\{\sigma^1, \ldots, \sigma^\tau | \sigma^0\} = \prod_{t=0}^{\tau-1} p_{\sigma^t \sigma^{t+1}}$$

for the initial value problem. Actually, to avoid confusion with the official definition of Markov chain, which specifies probabilities for all sequences not just conditional on initial conditions, I should give this a different name, like "pre-Markov chain". The distinction is important because we are trying to decide when specification of the transition probabilities implies a unique probability distribution on doubly infinite sequences. Nevertheless, I'll follow many physicists and call it a "Markov chain" except in the Appendix.

Even though they do not determine the unconditional probabilities of sequences, the transition probabilities do define the probability distributions for the two-point boundary value problems with σ^m, σ^n given (assuming there exists a path with positive probability):

$$P\{\sigma^{m+1}, \ldots, \sigma^{n-1} | \sigma^m, \sigma^n\} = \frac{1}{Z} \prod_{t=m}^{n-1} p_{\sigma^t \sigma^{t+1}},$$

where $Z(\sigma^m, \sigma^n)$ is a normalisation constant. This can be rewritten as

$$P\{\sigma^{m+1}, \ldots, \sigma^{n-1} | \sigma^m, \sigma^n\} = \frac{1}{Z} \exp\left(- \sum_{t=m}^{n-1} \log \frac{1}{p_{\sigma^t \sigma^{t+1}}} \right).$$

Thus $\log \frac{1}{p_{\sigma\sigma'}}$ can be interpreted as an "energy" contribution from the transition $\sigma \to \sigma'$ of state along a 1D statistical mechanics chain at temperature 1, and the above equation for the probability distributions for two-point boundary value problems is precisely that defining (Boltzmann-)Gibbs phases of such a statistical mechanics model.

A *Gibbs phase* for a Markov chain is a probability distribution on doubly infinite sequences of states such that the conditional probability distribution for all two-point boundary value problems is given by this formula. Of course, for an aperiodic irreducible Markov chain there is a unique Gibbs phase, generated by its stationary measure and the transition probabilities.

3.4 Gibbs phases for Expanding maps

For expanding maps with symbolic dynamics and Hölder continuous derivative on each branch there is an analogous way to determine the possible long-time statistics of orbits, but in general we have to consider boundary value problems where all the symbols before some time and after some other time are given, rather than just two-point.

A map f is *expanding* if there is a $\tau > 0$ and $\lambda > 1$ such that f^τ expands the distance between any two nearby points by at least λ. Note that such maps are necessarily noninvertible, but I'll consider invertible maps later.

If we start from an absolutely continuous measure with density ϱ at some large negative time M, and we observe the symbol sequence σ from time M to some large positive N, then we could ask what is the probability of the subsequence $\sigma^m, \ldots, \sigma^n$ for some fixed times m and n between M and N, given the remaining symbols. I will find it useful to assume that $\log \varrho$ is Hölder continuous on each subset with a given long enough symbol sequence, though this should not really be necessary.

The answer is proportional to $\int \varrho(x^M)\, dx^M$ over the set of x^M with symbol sequence $\sigma^M, \ldots, \sigma^N$, this being the measure of the set of initial conditions which follows the given symbol sequence. Using the map f, which is invertible in the domain where the given symbol sequence occurs, we can change variable of integration from x^M to $x^{M+1} = f(x^M)$, obtaining

$$\int \varrho(x^M)\ dx^M = \int \frac{\varrho(x^M)}{|\det Df_{x^M}|}\ dx^{M+1}$$

over the set with symbol sequence $\sigma^{M+1}, \ldots, \sigma^N$, where x^M is regarded as a function of x^{M+1} by using the appropriate inverse branch of f. Iterating this procedure up to any time $K \leq N$, we find that the probabilities of symbol sequences $\sigma^m, \ldots, \sigma^n$ given the rest are proportional to

$$\int \frac{\varrho(x^M)}{\prod_{t\in\{M,\ldots,K-1\}} |\det Df_{x^t}|}\ dx^K$$

over the set with symbol sequence $\sigma^K, \ldots, \sigma^N$.

For $K > n$, the range of integration is independent of $\sigma^m, \ldots, \sigma^n$. For M large (negative), the set of x^M with given symbol sequence $\sigma^M, \ldots, \sigma^{m-1}$ has exponentially small diameter in $m - M$, so by the assumed continuity of $\log \varrho$, ϱ can be approximated by constant density on it. If we choose K to be roughly halfway between n and N, then for N large, $\prod_{t\in\{M,\ldots,K-1\}} |\det Df_{x^t}|$ is close to constant for each symbol sequence (because Df is Hölder continuous and the sets over which the x^t can range are exponentially small in $N - t$). Thus as $M \to -\infty$, $N \to +\infty$, the conditional probabilities of $\sigma^m, \ldots, \sigma^n$ are proportional to $1/\prod_{t\in\{M,\ldots,K-1\}} |\det Df_{x^t}|$, which acquires a well defined meaning because the effect of the terms far from times $t \in \{m, \ldots, n\}$ depends exponentially weakly in $K - t, t - M$ on $\sigma^m, \ldots, \sigma^n$.

It is common to rewrite this relative conditional probability as

$$\exp\left(-\sum_{t\in\{M,\ldots,K-1\}} \log|\det Df_{x^t}|\right),$$

to make it look like the condition for a Gibbs measure in statistical mechanics. Thus we see that the possible probabilistic behaviours for expanding maps with symbolic dynamics started far in the past from an absolutely continuous measure correspond to the Gibbs phases for a 1D spin chain with energy contribution from site t given by $\log|\det Df_{x^t}|$, where x^t is the state at time t corresponding to a given symbol sequence. Since x^t depends exponentially weakly on symbols distant from t and Df is Hölder continuous, this gives an exponentially decaying interaction between sites in the spin chain.

3.5 Gibbs phases for PCA

For a PCA with space S, if we wish to know the probability measures on all possible realisations on $S \times T$ (with $T = \mathbb{Z}$) which are consistent with the transition probabilities, including arbitrarily far in the past, then we can again write them as Gibbs phases.

As for general Markov chains, we consider boundary (rather than initial) value problems, but now in space-time rather than just time. For each bounded subset Λ of space-time $S \times T$ and state η on the complement of Λ, we can ask for the conditional probability distribution for the state σ on Λ. Define the *templates* for the PCA to be the subsets of $S \times T$ of the form $\{(s, t+1)\} \cup \{(r,t) : r \in N(s)\}$, where $N(s)$ is the neighbourhood of dependence for s, and for state σ and template τ define the *weight* $w_\tau(\sigma)$ to be the transition probability $p_\tau(\sigma)$ for the transition represented by the restriction of σ to τ. Then

$$P(\sigma \text{ on } \Lambda | \eta \text{ off } \Lambda) = \frac{1}{Z} \prod_{\tau \cap \Lambda} p_\tau(\sigma \vee \eta),$$

where $\tau \cap \Lambda$ means the set of templates having nonempty intersection with Λ, $\sigma \vee \eta$ is the state on $S \times T$ given by σ on Λ and η off Λ, and Z is chosen to normalise the distribution (provided that there is at least one σ compatible with η). Thus we view PCA on S as a class of Markov random field on $S \times T$ (with the special features that the templates have the shape above and the weights on a template with given state on $N(s)$ sum to 1 over the state at the "apex"). This conditional probability can be rewritten as

$$P(\sigma \text{ on } \Lambda | \mu \text{ off } \Lambda) = \frac{1}{Z} \exp\left(- \sum_{\tau \cap \Lambda} \log \frac{1}{p_\tau(\sigma \vee \eta)} \right),$$

which looks like the condition for a Gibbs distribution for a statistical mechanics spin lattice on $S \times T$ where each template τ contributes an energy $\log \frac{1}{p_\tau(\sigma)}$ (in temperature units).

3.6 CML with Symbolic dynamics

The concept of symbolic dynamics extends to some CML, in particular those we constructed to simulate PCA, and also uniformly hyperbolic CML as I'll indicate below. For a CML, a trajectory x is an array of x_s^t with $s \in S, t \in T$ ($T = \mathbb{Z}_+$ or $\mathbb{Z}$, for the expanding or invertible case, respectively), and the symbol sequence is replaced by a symbol table $\sigma = (\sigma_s^t)_{s \in S, t \in T}$.

Our examples have constant derivative on each piece corresponding to a given symbolic transition, but for general expanding CML with symbolic dynamics, the derivative depends on the exact state making a given symbolic transition, which is determined by more distant symbols in space and time. So to justify the concept of Gibbs states for general expanding CML with symbolic dynamics, we need to examine the dependence of the derivative on the full symbol table. Following [PS, Ma], here is a sketch (which applies equally well in the invertible case).

For CML F on M^S, the orbits $x^{t+1} = F(x^t)$, $t \in T$, correspond to the fixed points of the map $\mathcal{F}$ on $M^{S \times T}$ defined by

$$(\mathcal{F}x)_s^t = F_s(x^{t-1}).$$

If F is differentiable (in uniform norm) with bounded derivative, then so is $\mathcal{F}$. I say an orbit x is *uniformly hyperbolic* if it is a non-degenerate fixed point of $\mathcal{F}$, i.e. $I - D\mathcal{F}_x$ is invertible; since $D\mathcal{F}$ is bounded it follows that $(I - D\mathcal{F}_x)^{-1}$ is bounded. I say a set U of uniformly hyperbolic orbits is *uniformly hyperbolic* if there exists $C \in \mathbb{R}$ such that $||(I - D\mathcal{F}_x)^{-1}|| \leq C$ for all $x \in U$.

If the coupling in F is local (or exponentially decaying in a suitable sense) then it can be shown that this definition of uniform hyperbolicity is equivalent to the usual one in terms of a splitting of the tangent space into a contracting subspace and a backwards contracting subspace with uniform exponential contraction estimates and angle of splitting bounded away from zero. In addition, however, we obtain exponential decorrelation in space. The point is that uniform hyperbolicity plus exponentially decaying coupling implies that the matrix elements of $(I - D\mathcal{F}_x)^{-1}$ decay exponentially from the diagonal in both space and time. In particular, this implies that the dependence of x_s^t on symbol $\sigma_{s'}^{t'}$ decays exponentially with $d((s,t),(s',t'))$, which will be required in the next section.

This viewpoint on uniform hyperbolicity for CML has many consequences for the dynamics at the topological level as well as probabilistic. In particular, it allows one to construct symbolic dynamics for weakly coupled uniformly hyperbolic maps [PS] and to obtain many of the results of Afraimovich's lectures. The idea is that given a uniformly hyperbolic set U of orbits for CML F then all C^1-small perturbations of F possess a unique continuation of U, with topologically equivalent dynamics. Specifically, if F depends on parameters λ, the continuation $x(\lambda)$ of an orbit $x(0) \in U$ is given by

$$\frac{\partial x}{\partial \lambda} = (I - D\mathcal{F}_x)^{-1} \frac{\partial \mathcal{F}}{\partial \lambda}(x).$$

This can be used to make further examples of interesting topological dynamics in CML. For example, one can make CML with localised chaotic attractors, by taking the on-site dynamics to have an attracting fixed point and a uniformly hyperbolic chaotic attractor (e.g. Plykin or solenoid), choosing one site to be on the chaotic attractor and the rest on the fixed point, and adding any smooth weak coupling [Ma].

Furthermore, the shadowing theorem extends to this space-time context, so the approach allows to construct symbolic dynamics for any locally maximal uniformly hyperbolic set for a CML (choose a sufficiently fine set of time-periodic orbits and then use shadowing to associate to any orbit at each $(s,t) \in S \times T$ one of the periodic orbits, thus constructing its symbol table).

3.7 Gibbs phases for expanding CML with symbolic dynamics

For a CML F on M^S, the phases should describe the statistics of orbits from all nice enough initial probability distributions on the state space, started far enough in the past. It is not yet clear what is the sensible class of initial probability distributions to consider: see Jarvenpää's lectures. For safety, I will restrict attention to those with a C and a $\delta > 0$ such that their marginals on any finite subset $L \subset S$ have an L^1 density ϱ_L and

$$|\log \varrho_L(x) - \log \varrho_L(y)| \leq C \sum_{s \in L} d(x_s, y_s)^\delta,$$

as in [BK] and [GM], but in fact I will require it only when x and y have the same symbol table on some neighbourhood of $L \times \{0\}$ in $S \times T$. I believe this Hölder condition on the logarithm of the densities is unnecessarily strong, but I have not yet formulated a better one.

Given an initial measure satisfying the Hölder condition at a large negative time M, we ask for the probabilities of symbol tables on a given $\Lambda \subset S \times T$, conditional on the symbol table outside Λ. Fatten Λ to some larger patch Λ', in both space and time, to allow for the fact that in general the symbols on Λ depend on the state off Λ as well as on Λ, but exponentially weakly as the distance away from Λ increases. Because the CML has finite range coupling, disturbances can propagate at most at some speed. Let S^M be the subset of S at time M capable of influencing the state on Λ'. Fatten S^M in space to a set J. Then the conditional probabilities P of symbols on Λ are proportional to $\int \varrho_J(x^M) \prod_{s \in J} dx_s^M$ over the subset of states x^M on J giving rise to the given symbol table (modulo a slight dependence on the state outside J in general, which one should bound).

Change variable from x_s^M to x_s^{M+1} at all sites $s \in S^M$ such that $N(s) \subset S^M$, using the map F. Call this set of sites S^{M+1}. The volume change factor is the determinant of the block of DF corresponding to S^{M+1}, which I'll denote $DF_{S^{M+1}}$, evaluated at the state x^M at time M. Then

$$P = \frac{1}{Z} \int \frac{\varrho_{S^M}(x^M)}{|\det DF_{S^{M+1}}(x^M)|} \prod dx_s^t,$$

where (s,t) ranges over S^{M+1} at time $M+1$ and the remaining sites of J at time M.

It is convenient to rewrite

$$|\det DF_{S^{M+1}}(x^M)|^{-1} = \exp\left(-tr \log DF_{S^{M+1}}(x^M)\right).$$

For any invertible operator A, $tr \log A$ can be calculated in an arbitrary bounded connected neighbourhood of a given A_0, up to a constant, by writing $A = A_0 \prod_i (I - B_i)$ for a suitable finite product of operators with each $||B_i|| < 1$, using $\log(I - B) = -\sum_{n \geq 1} \frac{B^n}{n}$ so $tr \log(I - B) = -\sum_{n \geq 1} \frac{tr B^n}{n}$, and letting $tr \log A = \sum_i tr \log(I - B_i)$ (it is not necessary to add a value for $tr \log A_0$ because the choice would change only the normalisation constant Z). Alternatively, it can be calculated as $\int_0^1 tr\left(A^{-1}(\lambda) A'(\lambda)\right) \, d\lambda$ along any differentiable path of invertible operators to A from a convenient reference point

A_0. Note that neither approach defines $\log A$ itself if the B_i do not commute or $A'(\lambda)$ does not commute with $A(\lambda)$ along the path (one has to use the Campbell-Baker-Hausdorff formula; we neglected to mention this in [GM]), but both always give the correct trace.

Next, we can expand

$$\exp\left(-tr \log DF_{S^{M+1}}(x^M)\right) = \exp\left(-\sum_{s \in S^{M+1}} tr(\log DF)_{ss}(x^M)\right),$$

where $(\log DF)_{ss}$ denotes the diagonal block of $\log DF$ corresponding to site s (of dimension equal to that of the local state space, so 1 for all our expanding examples). Since all that matters is the sum over s, one can again replace the true $\log DF$ by either of the surrogates of the previous paragraph.

Now iterate the change of variables to a nested sequence of subsets S^t at times $t = M, \ldots, \tau$, where τ is the largest time in Λ'. Then

$$P = \frac{1}{Z} \int \varrho_{S^M}(x^M) \exp\left(-\sum_{(s,t) \in K} tr(\log DF(x^t))_{ss}\right) \prod_{(s,t) \in L} dx_s^t,$$

where K is the subset of $S \times T$ with $t \geq M$ that can influence the state on S^τ at time τ and L is the union of its "boundary" and the sites of J at time M not in S^{M+1}. See Fig. 7 for an illustration for the case of a 1D nearest neighbour CML.

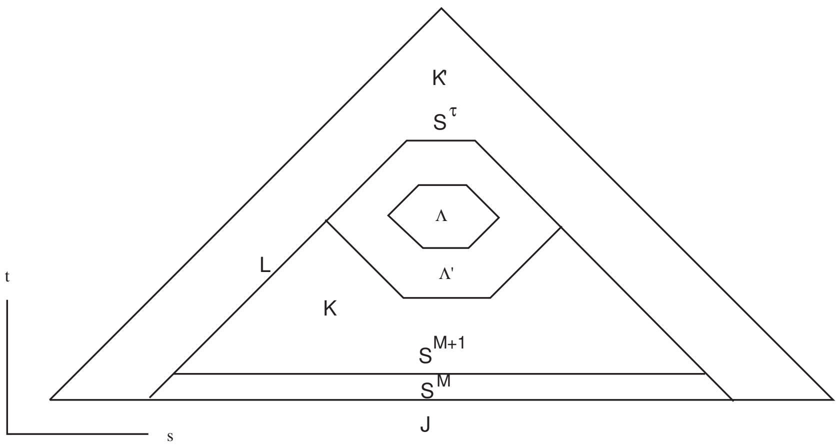

Fig. 7. Illustration of the change of variables for a 1D nearest neighbour CML.

Since the dynamics is expanding and Λ is contained in the subset of $S \times T$ which can influence the state on L, the domain of integration is the same for all σ (given η). Now x on J determines x on a fattening K' of K. The exponent is almost independent of x in the domain of integration for each fixed σ and η, because the symbolic state determines the actual state with exponential accuracy and DF is assumed to be Hölder continuous. Finally, ϱ_{SM} can be approximated by a constant, independent of σ, because of the Hölder condition on initial measures.

Thus, taking appropriate limits,

$$P(\sigma \text{ on } \Lambda | \eta \text{ off } \Lambda) = \frac{1}{Z'(\Lambda, \eta)} \exp\left(- \sum_{(s,t)} tr \log DF(x^t)_{ss} \right) .$$

I came up with this formula in February 1995 (unpublished notes); the same idea was presented independently in [BK] and used again in [JP].

3.8 Gibbs phases for Invertible dynamical systems

For invertible systems (whether ordinary or CML), symbolic dynamics uses backward-time symbols as well as forward, and the appropriate formulae for Gibbs phases use $\det Df^u$ and $\log Df^u$, where Df^u is the restriction of the derivative of f to the unstable subspace, and some choice of way of comparing volumes on the unstable subspaces of different points has to be made.

For an invertible system, however, it is also possible to ask for the probabilistic behaviour at large negative times. This is given by the backwards Gibbs phases, which use $(Df^s)^{-1}$ instead of Df^u (where Df^s is the restriction of Df to the stable subspace). For systems preserving a volume form, these give the same result, but in general they give different results.

An example where this can be seen by simple simulation (as well as proved) is the torus map

$$\begin{aligned} x' &= 2x + y + \frac{k}{2\pi} \sin 2\pi x \\ y' &= x + y \end{aligned}$$

for $|k| \in (0, 1)$. It is uniformly hyperbolic and topologically mixing and has unique forwards and backwards Gibbs phases, but they are different (and singular with respect to each other and Lebesgue measure).

4 Some challenges

How robust are the examples of indecomposable CML with non-unique phase given here? By Toom's work, Examples 1–4 (and probably 5) are robust to small changes in the slopes provided they remain constant, but presumably one can also make the slopes slightly non-constant and still keep the phenomena, because the resulting Gibbsian random fields on space-time should not be very different from the Markovian ones of the PCA.

Can you make continuous (or smooth) examples of indecomposable CML with non-unique phase? If our examples are very robust then one might be able to make the local map depend continuously on the state of neighbours, rather than just on their symbolic state. Alternatively, to simulate a PCA we chose to use piecewise constant expansion factors, but really we needed this only modulo a coboundary. Does the freedom to choose a coboundary allow to make continuous examples? As a third alternative, might it be possible to make a CML version of Toom's proof for the majority voter PCA? the idea would be to make a continuous CML which is monotone in the usual dynamical systems sense within each local symbol.

Can you explain the "stable chaos" of [GLP]?

What do Gibbs phases for CML tell us about "typical" orbits? For ordinary dynamical systems, they give the statistics for the orbit of Lebesgue almost every point in the basin of attraction. The problem for spatially extended systems is to make sense of "typical".

What is the most general notion of "nice enough" initial probability distribution which would lead to selection of Gibbs phases?

Can you find examples of non-unique phase without using symbolic dynamics? Unique phase can be proved for some classes of system without symbolic dynamics (e.g. [Ba] and Keller's lectures), but the methods do not yet suggest how to make examples with non-unique phase. Alternatively, the concept of "equilibrium state" does not require symbolic dynamics, though would exclude nontrivial collective behaviour.

Can you make symplectic examples of reversible CML with non-reversible phase? This would connect more closely to Hamiltonian dynamics and hence to fundamental models of the universe. The symplectic form need not be the standard one.

Can you make a reversible CML exhibiting phases with a second law of thermodynamics, or radiation conditions, rather than just being non-reversible? For example, one might hope to make a form of nontrivial collective behaviour which is not recurrent, or a skew product over the kinetic Ising model which generates outgoing waves in forward time when the underlying state is $+$ and backward time when it is $-$. Perhaps a 2D version of Stavskaya's PCA has a phase with an interface zone across which the density of $-$ increases from 0 to its value for the eternal transient phase, and a reversible version could be made?

Appendix: Reversibility of Markov chains and PCA

It is common to regard Markov chains as the epitome of irreversible behaviour: a preferred direction of time seems to be built in because we usually characterise them by their forward transition probabilities, and memory of the initial state is lost as time goes forwards (completely if the chain is irreducible and aperiodic).

Nevertheless, the Markov property does not distinguish between backward and forward time. A Markov chain (e.g. [Ke]) is a discrete-time stochastic process on a countable state space such that the sequences of states in the past and future of any moment are independent given the state at that moment. From this it follows that the probability of any given finite sequence of states can be factorised as a probability of the initial state multiplied by a sequence of forward transition probabilities, but it could be written equally well as a probability for the final state multiplied by a sequence of backward transition probabilities (or even starting at any intermediate time and working out both ways in time). Note that this is true regardless of whether the Markov chain is stationary (a stochastic process is *stationary* if the probability distribution for realisations is invariant under time-shifts).

A stochastic process $(X_t)_{t\in\mathbb{Z}}$ is said to be *reversible* if the probability distribution for $(X_t,\ldots,X_t')$ is the same as for $(X_{\tau-t},\ldots,X_{\tau-t'})$ for all t,t',τ. Note that "reversible" is not the opposite of "irreversible". Most reversible Markov chains exhibit irreversible behaviour, in the sense of loss of dependence of the future state on the state at time 0, but they have identical loss of dependence in reverse time.

I consider it appropriate, even essential, to extend the concept of reversibility of stochastic processes to something which is a closer analogue of the notion in deterministic dynamics. If R is an *involution* of the state space (a map whose square is the identity I), I say a stochastic process $(X_t)_{t\in\mathbb{Z}}$ is *R-reversible* if $(R(X_{\tau-t}),\ldots,R(X_{\tau-t'}))$ has the same distribution as $(X_t,\ldots,X_t')$ for all t,t',τ. Just as is well known for the standard case of $R=I$, an R-reversible stochastic process is automatically stationary (just choose two different values for τ). Also if π is the probability distribution at time 0 (and hence at any time, by stationarity) for an R-reversible Markov chain with transition probabilities p_{ij} then by taking $t=t'=\tau=0$ we obtain $\pi_{R(i)}=\pi_i$ for all states i and by taking $t=0, t'=\tau=1$ we obtain

$$\pi_i p_{ij} = \pi_{R(j)} p_{R(j)R(i)}$$

for all states i,j (in the standard case $R=I$, the first relation is vacuous and the second is known as "detailed balance"). Furthermore, if a stationary Markov chain satisfies these two conditions then it is easy to prove it is R-reversible (generalising another standard result).

Often, however, under the name "Markov chain" we are presented just with forward transition probabilities p_{ij}, specifying conditional probabilities like $P(i_1,\ldots,i_t|i_0)$ but not $P(i_0)$ nor $P(i_0|i_1)$, for example. This is not a full prescription of a Markov chain, which officially specifies the probabilities of all events in the space of sequences, so to avoid confusion I'll call it a *pre-Markov chain*. It is precisely our situation: we wish to know the possible stochastic processes compatible with the transition probabilities, including arbitrarily far in the past, and decide if there is a unique one or not.

A pre-Markov chain does not define a backwards pre-Markov chain, so is not invertible, yet we'd like to be able to say if a pre-Markov chain is reversible or not. I propose a definition, based on considering the transition probabilities as defining probabilities for two-point boundary value problems rather than initial value problems (as in Section 3.3).

To simplify exposition, let's first restrict attention to *autonomous* pre-Markov chains (meaning the transition probabilities do not depend on time). A pre-Markov chain (p_{ij}) defines the probabilities of sequences $i_0,\dots,i_\tau$ given i_0, i_τ (provided a path of positive probability exists) to be

$$P(i_0,\dots,i_\tau|i_0,i_\tau) = \frac{1}{Z}\prod_{t=0}^{\tau-1} p_{i_t i_{t+1}},$$

with Z a normalisation factor. Given involution R, say the pre-Markov chain is *R-reversible* if for each two-point boundary value problem the probability distribution is the same as for that given by reflecting all sequences in time and in state using R, equivalently,

$$\frac{1}{Z}\prod_{t=0}^{\tau-1} p_{i_t i_{t+1}} = \frac{1}{Z'}\prod_{t=0}^{\tau-1} p_{R(i_{t+1})R(i_t)}.$$

This holds if there exists a probability distribution π such that

$$p_{R(j)R(i)} = \pi_i p_{ij}/\pi_j$$

for all i, j, and for a large class of pre-Markov chains this is also a necessary condition (see next four paragraphs). Such a π is an equilibrium measure for the pre-Markov chain ($\sum_i \pi_i p_{ij} = \pi_j$ for all j), so (p,π) defines a Markov chain. If $R = I$ it is reversible, but for $R \neq I$ if $\pi_{R(i)} \neq \pi_i$ for some i then it is not R-reversible. Thus the notion of R-reversibility allows the possibility of making R-reversible pre-Markov chains with non-R-reversible equilibrium measure.

To see that the above sufficient condition is often also necessary, given p for which the graph of allowed transitions is connected and each state can be both preceded and followed by a cycle (and a hypothesis on the highest common factors of lengths of cycles, to be formulated later), the general solution for non-negative q of

$$\frac{1}{Z}\prod_{t=0}^{\tau-1} p_{i_t i_{t+1}} = \frac{1}{Z'}\prod_{t=0}^{\tau-1} q_{i_t t_{t+1}}$$

for all i_0, i_τ for which there exists an allowed path is

$$q_{ij} = DV_j p_{ij}/V_i$$

for some $D > 0$ and positive function V.

To prove this, let $z_{ij} = q_{ij}/p_{ij}$ for allowed transitions, so the condition is that the product of $z_{i_t i_{t+1}}$ along all paths from i_0 to i_τ in the same time τ is the same. Choose a cycle $i_0,\dots,i_N = i_0$ and let D be the N^{th} root of the product of $z_{i_t i_{t+1}}$ around it. Any other cycle gives the same value because if γ_1, γ_2 are two cycles of lengths N_1, N_2, and one can get from γ_1 to γ_2 by a path γ then the two paths $\gamma_1^{N_2}\gamma$ and $\gamma\gamma_2^{N_1}$ connect the same two states in the same time so have the same product of z. Similarly, if γ_1 and γ_2 can be reached from a common point then there exists a preceding cycle γ_3, so γ_1, γ_2 have the same value of D as γ_3 (and similarly if they can lead to a common point). Iterating this procedure for cycles which are linked only by a chain of segments in alternating directions, we show that all cycles have the same value of D.

Now choose a reference state o, let $V_o = 1$ and for each j reachable from o let V_j be the product of the z along a path from o to j divided by D^τ where τ is the length of the path. There might be more than one path from o to j. If they have the length then the product of z is the same, so no ambiguity results. If path γ_1 from o to j is longer than γ_2 by $K > 0$ then j is followed by some path γ_3 ending in a cycle γ. If the length of γ divides K, say N times, then $\gamma_1\gamma_3$ and $\gamma_2\gamma_3\gamma^N$ joins the same points in the same time, so the value of V_j is the same using γ_1 or γ_2. Similarly, if o is preceded by a cycle of length dividing K, or if there are cycles γ_4 preceding o and γ_5 following j whose lengths have highest common factor dividing K, one can construct two paths joining two points in the same time and hence see that V_j is well defined. Similarly if γ_2 contains a state on a cycle with length dividing K, or if both contain states on cycles with lengths whose highest common factor divides K, then we find the same value for V_j. For j that can lead to o, one can define V_j to be D^τ divided by the product (which is non-zero) of z for backward transitions from o, and again argue that this produces no contradiction. Finally, one can alternate forwards and backwards paths to define V_j for all j, because we assumed the graph is connected.

Thus under these assumptions, for an R-reversible pre-Markov chain,

$$p_{R(j)R(i)} = DV_j p_{ij}/V_i$$

for some D and V. Summing over i yields $1 = DV_j \sum_i \frac{p_{ij}}{V_i}$, so $D\sum_i \frac{p_{ij}}{V_i} = \frac{1}{V_j}$. Now assume $Z = \sum_i \frac{1}{V_i} < \infty$ and sum over j, to obtain $DZ = Z$, so $D = 1$. Then $\pi_i = \frac{1}{ZV_i}$ is a probability distribution such that for all i, j,

$$p_{R(j)R(i)} = \pi_i p_{ij}/\pi_j.$$

The considerations of this Appendix extend from Markov chains to PCA. When we talk about PCA, we usually mean a pre-PCA, specifying conditional probabilities for initial value problems, rather than a full stochastic process. This suffices to determine the probabilities for boundary value problems, however, as in Section 3.5. Hence given an involution R of the local state space we can define a pre-PCA to be *R-reversible* if the conditional probabilities for all time and state reflected boundary value problems are the same as for the originals. A sufficient condition for this is that the associated Markov random field model be equivalent to one with time-symmetric templates such that the weights for state-time-reflected configurations differ by a (multiplicative) coboundary, i.e. a function of configuration on the template such the product over any bounded subset Λ of space-time given the state outside is independent of the state on Λ (such as V_j/V_i for the case of a pre-Markov chain p_{ij}). One could also allow involutions of the product state space not of product form, but I won't pursue that here.

Now for the kinetic Ising model, which is not autonomous but period-2, we need to modify the definition of reversibility again. I'll say a period-2 Markov chain (meaning the distribution is invariant under shifts by even times) is R-reversible if the distribution of sequences reflected about half-odd integer times and with states reflected by R is unchanged. Similarly, a period-2 pre-Markov chain is R-reversible if the probabilities for all two-point boundary value problems are invariant under reflection in a half-odd integer time and reflection of the states by R. The same goes for period-2 PCA.

For example, the kinetic Ising model is reversible with $R = I$, because an equivalent Markov random field (i.e. after multiplication by a coboundary) is that with templates τ of the form $\{(s,t),(s \pm (1,0),t),(s \pm (0,1),t),(s,t+1),(s \pm (1,0),t+1),(s \pm (0,1),t+1)\}$ for (s,t) even, and weights

$$w_\tau = \frac{\exp \frac{1}{2}(\sigma_s^t + \sigma_s^{t+1})Jm}{\exp Jm + \exp -Jm}$$

if $\sigma_r^{t+1} = \sigma_r^t$ for all $r \in N(s)$ (where $m = \sum_{r \in N(s)} \sigma_r^t$), 0 otherwise, which are invariant under time-reflection. The coboundary consists of factors $\exp \frac{1}{2} J \sigma_s^{t+1} \sigma_r^{t+1}$ and $\exp \frac{1}{2} J \sigma_s^t \sigma_r^t$ for each $r \in N(s)$, each of which cancel with corresponding factors in overlapping templates.

The kinetic Ising PCA is also reversible with respect to interchange of $\pm$.

References

[Ba] Bardet J-B, Limit theorems for coupled analytic maps, Probab Theory Related Fields 124 (2002) 151–177.

[Be] Beck C, Spatio-temporal chaos and vacuum fluctuations of quantized fields (World Sci, 2002).

[BG] Bennett C, Grinstein G, Role of irreversibility in stabilizing complex and nonergodic behaviour in locally interacting discrete systems, Phys Rev Lett 55 (1985) 657–660.

[BK] Bricmont J, Kupiainen A, High temperature expansions and dynamical systems, Commun Math Phys 178 (1996) 703–32.

[BuS] Bunimovich LA, Sinai Ya G, Space-time chaos in coupled map lattices, Nonlinearity 1 (1988) 491–516.

[BS] Burton R, Steif JE, Non-uniqueness of measures of maximal entropy for subshifts of finite type, Erg Th Dyn Sys 14 (1994) 213–35.

[CM] Chaté H, Manneville P, Collective behaviors in spatially extended systems with local interactions and synchronous updating, Prog Theor Phys 87 (1992) 1–60.

[CK] Crutchfield JP, Kaneko K, Are attractors relevant to turbulence? Phys Rev Lett 60 (1988) 2715–8.

[Ga] Gacs P, Reliable cellular automata with self-organization, J Stat Phys 103 (2001) 45–267.

[GM] Gielis G, MacKay RS, Coupled map lattices with phase transition, Nonlinearity 13 (2000) 867–888.

[GLP] Ginelli F, Livi R, Politi A, Emergence of chaotic behaviour in linearly stable systems, J Phys A 35 (2002) 499–516.

[Gr] Gray LF, A reader's guide to Gacs' "positive rates" paper, J Stat Phys 103 (2001) 1–44.

[Ha] Haken H, Cooperative phenomena in systems far from thermal equilibrium and in nonphysical systems, Rev Mod Phys 47 (1975) 67–121.

[JP] Jiang M, Pesin YaB, Equilibrium measures for coupled map lattices: existence, uniqueness and finite dimensional approximations, Commun Math Phys 193 (1998) 677–711.

[Ka84] Kaneko K, Period-doubling of kink-antikink patterns, quasiperiodicity in antiferro-like structures and spatial intermittency in coupled map lattices, Prog Th Phys 72 (1984) 480–486.

[Ka90] Kaneko K, Globally coupled chaos violates the law of large numbers, Phys Rev Lett 65 (1990) 1391–4.

[KH] Katok A, Hasselblatt B, Introduction to the modern theory of dynamical systems (Cambridge Univ Press, 1995).

[Ke] Kelly FP, Reversibility and stochastic networks (Wiley, 1979).

[Ku] Kuramoto Y, Cooperative dynamics of oscillator community, Prog Theor Phys 79 (1984) 223–240.

[LMS] Lebowitz JL, Maes C, Speer E, Statistical mechanics of probabilistic cellular automata, J Stat Phys 59 (1990) 117–70.

[Lig] Liggett TM, Interacting particle systems (Springer, 1985).

[LM] Losson J, Mackey MC, Thermodynamic properties of coupled map lattices, in: Stochastic and spatial structures of dynamical systems, eds Strien S van, Verduyn Lunel SM (N Holland, 1996), 41–69.

[Ma] MacKay RS, Dynamics of networks: features that persist from the uncoupled limit, in "Stochastic and spatial structures of dynamical systems", eds SJ van Strien, SM Verduyn Lunel (N Holland, 1996), 81–104.

[MH] Miller J, Huse DA, Macroscopic equilibrium from microscopic irreversibility in a chaotic coupled map lattice, Phys Rev E 48 (1993) 2528–35.

[Pe] Penrose R, The road to reality (Random House, 2004)

[PS] Pesin YaB, Sinai YaG, Space-time chaos in chains of weakly coupled interacting hyperbolic mappings, Adv Sov Math 3 (1991) 165–198.

[Sa] Sakaguchi H, Phase transition in coupled Bernoulli maps, Prog Theor Phys 80 (1988) 7–12.

[Si] Sinai YaG, Gibbs measures in ergodic theory, Russ Math Surv 27:4 (1972) 21–69.

[T+] Toom AL, Vasilyev NB, Stavskaya ON, Mityushin LG, Kurdyumov GL, Pirogov SA, Discrete local Markov systems, in: Stochastic cellular systems, ergodicity, memory, morphogenesis, eds Dobrushin RL, Kryukov VI, Toom AL (Manchester Univ Press, 1990), 1–182.

[Vi] Viana M, Stochastic dynamics of deterministic systems (Instituto de Matematica Pura e Aplicada, 1997).