

Overdetermined ODEs and rigid periodic states in network dynamics

Ian Stewart

Abstract. We consider four long-standing Rigidity Conjectures about synchrony and phase patterns for hyperbolic periodic orbits of admissible ODEs for networks. Proofs of stronger local versions of these conjectures, published in 2010–12, are now known to have a gap, but remain valid for a broad class of networks. Using different methods we prove local versions of the conjectures under a stronger condition, ‘strong hyperbolicity’, which is related to a network analogue of the Kupka–Smale Theorem. Under this condition we also deduce global versions of the conjectures and an analogue of the H/K Theorem in equivariant dynamics. We prove the Rigidity Conjectures for all 1- and 2-colourings and all 2- and 3-node networks by showing that strong hyperbolicity is generic in these cases.

1. Introduction

The phenomenon of synchrony in network dynamics has been widely studied for decades; see for example Boccaletti et al. [12] and Wang [76]. In the network context, sets of synchronous nodes are often called *clusters*. Another term is ‘partial synchrony’; see Belykh et al. [9, 11], Belykh and Hasler [10], Pecora et al. [61], Pogromsky [64], Pogromsky et al. [63]. In neurobiology, neurons are synchronised if they ‘fire together’, a relationship that is significant for neural processing and the architecture of the brain (Kopell and LeMasson [54], Singer [65], Uhlhaas et al. [74]). Manrubia et al. [59] discuss synchronisation in neural networks. Van Vreeswijk and Hansel [75] analyse several models, including a coupled system of Hodgkin–Huxley neurons that can produce spikes and bursts. Mosekilde et al. [60] describe a model of synchronisation in nephrons, structures in the kidneys that help regulate blood pressure.

A closely associated phenomenon is the occurrence of *phase patterns*: specific phase shifts (as fractions of the period) between nodes with otherwise identical time-periodic waveforms. Such phenomena occur in models of animal locomotion (Buono

2020 Mathematics Subject Classification. Primary 34C25; Secondary 05C99, 34C15.

Keywords. Network, periodic orbit, rigid, synchrony, phase shift, balanced, hyperbolic, strongly hyperbolic, overdetermined ODE.

and Golubitsky [14], Golubitsky et al. [40, 41]), peristalsis (Chambers et al. [18], Gjorgjieva et al. [29]), respiration (Butera et al. [15, 16]), and binocular rivalry and visual illusions (Curtu [19], Diekman et al. [23–25, 70]). There are also applications in the physical sciences, for instance to robotics (Campos et al. [17], Liu et al. [58]) and coupled lasers (Glova [30], Zhang et al. [77]).

The common occurrence of such patterns suggests that a unified theory, providing a conceptual framework applicable to arbitrary networks, could prove useful. One such framework is the ‘coupled cell’ formalism of [36, 39, 43, 71], which takes its inspiration from the topological approach to nonlinear dynamics of Arnold [7], Smale [67], and many others, and the analogous theory of equivariant dynamics and bifurcation of Golubitsky et al. [35, 42]. In this formalism, a network determines a class of admissible ODEs, and the primary aim is to relate the dynamics of such equations to the network architecture. Synchrony and phase patterns are closely associated with balanced colourings, which are combinatorial features of the network, and the associated quotient networks, whose admissible ODEs prescribe the dynamics of synchronous clusters. There are numerous existence theorems for steady and periodic states with prescribed synchrony and phase patterns, and some ideas extend to synchronised chaos. We summarise this formalism in Section 4.

However, the theory of synchrony and phase patterns for periodic states remains incomplete, because several key results are still conjectural. The main ones are the Rigidity Conjectures, discussed in Section 1.3. They have been proved for a broad class of networks, but it has recently been realised that the published proofs make a tacit assumption that fails for some networks; see Section 1.4. The aim of this paper is to prove these conjectures without that assumption, but under an extra technical hypothesis: ‘strong hyperbolicity’ of the periodic orbit. In some cases it is possible to dispense with this condition; in particular we prove the Rigidity Conjectures unconditionally for networks with up to 3 nodes and for all 1- and 2-colour patterns on any network. The general case, without assuming strong hyperbolicity, remains open, but could be dealt with using similar methods if it is possible to prove suitable network analogues of the Kupka–Smale Theorem, Section 1.5. The method shows that if counterexamples to the Rigidity Conjectures exist, they must have extremely degenerate systems of periodic orbits, Section 14.5.

For technical reasons indicated in Section 4.3, the analysis is carried out within a mild generalisation of the standard coupled cell formalism, developed in detail in [39]. We focus on synchrony and phase relations between nodes (formerly called ‘cells’), defined as follows. Consider an admissible ODE for the network (one that respects the network structure)

$$\dot{x} = f(x), \quad x_c \in P_c \quad (1.1)$$

where c runs through the set $\mathcal{C} = \{1, \dots, n\}$ of nodes, $x = (x_1, \dots, x_n)$, $f = (f_1, \dots, f_n)$, and the node spaces P_c are finite-dimensional real vector spaces. An orbit $\mathbb{X} = \{x(t)\}$ of (1.1) is *synchronous* for nodes c, d if the corresponding node states are equal for all time

$$x_c(t) \equiv x_d(t). \quad (1.2)$$

The orbit \mathbb{X} is *phase-related* for nodes c, d if the corresponding node states are equal for all times, up to a phase shift θ

$$x_c(t) \equiv x_d(t + \theta). \quad (1.3)$$

The *synchrony pattern* for $x(t)$ is the equivalence relation determined by all synchronous pairs. It can also be viewed as a partition of the nodes into synchronous clusters, or as a colouring in which synchronous nodes are given the same colour. We pass between these three interpretations without further comment.

For technical reasons, explained in Sections 7.6 and 7.7, we prefer to use local versions of these relations. Let $t_0 \in \mathbb{R}$. Then \mathbb{X} is *locally synchronous* for nodes c, d at $x(t_0)$ if there exists an open set $J \subseteq \mathbb{R}$ containing t_0 such that (1.2) holds for all $t \in J$, and *locally phase-related* for nodes c, d at $x(t_0)$ if (1.3) holds for all $t \in J$. The *local synchrony pattern* at $x(t_0)$ is the equivalence relation determined by all synchronous pairs.

The *local phase pattern* at a point $x(t_0) \in \mathbb{X}$ is slightly more complicated. Phase relations between pairs of nodes need not be unique, because the minimal period of $x_c(t)$ at a node c may differ from the minimal period of the entire orbit \mathbb{X} . For an oscillating node this can happen for a multirhythm state [37]; in fact it is possible for all minimal periods at nodes to differ from the overall minimal period. Moreover, the state of a steady node is fixed by all phase shifts. The *phase pattern* for \mathbb{X} encodes phase relations between pairs of nodes as sets $\Theta(c, d)$ containing all θ such that (1.3) is valid. The set $\Theta = \{\Theta(c, d)\}$ has a natural groupoid structure [72]. This definition also has a local version where again we consider only $t \in J$.

1.1. Rigidity

The notion of rigidity is central to the general theory of synchrony and phase patterns, because it excludes artificial examples where these patterns arise for non-generic reasons, such as couplings that are generically nonzero but vanish on the states concerned. A prerequisite for defining rigidity is that the relevant state (which we take to be either an equilibrium or a periodic orbit) should be hyperbolic. For an equilibrium this means that the Jacobian has no zero or purely imaginary eigenvalues, Guckenheimer and Holmes [44, Section 1.4]. For periodic orbits, the Floquet multipliers should not lie on the unit circle except for a simple eigenvalue 1 associated with the periodic orbit,

Hassard et al. [45, Section 1.4]. Equivalently, the derivative of a Poincaré return map should have no eigenvalues on the unit circle [44, Section 1.5].

In a general dynamical system, hyperbolicity of a given periodic orbit or equilibrium is a generic property. That is, it is:

- *Dense*: Every periodic orbit becomes hyperbolic after an arbitrarily small perturbation (if necessary).
- *Open*: After any sufficiently small perturbation, a hyperbolic periodic orbit remains hyperbolic.

The density property follows from the Kupka–Smale Theorem, a considerably stronger statement; see Kupka [55], Smale [66], and Peixoto [62]. Openness is obvious because eigenvalues (of the Jacobian at an equilibrium or the derivative of a Poincaré map at a periodic orbit) perturb continuously, Lancaster and Tismenetsky [56].

We describe the situation for periodic orbits; there are simpler analogous statements for equilibria. Consider an ODE

$$\dot{x} = f(x), \quad x \in \mathbb{R}^m \quad (1.4)$$

for a smooth (that is, C^∞) vector field f on a finite-dimensional Euclidean space \mathbb{R}^m . In the theory of general dynamical systems, any hyperbolic periodic state \mathbb{X} of this ODE with period T persists when f is replaced by any sufficiently small perturbation \tilde{f} ; see Hirsch et al. [47]. That is, locally there is a unique perturbed periodic state $\tilde{x}(t)$ near $x(t)$ with period \tilde{T} near T . Throughout this paper, ‘small’ refers to the C^1 norm, Section 6.1.

Definition 1.1. Suppose that (1.4) is an admissible ODE for a network. A property \mathcal{P} of a hyperbolic periodic orbit \mathbb{X} is *rigid* if, for any admissible perturbation $\tilde{f} = f + p$ of the vector field f , where p is sufficiently small, the perturbed periodic orbit $\tilde{\mathbb{X}}$ also has property \mathcal{P} .

Hyperbolicity ensures that a locally unique perturbed periodic orbit exists, so this definition makes sense. We allow property \mathcal{P} to depend on the period T , which is replaced by \tilde{T} in the perturbed ODE. So, for example, ‘nodes 1 and 2 are out of phase by half a period’ might be a rigid property. Rigidity is an ‘openness’ condition: in a suitable topology, the set of \mathbb{X} with property \mathcal{P} is open.

1.2. Motivation from equivariant dynamics

In this paper we focus on rigid synchrony and phase patterns of orbits of admissible ODEs for networks. A natural setting for these patterns, and an important source of motivation for the network theory, is equivariant dynamics [36, 42]. Here the map f

in (1.4) is equivariant for the action of a group Γ on \mathbb{R}^k ; that is,

$$f(\gamma x) = \gamma f(x) \quad \forall \gamma \in \Gamma.$$

Associated with any periodic orbit $\mathbb{X} = \{x(t)\}$ are two subgroups of Γ

$$K = \{\gamma : \gamma x(t) = x(t) \quad \forall t \in \mathbb{R}\},$$

$$H = \{\gamma : \gamma \mathbb{X} = \mathbb{X} \quad \forall t \in \mathbb{R}\}.$$

We call K the group of pointwise symmetries of \mathbb{X} , and H the group of setwise symmetries. The possible pairs (H, K) are classified by the H/K Theorem of Buono and Golubitsky [14]. For finite Γ the main conditions are that $K \triangleleft H$ and the quotient group $H/K \cong \mathbb{Z}_r$ is cyclic. There are other technical conditions if the state space has low dimension. It can then be shown that \mathbb{X} is a discrete rotating wave

$$\gamma x(t) = x(t + \theta), \quad \theta = mT/r$$

for all $t \in \mathbb{R}$ and for some integer m . The group K plays the role of a (global) synchrony pattern ($\theta = 0$), while H (or H/K) plays the role of a (global) phase pattern. By [36, Corollary 3.7], both H and K are rigid properties of \mathbb{X} . (The term used there is ‘robust’.) So the synchrony and phase patterns for (H, K) are rigid. It is easy to prove that if \mathcal{G} is a network with symmetry group Γ then every admissible map is Γ -equivariant. Therefore any synchrony or phase pattern arising from a pair (H, K) is rigid. However, equivariant maps need not be admissible [5, Section 3.1]. Examples of synchrony and phase patterns of these kinds can be found in many papers, for instance [4–6], Buono and Golubitsky [14], Golubitsky et al. [31, 32, 37, 57, 70].

It is well known that rigid synchrony patterns in networks can arise for reasons more general than symmetry. In particular, any balanced colouring of the nodes determines a rigid synchrony pattern [43, 71]. If the Rigid Synchrony Conjecture holds for the network, the converse is true. Moreover, if the Rigid Phase Conjecture also holds for the network, its rigid phase patterns arise from cyclic symmetries of quotient networks by balanced colourings [73].

Another source of motivation for the Rigidity Conjectures is their analogues for equilibria. It is proved in [43] that if a synchrony pattern of a steady state is rigid, then the corresponding colouring is balanced. That is, synchronous nodes have synchronous inputs, up to input isomorphism. Aldis [3, Chapter 7] gives another proof using transversality. A third proof, using methods along the lines of this paper, is in [69].

1.3. Rigidity Conjectures

In this paper, an orbit $\mathbb{X} = \{x(t)\}$ is defined to be *periodic* if there exists $T > 0$ such that $x(t + T) \equiv x(t)$, and $x(t)$ is not constant (which would be a steady state). A

periodic orbit of a network may be steady at some nodes; that is, some components $x_c(t)$ can be constant as t varies. However, $x_d(t)$ must *oscillate* (not be constant) for some node d . This can happen, for example, in feedforward networks, where a steady node is an input to oscillating nodes; see for example Golubitsky et al. [32].

The Rigidity Conjectures, stated about 15 years ago, comprise:

- (a) *Rigid Input Conjecture*: For any rigid synchrony or phase pattern, synchronous or phase-related nodes are input isomorphic. That is, they have the same number of input arrows for any given arrow-type.
- (b) *Rigid Synchrony Conjecture*: For any rigid synchrony pattern, corresponding input nodes inherit the same synchrony pattern, if suitably identified.
- (c) *Rigid Phase Conjecture*: For any rigid phase pattern, corresponding input nodes inherit the same phase pattern, if suitably identified.
- (d) *Full Oscillation Conjecture*: If a transitive network has a hyperbolic periodic state, there exist arbitrarily small admissible perturbations for which every node oscillates.

The identifications in (b) and (c) must be made using input isomorphisms, which preserve the node dynamics and the numbers and types of couplings. Although rigidity is not mentioned specifically in (d), it is implicit: a node that is steady after any small perturbation is rigidly steady. Moreover, (d) is a simple consequence of (c).

All four conjectures can also be stated as local versions, in which the relevant hypotheses are assumed to hold only for some non-empty open interval of time. We append the word ‘local’ to distinguish these. The local versions are not just generalisations: they have technical advantages and are essential to this paper. As it happens, the local versions imply the global ones, but this is not immediate from the definitions. The Rigidity Conjectures are all motivated by the same intuition: if two nodes have the same dynamics, except perhaps for a phase shift, then the same should be true for the nodes that input to them, up to some bijection and for the same phase shift. In (a) the common feature is having the same number of inputs of each arrow type; in (b) it is synchrony, in (c) it is a phase relation, and in (d) it is the node not oscillating.

The condition of rigidity is required for (a), (b), and (c), because without it, these conjectures are false [72, Section 7]. All known counterexamples are ‘non-generic’, having very special features that can be destroyed by small admissible perturbations of the ODE. It therefore makes sense to impose suitable genericity conditions. The natural choice is rigidity: the property concerned persists under small admissible perturbations. As mentioned in Section 1.1, the periodic orbit must be hyperbolic to ensure that a locally unique perturbed periodic orbit exists. See Section 6.2 for details.

For (d) a network is transitive (strongly connected, path-connected) if, for any two nodes c, d , there is a directed path from c to d . Statement (d) can be false if the

network is not transitive; for instance, in feedforward networks nodes upstream from an oscillating node can be rigidly steady.

1.4. Previous results

We summarise the current state of play for the Rigidity Conjectures. Conjectures (a), (b), (c) were stated in 2006 in Golubitsky and Stewart [37, Section 10], with a claim that (a) had been proved using strongly admissible coordinate changes. At that time, all four conjectures had been ‘folklore’ for some years. Fully oscillatory states are mentioned by Josič and Török [52] as hypotheses of existence theorems for periodic states of symmetric networks. Conjecture (d) is stated explicitly in [73]. A proof of conjecture (a) is presented in [72]. Proofs of conjectures (a,b,d) are presented in Golubitsky et al. [33] for stronger local versions. Similar methods were applied in Golubitsky et al. [34] to prove (c), again in a local version. Since the local versions have weaker hypotheses, they are stronger than the global ones. Curiously, the local versions are technically more tractable, Sections 7.2 and 7.6. Joly [51] proves the Full Oscillation Conjecture, but only for fully inhomogeneous networks. This proof uses transversality methods.

Readers familiar with the literature may wonder why we refer to the above statements as conjectures. The reason is that it has recently been noticed that there is a gap in the proofs in [33, 34, 72], which is related to the coordinate changes employed; see [69, Appendix]. Specifically, it is assumed that certain coordinate changes are ‘strongly admissible’, when sometimes they are not. This gap can be repaired by requiring the network to be *semihomogeneous*: such that input equivalence is the same as state (previously cell) equivalence. This class includes all homogeneous networks (all nodes have the same number of input arrows of each arrow type) and all fully inhomogeneous networks (all nodes and all arrows have different types and no multiple arrows or self-loops occur); see Golubitsky et al. [28, 38]. However, not all networks are semihomogeneous; a simple example is discussed in Section 3.

As already remarked, in this paper we prove all four conjectures, for any finite network, but only by making an extra assumption on the periodic orbit concerned, which we call ‘strong hyperbolicity’. A weaker property, ‘stable isolation’, suffices, but existing proofs of this property for specific networks establish strong hyperbolicity in any case. Both properties are closely related to the Kupka–Smale Theorem, which (among other things) asserts that hyperbolicity of all equilibria and periodic orbits is generic in a general dynamical system; see Section 1.5. A Kupka–Smale analogue for admissible ODEs would imply strong hyperbolicity; see Section 6.10.

The analysis shows that any hypothetical counterexample to any of the Rigidity Conjectures must have at least 4 nodes and involve a synchrony pattern with at least 3 colours. Moreover, the periodic orbit concerned must have properties that are highly

non-generic in a general dynamical system, and seem unlikely even when network constraints are imposed. Thus the results proved here, although requiring stronger hypotheses, represent a significant strengthening of the evidence in support of the Rigidity Conjectures.

1.5. Kupka–Smale Theorem

The Kupka–Smale Theorem plays a central role in this paper. It asserts a form of genericity (open and dense) for discrete dynamical systems (diffeomorphisms) or continuous ones (flows). We paraphrase the result:

Theorem 1.2 (Kupka–Smale Theorem). *Let M be a C^∞ manifold and let \mathcal{V} be the space of all C^r vector fields on M with the C^r topology, $r \geq 1$. For a general dynamical system on M , the following three properties are generic in \mathcal{V} ; that is open and dense – indeed, residual:*

- (a) *Every equilibrium point is hyperbolic.*
- (b) *Every periodic orbit is hyperbolic.*
- (c) *The stable and unstable manifolds of all equilibria and periodic orbits intersect transversely.*

Proof. For precise statements and definitions see Kupka [55], Smale [66], and the simplified proof in Peixoto [62, Section 2]. ■

There is a corresponding theorem for discrete dynamical systems (diffeomorphisms). In this paper we require both the continuous and the discrete versions, but the more delicate property (c) is not needed. Field [27] proves equivariant versions for discrete and continuous dynamics, which we apply in Section 14 to prove the Rigidity Conjectures for all 1- and 2-colourings and all 2-node and 3-node networks.

The original theorem of Kupka and Smale was restricted to compact manifolds, where the C^r topology is the usual metric topology defined by bounding norms of derivatives of order up to r . This restriction was removed by Peixoto, assuming the Whitney C^r topology. In this paper state spaces are \mathbb{R}^k , hence noncompact, but we can effectively reduce to the case of a compact manifold by using bump functions to make the admissible vector field vanish outside a compact set whose interior contains the periodic orbit under consideration. Therefore we can work with the metric C^r topology. (The space of C^∞ maps is not closed in this topology.) We make these statements precise in Section 6.

1.6. Is hyperbolicity generic for networks?

For a general dynamical system (equivalent to an all-to-all coupled fully inhomogeneous network) the Kupka–Smale Theorem proves that hyperbolicity of all periodic orbits is a generic property. Field [27] proves an equivariant analogue, which applies to some symmetric networks. Not all, because in this case admissible maps are always equivariant, but equivariant maps for networks need not be admissible [4].

Hyperbolicity and strong hyperbolicity are probably generic properties for all networks, except for one class where this statement is known to be false by Josić and Török [52, Remark 1]; see Section 6.5. Indeed, the Kupka–Smale Theorem fails for such networks. We suspect that this class of networks contains all exceptions, and that it can be avoided in the proof of the Rigidity Conjectures below, but currently we are unable to prove either of these statements.

1.7. Implications between the conjectures

In principle, any of the Rigidity Conjectures might be valid for some networks, or some periodic orbits, but false for others. It is therefore convenient to phrase the results of the conjectures as positive properties of the network and the periodic orbit:

Definition 1.3. Let \mathcal{G} be a network and let \mathbb{X} be a hyperbolic periodic orbit of an admissible ODE. Then the pair $(\mathcal{G}, \mathbb{X})$ has the following properties if the stated conditions hold:

- (a) *Rigid Input Property (RIP)*: For any rigid synchrony or phase pattern of \mathbb{X} , synchronous or phase-related nodes are input equivalent.
- (b) *Rigid Synchrony Property (RSP)*: For any rigid synchrony pattern of \mathbb{X} , corresponding input nodes inherit the same synchrony pattern, if suitably identified.
- (c) *Rigid Phase Property (RPP)*: For any rigid pattern of phase relations of \mathbb{X} , corresponding input nodes inherit the same phase pattern, if suitably identified.
- (d) *Full Oscillation Property (FOP)*: If \mathcal{G} is transitive and \mathbb{X} is a hyperbolic periodic state, there exist arbitrarily small perturbations for which every node oscillates.

If any of the above statements holds for all hyperbolic \mathbb{X} , we say that \mathcal{G} has the property concerned.

Each conjecture states that all networks have the corresponding property. There are local versions, where the stated conditions hold on a non-empty open interval of time, which we call the Local Rigid Synchrony Property (LRSP), and so on.

The properties are closely related. Known implications among the corresponding properties, for any specific \mathcal{G} , are

$$\text{FOP} \Leftarrow \text{RPP} \iff \text{RSP} \implies \text{RIP}.$$

Three implications are trivial: $\text{RPP} \implies \text{RSP}$, $\text{RSP} \implies \text{RIP}$, and $\text{RPP} \implies \text{RIP}$. It is also clear that $\text{RPP} \implies \text{FOP}$, because a node variable x_c is in equilibrium if and only if $x_c(t) \equiv x_c(t + \theta)$ for all $\theta \in \mathbf{S}^1$ (see Section 12 for a full discussion and a more general version of the result).

The most surprising implication is that $\text{RSP} \implies \text{RPP}$. This follows using the ‘doubling’ trick of Golubitsky et al. [34], which converts a phase relation on \mathcal{G} into a synchrony relation on two disjoint copies $2\mathcal{G}$, for a special periodic orbit on a 2-torus foliated by periodic orbits. Some care is needed to show that $2\mathcal{G}$ has suitable versions of the required properties, see Section 11.

2. Summary of paper

In [33, 34] the conjectures are proved in the order

$$\text{Rigid Input} \implies \text{Full Oscillation} \implies \text{Rigid Synchrony} \implies \text{Rigid Phase}.$$

The viewpoint there is local: local versions are more tractable and lead to stronger results while avoiding technical obstacles. Here we also consider local versions, for similar reasons, but we employ a different strategy:

$$\begin{array}{c} \text{Rigid Synchrony} \implies \text{Rigid Phase} \implies \text{Full Oscillation} \\ \Downarrow \\ \text{Rigid Input} \end{array}$$

The key result for the method employed in this paper is therefore the local form of (b), the Local Rigid Synchrony Property.

Consider an admissible ODE $\dot{x} = f(x)$, where $x = (x_1, \dots, x_n)$, and write it in components as $\dot{x}_c = f_c(x)$, where $1 \leq c \leq n$. Assume, for a contradiction, that this ODE has a hyperbolic periodic orbit $\mathbb{X} = \{x(t)\}$ with a rigid synchrony pattern \bowtie . Any such orbit satisfies the condition

$$c \bowtie d \implies x_c(t) = x_d(t) \quad \forall t \in \mathbb{R}$$

which implies that x_c satisfies two equations

$$\dot{x}_c = f_c(x), \quad \dot{x}_c = f_d(x).$$

If the colouring is balanced, these equations are identical, but if it is unbalanced, they are formally inconsistent – they involve different functions or the same function evaluated at different points. This does not of itself create a contradiction, because f_c and f_d might agree on \mathbb{X} while being distinct elsewhere. It seems highly implausible that such a situation can persist under all small admissible perturbations, but proving that is another matter. In the unbalanced case, the result of this substitution is an ‘overdetermined ODE’ (OODE), with more equations than unknowns, and the method of proof that we employ is to construct perturbations that exploit the formal inconsistency of the OODE to derive a contradiction.

Our method requires an additional assumption on the orbit \mathbb{X} : ‘strong hyperbolicity’ (Section 6.7), or the weaker property of ‘stable isolation’ (Section 6.9). We can establish this property rigorously for some networks and colourings. Heuristically, this extra property is plausible, and hyperbolic periodic orbits that lack it ‘ought to be’ of infinite codimension, hence highly non-generic. However, we are unable to prove this in full generality with current methods. For further discussion see Section 6.7 and Section 6.9.

As motivation, and to provide a simple example of the proof technique, Section 3 considers the special case of a 3-node directed ring with two arrow types. In this network, admissible diagonal maps are not strongly admissible, so the results of Golubitsky et al. [33, 34] do not apply. Using the relevant OODEs, we prove the Rigid Synchrony Conjecture for this network. This example emphasises the central role of the Kupka–Smale Theorem.

Section 4 recalls the relevant features of the basic formalism of coupled cell networks (henceforth just ‘networks’) from [43, 71], including properties of the quotient network by a balanced colouring. An important point is a simple characterisation of admissible maps in Proposition 4.7, originally proved in [71, Proposition 4.6], which reduces admissibility to invariance under the appropriate vertex group.

Section 5 sets up a generalisation of the usual quotient network construction for a balanced colouring \bowtie , by throwing away the balance condition. The resulting ‘quasi-quotient’ depends on a choice of representatives \mathcal{R} for the colouring, and fails to have most of the useful properties of the quotient. However, it retains two key properties in relation to states with synchrony pattern \bowtie : such states induce solutions for the induced ODE on the quasi-quotient, and solutions of the induced equation that also satisfy a set of constraint equations lift to solutions of the original ODE with synchrony pattern \bowtie .

Section 6 reviews standard results concerning hyperbolicity and genericity in general dynamical systems, provides rigorous definitions of the C^1 norm and rigidity, relates these concepts to network dynamics, and introduces two properties that are central to the methods of this paper: stable isolation and strong hyperbolicity. It ends with a discussion of the Kupka–Smale Theorem and network analogues.

Section 7 begins the general programme to prove the Local Rigid Synchrony Conjecture for any network under the assumption of strong hyperbolicity. Here we define a local version of rigid synchrony, needed to set up the proof. We discuss lower semi-continuity of colourings, the existence of generic points for local rigid synchrony, and technical obstacles that arise if instead we try to work with global rigid synchrony.

Section 8 describes a method for constructing admissible perturbations with small support and small C^1 norm. The construction uses bump functions and a symmetrisation technique.

The Local Rigid Synchrony Conjecture, which lies at the heart of this paper, is proved in Section 9 under the hypothesis of strong hyperbolicity. The main obstacle to finding such a proof has always been to gain enough control over how the perturbed periodic orbit $\tilde{\mathbb{X}}$ moves when the vector field f is perturbed. Previous attacks on the conjecture employ various strategies to do this, such as constructing flows geometrically [72] or using perturbations related to adjacency matrices and delicate estimates for integrals along the periodic orbit [33]. Instead, we analyse the quasi-quotient $\mathcal{G}^{\mathcal{R}}$ for a set of representatives \mathcal{R} of the synchrony colouring and the structure of the resulting ODE. This leads to an induced ODE for $\mathcal{G}^{\mathcal{R}}$ on the synchrony space $P^{\mathcal{R}}$, together with constraint equations. We retain control of $\tilde{\mathbb{X}}$ by constructing perturbations that leave the corresponding orbit in $P^{\mathcal{R}}$ unchanged, assuming local rigidity. Indeed, the constraint equations created by local rigidity imply that $\tilde{\mathbb{X}} = \mathbb{X}$ for sufficiently small perturbations of that kind, so the periodic orbit in P is also unchanged. The constraint equations then lead to a contradiction when the synchrony colouring is unbalanced. The Rigid Input Property follows immediately because balanced colourings refine input equivalence.

Section 10 parlays this local theorem into a proof of the standard global version of the Rigid Synchrony Property, using the lattice of colourings [68].

The Local Rigid Phase Property is deduced in Section 11 using the same trick that underlies the proof in [34], namely, form two disjoint copies of the network and convert a phase relation into a synchrony relation. The Local Rigid Synchrony Property does not apply directly, because doubling the network destroys hyperbolicity and strong hyperbolicity; instead of being isolated, periodic orbits defined by phase-shifted pairs foliate a 2-torus. However, rigidity of the phase shift gives a *canonical* choice of perturbed periodic orbit among those on the torus. The proof of the Local Rigid Synchrony Conjecture generalises directly to this situation, and the conclusion that any rigid synchrony relation is balanced implies the Local Rigid Phase Conjecture.

We deduce a local version of the Full Oscillation Conjecture in Section 12. The main observation is that if a periodic state of the network is steady at some node, for some time interval J , then on J that node is phase-related to itself by all possible phase shifts.

Section 13 provides a brief discussion of implications for analogues of the H/K Theorem in equivariant dynamics, whose proofs, given in [34,73], are contingent upon the four Rigidity Properties. These results are striking consequences of those properties because they classify all possible locally or globally rigid synchrony and phase patterns. In particular, they show that rigid phase patterns arise from cyclic group symmetries of a quotient network (or its completion if the quotient is not transitive).

In special cases, strong hyperbolicity of the periodic orbit can be proved to be generic using current techniques. In Section 14 we carry out this programme for 1- and 2-colourings, and deduce the Rigidity Properties for all networks with at most 3 nodes. Here a key role is played by the notion of ODE-equivalence: distinct networks with the same space of admissible maps. We classify 2-node networks up to ODE-equivalence.

Finally, we remark on the strong link that emerges from the methods of this paper, relating the four Rigidity Properties to results of Kupka–Smale type for networks, and point out that if a counterexample to any of the Rigidity Conjectures exists, the periodic orbit \mathbb{X} must satisfy conditions that would be extremely degenerate in a general dynamical system.

3. 3-node example

We set the scene by analysing a simple example, used later to motivate the proof of the Rigid Synchrony Conjecture for strongly hyperbolic periodic orbits. (For this network, strong hyperbolicity can be proved.) To focus on the main idea we postpone formal definitions of the terminology to Section 4. The discussion should be clear without these. The main feature of the general case that does not arise for this example is invariance under vertex groups. This step requires a straightforward symmetrisation, Section 8.1.

Figure 1 shows a network \mathcal{G} with three nodes, $\mathcal{C} = \{1, 2, 3\}$. There are two arrow types, solid and dashed. The shading on the nodes indicates ‘input type’, formalised below. Each of nodes 1 and 2 receives its input from a single solid arrow, so they have




Figure 1. A 3-node directed ring \mathcal{G} , with two arrow types, that is not semihomogeneous.

the same input type. Node 3 receives its input from a single dashed arrow, so it has a different input type.

A network is *semihomogeneous* if state equivalent nodes are input equivalent. This class includes all homogeneous networks and all fully inhomogeneous ones, together with many others. The proofs in [33,34] are valid for all semihomogeneous networks. However, there exist networks that are not semihomogeneous, the simplest example being \mathcal{G} . Indeed, all nodes of \mathcal{G} are cell (or state) equivalent, but there are two distinct input equivalence classes, $\{1, 2\}$ and $\{3\}$. Therefore the results of [33,34] do not apply to \mathcal{G} . However, by constructing suitable admissible perturbations, we prove that in fact it does have all four Rigidity Properties. The central role of the Kupka–Smale Theorem arises naturally from the method.

Admissible ODEs for \mathcal{G} (defined in Section 4.3) have the form

$$\begin{aligned}\dot{x}_1 &= f(x_1, x_3), \\ \dot{x}_2 &= f(x_2, x_1), \\ \dot{x}_3 &= g(x_3, x_2).\end{aligned}\tag{3.1}$$

Here the variables x_c lie in *node spaces* P_c , which we take to be real vector spaces \mathbb{R}^{k_c} . In order for (3.1) to respect the network architecture, the domains and ranges of these functions must be

$$\begin{aligned}f &: P_1 \times P_3 \rightarrow P_1, \\ f &: P_2 \times P_1 \rightarrow P_2, \\ g &: P_3 \times P_2 \rightarrow P_3.\end{aligned}$$

The function f occurs twice, so the domains and ranges in the two cases must coincide. That is, $P_1 = P_2$ and $P_3 = P_1$. In other words, the network structure and the definition of admissible ODEs requires $P_1 = P_2 = P_3$. This is an example of ‘state equivalence’, which replaces the usual ‘compatibility conditions’ on head and tail nodes of arrows in [43,71]. The reasons for this change are discussed in Section 4.3, and in greater detail in [39].

3.1. Strong admissibility

We find the strongly admissible maps for this network and verify their composition properties directly. Some notation defined in Section 4 is convenient. To avoid complications concerning state equivalence, we use the previous notation \sim_C for cell equivalence and \sim_I for input equivalence. (In this case, cell equivalence with the previous compatibility conditions turns out to be the same as state equivalence.)

Strongly admissible maps are defined in [43] as ‘diagonal’ maps

$$g(x) = (g_1(x_1), \dots, g_n(x_n))\tag{3.2}$$

such that $g_i = g_j$ whenever $i \sim_C j$. It is proved there that if f is admissible and g is strongly admissible, then both $f \circ g$ and $g \circ f$ are admissible.

In [33, 34, 72] it was tacitly assumed that the same composition properties hold if $g_i = g_j$ whenever $i \sim_I j$; that is, if g is diagonal and admissible. We show that this statement is false for \mathcal{G} , implying that new methods are required to prove any of the Rigidity Conjectures for this network.

Since the two solid arrows have the same type, the previous compatibility condition requires all three nodes to have the same cell type. The network \mathcal{G} has two different input types $\{1, 2\}$ and $\{3\}$, so \sim_C is different from \sim_I . Admissible diagonal maps (3.2) have $g_1 = g_2 \neq g_3$. Strongly admissible maps, as defined in [43, Definition 7.2], have $g_1 = g_2 = g_3$.

First, we show that the only diagonal maps g that compose on the right with admissible maps f to give admissible maps are maps g where $g_1 = g_2 = g_3$. That is, $g_i = g_j$ whenever $i \sim_C j$, in accordance with [43, Lemma 7.3]. We have

$$f(x) = \begin{bmatrix} f_1(x_1, x_3) \\ f_1(x_2, x_1) \\ f_2(x_3, x_2) \end{bmatrix}, \quad f \circ g(x) = \begin{bmatrix} f_1(g_1(x_1), g_3(x_3)) \\ f_1(g_2(x_2), g_1(x_1)) \\ f_2(g_3(x_3), g_2(x_2)) \end{bmatrix}.$$

If $f \circ g$ is admissible for all f , then the first two components yield

$$f_1(g_2(x_2), g_1(x_1)) = f_1(g_1(x_2), g_3(x_1)).$$

Take $f_1(u, v) = u$ to give $g_2(x_2) = g_1(x_2)$, and then $f_1(u, v) = v$ to give $g_1(x_1) = g_3(x_1)$. Therefore $g_1 = g_2 = g_3$.

Conversely, any map of this form composes on the right to give an admissible map.

In contrast, we now show that the maps g that compose on the left with admissible maps f to give admissible maps are maps g where $g_1 = g_2$. That is, $g_i = g_j$ whenever $i \sim_I j$. These are precisely the admissible diagonal maps. Now

$$g \circ f(x) = \begin{bmatrix} g_1(f(x_1, x_3)) \\ g_2(f(x_2, x_1)) \\ g_3(f(x_3, x_2)) \end{bmatrix}.$$

If $g \circ f$ is admissible for all f , then the first two components yield

$$g_2(f(x_1, x_3)) = g_1(f(x_1, x_3)).$$

Take $f(u, v) = u$, obtaining $g_2(u) = g_1(u)$. Therefore $g_1 = g_2$. Conversely, any map of this form composes on the left to give an admissible map.

3.2. Construction of suitable perturbations

Since the network \mathcal{G} of Figure 1 is not semihomogeneous, the results of [33, 34, 73] do not apply. Nevertheless, we now prove by a different method that \mathcal{G} has the Local Rigid Synchrony Property. As mentioned in Section 1.7 and proved in Sections 11 and 12, it therefore has the other three Local Rigidity Properties as well. The method used for this example motivates the subsequent approach to synchrony patterns on arbitrary networks. For this network we obtain a complete proof, because we can apply the standard Kupka–Smale Theorem and the equivariant version of Field [27] for the symmetry group \mathbb{Z}_2 . In the general case some network version, not yet proved, is required: this is why we impose strong hyperbolicity in the bulk of this paper. We discuss the first case in detail, to establish the logic, and provide less detail for the other cases.

Admissible ODEs take the form $\dot{x} = f(x)$ with certain conditions on the components f_c . We consider an arbitrary 1-parameter family of perturbations $\dot{x} = f(x) + \varepsilon p(x)$, where p is also admissible. Explicitly, admissibility requires

$$\begin{aligned}\dot{x}_1 &= f(x_1, x_3) + \varepsilon p(x_1, x_3), \\ \dot{x}_2 &= f(x_2, x_1) + \varepsilon p(x_2, x_1), \\ \dot{x}_3 &= g(x_3, x_2) + \varepsilon q(x_3, x_2),\end{aligned}\tag{3.3}$$

where p, q are arbitrary smooth functions because vertex symmetries are trivial. We can choose p, q to be bounded using bump functions, so the perturbation is C^1 -small when $\varepsilon \ll 1$. See Section 8.2.

The only balanced colouring is the trivial one with all nodes of different colours. We show that for every other colouring \bowtie , rigid synchrony leads to a contradiction. This is obtained by applying the Kupka–Smale Theorem (or Field’s equivariant version) for certain 1- and 2-node networks, when p is constructed to have certain properties that depend on the colouring \bowtie . Throughout we assume only that the synchrony pattern \bowtie is valid for t in some non-empty open interval J , and choose a point $t_0 \in J$.

Case (A): $\bowtie = \{\{1, 2, 3\}\}$. (Here and elsewhere we write \bowtie as a partition of \mathcal{C} .) For given ε , any (periodic) orbit $\mathbb{X}^\varepsilon = (x_1^\varepsilon, x_2^\varepsilon, x_3^\varepsilon)$ with this synchrony pattern has the fully synchronous form

$$(u^\varepsilon(t), u^\varepsilon(t), u^\varepsilon(t)).$$

We assume that \mathbb{X}^0 is hyperbolic and $0 < \varepsilon \ll 1$. Taking a suitable Poincaré section Σ at $x(t_0)$ and setting initial conditions by requiring $x^\varepsilon(t_0) \in \Sigma$, we can assume that $x^\varepsilon(t)$ varies continuously (indeed, by the implicit function theorem applied to a first-return map, smoothly) with ε and t .

Substituting $x_1^\varepsilon = x_2^\varepsilon = x_3^\varepsilon = u^\varepsilon$ in (3.3), this state must satisfy the conditions

$$\begin{aligned} \dot{u}^\varepsilon &= f(u^\varepsilon, u^\varepsilon) + \varepsilon p(u^\varepsilon, u^\varepsilon), \\ \dot{u}^\varepsilon &= f(u^\varepsilon, u^\varepsilon) + \varepsilon p(u^\varepsilon, u^\varepsilon), \\ \dot{u}^\varepsilon &= g(u^\varepsilon, u^\varepsilon) + \varepsilon q(u^\varepsilon, u^\varepsilon). \end{aligned} \quad (3.4)$$

The first component determines u^ε uniquely, for given initial conditions. The second is the same as the first. The third is different, and potentially contradictory; we use it to derive a contradiction.

The projection $\mathbb{U}^0 = \{(u^0(t))\}$ of \mathbb{X}^0 into $P_1 = P_2 = P_3$ is a periodic orbit of the ‘induced ODE’

$$\dot{y} = f(y, y). \quad (3.5)$$

We pre-prepare f so that \mathbb{U}^0 is hyperbolic on P_1 . This follows from the Kupka–Smale Theorem, since (3.5) is an ODE on P_1 and any perturbation $p(y)$ can be expressed in the form $g(y, y, y)$. By rigidity, the local synchrony pattern $\{\{1, 2, 3\}\}$ applies to this perturbed ODE provided we make the perturbation small enough. The open interval J may have to be replaced by a smaller open interval J' where $t_0 \in J' \subseteq J$.

Having pre-prepared f and \mathbb{X}^0 to make \mathbb{U}^0 hyperbolic on P_1 , we can realise the contradiction as follows. To simplify notation, write

$$u^* = u^0(t_0).$$

Define $p(x) \equiv 0$ for all x , and define q so that $q(u^*, u^*) \neq 0$. This is possible because p, q are arbitrary independent smooth maps. The first equation now becomes

$$\dot{u}^\varepsilon = f(u^\varepsilon, u^\varepsilon)$$

which is the same as the unperturbed equation. Therefore, near (u^*, u^*) , the periodic orbits u^ε and u^0 satisfy the same ODE, and $u^\varepsilon \rightarrow u^0$ as $\varepsilon \rightarrow 0$.

Since \mathbb{U}^0 is hyperbolic on P_1 , there is a locally unique periodic orbit near $u^0(t)$. But $u^\varepsilon \rightarrow u^0$ as $\varepsilon \rightarrow 0$. Therefore, for small $\varepsilon > 0$, we have $u^\varepsilon(t) \equiv u^0(t)$ for all t near t_0 (which implies equality for all t by uniqueness of solutions to ODEs). Therefore $\dot{u}^\varepsilon(t) \equiv \dot{u}^0(t)$. Set $t = t_0$ to obtain

$$g(u^*, u^*) + \varepsilon q(u^*, u^*) = g(u^*, u^*) + 0 \cdot q(u^*, u^*) = g(u^*, u^*).$$

This implies that $\varepsilon q(u^*, u^*) = 0$ for some $\varepsilon > 0$. However, we chose q so that this is false. Therefore \mathbb{X} cannot have local rigid synchrony pattern $\{1, 2, 3\}$.

Case (B): $\bowtie = \{\{1, 2\}, \{3\}\}$. By Case (A) we may assume that \bowtie is the finest colouring such that \mathbb{X} has local synchrony pattern \bowtie at t_0 . We follow similar reasoning, and omit routine details.

For given ε , any (periodic) orbit $\mathbb{X}^\varepsilon = (x_1^\varepsilon, x_2^\varepsilon, x_3^\varepsilon)$ with this local synchrony pattern has the form

$$(u^\varepsilon(t), u^\varepsilon(t), v^\varepsilon(t)).$$

Substitute $x_1^\varepsilon = x_2^\varepsilon = u^\varepsilon$ and $x_3^\varepsilon = v^\varepsilon$ in (3.3) to obtain

$$\dot{u}^\varepsilon = f(u^\varepsilon, v^\varepsilon) + \varepsilon p(u^\varepsilon, v^\varepsilon), \quad (3.6)$$

$$\dot{u}^\varepsilon = f(u^\varepsilon, u^\varepsilon) + \varepsilon p(u^\varepsilon, u^\varepsilon), \quad (3.7)$$

$$\dot{v}^\varepsilon = g(v^\varepsilon, u^\varepsilon) + \varepsilon q(v^\varepsilon, u^\varepsilon). \quad (3.8)$$

Components (3.6) and (3.8) determine u^ε and v^ε uniquely, for given initial conditions. Equation (3.7) is formally different from (3.6), and potentially contradictory.

The perturbation terms in (3.6) and (3.8) have the form $(p(u^\varepsilon, v^\varepsilon), q(v^\varepsilon, u^\varepsilon))$, which is a general vector field on $P_1 \times P_3$ with variables $u^\varepsilon, v^\varepsilon$. We can therefore apply the Kupka–Smale Theorem (for a general dynamical system) to pre-prepare f, \mathbb{X} so that the projected orbit $\mathbb{U}^0 = \{(u^0(t), v^0(t))\}$ is hyperbolic on $P_1 \times P_3$. We retain the same notation.

Choose a time $t_1 \in J$ so that if $u^* = u^0(t_1)$ and $v^* = v^0(t_1)$ then $u^* \neq v^*$. If this is not possible then we are in Case (A), already dealt with; this is also contrary to \bowtie being the finest local synchrony pattern.

Define $q \equiv 0$, and define p so that $p(u, v) \equiv 0$ in a neighbourhood of (u^*, v^*) , but $p(u^*, u^*) \neq 0$. This is possible since $u^* \neq v^*$, so $(u^*, u^*) \neq (u^*, v^*)$. Indeed, we can use a bump function to make p vanish outside a small neighbourhood of (u^*, u^*) , but be nonzero near (u^*, u^*) .

When (u, v) is near (u^*, v^*) , equations (3.6) and (3.8) reduce to

$$\dot{u}^\varepsilon = f(u^\varepsilon, v^\varepsilon), \quad (3.9)$$

$$\dot{v}^\varepsilon = g(v^\varepsilon, u^\varepsilon), \quad (3.10)$$

which is the same ODE as the unperturbed equation, but with variables $u^\varepsilon, v^\varepsilon$ in place of u^0, v^0 .

As before, the pre-preparation guarantees local uniqueness of perturbed periodic orbits on $P_1 \times P_3$, so this implies that $p(u^*, u^*) = 0$, a contradiction.

Case (C): $\bowtie = \{\{1, 3\}, \{2\}\}$. The argument has a similar structure. Synchronous orbits have the form

$$(u^\varepsilon(t), v^\varepsilon(t), u^\varepsilon(t)).$$

Substitute $x_1^\varepsilon = x_3^\varepsilon = u^\varepsilon$ and $x_2^\varepsilon = v^\varepsilon$ in (3.3) to obtain

$$\dot{u}^\varepsilon = f(u^\varepsilon, u^\varepsilon) + \varepsilon p(u^\varepsilon, u^\varepsilon), \quad (3.11)$$

$$\dot{v}^\varepsilon = f(v^\varepsilon, u^\varepsilon) + \varepsilon q(v^\varepsilon, u^\varepsilon), \quad (3.12)$$

$$\dot{u}^\varepsilon = g(u^\varepsilon, v^\varepsilon) + \varepsilon p(u^\varepsilon, v^\varepsilon). \quad (3.13)$$

Components (3.12) and (3.13) determine u^ε and v^ε uniquely, for given initial conditions. Equation (3.11) is formally different from (3.11). When $\varepsilon = 0$ equations (3.12) and (3.13) are a general ODE on $P_1 \times P_2$, and (p, q) is an arbitrary vector field on $P_1 \times P_2$. We can use the Kupka–Smale Theorem to pre-prepare f, g so that the periodic orbit $\mathbb{U}^0 = \{u^0(t), v^0(t)\}$ is hyperbolic on $P_1 \times P_2$.

Choose the perturbation so that $q \equiv 0, p(v, u) \equiv 0$ near (v^*, u^*) , but $p(u^*, u^*) \neq 0$. When (u, v) is near (u^*, v^*) , the orbit $(u^\varepsilon, v^\varepsilon)$ satisfies the same ODE as (u^0, v^0) , and hyperbolicity implies local uniqueness, so $(u^\varepsilon(t), v^\varepsilon(t)) = (u^0(t), v^0(t))$ near t_1 . As before, this implies that $p(u^*, u^*) = 0$, a contradiction.

Case (D): $\bowtie = \{\{2, 3\}, \{1\}\}$. Again the argument has a similar structure. Synchronous orbits have the form

$$(u^\varepsilon(t), v^\varepsilon(t), v^\varepsilon(t)).$$

Substitute $x_1^\varepsilon = u^\varepsilon$ and $x_2^\varepsilon = x_3^\varepsilon = v^\varepsilon$ in (3.3) to obtain

$$\dot{u}^\varepsilon = f(u^\varepsilon, v^\varepsilon) + \varepsilon p(u^\varepsilon, v^\varepsilon), \quad (3.14)$$

$$\dot{v}^\varepsilon = f(v^\varepsilon, u^\varepsilon) + \varepsilon q(v^\varepsilon, u^\varepsilon), \quad (3.15)$$

$$\dot{v}^\varepsilon = g(v^\varepsilon, v^\varepsilon) + \varepsilon p(v^\varepsilon, v^\varepsilon). \quad (3.16)$$

Components (3.14) and (3.15) determine u^ε and v^ε uniquely, for given initial conditions. Equation (3.16) is formally different from (3.15).

When $\varepsilon = 0$ equations (3.14) and (3.15) are a general \mathbb{Z}_2 -equivariant ODE on $P_1 \times P_2$, where \mathbb{Z}_2 swaps u^ε and v^ε , and (p, q) is an arbitrary \mathbb{Z}_2 -equivariant vector field on $P_1 \times P_2$. We can use Field's equivariant Kupka–Smale Theorem to pre-prepare f, g so that the periodic orbit $\mathbb{U}^0 = \{u^0(t), v^0(t)\}$ is hyperbolic on $P_1 \times P_2$.

Choose the perturbation so that $q \equiv 0, p(u, v) \equiv 0$ near (u^*, v^*) , but $p(v^*, v^*) \neq 0$. When (u, v) is near (u^*, v^*) , the orbit $(u^\varepsilon, v^\varepsilon)$ satisfies the same ODE as (u^0, v^0) , and hyperbolicity implies local uniqueness, so $(u^\varepsilon(t), v^\varepsilon(t)) = (u^0(t), v^0(t))$ near t_1 . As before, this implies that $p(v^*, v^*) = 0$, a contradiction.

We conclude that the only local rigid synchrony pattern is trivial, verifying the Local Rigid Synchrony Property for \mathcal{G} . The other three Local Rigidity Properties follow, as outlined in Section 1.7 and discussed in detail in Sections 11 and 12.

Remark 3.1. Since the Rigidity Conjectures were first stated it has been clear that the main obstacle to proving them is to retain enough control over the behaviour of the perturbed periodic orbit. In [33, 34] this is achieved by delicate estimates. The method employed above controls the perturbed periodic orbit by *not changing it*. Obviously a zero perturbation has this property, but the admissible perturbation that we construct changes the constraint equations. This construction leads to a contradiction when the local synchrony colouring is rigid but not balanced.

This example suggests that a similar type of perturbation of the induced ODE on the synchrony space might be used for an arbitrary network, and that the main obstacle is to prove a suitable version of the Kupka–Smale Theorem, so that (assuming rigidity) the perturbed periodic orbit is the same as the unperturbed one, but the constraints of synchrony lead to a contradiction. In the rest of the paper we show that this approach succeeds, modulo a version of Kupka–Smale for networks.

4. Formal definition of a network and standard properties

We now proceed to the general case. First, we briefly recall some basic concepts of the ‘coupled cell’ network formalism introduced in [71] and generalised in [43], and state some standard notations, definitions, and results. For further details, see [33, 37, 39, 43, 69]. We introduce a further slight generalisation, which resolves the dual role of ‘cell equivalence’ in the previous formalism. All of the standard theory extends to this more general setting, which applies to a wider range of ODEs with network structure. Full details are presented in [39]; everything in this paper is valid in this more general setting.

We begin with the formal setting for networks.

Definition 4.1. A network $\mathcal{G} = (\mathcal{C}, \sim_{\mathcal{C}}, \mathcal{A}, \sim_{\mathcal{A}}, \mathcal{H}, \mathcal{T})$ consists of:

- (a) A finite set of *nodes* \mathcal{C} and a *node-type* assigned to each node. Write

$$c \sim_{\mathcal{C}} d$$

if $c, d \in \mathcal{C}$ have the same node-type.

- (b) A finite set of *arrows* \mathcal{A} and an *arrow-type* assigned to each arrow. Write

$$a \sim_{\mathcal{A}} b$$

if $a, b \in \mathcal{A}$ have the same arrow-type. (The previous notation uses \mathcal{E} for \mathcal{A} and \sim_E for $\sim_{\mathcal{A}}$.)

The node type can be viewed as a distinguished ‘internal’ arrow-type.

- (c) Each $a \in \mathcal{A}$ has a *head node* $\mathcal{H}(a)$ and a *tail node* $\mathcal{T}(a)$ in \mathcal{C} . When viewing a node $c \in \mathcal{C}$ as an internal arrow, we define $\mathcal{H}(c) = c = \mathcal{T}(c)$.

Remark 4.2. Readers familiar with the literature will observe that we have omitted from this definition the standard ‘compatibility condition’ that arrow-equivalent arrows have node-equivalent heads and node-equivalent tails. This condition combines two roles for cell-equivalence that are better kept distinct, namely equality of state spaces (which we call \sim_S below) and equality of the distinguished ‘internal

arrows' on nodes (where we retain the notation \sim_C). In its place, we impose a natural condition on the state spaces (or phase spaces) assigned to nodes, see Definition 4.5 (a).

4.1. Input sets and tuples

Definition 4.3. Let c, d be nodes in \mathcal{C} .

- (a) The *input set* of c is the set $I(c)$ of all arrows $a \in \mathcal{A}$ such that $\mathcal{H}(a) = c$.
- (b) An *input isomorphism* is an arrow-type preserving bijection $\beta : I(c) \rightarrow I(d)$. That is, $c \sim_C d$ and $a \sim_A \beta(a)$ for all $a \in I(c)$.
- (c) Two nodes c and d are *input isomorphic* or *input equivalent* if there exists an input isomorphism from $I(c)$ to $I(d)$. In this case we write

$$c \sim_I d.$$

The set of input isomorphisms from c to d is denoted by $B(c, d)$. The disjoint union of these sets

$$\mathcal{B} = \bigcup_{c, d \in \mathcal{C}} B(c, d) \quad (4.1)$$

is a groupoid; see Brown [13], Higgins [46], and [37, 43, 71].

4.2. Redundancy

The definitions of node- and arrow-types, as stated, allow nodes or arrows to be assigned the same type even when they are not related by an input isomorphism – that is, they are in different groupoid orbits. This *redundancy* is often convenient, especially when drawing network diagrams. It does not affect the class of admissible maps, which depends only on the input isomorphisms, but it can cause problems in some constructions and introduces an ambiguity into the definition of the adjacency matrix for a given arrow type. Redundancy can be avoided by requiring the types to be the same if and only if the nodes or arrows are related by an input isomorphism. The resulting network is said to be *irredundant*, and we assume this throughout.

4.3. Admissible maps and ODEs

We now define admissible maps and ODEs, and state an equivalent property that is central to this paper.

Assign to each node $c \in \mathcal{C}$ a *node space* P_c . This is usually taken to be a real vector space \mathbb{R}^{k_c} , and we make this assumption throughout the paper. The overall

state space of the *network system* (or *coupled cell system*) of ODEs is

$$P = P_1 \times \cdots \times P_n.$$

In node coordinates, a map $f : P \rightarrow P$ has components f_c for $c \in \mathcal{C}$ such that

$$f_c : P \rightarrow P_c.$$

Remark 4.4. More generally, node spaces can be C^∞ manifolds, Field [26]. In phase oscillator models (Izhikevich and Ermentrout [50]) all node spaces are the circle, so $P_c = \mathbf{S}^1$. The methods employed in this paper probably generalise to manifolds. However, Golubitsky et al. [31] show that the topology of node spaces can change the list of possible phase patterns in the H/K Theorem, so it should not be assumed that all of the results proved here automatically remain valid when node spaces are manifolds, or that they are independent of their topology.

In Example 3 we noted that in order for admissible ODEs to make sense, certain equalities are forced on node spaces. These equalities arise whenever nodes $c \neq d$ are input isomorphic. For any input arrow $e \in I(c)$, and to $e' = \beta(e) \in I(d)$ where $\beta \in B(c, d)$, we require

$$P_{\mathcal{H}(e)} = P_{\mathcal{H}(e')}, \quad P_{\mathcal{T}(e)} = P_{\mathcal{T}(e')}. \quad (4.2)$$

The first equation reduces to $P_c = P_d$, so input isomorphic nodes must have the same state space. However, the second equation can impose further equalities. We say that i, j are *state-equivalent*, written $i \sim_S j$, if the above equations, taken over all c, d , require $P_i = P_j$. This is the transitive closure of the relation on \mathcal{C} defined by (4.2). It resolves an ambiguity in the usual concept of cell equivalence by distinguishing between having the same node space, and having the same node dynamic. It also extends the possible types of network without changing any of the basic theorems or proofs [39].

For any tuple of nodes $\mathbf{c} = (c_1, \dots, c_m)$ we write

$$\begin{aligned} P_{\mathbf{c}} &= P_{c_1} \times \cdots \times P_{c_m}, \\ x_{\mathbf{c}} &= (x_{c_1}, \dots, x_{c_m}). \end{aligned}$$

The input set of node c defines an *input tuple* $I(c) = (c, \mathcal{T}(i_1), \dots, \mathcal{T}(i_\nu))$, where the i_j are the arrows satisfying $\mathcal{H}(i_j) = c$. For brevity we follow [72, 73] and define the ν -tuple of tail nodes $x_{\mathcal{T}(c)}$ and the space $P_{\mathcal{T}(c)}$ by

$$x_{\mathcal{T}(c)} = x_{\mathcal{T}(I(c))} = (x_{\mathcal{T}(i_1)}, \dots, x_{\mathcal{T}(i_\nu)}) \in P_{\mathcal{T}(I(c))} = P_{\mathcal{T}(c)}.$$

Definition 4.5. Let \mathcal{G} be a network. A map $f : P \rightarrow P$ is \mathcal{G} -admissible if:

- (a) *Node compatibility:* The node spaces satisfy $P_c = P_d$ whenever $c \sim_S d$.
- (b) *Domain condition:* For every node c , there exists a function $\hat{f}_c : P_c \times P_{T(c)} \rightarrow P_c$ such that

$$f_c(x) \equiv \hat{f}_c(x_c, x_{T(c)}).$$

In particular, the domain of \hat{f}_c (which, in effect, is the relevant domain of f_c) is $P_c \times P_{T(c)}$.

- (c) *Pullback condition:* If nodes c, d are input equivalent, then for every $\beta \in B(c, d)$:

$$\hat{f}_d(x_d, x_{T(d)}) \equiv \hat{f}_c(x_d, \beta^* x_{T(d)}) \quad (4.3)$$

where the *pullback map* is defined by

$$\beta^* x_{T(d)} = (x_{\mathcal{T}(\beta(i_1))}, \dots, x_{\mathcal{T}(\beta(i_v))}) \in P_{T(c)}. \quad (4.4)$$

In particular, we can apply (4.3) when $c = d$. This shows that

$$\hat{f}_c(x_c, x_{T(c)}) \text{ is } B(c, c)\text{-invariant}$$

where the *vertex group* $B(c, c)$ acts trivially on the first coordinate x_c and permutes the coordinates of $x_{T(c)}$ according to the pullback maps (4.4). That is, the action of β is

$$(x_c, x_{T(c)}) \mapsto (x_c, \beta^* x_{T(c)}). \quad (4.5)$$

Triviality of this action on the first coordinate, and the distinguished nature of that coordinate, are crucial to this paper.

Remarks 4.6. (a) The group $B(c, c)$ is finite and is a direct product of symmetric groups, one for each arrow-type.

- (b) From now on it is convenient to omit the hat on \hat{f}_c and consider f_c as a map $f_c : P_c \times P_{T(c)} \rightarrow P_c$.

4.4. Alternative characterisation of admissibility

The definition of pullback maps provides a ‘coordinate-free’ definition of admissible ODEs. We now deduce a standard characterisation of admissible maps, based on a specific choice of coordinates in the domains of component maps f_c , which is more convenient for the purposes of this paper.

Choose an ordering on arrow types, so that arrows of a given type occur as a block; then order arrows arbitrarily within each block. Call this a *standard ordering* of arrows. It is easy to prove that the groupoid \mathcal{B} is generated by all vertex symmetry

groups $B(c, c)$ together with a single *transitional map* $\beta_{cd} : I(c) \rightarrow I(d)$ for each $c \neq d$ with $c \sim_I d$. Moreover, if input variables for input equivalent nodes are listed in standard order, the natural transitional map is the identity. This is why the usual way to represent symmetries of components using an overline on the relevant input variables is possible [43, 71]. The overlines correspond to the blocks of arrows with a given arrow type, and substitution of corresponding variables gives the identity transition map.

The group $B(c, c)$ acts on the input set $I(c)$ by permuting arrows and preserving arrow-type, so it preserves blocks of arrows in standard order. We can now characterise admissible maps in terms of $B(c, c)$ -invariance, avoiding explicit reference to pullback maps.

Proposition 4.7. *A map $f : P \rightarrow P$ is admissible if and only if, in standard order:*

- (1) f_c is invariant under $B(c, c)$ for each c in a set of representatives of the input equivalence classes.
- (2) $c \sim_I d \Rightarrow f_c = f_d$.

Proof. This follows from [71, Lemma 4.5 and Proposition 4.6], with the extra observation that when the inputs are in standard order the β_{cd} can be taken to be the identity. ■

Proposition 4.7 implies that admissible maps can be constructed as follows. Choose a set of representatives \mathcal{S} for input equivalence. For each $s \in \mathcal{S}$ let p_s be any smooth $B(s, s)$ -invariant map $p_s : P_s \times P_{T(s)} \rightarrow P_s$. The maps p_s can be chosen independently for each s . In standard order, for all $c \in \mathcal{C}$, define

$$p_c = p_s \quad \text{where } s \sim_I c \text{ and } s \in \mathcal{S}.$$

The resulting map p is admissible because

$$p_d(x_d, x_{T(d)}) = p_c(x_d, \beta^* x_{T(d)}) = p_c(x_d, x_{T(d)})$$

when $\beta^* = \text{id}$.

4.5. Balanced colourings

A *colouring* of a network is a partition of the nodes into disjoint subsets, the *parts*

$$\mathcal{C} = \mathcal{C}_1 \dot{\cup} \mathcal{C}_2 \dot{\cup} \cdots \dot{\cup} \mathcal{C}_m.$$

The *colour* $[c]$ of node c is the unique i such that $c \in \mathcal{C}_i$. A colouring can also be viewed as an equivalence relation ‘in same part’ or ‘same colour’. We pass without

comment between these three interpretations, but mainly refer to colourings. We use the same symbol \bowtie for all three, and often specify it as a partition.

Associated with any colouring \bowtie is the *polydiagonal* (or *synchrony space*)

$$\Delta_{\bowtie} = \{x \in P : x_c = x_d \Leftrightarrow c \bowtie d\}.$$

The name indicates that this notion is a generalisation of the usual diagonal subspace $\{(y, y, y, \dots, y)\}$. Another common term is *synchrony space*.

Definition 4.8. A colouring is *balanced* if whenever c and d have the same colour, there is a colour-preserving input isomorphism $\beta : I(c) \rightarrow I(d)$. That is, $\mathcal{T}(e)$ and $\beta(\mathcal{T}(e))$ have the same colour for all arrows $e \in I(d)$. Symbolically,

$$c \bowtie d \implies \mathcal{T}(e) \bowtie \beta(\mathcal{T}(e)) \quad \forall e \in I(c).$$

This concept is central to network dynamics because Δ_{\bowtie} is *flow-invariant*, that is, invariant under *any* admissible map, if and only if \bowtie is balanced; see [43, Theorem 4.3] or [71, Theorem 6.5]. The space Δ_{\bowtie} is defined even when \bowtie is unbalanced, but is no longer flow-invariant.

Associated with any balanced colouring of \mathcal{G} is a *quotient network* \mathcal{G}/\bowtie whose admissible maps are precisely the restrictions to Δ_{\bowtie} of the admissible maps of the original network when Δ_{\bowtie} is canonically identified with $\prod_{s \in \mathcal{S}} P_s$ for a set of representatives \mathcal{S} of \bowtie ; see [43, Section 5]. The validity of this theorem requires the multiarrow formalism introduced in that paper; the differences that occur in the single-arrow formalism are described in [21].

4.6. Synchrony and phase relations: sufficient conditions

The calculations that motivate the Rigid Synchrony and Rigid Phase Conjectures combine the pullback condition (4.3) for admissibility with equations (1.2) and (1.3), as follows. From (1.2) we obtain $\dot{x}_c(t) \equiv \dot{x}_d(t)$, so

$$f_c(x_c, x_{T(c)}) \equiv f_d(x_d, x_{T(d)}) \equiv f_c(x_d, \beta^* x_{T(d)}).$$

Therefore, a sufficient condition for synchrony of nodes c, d is

$$x_{T(c)}(t) \equiv \beta^* x_{T(d)}(t) \quad \forall t \in \mathbb{R} \text{ whenever } x_c(t) \equiv x_d(t). \quad (4.6)$$

The Rigid Synchrony Conjecture states that with the additional hypothesis of rigidity, condition (4.6) is also necessary. This condition is equivalent to the relation of synchrony being balanced. Similar reasoning for (1.3) leads to the sufficient condition

$$x_{T(c)}(t) \equiv \beta^* x_{T(d)}(t + \theta) \quad \forall t \in \mathbb{R} \text{ whenever } x_c(t) \equiv x_d(t + \theta). \quad (4.7)$$

The Rigid Phase Conjecture states that condition (4.7) is also necessary for (1.3) to hold, with the additional hypothesis of rigidity.

5. Quasi-quotients

Each induced ODE obtained in Section 3 can be characterised as an admissible ODE for a smaller network whose nodes correspond to the colours. If a colouring \bowtie is balanced, the smaller network is the usual quotient network. If \bowtie is not balanced, we can still construct a smaller network as a ‘quasi-quotient’ $\mathcal{G}^{\mathcal{R}}$ for any set of representatives \mathcal{R} . Uniqueness now fails: different choices of \mathcal{R} can give different quasi-quotients. Dynamics with synchrony pattern \bowtie projects to give dynamics on $\mathcal{G}^{\mathcal{R}}$, but the converse fails because the discarded ‘constraint equations’ need not be satisfied. For these reasons, quasi-quotients seem not to have been considered previously. However, they arise naturally from the methods of this paper, and they have one very useful property: all admissible maps for $\mathcal{G}^{\mathcal{R}}$ lift to (that is, are induced from) admissible maps for \mathcal{G} . This property allows us to construct $\mathcal{G}^{\mathcal{R}}$ -admissible perturbations using only the topology of $\mathcal{G}^{\mathcal{R}}$, and then lifting them to admissible perturbations on \mathcal{G} . We therefore develop the basic properties of quasi-quotients required in later proofs.

5.1. Definition of quasi-quotient

Let \mathcal{G} be a network with nodes \mathcal{C} , let \bowtie be a colouring of \mathcal{G} (which need not be balanced), and choose a set \mathcal{R} of representatives for \bowtie .

Definition 5.1. If $c \in \mathcal{C}$, write $[c]$ for the unique element of \mathcal{R} such that $r \bowtie c$. In particular, $[r] = r$ if and only if $r \in \mathcal{R}$.

Definition 5.2. The *quasi-quotient* network $\mathcal{G}^{\mathcal{R}}$ has nodes $r \in \mathcal{R}$, whose node type is the same as the node-type of r in \mathcal{G} .

The arrows e for $\mathcal{A}^{\mathcal{R}}$ of $\mathcal{G}^{\mathcal{R}}$ are identified (via a bijection $e \mapsto e'$) with the arrows

$$e' \in \bigcup_{r \in \mathcal{R}} I(r).$$

Under this identification, head and tail nodes in $\mathcal{G}^{\mathcal{R}}$ are defined by

$$\mathcal{H}^{\mathcal{R}}(e) = \mathcal{H}(e'), \quad \mathcal{T}^{\mathcal{R}}(e) = [\mathcal{T}(e')].$$

Arrows e_1, e_2 in $\mathcal{A}^{\mathcal{R}}$ have the same arrow type if and only if e'_1, e'_2 have the same arrow type in \mathcal{G} .

Informally, we construct $\mathcal{G}^{\mathcal{R}}$ by taking all nodes in \mathcal{R} , together with their input arrows. Then the tail node of each input arrow is found by replacing its tail node in \mathcal{G} by the unique node in \mathcal{R} that has the same colour.

For example, let \mathcal{G} be the 3-node network of Figure 1. The corresponding quasi-quotients for all nontrivial (\bowtie, \mathcal{R}) are shown in Figure 2.




Figure 2. Quasi-quotients of the 3-node ring.

5.2. Admissible maps for quasi-quotients

Given node spaces P_c for $c \in \mathcal{C}$, define the node space for $r \in \mathcal{R}$ to be P_r . The state space of $\mathcal{G}^{\mathcal{R}}$ is then

$$P^{\mathcal{R}} = \prod_{r \in \mathcal{R}} P_r.$$

For any tuple (c_1, \dots, c_k) of nodes $c_i \in \mathcal{C}$, define the corresponding tuple of nodes of \mathcal{R} by

$$[(c_1, \dots, c_k)] = ([c_1], \dots, [c_k]). \tag{5.1}$$

We now show that any \mathcal{G} -admissible map on P defines a unique $\mathcal{G}^{\mathcal{R}}$ -admissible map on $P^{\mathcal{R}}$. Conversely, every $\mathcal{G}^{\mathcal{R}}$ -admissible map on $P^{\mathcal{R}}$ lifts to a \mathcal{G} -admissible map on P , but this need not be unique. The key observation is:

- Proposition 5.3.** (a) Any input isomorphism from $I(r)$ to $I(s)$ in $\mathcal{G}^{\mathcal{R}}$ identifies naturally with an input isomorphism from $I(r)$ to $I(s)$ in \mathcal{G} .
 (b) Nodes $r, s \in \mathcal{R}$ have the same input type in $\mathcal{G}^{\mathcal{R}}$ if and only if they have the same input type in \mathcal{G} .

Proof. (a) This is immediate from the definition of arrows and arrow-types in Definition 5.2. In detail: we have identified arrows in $\mathcal{A}^{\mathcal{R}}$ with arrows in \mathcal{A} via the bijection $e \mapsto e'$. The definition of head nodes of these arrows implies that the input arrows of r in $\mathcal{G}^{\mathcal{R}}$ correspond bijectively to the input arrows of r in \mathcal{G} , preserving arrow types. The same goes for input arrows of s . Therefore, any input isomorphism $I(r) \rightarrow I(s)$ in $\mathcal{G}^{\mathcal{R}}$ corresponds to an input isomorphism $I(r) \rightarrow I(s)$ in \mathcal{G} , and conversely.

(b) This is now immediate. ■

5.3. Properties of quasi-quotients

Most features of the usual quotient network construction do not carry over to quasi-quotients, but a few useful ones do. We exploit these features in the proof of the main theorem, Theorem 9.2.

Theorem 5.4. (a) Restriction: Every \mathcal{G} -admissible map $f : P \rightarrow P$ determines a $\mathcal{G}^{\mathcal{R}}$ -admissible map $f^{\mathcal{R}} : P^{\mathcal{R}} \rightarrow P^{\mathcal{R}}$ defined by

$$f_r^{\mathcal{R}}(x_r, x_{T(r)}) = f_r(x_r, x_{T(r)}). \quad (5.2)$$

(b) Lifting: For every $\mathcal{G}^{\mathcal{R}}$ -admissible map $g : P^{\mathcal{R}} \rightarrow P^{\mathcal{R}}$, there exists a \mathcal{G} -admissible map $f : P \rightarrow P$ such that $f^{\mathcal{R}} = g$.

(c) Smallness of Lift: If $\|g\|_1 < \varepsilon$, we can define f so that $\|f\|_1 < \varepsilon$.

Proof. (a) By Definition 5.2, the domain condition for $f^{\mathcal{R}}$ is that there exists $\hat{f}^{\mathcal{R}} : P_r \times P_{[T(r)]}$ such that $f^{\mathcal{R}}(x) = \hat{f}^{\mathcal{R}}(x_r, x_{[T(r)]})$. This is consistent with (5.2).

The identifications in Definition 5.2 imply that if $r, s \in \mathcal{R}$ and $\beta \in B(r, s)$ (for \mathcal{G}) then β identifies with an input isomorphism in $\mathcal{G}^{\mathcal{R}}$, which we also denote by β . The pullback condition for f is

$$f_d(x_d, x_{T(d)}) = f_c(x_d, \beta^* x_{T(d)})$$

for all input isomorphisms $\beta : I(c) \rightarrow I(d)$. Therefore

$$f_s^{\mathcal{R}}(x_s, x_{[T(s)]}) = f_s(x_s, x_{T(s)}) = f_r(x_r, \beta^* x_{T(r)}) = f_r^{\mathcal{R}}(x_s, \beta^* x_{[T(s)]})$$

which is the pullback condition for $f^{\mathcal{R}}$.

(b) Here it is convenient to use the alternative characterisation of admissible maps in Proposition 4.7. Let $g : P^{\mathcal{R}} \rightarrow P^{\mathcal{R}}$ be a $\mathcal{G}^{\mathcal{R}}$ -admissible map. We must construct a \mathcal{G} -admissible map $f : P \rightarrow P$ such that $f^{\mathcal{R}} = g$. The nodes \mathcal{C} split into two disjoint subsets $\mathcal{C}_1, \mathcal{C}_2$ defined by

$$\mathcal{C}_1 = \{c \in \mathcal{C} : \exists r \in \mathcal{R}, c \sim_I r\}, \quad \mathcal{C}_2 = \mathcal{C} \setminus \mathcal{C}_1.$$

With arrows in standard order, define f by

$$f_c(x_c, x_{T(c)}) = \begin{cases} f_{[c]}(x_{[c]}, x_{[T(c)]}) & \text{if } c \in \mathcal{C}_1, \\ 0 & \text{if } c \in \mathcal{C}_2. \end{cases} \quad (5.3)$$

Clearly f satisfies the required domain conditions to be \mathcal{G} -admissible. To verify the pullback conditions, we must show that f_c is $B(c, c)$ -invariant on $P_{T(c)}$. The rest follows because transition maps are now the identity. Since $c \in \mathcal{C}_1$, invariance under $B(c, c)$ follows from Proposition 5.3 and $B([c], [c])$ -invariance of $f_{[c]}$ on $P_{[T(c)]}$.

(c) In (5.3), we have $f_{[c]}(x_{[c]}, x_{[T(c)]}) = 0$ when $c \in \mathcal{C}_2$. With the C^1 norm defined as in Definition 6.2, $\|g\|_1 < \varepsilon$ implies that $\|f\|_1 < \varepsilon$. ■

Remark 5.5. For (b), the choice of f_c for $c \in \mathcal{C}_2$ can be replaced by the corresponding components g_c of any \mathcal{G} -admissible map g , by Proposition 4.7. The choice on \mathcal{C}_1 is unique.

In the proof of the key Lemma 9.4 below, we make the lift have small compact support, which implies that it is C^1 -bounded, but we do not want it to vanish identically for $c \in \mathcal{C}_2$. This can be done by making g have small compact support but not requiring $g \equiv 0$.

5.4. Induced ODE

Associated with any quasi-quotient is a version of the usual restricted ODE for a quotient network. Because Δ_{\bowtie} need not to be flow-invariant, the domain of the ODE is restricted to $P^{\mathcal{R}}$ and its codomain is projected onto $P^{\mathcal{R}}$.

Definition 5.6. Consider an admissible ODE (6.3) and write it in components as

$$\dot{x}_c = f_c(x_c, x_{T(c)}), \quad x \in P, 1 \leq c \leq n. \quad (5.4)$$

The *induced ODE* for the pair (\bowtie, \mathcal{R}) is the ODE

$$\dot{x}_r = f_r(x_r, x_{[T(r)]}), \quad x \in P, r \in \mathcal{R} \quad (5.5)$$

where $[T(r)]$ is defined by (5.1). The *constraint equations* are the corresponding equations on the other components, implied by the synchrony relations $x_c(t) \equiv x_{[c]}(t)$,

$$\dot{x}_{[c]} = f_c(x_{[c]}, x_{[T(c)]}) \quad x \in P, c \in \mathcal{C} \setminus \mathcal{R}. \quad (5.6)$$

Example 5.7. Again, consider the 3-node network of Figure 1, with admissible ODEs (3.1). Consider Case (B) of Section 3.2 with colouring $\bowtie = \{\{1, 2\}, \{3\}\}$. There are two choices of \mathcal{R} .

If $\mathcal{R} = \{1, 3\}$, the induced ODE is

$$\dot{x}_1 = f(x_1, x_3), \quad \dot{x}_3 = g(x_3, x_1),$$

with constraint

$$\dot{x}_1 = f(x_1, x_1).$$

If $\mathcal{R} = \{2, 3\}$, the induced ODE is

$$\dot{x}_2 = f(x_2, x_2), \quad \dot{x}_3 = g(x_3, x_2),$$

with constraint

$$\dot{x}_2 = f(x_2, x_3).$$

The induced ODEs are the admissible ODEs for the corresponding quasi-quotients in Figure 2, in accordance with Theorem 5.4. The same remark holds for all other choices of (\bowtie, \mathcal{R}) in the figure.

The induced ODE depends on the choice of \bowtie and \mathcal{R} , and in general solutions need not lift back to (5.4). More precisely:

- Theorem 5.8.** (a) Any orbit $\mathbb{X} = \{x(t)\}$ of (5.4) defines an orbit $\mathbb{X}^{\mathcal{R}} = \{x^{\mathcal{R}}(t)\}$ of (5.5).
 (b) The orbit $x(t) \in \Delta_{\bowtie}$ if and only if it satisfies the additional constraints (5.6).
 (c) Solutions of (5.5) lift uniquely to a solution of (5.4), provided that such a solution satisfies the constraints. ■

The proof is obvious. We emphasise that (c) requires the constraint equations to be satisfied as well as the induced ODE. This can, for example, be implied by rigidity. Our aim in this paper is to obtain a contradiction to this property in suitable circumstances.

When \bowtie is not balanced, then for any choice of \mathcal{R} the constraints include at least one component that differs formally from the corresponding component of the induced equation. We exploit this formal difference to obtain a contradiction to rigidity.

When \bowtie is balanced, the constraints just repeat the corresponding components of the induced ODE, and this is the same as the usual restricted ODE. In this case no contradiction occurs.

5.5. Perturbations

Suppose that an ODE $\dot{y} = g(y)$ has a non-hyperbolic periodic orbit \mathbb{Y} . If g is perturbed to a nearby map \tilde{g} , there may be no periodic orbits near \mathbb{Y} , or more than one. Thus, we cannot talk of ‘the’ perturbed periodic orbit $\tilde{\mathbb{Y}}$.

If f is \mathcal{G} -admissible, with a hyperbolic periodic orbit \mathbb{X} , and we consider the induced ODE $\dot{y} = g(y)$ for $\mathcal{G}^{\mathcal{R}}$, the same remark applies to $\mathbb{Y} = \mathbb{X}^{\mathcal{R}}$, because \mathbb{Y} need not be hyperbolic. However:

Lemma 5.9. *If \mathbb{X} is hyperbolic with rigid synchrony pattern \bowtie , and \mathcal{R} is a set of representatives, then for any small $\mathcal{G}^{\mathcal{R}}$ -admissible perturbation $\tilde{f}^{\mathcal{R}}$ of $f^{\mathcal{R}}$ the induced orbit $\mathbb{X}^{\mathcal{R}}$ has a uniquely defined canonical perturbed periodic orbit $\tilde{\mathbb{X}}^{\mathcal{R}}$.*

Proof. Let $p^{\mathcal{R}} = \tilde{f}^{\mathcal{R}} - f^{\mathcal{R}}$ where $\|p^{\mathcal{R}}\|_1$ is small. By Theorem 5.8 (c) we can lift $p^{\mathcal{R}}$ to a \mathcal{G} -admissible map p whose norm is equally small. If $\tilde{f} = f + p$ then $\tilde{f}^{\mathcal{R}} = f^{\mathcal{R}} + p^{\mathcal{R}} = \tilde{f}^{\mathcal{R}}$. Let $\tilde{\mathbb{X}}$ be the unique perturbed periodic orbit near \mathbb{X} for \mathcal{G} . By rigidity, $\tilde{\mathbb{X}}^{\mathcal{R}}$ is a periodic orbit of $f^{\mathcal{R}}$, and is near $\mathbb{X}^{\mathcal{R}}$. This procedure defines $\tilde{\mathbb{X}}^{\mathcal{R}}$ uniquely. ■

Corollary 5.10. *If all periodic orbits of $\tilde{f}^{\mathcal{R}}$ near $\mathbb{X}^{\mathcal{R}}$ are hyperbolic, then $\tilde{\mathbb{X}}^{\mathcal{R}}$ is hyperbolic. ■*

6. Properties related to hyperbolicity

In this section we recall background results and concepts that are needed for the proofs of the main theorems, and provide rigorous definitions for concepts that until now have been treated informally for illustrative purposes.

6.1. C^1 norm

We begin by clarifying the sense in which a perturbation is to be considered ‘small’, a technical point that we have hitherto slid over. In order for hyperbolicity to imply the existence of a locally unique perturbed periodic orbit, we use the C^1 topology. It is also convenient to define this in a way that is tailored to the network setting, with distinguished node spaces, as follows.

Choose a fixed state space $P = P_1 \times \cdots \times P_n$ where $P_c = \mathbb{R}^{k_c}$ for finite k_c and $1 \leq c \leq n$. Let $C^1(P, P)$ be the Banach space of admissible C^1 -bounded C^1 maps $f : P \rightarrow P$ with the C^1 norm

$$\|f\|_1 = \sup_{x \in P} (\|f(x)\|, \|Df(x)\|) \quad (6.1)$$

where Df is the derivative. In the context of this paper it is convenient to define the norm on state space P by

$$\|(x_c)_{c \in \mathcal{C}}\| = \max\{\|x_c\|_E : c \in \mathcal{C}\} \quad (6.2)$$

where $\|v\|_E$ is the Euclidean norm.

By Abraham et al. [2, Proposition 2.1.10 (ii)], all norms on a finite-dimensional real vector space are equivalent, so this definition is equivalent to the usual C^1 norm.

Remark 6.1. In order for locally unique perturbed periodic orbits to exist, we must work with perturbations p for which the C^1 norm is bounded. This condition can always be arranged using a bump function, Abraham et al. [2, Lemma 4.2.13], to modify any admissible map f so that it vanishes outside some large compact set K that contains \mathbb{X} , while leaving \mathbb{X} unchanged and f unchanged in a neighbourhood of \mathbb{X} . However, we require C^1 -boundedness only for perturbations p of f , not for f itself, and p will always be defined in a manner that ensures it is bounded, so this modification of f is not required in this paper.

6.2. Hyperbolic periodic orbits

Hyperbolicity (for equilibria, periodic orbits, or more generally for invariant submanifolds) is defined in many sources, for example Abraham and Mardsen [1], Arrowsmith and Place [8], Hirsch and Smale [48], Guckenheimer and Holmes [44], and Katok and Hasselblatt [53]. An equilibrium x^0 of (1.1) is hyperbolic if and only if the derivative (Jacobian) $Df|_{x^0}$ has no eigenvalues on the imaginary axis (zero included). A periodic orbit is hyperbolic if its linearised Poincaré return map, for some (hence any) Poincaré section, has no eigenvalues on the unit circle. That is, exactly one Floquet multiplier (equal to 1) lies on the unit circle; equivalently, exactly one Floquet exponent (equal to 0) lies on the imaginary axis [45].

The following result is standard, and can be proved by applying the implicit function theorem to a Poincaré map. A more general proof for invariant submanifolds can be found in Hirsch et al. [47, Theorem 4.1 (f)]. For the purposes of this paper it is convenient to state it for 1-parameter families of perturbations.

Lemma 6.2. *Let $\mathbb{X} = \{x(t)\}$ be a hyperbolic periodic orbit of a smooth ODE $\dot{x} = f(x)$ on \mathbb{R}^n . Let $f + \varepsilon p$ be any 1-parameter family of perturbations, with $\|p\|_1$ bounded. Then for $\varepsilon \ll 1$ there exists, near $x(t)$, a locally unique periodic orbit $\mathbb{X}^\varepsilon = \{x^\varepsilon(t)\}$ of the perturbed ODE $\dot{x} = f(x) + \varepsilon p(x)$. ■*

The perturbed periodic orbit is locally unique in the sense that, for any 1-parameter family of sufficiently small perturbations, there is a locally unique path of periodic orbits that includes the unperturbed one. Here and elsewhere, $\varepsilon \ll 1$ means $\varepsilon < \varepsilon^*$ for some $\varepsilon^* > 0$ with specified properties.

6.3. Open properties

Again, choose a fixed state space $P = P_1 \times \cdots \times P_n$, and let $C^1(P, P)$ be the Banach space of admissible C^1 -bounded C^1 maps $f : P \rightarrow P$ with the C^1 norm (6.1). Let $C^\infty(P, P)$ be the space of admissible C^∞ maps $f : P \rightarrow P$, which is not a Banach space. We use $\|\cdot\|_1$ to put a topology on the space

$$\mathcal{F}(P) = C^1(P, P) \cap C^\infty(P, P).$$

This topology is applied only to ‘small perturbations’ p of C^∞ maps f , because smooth maps need not be C^1 -bounded. We use the same notation when all maps are required to be \mathcal{G} -admissible for a network \mathcal{G} , indicating this condition by context.

If we focus only on a suitable compact subset K of a state space, we can replace any admissible map f by a bounded one that agrees with f on K , see Remark 6.1. However, the space of admissible C^1 -bounded C^∞ maps $f : P \rightarrow P$ is still not a Banach space. Nevertheless, we can state:

Definition 6.3. A property \mathcal{O} of maps $f \in C^\infty(P, P)$ is *open* if whenever f has property \mathcal{O} , there exists $\varepsilon > 0$ such that, for all $q \in \mathcal{F}(P)$ with $\|q\|_1 < \varepsilon$, the map $f + q$ has property \mathcal{O} . Equivalently, the set of all f with property \mathcal{O} is open in the C^1 norm, and is preserved by all C^1 -small perturbations of f .

We use the same terminology for network dynamics, requiring the maps involved to be admissible.

In the sequel we use a 1-parameter family of perturbations $q = \varepsilon p$ for a fixed $p \in \mathcal{F}(P)$, that is, we consider the family $f + \varepsilon p$ for $\varepsilon \ll 1$. We use the weaker condition that \mathcal{O} holds for all p and all $\varepsilon \ll 1$. That is, we do not require the upper bound on ε to be uniform in p .

6.4. Rigidity

We extend Definition 6.3 to properties of a hyperbolic periodic orbit of an admissible ODE, replacing ‘open’ by ‘rigid’ to preserve traditional terminology. Restating (1.1) for convenience, let the admissible ODE be

$$\dot{x} = f(x), \quad x \in P \tag{6.3}$$

and let $\mathbb{X} = \{x(t) : t \in \mathbb{R}\}$ be a hyperbolic periodic orbit with period T . By Lemma 6.2, if $p \in \mathcal{F}(P)$ is any admissible map and ε is sufficiently small then the perturbed ODE

$$\dot{x}^\varepsilon = f(x^\varepsilon) + \varepsilon p(x^\varepsilon), \quad x^\varepsilon \in P \tag{6.4}$$

has a *unique* perturbed periodic orbit $\mathbb{X}^\varepsilon = \{x^\varepsilon(t) : t \in \mathbb{R}\}$ that is near \mathbb{X} in the Hausdorff metric for the C^1 topology, with period T^ε near T . This equation becomes (6.3) when $\varepsilon = 0$. In particular, $\mathbb{X}^0 = \mathbb{X}$ and $T^0 = T$.

As is customary, we use the same symbol x^ε to denote an arbitrary variable in P and a specific solution (orbit, trajectory) of the ODE. The alternative is to introduce cumbersome notation to distinguish the two meanings.

Using a fixed Poincaré section Σ to \mathbb{X} to define the initial condition by $x^\varepsilon(t_0) \in \Sigma$, and considering the Poincaré map and hyperbolicity, we can assume that for $\varepsilon \geq 0$ and any $t \in \mathbb{R}$ the point $x^\varepsilon(t)$ varies smoothly with ε , and so does T^ε .

Definition 6.4. A property \mathcal{O} of \mathbb{X} relative to T is *rigid* if \mathbb{X}^ε has property \mathcal{O} relative to T^ε for $\varepsilon \ll 1$.

Hyperbolicity of a given \mathbb{X} for f is an open property of f , and also a rigid property of \mathbb{X} . Rigid synchrony and phase patterns of \mathbb{X} are obviously rigid properties of \mathbb{X} . So are local rigid synchrony and phase patterns, defined in Section 1.

If \mathbb{X} has a balanced synchrony pattern (or local synchrony pattern) \bowtie , the corresponding periodic orbit on the quotient \mathcal{G}/\bowtie is also hyperbolic, because Δ_{\bowtie} is

flow-invariant. However, this result no longer applies if \bowtie is not balanced. This is why we require \mathbb{X} to be stably isolated or strongly hyperbolic, Sections 6.9 and 6.7.

6.5. Failure of hyperbolicity

There is one class of networks for which an analogue of the Kupka–Smale Theorem cannot hold, for feedforward reasons. To discuss it, we use standard ideas from Floquet theory [45].

Every network decomposes into *transitive components* [28]¹, and the set \mathcal{T} of transitive components has a natural partial ordering induced by directed paths. Dynamically, this ordering gives admissible ODEs a feedforward structure. If \mathcal{T} has more than one maximal element, and the periodic state \mathbb{X} oscillates on at least two maximal components, this state is not hyperbolic. This follows because the Floquet operator has block-triangular form induced by the partial ordering.

In more detail, Josić and Török [52, Remark 1] observe that for the network of Figure 3, periodic orbits cannot be hyperbolic unless node 1 or node 2 is steady. To prove this, observe that admissible ODEs for this network have the form

$$\begin{aligned}\dot{x}_1 &= f(x_1), \\ \dot{x}_2 &= g(x_2), \\ \dot{x}_3 &= h(x_1, x_2, x_3).\end{aligned}\tag{6.5}$$

We must show that if such an ODE has a hyperbolic periodic orbit

$$\mathbb{X} = \{(x_1(t), x_2(t), x_3(t)) : t \in \mathbb{R}\}$$

then either $x_1(t)$ or $x_2(t)$ is steady. This follows because, setting $F = (f, g, h)$, the Floquet multipliers are the eigenvalues of $D_{x(t)}F$, which is lower triangular. Therefore, the evolution operator is also lower triangular, and has two eigenvalues equal to 1. These eigenvalues correspond to the two diagonal blocks describing the evolution in the spaces of the variables x_1 and x_2 . (A periodic orbit always has a Floquet multiplier 1 for an eigenvector along the orbit; see [45, Chapter 1 Note 5].)




Figure 3. Connected network with two distinct maximal transitive components.

¹See also en.wikipedia.org/wiki/Strongly_connected_component.

As Josić and Török observe, similar remarks apply if we replace nodes 1 and 2 by two disjoint transitive components that are maximal in the partial ordering: that is, they force (feed forward into) the rest of the network, which replaces node 3. In (6.5) let x_1, x_2, x_3 be coordinates on, respectively, the two maximal components and the rest of the network. The same argument then applies.

6.6. Implications for quasi-quotients

In the present context, such networks cannot occur as the overall network \mathcal{G} , since we assume \mathbb{X} hyperbolic. However, we also require hyperbolicity (or a similar property) for certain quasi-quotients $\mathcal{G}^{\mathcal{R}}$. A ‘bad’ choice of representatives can create more than one maximal transitive component in $\mathcal{G}^{\mathcal{R}}$. Figure 4 (left) shows the simplest (connected) example of this kind. The ‘good’ choice $\mathcal{R} = \{1, 3\}$ yields an induced network with feedforward structure and a single maximal component, Figure 4 (middle). In contrast, the ‘bad’ choice $\mathcal{R} = \{1, 4\}$ yields an induced network with two disconnected maximal components, Figure 4 (right). If nodes 1 and 3 force the same component of some larger network, the corresponding quasi-quotient has no hyperbolic periodic orbits.

We do not know whether there always exists such an \mathcal{R} when \mathcal{G} has only one maximal component.



Figure 4. *Left:* 2-coloured regular network. *Right:* Two distinct induced networks, one ‘good’, one ‘bad’.

6.7. Strong hyperbolicity

This paper relies on the following concept:

Definition 6.5. A periodic orbit \mathbb{X} is *strongly hyperbolic* for a colouring \bowtie if \mathbb{X} is hyperbolic and there exists a set of representatives \mathcal{R} such that, if necessary after an arbitrarily small perturbation, the induced orbit $\mathbb{X}^{\mathcal{R}}$ is a hyperbolic periodic orbit of the induced ODE.

The orbit \mathbb{X} is *strongly hyperbolic* if it is strongly hyperbolic for every colouring \bowtie .

Suppose that \mathbb{X} is hyperbolic, and perturb f to a nearby admissible map \tilde{f} . There is a locally unique perturbed periodic orbit $\tilde{\mathbb{X}}$ near \mathbb{X} . Lemma 5.9 implies that there

is a unique canonical choice for ‘the’ perturbed induced periodic orbit $\tilde{\mathbb{X}}^{\mathcal{R}}$, which is near $\mathbb{X}^{\mathcal{R}}$. Thus when \mathbb{X} has a rigid synchrony pattern we can keep track of each $\mathbb{X}^{\mathcal{R}}$ in a meaningful manner when f is perturbed, even when $\mathbb{X}^{\mathcal{R}}$ is not known to be hyperbolic on $P^{\mathcal{R}}$.

The condition of strong hyperbolicity lets us ‘pre-perturb’ the admissible map f to \tilde{f} , thereby ensuring that without loss of generality the uniquely defined periodic orbit $\mathbb{X}^{\mathcal{R}}$ is hyperbolic on $P^{\mathcal{R}}$. This step is crucial for the proof of the Local Rigid Synchrony Property.

6.8. Local Kupka–Smale Theorem

In Section 14.4 we infer strong hyperbolicity from a local version of the Kupka–Smale Theorem:

Lemma 6.6. *Let f be any smooth vector field on R^k , with a periodic orbit \mathbb{X} of period T . Then:*

- (a) *For all sufficiently small $\delta_1 > 0$ there exists δ_2 with $0 < \delta_2 < \delta_1$, such that for all t with $0 \leq t \leq 2T$ we have*

$$\psi^t(N_{\delta_2}(\mathbb{X})) \subseteq N_{\delta_1}(\mathbb{X})$$

where ψ^t is the flow of f and N_δ is the tubular neighbourhood of \mathbb{X} of radius δ .

- (b) *There exists a smooth perturbation \tilde{f} of f and $\delta_3, \delta_4 > 0$ with $\delta_3 \leq \delta_1$ and $\delta_4 \leq \delta_2$ such that*

$$\tilde{\psi}^t(N_{\delta_4}(\mathbb{X})) \subseteq N_{\delta_3}(\mathbb{X})$$

where $\tilde{\psi}^t$ is the flow of \tilde{f} , and every periodic orbit of \tilde{f} inside $N_{\delta_4}(\mathbb{X})$ is hyperbolic.

Proof. This is a restatement of Peixoto [62, Lemma 3]. The bound $2T$ on t is introduced because the perturbed period \tilde{T} may be (slightly) larger than T . ■

Definition 6.7. Let \mathcal{G} be a network and let \mathbb{X} be a hyperbolic periodic orbit of an admissible vector field f . If for every colouring \bowtie there exists a set of representatives \mathcal{R} such that statement (b) holds for all $\mathbb{X}^{\mathcal{R}}$, we say that \mathcal{G} is *locally Kupka–Smale* for \mathbb{X} .

Lemma 6.8. *If \mathcal{G} is locally Kupka–Smale for \mathbb{X} , then for any \bowtie there is a choice of \mathcal{R} such that the canonical induced periodic orbit $\mathbb{X}^{\mathcal{R}}$ can be made hyperbolic by an arbitrarily small admissible perturbation.*

Proof. The perturbed induced periodic orbit $\tilde{\mathbb{X}}^{\mathcal{R}}$ lies inside $N_{\delta_4}(\mathbb{X})$. ■

6.9. Stable isolation

A rigid property that is weaker than strong hyperbolicity but could play the same role in the method used here is:

Definition 6.9. A hyperbolic periodic orbit \mathbb{X} is *stably isolated* for \triangleright if there exists \mathcal{R} such that, if necessary after an arbitrarily small perturbation:

- (a) The induced orbit $\mathbb{X}^{\mathcal{R}}$ is an isolated periodic orbit of the induced ODE.
- (b) This property is preserved by any further sufficiently small perturbation.

The orbit \mathbb{X} is *stably isolated* if it is stably isolated for every colouring \triangleright .

Remarks 6.10. (a) Condition (b) ensures that ‘having a stably isolated periodic orbit’ is an open property of f . Here ‘stable’ refers to structural stability, not stability to perturbations of initial conditions such as asymptotic or linear stability.

- (b) If \triangleright is balanced, any hyperbolic \mathbb{X} is stably isolated, because Δ_{\triangleright} is flow-invariant, so $\mathbb{X}^{\mathcal{R}}$ is hyperbolic.

Proposition 6.11. *If \mathbb{X} is strongly hyperbolic for \triangleright then it is stably isolated for \triangleright . If \mathbb{X} is strongly hyperbolic then it is stably isolated.*

Proof. Hyperbolic periodic orbits are isolated by the local uniqueness assertion in Lemma 6.2. ■

6.10. Kupka–Smale networks

There is a connection between strong hyperbolicity and the Kupka–Smale Theorem, requiring only analogues of (a) and (b) in Theorem 1.2. This motivates:

Definition 6.12. A network \mathcal{G} is a *Kupka–Smale network* if properties (a) and (b) in Theorem 1.2 hold for a generic set of admissible vector fields.

Theorem 6.13. *Let \mathbb{X} be a hyperbolic periodic orbit and let \triangleright be a colouring.*

- (a) *If $\mathcal{G}^{\mathcal{R}}$ is a Kupka–Smale network for some set of representatives \mathcal{R} for \triangleright then \mathbb{X} is stably isolated for \triangleright .*
- (b) *If $\mathcal{G}^{\mathcal{R}}$ is locally Kupka–Smale near $\mathbb{X}^{\mathcal{R}}$ for some \mathcal{R} , then \mathbb{X} is stably isolated for \triangleright .*
- (c) *If \mathbb{X} is stably isolated for \triangleright then it is strongly hyperbolic for \triangleright .*
- (d) *If \mathbb{X} is stably isolated for \triangleright then \mathbb{X} is not a limit of a continuum of periodic orbits, whose members do not equal \mathbb{X} on some open neighbourhood.*
- (e) *The implications (a)–(d) hold when ‘for \triangleright ’ is deleted from all of them.*

Proof. All implications are trivial. ■

In the remainder of the paper we prove that when the local synchrony pattern \bowtie of \mathbb{X} is rigid, the property stated in (d) implies that \bowtie is balanced. That is, \mathcal{G} has the Rigid Synchrony Property for \mathbb{X} . The implications (a)–(d) provide alternative strategies for proving the Local Rigid Synchrony Property in specific cases.

7. Local rigidity

In this section we compare and contrast local and global versions of rigidity, discuss the existence of generic points, and motivate the focus on local versions of the Rigidity Conjectures.

7.1. Synchrony patterns on subsets

Let $\mathbb{X} = \{x(t) : t \in \mathbb{R}\}$ as usual. Let $J \subseteq \mathbb{R}$ be any subset. Define the *local synchrony pattern of \mathbb{X} on J* to be the colouring \bowtie_J defined by

$$c \bowtie_J d \iff x_c(t) \equiv x_d(t) \quad \forall t \in J. \quad (7.1)$$

If $J = \{x\}$ is a singleton, we write \bowtie_x instead of $\bowtie_{\{x\}}$.

In this notation the *global synchrony pattern* of \mathbb{X} is the colouring $\bowtie_{\mathbb{R}}$. However, it is just as valid, and more intuitive to denote it by $\bowtie_{\mathbb{X}}$, which we do from now on. The corresponding polydiagonal is accordingly denoted $\Delta_{\mathbb{X}}$. This pattern is *globally rigid* if, for small enough perturbations, the perturbed periodic orbit has the same global synchrony pattern. The usual form of the Rigid Synchrony Conjecture (see [37, Section 10]) states that any rigid global synchrony pattern $\bowtie_{\mathbb{X}}$ is balanced.

It is easy to establish an alternative characterisation of $\bowtie_{\mathbb{X}}$ in terms of polydiagonals:

Proposition 7.1. *The global synchrony pattern of \mathbb{X} is the unique colouring \bowtie such that*

- (a) $\mathbb{X} \subseteq \Delta_{\bowtie}$,
- (b) \mathbb{X} is not contained in any polydiagonal strictly smaller than Δ_{\bowtie} . ■

Uniqueness follows from the lattice structure.

7.2. Locally rigid synchrony

For technical reasons, discussed in Section 7.6, it is better to work with a local version of rigidity. (The same point was made in [33, 34].) We modify Definition 6.4:

Definition 7.2. A property \mathcal{P} of a hyperbolic periodic orbit \mathbb{X} relative to its period T is *locally rigid near* $t_0 \in \mathbb{R}$ if

- (a) There exists an open interval $J \subseteq \mathbb{R}$ with $t_0 \in J$ such that $x^0(t)$ has property \mathcal{P} for all $t \in J$.
- (b) There exists an open interval $J_1 \subseteq \mathbb{R}$ with $t_0 \in J_1$ such that if ε is small enough, $x^\varepsilon(t)$ has property \mathcal{P} (relative to T^ε) for all $t \in J_1$.

Remark 7.3. This 1-parameter version is sufficient for the purposes of this paper. We do not need the bound on ε to be uniform in p .

7.3. The lattice of colourings

In order to compare $x_c(t)$ with $x_d(t)$, they must both belong to the same space, so we require $P_c = P_d$. Any colouring defined by synchrony properties satisfies this condition, so from now on we do not refer to it explicitly. Formally, for any colouring \bowtie , the condition $c \bowtie d$ requires c and d to be state equivalent, denoted by $c \sim_S d$.

Recall that every colouring \bowtie defines a polydiagonal

$$\Delta_{\bowtie} = \{x \in P : c \bowtie d \Rightarrow x_c = x_d\}$$

and it also defines a partition whose parts correspond to the colours. There is a natural partial order \preceq on colourings:

Definition 7.4. A colouring \bowtie_1 is *finer than* a colouring \bowtie_2 , written $\bowtie_1 \preceq \bowtie_2$, if

$$c \bowtie_1 d \implies c \bowtie_2 d.$$

Contrary to normal English, this includes the possibility that the colourings are the same, up to a permutation of the colours. We also say that \bowtie_2 is *coarser than* \bowtie_1 . To remove the possibility of equality we use the terms *strictly finer* and *strictly coarser*.

The following proposition is obvious:

Proposition 7.5. *The following properties are equivalent:*

- (a) *The colouring \bowtie_1 is finer than \bowtie_2 .*
- (b) *Every part of the partition defined by \bowtie_1 is contained in some part of the partition defined by \bowtie_2 .*
- (c) $\Delta_{\bowtie_1} \supseteq \Delta_{\bowtie_2}$.

Section 4 of [68] proves, in slightly different terminology, that with this partial ordering the set of all colourings is a lattice in the sense of partially ordered sets, Davey and Priestley [20]. Lemma 4.3 of that paper proves that this lattice is dual

to the lattice of polydiagonals under inclusion, property (c) of Proposition 7.5. The balanced colourings form a sublattice, as do the balanced polydiagonals. Since the number of polydiagonals is finite, these are finite lattices.

7.4. Semicontinuity of colourings

The next proposition, which is well known and easy to prove, states that sufficiently small changes to $x \in P$ can make the colouring finer, but not strictly coarser. (Recall that in Definition 7.4 the terms ‘finer’ and ‘coarser’ permit equality.)

Proposition 7.6. *If $x \in P$, with synchrony pattern \bowtie_x , and $y \in P$ with $\|y - x\|$ sufficiently small, then \bowtie_x is finer than \bowtie_y .*

Proof. Suppose that $x_c \neq x_d$, and let $\|x_c - x_d\| = \delta > 0$. Suppose that $\|x - y\| \leq \delta/3$. Then

$$\begin{aligned} \delta = \|x_c - x_d\| &\leq \|x_c - y_c\| + \|y_c - y_d\| + \|y_d - x_d\| \\ &\leq \delta/3 + \|y_c - y_d\| + \delta/3 \end{aligned}$$

so $\|y_c - y_d\| \geq \delta/3 > 0$. ■

Thus, one implication of local rigidity is that the synchrony pattern at $x(t_0) \in \mathbb{X}$ does not become coarser at sufficiently close points $\tilde{x}(t) \in \tilde{\mathbb{X}}$.

Intuitively, Proposition 7.6 states that small perturbations cannot create new equalities of coordinates, but they might break up existing equalities. Equivalently, small perturbations cannot make the synchrony colouring coarser, but they might make it finer. Another equivalent condition is that small perturbations cannot make the polydiagonal Δ_{\bowtie_x} smaller, but they might make it larger. Technically, the number of colours is *lower semicontinuous* with respect to small perturbations, Itô [49].

7.5. Generic points

Given \mathbb{X} , the synchrony pattern $\bowtie_{x(t_0)}$ at a point $x(t_0) \in \mathbb{X}$ is defined by

$$c \bowtie_{x(t_0)} d \iff x_c(t_0) = x_d(t_0).$$

In general this pattern may vary with t_0 . By Proposition 7.6, synchrony patterns are lower semicontinuous. Intuitively, under sufficiently small perturbations colour clusters can break up, but not merge.

In Section 7.1 we defined the local synchrony pattern on an interval J . Depending on J , this pattern might change as t runs over J . Semicontinuity implies that, by shrinking J if necessary, we can assume that

$$\bowtie_{x(t_0)} = \bowtie_{x(t_1)} \quad \forall t_0, t_1 \in J.$$

We assume this from now on.

When this condition holds, any point in $x(J) = \{x(t) : t \in J\}$ is a *generic point*. That is, it has the same synchrony pattern as the local pattern on J . This result is a trivial consequence of Definition 7.8, but we state it explicitly because it is central to the proof of the Rigid Synchrony Conjecture. It is valid for a local synchrony pattern, but need not be for the global synchrony pattern. Although the proof is trivial, this observation is crucial to the construction of suitable perturbations, and is why we consider locally rigid synchrony instead of globally rigid synchrony.

Proposition 7.7. *Let $\triangleright_{x(t_0)}$ be the local rigid synchrony pattern for \mathbb{X} at $x(t_0)$. If \mathcal{R} is a set of representatives for $\triangleright_{x(t_0)}$ then*

$$i \neq j \in \mathcal{R} \implies x_i(t_0) \neq x_j(t_0).$$

Proof. If $x_i(t_0) = x_j(t_0)$ then $i \triangleright_{x(t_0)} j$, so $i = j$ since \mathcal{R} is a set of representatives. ■

Local synchrony is technically more tractable than global synchrony for several reasons. One is that generic points need not exist for global synchrony patterns; see Section 7.6. (Existence of a generic point is claimed in [72], but the proof there is fallacious.) Another is that in the global case, the perturbation technique requires admissible perturbations to vanish near the whole of \mathbb{X} , but not at certain other points. In certain circumstances this can conflict with admissibility. Local perturbations – those with small support – avoid this problem.

Ironically, once we have used local rigid synchrony to prove that (subject to strong hyperbolicity) a network has the Local Rigid Synchrony Property, we can deduce the Global Rigid Synchrony Property and show that a rigid global synchrony pattern is the same as any local one. But working from the beginning with rigid global synchrony runs into the technical difficulties mentioned above. See Section 10.

7.6. Changes in local synchrony

In general dynamical system, different points on a periodic orbit can have different synchrony patterns. Moreover, a general dynamical system can be viewed as a fully inhomogeneous network ODE with all-to-all coupling. Therefore we must deal with the possibility of a periodic orbit akin to Figure 5.

Here Δ_1, Δ_2 are distinct polydiagonals. The periodic orbit $\mathbb{X} = \{x(t)\}$ lies in Δ_1 for some intervals of time t , but in Δ_2 for another interval. The transitions occur smoothly via intervals of time in which $x(t) \in \Delta_1 \cap \Delta_2$. Considerably more complicated changes in the local synchrony pattern are possible; for example it might change infinitely many times. If this or anything similar occurs, then the polydiagonal $\Delta_{\mathbb{X}}$ for

the global synchrony pattern $\bowtie_{\mathbb{X}}$ of \mathbb{X} is different from both Δ_1 and Δ_2 . Indeed,

$$\Delta_1 \subsetneq \Delta_{\mathbb{X}}, \quad \Delta_2 \subsetneq \Delta_{\mathbb{X}}$$

and no point or small interval J on \mathbb{X} has synchrony pattern $\bowtie_J = \bowtie_{\mathbb{X}}$, the global synchrony pattern of \mathbb{X} . We cannot infer local rigidity directly from global rigidity, because perturbed orbits need not remain locally on either Δ_1 or Δ_2 . All we know is that they stay inside the strictly larger polydiagonal $\Delta_{\mathbb{X}}$.

7.7. Local rigid synchrony and generic points

If \mathbb{X} is as in Figure 5, there is no generic point for the global synchrony pattern, because no point on the orbit lies in $\Delta_{\mathbb{X}} \setminus (\Delta_1 \cup \Delta_2)$. However, there exist generic points for the local synchrony patterns.

Such behaviour is easily perturbed away in a general dynamical system, but it is not clear whether this can always be done for network dynamics. We get round this issue by considering a local version of rigid synchrony. This property is local in time, but still requires a global condition: existence of a hyperbolic periodic orbit. A local version of rigidity was introduced by Golubitsky et al. [33, Section 2] for similar reasons; see in particular their Definition 2.9 (c).

The natural definition of the global pattern of rigid synchrony $\bowtie_{\mathbb{X}}$ on \mathbb{X} is to take $J = \mathbb{R}$ in (7.1); which leads to

$$c \bowtie_{\mathbb{X}} d \iff \tilde{x}_c(t) \equiv \tilde{x}_d(t) \quad \forall t \in \mathbb{R}$$

where $\tilde{x}(t)$ is any perturbed periodic orbit for a sufficiently small admissible perturbation. This colouring is the same as \approx^{rig} in [72]. However, because of the possible non-existence of generic points, we work locally:

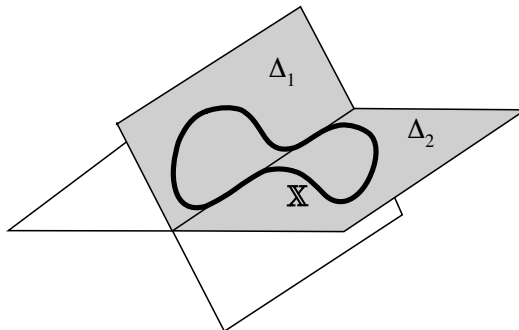


Figure 5. A periodic orbit with more than one local synchrony pattern, which potentially might be rigid.

Definition 7.8. A colouring \bowtie is a *locally rigid synchrony pattern* or *rigid local synchrony pattern* for \mathbb{X} at $x(t_0)$ if and only if there exists $\varepsilon > 0$ and $t^* > 0$ such that

- (a) \mathbb{X} remains hyperbolic after a perturbation p with $\|p\|_1 < \varepsilon$,
- (b) the condition

$$\tilde{x}_c^p(t) = \tilde{x}_d^p(t) \iff c \bowtie d$$

is valid for all perturbations p with $\|p\|_1 < \varepsilon$ and all t such that $|t - t_0| < t^*$.

That is, the synchrony pattern for $\tilde{x}^p(t)$ is the same for all small enough p and for all t close enough to t_0 . We define this to be *the* local synchrony pattern of \mathbb{X} near t_0 , and denote it by $\bowtie_{x(t_0)}$.

8. Construction of admissible perturbations

We now set up a general method for defining admissible perturbations with arbitrarily small support and arbitrarily small C^1 norm. We describe the method for a general network \mathcal{G} , but apply it below to quasi-quotients $\mathcal{G}^{\mathcal{R}}$.

8.1. Symmetrisation

We will use Proposition 4.7 to construct admissible perturbations, supported on (usually small) sets, by defining suitable bump functions and then symmetrising, as explained below in Definition 8.1. Let

$$\begin{aligned} g_c &: P_c \times P_{T(c)} \rightarrow P_c, \\ (x_c, x_{T(c)}) &\mapsto g_c(x_c, x_{T(c)}) \end{aligned}$$

be a smooth admissible map. By (4.5) the vertex group $\Gamma = B(c, c)$ acts on such maps by

$$\gamma g_c(x_c, x_{T(c)}) = g_c(x_c, \gamma^* x_{T(c)})$$

where γ^* is the pullback of γ .

Definition 8.1. The *symmetrisation* of g_c is g_c^Γ , where

$$g_c^\Gamma(x_c, x_{T(c)}) = \frac{1}{|\Gamma|} \sum_{\gamma \in \Gamma} (x_c, \gamma^* x_{T(c)}).$$

Symmetrisation on each input class, combined with standard-order identifications, yields the groupoid symmetrisation construction of [72] up to constant factors. The normalisation factor $1/|\Gamma|$ ensures that

$$\|g_c^\Gamma\|_1 \leq \|g_c\|_1 \tag{8.1}$$

bearing in mind that $\|g_c^\gamma\|_1 = \|g_c\|_1$ for all $c \in \mathcal{C}$ because Γ acts by permuting coordinates. Clearly, g_c^Γ is Γ -invariant, that is, $B(c, c)$ -invariant. Moreover,

$$\text{supp}(g_c^\Gamma) = \bigcup_{\gamma \in \Gamma} \gamma \text{supp}(g_c) \quad (8.2)$$

where supp is the support

$$\text{supp}(h) = \text{cl}\{x : h(x) \neq 0\}$$

and cl indicates the closure.

In this paper we apply symmetrisation to components of admissible maps for quasi-quotients $\mathcal{G}^{\mathcal{R}}$. The following simple lemma lets us perform symmetrisations on $P^{\mathcal{R}}$ and lift them to symmetrisations on P .

Lemma 8.2. *With the above notation,*

$$[\beta^* x_{T(c)}] = \beta^* x_{[T(c)]}.$$

Proof. Permuting entries, and replacing c by $[c]$, commute. In more detail, let $T(c) = (i_1, \dots, i_k)$, so that $[T(c)] = ([i_1], \dots, [i_k])$. Then

$$\begin{aligned} [\beta^* x_{T(c)}] &= (x_{[\beta(i_1)]}, \dots, x_{[\beta(i_k)]}) = (x_{\beta([i_1])}, \dots, x_{\beta([i_k])}) \\ &= \beta^*(x_{[i_1]}, \dots, x_{[i_k]}) = \beta^* x_{[T(c)]}. \quad \blacksquare \end{aligned}$$

8.2. Bump functions

We use bump functions to construct admissible C^∞ maps $p : P \rightarrow P$ with a specified compact support, taking a specific locally constant value \mathbf{w} on a neighbourhood of a specified point $z \in P$. We denote the open ball in the Euclidean norm of radius r centre x by $B_r(x)$.

Proposition 8.3. *Let $z \in P_c \times P_{T(c)}$ and suppose that $\delta > 0$. Let $\mathbf{w} \in P_c$. Then there exists a C^∞ map $\Psi_c : P_c \times P_{T(c)} \rightarrow P_c$ with compact support, hence C^1 -bounded, such that*

$$\Psi_c(y) = \begin{cases} \mathbf{w} & \text{if } \|y - z\| \leq \delta, \\ 0 & \text{if } \|y - z\| \geq 2\delta. \end{cases}$$

Proof. Let $x \in P_c \times P_{T(c)}$ and let $U = B_r(x)$, $V = B_s(x)$ where $r < s$. By Abraham et al. [2, Lemma 4.2.13] there is a C^∞ function $\phi : \mathbb{R}^k \rightarrow \mathbb{R}$ such that $\phi(x) = 1$ when $x \in U$ and $\phi(x) = 0$ when $x \notin V$. In particular, there exists a C^∞ bump function $\psi : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$\psi(x) = \begin{cases} 1 & \text{if } x \in (-1, 1), \\ 0 & \text{if } x \notin (-2, 2). \end{cases}$$

This has compact support. For any $\delta > 0$ let $\phi^\delta : \mathbb{R}^+ \rightarrow [0, \delta]$ be defined by

$$\phi^\delta(x) = \psi(x/\delta). \quad (8.3)$$

Then if $\|\cdot\|_E$ is the Euclidean norm, the function

$$\Phi(y) = \phi^\delta(\|y - z\|_E^2) \mathbf{w} \quad (8.4)$$

is smooth, supported on the compact set $\text{cl}(B_{2\delta}(z))$, and satisfies $\Phi(y) = \mathbf{w}$ if $\|y - z\| \leq \delta$. ■

8.3. Symmetrised bump functions

We construct admissible perturbations using symmetrised bump functions. In a general setting, let Γ be a finite group acting linearly and orthogonally on a real vector space V . Denote the Γ -orbit of $A \subseteq V$ by

$$\mathcal{O}(A) = \{\gamma.a : \gamma \in \Gamma, a \in A\}.$$

For finite subsets $Y, Z \subseteq V$ write the (Hausdorff) distance as

$$d(A, B) = \min\{\|y - z\| : y \in Y, z \in Z\}.$$

We state the following lemma for a general finite group action. It will be applied when the group is a vertex group.

Lemma 8.4. *Let V, W be real vector spaces, and let a finite group Γ act on V . Let $A, B \subseteq V$ be finite sets with disjoint group orbits*

$$\mathcal{O}(A) \cap \mathcal{O}(B) = \emptyset.$$

Let $0 \neq \mathbf{w} \in W$. Then there exists $\delta > 0$ and a Γ -invariant map $h : V \rightarrow W$ with compact support such that

$$h(y) = \begin{cases} 0 & \text{if } y \in \mathcal{O}(A), \\ \mathbf{w} & \text{if } d(y, \mathcal{O}(B)) \leq \delta, \\ 0 & \text{if } d(y, \mathcal{O}(B)) \geq 2\delta. \end{cases}$$

Proof. Since the group orbits $\mathcal{O}(A)$ and $\mathcal{O}(B)$ are disjoint, $d(\mathcal{O}(A), \mathcal{O}(B)) > 0$. Let

$$\delta < \frac{1}{2}d(\mathcal{O}(A), \mathcal{O}(B))$$

so that δ -neighbourhoods of $\mathcal{O}(A)$ and $\mathcal{O}(B)$ are disjoint. By Proposition 8.3 there is a bump function $\Phi : V \rightarrow W$ such that

$$\begin{aligned}\Phi(a) &= 0 & \text{if } a \in A, \\ \Phi(y) &= \mathbf{w} & \text{if } d(y, \mathcal{O}(B)) \leq \delta, \\ \Phi(y) &= 0 & \text{if } d(y, \mathcal{O}(B)) > 2\delta.\end{aligned}$$

Symmetrise, to obtain $h = \Phi^\Gamma$. Then h has compact support by definition, and the stated properties follow easily from Section 8.1. ■

9. Proof of Local Rigid Synchrony Property for strongly hyperbolic periodic orbits

We now adapt the method of Section 3 to prove the Local Rigid Synchrony Property for any network, under the extra condition that the periodic orbit is strongly hyperbolic (indeed, ‘stably isolated’ suffices).

If necessary, an arbitrarily small initial perturbation ensures that if \mathcal{R} is a set of representatives for local synchrony colouring then the induced periodic orbit $\mathbb{X}^{\mathcal{R}}$ is hyperbolic in $P^{\mathcal{R}}$, hence is an isolated periodic orbit in $P^{\mathcal{R}}$. The interval J and the size ε of a sufficiently small perturbation may have to be made smaller. Without loss of generality, we assume such a ‘pre-perturbation’ has been made, and retain the same notation for f, \mathbb{X} , and so on.

Recall that we consider a network \mathcal{G} with nodes $\mathcal{C} = \{1, 2, \dots, n\}$, with state spaces P_c for nodes c be \mathbb{R}^{k_c} , so that the total state space is $P = P_1 \times \dots \times P_n$. Consider a \mathcal{G} -admissible ODE (6.3), written in coordinates as

$$\dot{x}_c = f_c(x_c, x_{T(c)}), \quad c \in \mathcal{C}. \quad (9.1)$$

Let (6.4) be a family of small admissible perturbations. Let \mathbb{X} be a strongly hyperbolic periodic orbit of (6.3), with perturbed periodic orbit \mathbb{X}^ε for (6.4). Assume that at $t_0 \in \mathbb{R}$ the local synchrony pattern $\bowtie_{x(t_0)}$ is rigid.

We use Proposition 4.7 to define admissible perturbations p . This implies that we need to define p_c only for one node c from each input equivalence class, provided we ensure that p_c is $B(c, c)$ -invariant.

Remark 9.1. When $c \sim_I d$, Proposition 4.7 lets us identify P_c with P_d and $P_{T(c)}$ with $P_{T(d)}$. It also lets us identify the actions of $B(c, c)$ and $B(d, d)$, so $B(c, c)$ acts on $P_{T(d)}$ in the same way as it acts on $P_{T(c)}$.

We make these identifications without further comment from now on.

9.1. Statement of main theorem

As in [33, Sections 2 and 6], the issues raised in Section 7.6 make it necessary to prove a stronger local version of the Rigid Synchrony Property, namely:

Theorem 9.2. *If a strongly hyperbolic periodic orbit of a \mathcal{G} -admissible ODE has a locally rigid synchrony pattern on a non-empty open interval of time, then that pattern is balanced.*

The proof occupies the remainder of this section.

For any subset $J \subseteq \mathbb{R}$ write $x(J) = \{x(t) : t \in J\}$. In the sequel J is either a non-empty open interval or a point. Local rigidity of \mathbb{X} implies that there exists a non-empty open interval J such that the synchrony pattern $\bowtie_{x(t)}$ of any point $x(t) \in x(J)$ is the same for all $t \in J$, and is rigid. By Definition 7.8, this pattern includes *all* synchrony relations that hold on $x(J)$. By local rigidity, it is also the same colouring as that of $x^\varepsilon(t)$ for $\varepsilon \ll 1, t \in J$. (Here and elsewhere we may have to make J or ε smaller for such statements to be valid: we retain the same notation.)

Choose $t_0 \in J$, giving a point $x(t_0) \in \mathbb{X}$. As in Section 3.2, we choose a Poincaré section Σ transverse to \mathbb{X} at $x(t_0)$. Any perturbed periodic orbit meets Σ transversely in a unique point, provided the perturbation is sufficiently C^1 -small. We use this intersection point to fix initial conditions on perturbed periodic orbits $\tilde{\mathbb{X}}$ by requiring $\tilde{x}(t_0) \in \Sigma$. As just remarked, $\bowtie_{x(t_0)} = \bowtie_J$, and local rigidity similarly implies that $\bowtie_{\tilde{x}(t_0)} = \bowtie_{x(t_0)}$ for $\tilde{x}(t_0) \in \tilde{\mathbb{X}}$.

For a contradiction, assume that

$$\bowtie_{x(t_0)} \text{ is not balanced.} \tag{9.2}$$

The rest of this section proves that this cannot happen.

To simplify notation we write $\bowtie_{x(t_0)} = \bowtie^0$ from now on.

9.2. Induced OODE

The first step is to replace the unperturbed ODE (9.1) by the corresponding induced overdetermined ODE (or induced OODE) for the synchrony pattern with colouring \bowtie^0 . (We get an OODE because \bowtie^0 is not balanced.) Choose some set of representatives \mathcal{R} for the colours, and renumber nodes so that these are nodes $1, 2, \dots, m$. Denote the representative of node c in \mathcal{R} by $[c]$ as in Definition 5.1. Then $[c] = c$ when $c \in \mathcal{R}$, that is, for $1 \leq c \leq m$.

The set \mathcal{R} defines coordinates (u_1, \dots, u_m) on the polydiagonal $\Delta_{\bowtie^0} \subseteq P$ by setting $x_c = u_{[c]}$.

For example, if $|\mathcal{C}| = 5$ and the synchrony pattern is the partition $\{\{1, 2, 4\}, \{3, 5\}\}$, then (u_1, u_2) is identified with $(u_1, u_1, u_2, u_1, u_2) \in P$. Now u_r acts as the node coordinate for node r in $\mathcal{G}^{\mathcal{R}}$ when $1 \leq r \leq m$, and canonically identifies $P^{\mathcal{R}}$ with Δ_{\bowtie^0} .

By Proposition 7.7, local rigidity of \bowtie^0 implies

$$\text{the points } u_c^0(t_0) \text{ are distinct for } 1 \leq c \leq m. \quad (9.3)$$

Write input variables in standard order so that transitional pullback maps between distinct input equivalent nodes can be taken to be the identity. By Proposition 4.7, the only further requirement on the c -component for admissibility is $B(c, c)$ -invariance for each representative c of the input equivalence classes. Moreover, we can identify $B(d, d)$ with $B(c, c)$ whenever $c \sim_I d$, by Remark 9.1.

We have identified $\Delta_{\bowtie^0} \subseteq P$ with $P^{\mathcal{R}}$, with coordinates u_1, \dots, u_m . We now convert (9.1) into an OODE by substituting $u_{[c]}$ for x_c , for $1 \leq c \leq n$, as in Definition 5.6

$$\dot{u}_{[c]} = f_c(u_{[c]}, u_{[T(c)]}), \quad 1 \leq c \leq n. \quad (9.4)$$

Recall that $[T(c)]$ is the tuple of colours of tail nodes of input arrows $I(c)$. That is,

$$[(i_1, \dots, i_r)] = ([i_1], \dots, [i_r])$$

for the input tuple (i_1, \dots, i_r) . (The subscript of f_c remains c , not $[c]$, since we are not changing the map, just substituting different variables into it.) A solution \mathbb{X} of (9.1) satisfies the OODE (9.4) if and only if \mathbb{X} has synchrony pattern \bowtie^0 . Thus, \mathbb{X} determines a canonical periodic orbit $\mathbb{X}^{\mathcal{R}}$ for the $\mathcal{G}^{\mathcal{R}}$ -admissible induced ODE.

For all $u \in \mathbb{X}^{\mathcal{R}}$ and $c \in \{1, \dots, m\}$ we have

$$u_c = u_{[c]}, \quad u_{[T(c)]}(t) \in P_{[T(c)]}.$$

Say that the colouring is *unbalanced at nodes* c, d if distinct nodes c, d have the same colour, but their input sets are not colour-isomorphic. This is equivalent to

$$\text{either } c \not\sim_I d, \text{ or} \quad (9.5)$$

$$c \sim_I d \text{ and } x_{[T(c)]} \text{ and } x_{[T(d)]} \text{ lie in distinct orbits of } B(c, c). \quad (9.6)$$

Here we again use Remark 9.1, so that $B(c, c) = B(d, d)$ acts on *both* input tuples.

By (9.2), \bowtie^0 is not balanced at some $c \neq d$. Renumber the nodes so that $c = 1, d = m + 1$. Then: \bowtie^0 is unbalanced at nodes $1, m + 1$, and node 1 has the same colour as node $m + 1$.

Split (9.4) into the *induced OODE* for \mathcal{R} , which is

$$\dot{u}_c = f_c(u_c, u_{[T(c)]}), \quad 1 \leq c \leq m \quad (9.7)$$

together with the *conflicting equation*

$$\dot{u}_{m+1} = \dot{u}_1 = f_{m+1}(u_1, u_{[T(m+1)]}) \quad (9.8)$$

(bearing in mind that $[m+1] = 1$), with possible further constraint equations

$$\dot{u}_{[c]} = f_c(u_{[c]}, u_{[T(c)]}), \quad m+2 \leq c \leq n. \quad (9.9)$$

Because \bowtie^0 is unbalanced at nodes $1, m+1$, the two equations

$$\dot{u}_1 = f_1(u_1, u_{[T(1)]}), \quad \dot{u}_1 = f_{m+1}(u_1, u_{[T(m+1)]})$$

conflict, in the sense of (9.5) or (9.6): either nodes 1 and $m+1$ are not input equivalent, so that $f_{m+1} \neq f_1$, or they are input equivalent, so that $f_{m+1} = f_1$, but $u_{[T(m+1)]}$ and $u_{[T(1)]}$ do not lie in the same $B(1, 1)$ -orbit, using Lemma 8.2 to define the group action. (The same result holds if we replace x by u throughout.) We will use this conflict to deduce a contradiction. The remaining constraints (9.9) play no further role in the proof and are ignored. The various f_c may or may not be equal, and we distinguish these cases in the proof. Also, some input equivalence classes may not be represented among nodes $\{1, \dots, m\}$.

9.3. Perturbations

It is convenient to write

$$f^{\mathcal{R}} = (f_1, \dots, f_m)$$

for the map that appears in the induced ODE, to distinguish it from f . Now $f : P \rightarrow P$ and $f^{\mathcal{R}} : P^{\mathcal{R}} \rightarrow P^{\mathcal{R}}$. We use similar notation to distinguish a perturbation p on P from $p^{\mathcal{R}}$ on $P^{\mathcal{R}}$.

Theorem 5.4 implies that we can construct \mathcal{G} -admissible perturbations of f on P by first constructing $\mathcal{G}^{\mathcal{R}}$ -admissible perturbations $p^{\mathcal{R}}$ of $f^{\mathcal{R}}$ on $P^{\mathcal{R}}$ and then lifting to P . Moreover, lifting can be performed in a manner that does not increase the C^1 norm of the perturbation. We therefore work initially with $\mathcal{G}^{\mathcal{R}}$ -admissible maps on $P^{\mathcal{R}}$, using the previously defined coordinates u_1, \dots, u_m on $P^{\mathcal{R}}$.

Write the OODE for the perturbed map $f + \varepsilon p$ as

$$\dot{u}_c^\varepsilon = f_1(u_c^\varepsilon, u_{[T(c)]}^\varepsilon) + \varepsilon p_c(u_c^\varepsilon, u_{[T(c)]}^\varepsilon), \quad 1 \leq c \leq m \quad (9.10)$$

with conflicting equation

$$\dot{u}_{[m+1]}^\varepsilon = \dot{u}_1^\varepsilon = f_{m+1}(u_1^\varepsilon, u_{[T(m+1)]}^\varepsilon) + \varepsilon p_{m+1}(u_1^\varepsilon, u_{[T(m+1)]}^\varepsilon) \quad (9.11)$$

and possible further constraints

$$\dot{u}_c^\varepsilon = f_1(u_c^\varepsilon, u_{[T(c)]}^\varepsilon) + \varepsilon p_c(u_c^\varepsilon, u_{[T(c)]}^\varepsilon), \quad m+2 \leq c \leq n. \quad (9.12)$$

Assume, for a contradiction, that the OODE defined by (9.10), (9.11), and (9.12) has a solution for any admissible C^1 -bounded p . So although (9.8) conflicts with the first component of (9.7), they must agree on \mathbb{X} . Rigidity implies that the same statement holds for $\tilde{\mathbb{X}}$. Agreement of this kind is possible for a single admissible map f , but it is highly non-generic. We show that in the strongly hyperbolic case, it cannot remain valid for all small perturbations.

To avoid confusion, but at the expense of slightly more complicated notation, we write points on components of the induced unperturbed periodic orbit as $u_c^0(t)$ instead of $u_c(t)$. In particular, $u_c^0(t_0)$ is the c -component of the unique point (locally) at which the unperturbed periodic orbit meets the Poincaré section Σ . This notation is chosen to be consistent with the notation $u_c^\varepsilon(t)$ for the corresponding perturbed periodic orbit for a perturbation εp .

9.4. Proof strategy

The Local Rigid Synchrony Theorem 9.2 is an immediate consequence of the following two results, Lemma 9.3 and Lemma 9.4. The proof of Lemma 9.4 is deferred to Section 9.5.

Lemma 9.3. *Suppose there exists an admissible map p for \mathcal{G} such that for $1 \leq c \leq m$,*

$$p_c(x_c^\varepsilon(t), x_{[T(c)]}^\varepsilon(t)) = 0, \quad 1 \leq c \leq m, \quad \varepsilon \ll 1, |t - t_0| \ll 1 \quad (9.13)$$

and

$$p_{m+1}(x_{m+1}^0(t_0), x_{[T(m+1)]}^0(t_0)) \neq 0. \quad (9.14)$$

Then \bowtie^0 is not locally rigid.

Proof. The periodic orbit \mathbb{X} induces a canonically defined periodic orbit $\mathbb{X}^\mathcal{R}$ for (9.10) on $P^\mathcal{R}$. By (9.13), for all $\varepsilon \ll 1$ and t near t_0 , the induced equations for the perturbed $u_c^\varepsilon(t)$ are the same as those for the unperturbed coordinate $u_c^0(t)$. Now $\mathbb{X}^\varepsilon \rightarrow \mathbb{X}$ as $\varepsilon \rightarrow 0$. By strong hyperbolicity, if $\varepsilon \ll 1$ then $\mathbb{X}^\varepsilon = \mathbb{X}$. Since initial conditions are determined by the Poincaré section Σ , we have $x^\varepsilon(t) \equiv x(t)$ for $t \in J$. Thus

$$u_c^\varepsilon(t) \equiv u_c^0(t) \quad \forall t \in J.$$

(By uniqueness of solutions of ODEs, this identity is valid for all $t \in \mathbb{R}$.) The conflicting equation is

$$\dot{u}_1^0 = f_{m+1}(u_1^0, u_{[T(m+1)]}^0) + \varepsilon p_{m+1}(u_1^0, u_{[T(m+1)]}^0) \quad \forall \varepsilon \ll 1, t \in J.$$

Since all terms except ε are independent of ε ,

$$p_{m+1}(u_1^0(t), u_{[T(m+1)]}^0(t)) = 0 \quad \forall \varepsilon \ll 1, t \in J.$$

This is a contradiction since p_{m+1} is nonzero at $(u_1^0(t_0), u_{[T(m+1)]}^0(t_0))$, by (9.14). ■

Thus the proof of the Local Rigid Synchrony Property, Theorem 9.2, reduces to proving the second lemma:

Lemma 9.4. *There exists an admissible perturbation p , with support near $x^0(t_0)$ and its images under the vertex group $B(1, 1)$, that satisfies (9.13) and (9.14).*

We prove Lemma 9.4 in the next subsection. To do so, we need some additional observations.

9.5. Construction of the perturbation

It remains to construct a \mathcal{G} -admissible perturbation $p : P \rightarrow P$ that satisfies (9.13) and (9.14). We do this by constructing a $\mathcal{G}^{\mathcal{R}}$ -admissible perturbation $q : P^{\mathcal{R}} \rightarrow P^{\mathcal{R}}$ and lifting it to p using Theorem 5.4.

To construct q , we focus on the ODE (9.10) and the constraint (9.11), remembering that

$$m + 1 \bowtie^0 1.$$

To simplify notation, write

$$\begin{aligned} U_c &= (u_c^0(t_0), u_{[T(c)]}^0(t_0)), \\ \mathcal{O}_c &= \mathcal{O}(U_c) \quad \text{for } \Gamma = B(c, c), \\ U_c^\varepsilon &= (u_c^\varepsilon(t_0), u_{[T(c)]}^\varepsilon(t_0)), \\ \mathcal{O}_c^\varepsilon &= \mathcal{O}(U_c^\varepsilon) \quad \text{for } \Gamma = B(c, c). \end{aligned}$$

The identifications in Remark 9.1, transferred to $P^{\mathcal{R}}$, imply that the finite sets $\mathcal{O}_c, \mathcal{O}_d$ lie in the same space $P_c \times P_{[T(c)]}$ if $c \sim_I d$, and otherwise lie in distinct spaces. Moreover, when $c \sim_I d$ the groups $B(c, c), B(d, d)$ and their actions on this space are identified.

Lemma 9.5. *With the above notation,*

$$\mathcal{O}_c \cap \mathcal{O}_d = \emptyset \quad \forall c, d : 1 \leq c \neq d \leq m + 1, c \sim_I d.$$

Proof. Proposition 7.7 implies that the points $u_c^0(t_0)$ are distinct for $1 \leq c \leq m$. Therefore, by projection onto P_c , the points U_c are distinct for $1 \leq c \leq m$. Since Γ acts trivially on the ‘base point’ $u^0(t_0)$ of U_c , it follows that $\mathcal{O}_c \cap \mathcal{O}_d = \emptyset$ whenever $1 \leq c \neq d \leq m$.

Since $m + 1 \bowtie^0 1$, the only possible non-empty intersection occurs for $\mathcal{O}_1 \cap \mathcal{O}_{m+1}$. These two sets lie in the same space only when $m + 1 \sim_I 1$. But then, \mathcal{O}_1 and \mathcal{O}_{m+1} are disjoint by (9.6). \blacksquare

Proof of Lemma 9.4. The proof splits into three cases:

- (a) Node $m + 1$ is not input equivalent to any of nodes $1, \dots, m$.
- (b) Node $m + 1$ is input equivalent to node k where $2 \leq k \leq m$.
- (c) Node $m + 1$ is input equivalent to node 1.

The arguments are very similar in all cases, but differ in fine detail. We take the three cases in turn.

Case (a): Here we define the perturbation p directly for \mathcal{G} . We use a bump function and symmetrisation, as in Sections 8.2 and 8.1, to define p_{m+1} so that (9.14) holds. Specifically, in Lemma 8.4 make δ sufficiently small and take

$$\begin{aligned} \Gamma &= B(m + 1, m + 1), \\ V &= P_{m+1} \times P_{[T(m+1)]}, & W &= P_{m+1}, \\ A &= \emptyset, & B &= \{(x_{m+1}^0(t_0), x_{[T(m+1)]}^0(t_0))\}. \end{aligned}$$

The hypothesis $\mathcal{O}(A) \cap \mathcal{O}(B) = \emptyset$ clearly holds.

Define $p_{m+1} = h$ where h is as in Lemma 8.4. Then $p_{m+1}(B) \neq 0$, which is (9.14). Since variables are in standard order, the transition maps are the identity, so this determines all p_c with $c \sim_I m + 1$. On all remaining input equivalence classes, we set $p_c \equiv 0$. In particular, $p_c \equiv 0$ for $1 \leq c \leq m$, so (9.13) holds. Since h has compact support, so does p , and p is C^1 -bounded.

Case (b): Since $m + 1 \sim_I k$ and variables are in standard order, $g_{m+1} = g_k$ for any \mathcal{G} -admissible map g . Let

$$\mathcal{C}_1 = \{c : c \sim_I k \text{ \& } 1 \leq c \leq m\}. \quad (9.15)$$

We define a perturbation q for $\mathcal{G}^{\mathcal{R}}$ and then lift to obtain p using Theorem 5.4. We define q so that q has compact support on $P^{\mathcal{R}}$,

$$q_c(U_c^\varepsilon) = 0, \quad 1 \leq c \leq m, \quad \varepsilon \ll 1 \quad (9.16)$$

and

$$q_k(U_{m+1}) \neq 0. \quad (9.17)$$

To do so, we define q_c for $c \in \mathcal{C}_1$ and make it vanish on $\mathcal{R} \setminus \mathcal{C}_1$. Two conditions on q_k must be satisfied: (9.16) when $c = k$, and (9.17). If these conditions are satisfied, any lift p of q satisfies (9.13) and (9.14) and also has compact support, hence is C^1 -bounded. Therefore εp is C^1 -small for $\varepsilon \ll 1$. We are therefore finished once we show that (9.13) and (9.14) can be satisfied simultaneously.

To complete the proof, let \mathcal{C}_1 be as in (9.15). Let $q_d = 0$ for all $d \in \mathcal{R} \setminus \mathcal{C}_1$. For $c \in \mathcal{C}_1$, we define q_c using Lemma 8.4. Make δ sufficiently small and take

$$\begin{aligned}\Gamma &= B(m+1, m+1), \\ V &= P_k \times P_{[T(k)]}, & W &= P_k, \\ A &= \bigcup_{c \in \mathcal{C}_1} \mathcal{O}_c, & B &= \{U_{m+1}\}.\end{aligned}$$

The hypothesis $\mathcal{O}(A) \cap \mathcal{O}(B) = \emptyset$ holds by Lemma 9.5. This defines a $\mathcal{G}^{\mathcal{R}}$ -admissible perturbation $q : P^{\mathcal{R}} \rightarrow P^{\mathcal{R}}$. Lift q to a \mathcal{G} -admissible map $p : P \rightarrow P$, making it zero on all components that are not input equivalent to a node in $\mathcal{R} \cup \{m+1\}$; then p has the required properties. Again p is C^1 -bounded, so εp is C^1 -small for $\varepsilon \ll 1$ and case (b) is proved.

Case (c): Again we define a perturbation q for $\mathcal{G}^{\mathcal{R}}$ and then lift to obtain p using Theorem 5.4.

In this case $m+1 \sim_I 1$ so $q_{m+1} = q_1$. Let $\mathcal{C}_1 = \{c : c \sim_I 1 \text{ \& } 1 \leq c \leq m\}$. As before, we define q_c for $c \in \mathcal{C}_1$ and make it vanish on $\mathcal{R} \setminus \mathcal{C}_1$. The only potential obstacle is that both (9.13) and (9.14) impose conditions on p_1 , which might be contradictory. We want q_1 to vanish at U_1^ε for $\varepsilon \ll 1$, but to be nonzero at U_{m+1}^0 . Again we use Lemma 8.4. Make δ sufficiently small, and take

$$\begin{aligned}\Gamma &= B(1, 1) = B(m+1, m+1), \\ V &= P_1 \times P_{[T(1)]}, & W &= P_1, \\ A &= \bigcup_{c \in \mathcal{C}_1} \mathcal{O}_c, & B &= \{U_{m+1}\}.\end{aligned}$$

The hypothesis $\mathcal{O}(A) \cap \mathcal{O}(B) = \emptyset$ holds by Lemma 9.5. This defines a $\mathcal{G}^{\mathcal{R}}$ -admissible perturbation $q : P^{\mathcal{R}} \rightarrow P^{\mathcal{R}}$. Lift q to a \mathcal{G} -admissible map $p : P \rightarrow P$, making it zero on all components that are not input equivalent to a node in \mathcal{R} . Then p has the required properties. Again p is C^1 -bounded, so εp is C^1 -small for $\varepsilon \ll 1$ and case (b) is proved.

This completes the proof of Theorem 9.2. ■

Remarks 9.6. (a) The same strategy gives another proof of the Rigid Equilibrium Theorem, first proved in [43, Theorem 7.6]. This new proof (see [69]) is simpler than the periodic case, and we can appeal to Sard's Theorem instead of assuming strong hyperbolicity as an extra hypothesis. This approach has some similarities to a proof based on transversality arguments by Aldis [3, Theorem 7.2.3], but is simpler.

(b) The same proof works if node spaces are arbitrary C^∞ manifolds. Indeed, only C^1 -smoothness is required throughout. In particular, the above results and proofs remain valid for phase oscillators, where the node spaces are the circle \mathbf{S}^1 .

10. Global rigid synchrony and the Rigid Input Property

Having established the main result of this paper, we can deduce the usual global version of the Rigid Synchrony Property [33, 72], assuming as before that \mathbb{X} is strongly hyperbolic. As discussed in Section 7.6, the main obstacle is the possibility, in principle, that a local synchrony pattern on some interval of time need not be the same as the global synchrony pattern of the entire periodic orbit \mathbb{X} . Indeed, the global synchrony pattern $\bowtie_{\mathbb{X}}$ need not equal the local synchrony pattern $\bowtie_{x(t)}$ for any specific t . In fact, rigidity prevents this happening, but the proof requires a little care. To make the proof precise we require a number of technical definitions. We also use the lattice of colourings from Section 7.3.

10.1. Local rigidity implies global rigidity

We now appeal to the Local Rigid Synchrony Theorem 9.2 to show that the above change of local synchrony pattern cannot occur if \mathbb{X} is strongly hyperbolic and the global pattern $\bowtie_{\mathbb{X}}$ is rigid. The following remark is useful:

Remark 10.1. Any finite number of perturbations performed in turn can be made arbitrarily small by making successive sizes be $\varepsilon/2, \varepsilon/4, \varepsilon/8, \dots$. The triangle inequality then shows that the combined perturbation has size $< \varepsilon$. We can also reduce the size of ε or the interval J finitely many times if required.

Theorem 10.2. *Assume that \mathbb{X} is strongly hyperbolic, and suppose that the global synchrony pattern $\bowtie_{\mathbb{X}}$ on \mathbb{X} is rigid. Then it is balanced.*

Proof. As before, denote the local synchrony pattern at x_{t_0} by \bowtie^0 . Let $\bowtie_{\mathbb{X}}$ be the global synchrony pattern of \mathbb{X} and assume this is rigid.

Since $\mathbb{X} \subseteq \Delta_{\bowtie_{\mathbb{X}}}$, the colouring \bowtie^0 is coarser than $\bowtie_{\mathbb{X}}$. Proposition 7.6 implies that after any sufficiently small perturbation, in which \mathbb{X} becomes $\tilde{\mathbb{X}}$ and x_0 becomes \tilde{x}_0 , the colouring $\tilde{\bowtie}^0$ can become finer than \bowtie^0 , but not coarser. Therefore $\tilde{\bowtie}^0$ remains coarser than $\bowtie_{\mathbb{X}}$ since $\bowtie_{\mathbb{X}}$ is rigid; that is, $\Delta_{\tilde{\bowtie}^0} \subseteq \Delta_{\bowtie_{\mathbb{X}}}$.

Continue making perturbations until $\tilde{\bowtie}^0$ is as fine as possible. By Remark 10.1, the combined perturbation can be made as small as we wish. Now $\tilde{\bowtie}^0$ is locally rigid. By Theorem 9.2, $\tilde{\bowtie}^0$ is balanced. Therefore $\tilde{\mathbb{X}}$ intersects the synchrony space for $\tilde{\bowtie}^0$, and flow-invariance implies that $\mathbb{X} \subseteq \Delta_{\tilde{\bowtie}^0}$. By definition of the global synchrony pattern, $\Delta_{\bowtie_{\mathbb{X}}} \subseteq \Delta_{\tilde{\bowtie}^0}$. Thus the two are equal, so $\bowtie_{\mathbb{X}}$ is balanced. ■

Corollary 10.3. *Let \mathbb{X} be a strongly hyperbolic periodic orbit with a locally rigid synchrony pattern on some non-empty open time interval. Then the entire orbit has that synchrony pattern, and it is balanced.*

Proof. By Theorem 9.2 the rigid local synchrony pattern \bowtie^0 is balanced. Therefore Δ_{\bowtie^0} is flow-invariant. But $\mathbb{X} \cap \Delta_{\bowtie^0} \neq \emptyset$, so $\mathbb{X} \subseteq \Delta_{\bowtie^0}$. ■

Some points on the periodic orbit might have extra equalities among their components, compared to $\bowtie_{\mathbb{X}}$, but these cannot be balanced and cannot persist rigidly on any non-empty open interval of time. In particular, the scenario of Figure 5 cannot occur rigidly when \mathbb{X} is strongly hyperbolic.

10.2. Local Rigid Input Property

The Local (hence also the global) Rigid Input Property now follows trivially:

Corollary 10.4. *The Local Rigid Input Property holds for all strongly hyperbolic periodic orbits.*

Proof. The Rigid Synchrony Property implies the Rigid Input Property for synchronous nodes, because any colour-preserving input isomorphism is, in particular, an input isomorphism. ■

11. Rigid Phase Property

In this section we deduce the Local Rigid Phase Property from Theorem 9.2 using the ‘doubling’ trick of Golubitsky et al. [34]; see also Aldis [3, Chapter 10].

We state the Rigid Phase Property in the following local form. The global Rigid Phase Property is the case $J = \mathbb{R}$. We use p rather than εp and make $\|p\|_1$ small, and write $\mathbf{S}^1 = \mathbb{R}/\mathbb{Z}$ for the circle group, representing the phase as a proportion of the period.

Theorem 11.1 (Local Rigid Phase Property). *Let \mathbb{X} be a strongly hyperbolic periodic orbit of a \mathcal{G} -admissible ODE, and let $\tilde{\mathbb{X}}$ be the corresponding perturbed periodic orbit for an admissible perturbation p with $\|p\|_1 \ll 1$. Suppose that two nodes c, d in \mathcal{G} are rigidly phase-related on a time interval J ; that is*

$$\tilde{x}_c^h(t) \equiv \tilde{x}_d(t + \theta \tilde{T}) \quad \forall t \in J, \forall p \ll 1 \tag{11.1}$$

for a fixed proportion $\theta \in \mathbf{S}^1$ of the perturbed period \tilde{T} . Then there exists a vertex symmetry $\beta \in B(d, d)$ such that

$$x_{I(c)}(t) \equiv \beta^* x_{I(d)}(t + \theta T) \quad \forall t \in J. \tag{11.2}$$

Here $\tilde{\mathbb{X}}$ and \tilde{T} depend on p , but we suppress p in the notation. Informally, a locally rigid phase shift implies that input sets of phase-related nodes are related by the same phase shift, up to the action of a vertex group element.

The proof of Theorem 11.1 closely mimics that of Theorem 9.2, so we omit routine details. First, we state a simple corollary:

Corollary 11.2. *If nodes are related by a locally rigid phase shift, they are input equivalent.*

Proof. The map β introduced just before equation (11.2) is an input equivalence. ■

The central idea in the proof of Theorem 11.1 is a trick from [3, 34], namely: construct two isomorphic copies $\mathcal{G}_1, \mathcal{G}_2$ of \mathcal{G} and form the disjoint union

$$2\mathcal{G} = \mathcal{G}_1 \dot{\cup} \mathcal{G}_2.$$

If the state space for \mathcal{G} is P then that for $2\mathcal{G}$ is $P \times P$. Take coordinates (x, y) on $P \times P$ where $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$. We recall some simple properties of the *doubled network* $2\mathcal{G}$ that are proved in [34].

Lemma 11.3. *Let f be \mathcal{G} -admissible. Then (with obvious identifications) (f, f) is $2\mathcal{G}$ -admissible, and all $2\mathcal{G}$ -admissible maps are of this form.*

Proof. See [34, Lemma 4.3]. ■

Assume as usual an admissible ODE (6.3) on P for \mathcal{G} . This induces an admissible ODE on P for $2\mathcal{G}$ of the form

$$\dot{x} = f(x), \quad \dot{y} = f(y). \quad (11.3)$$

The dynamics of \mathcal{G}_1 and \mathcal{G}_2 are *decoupled*, so a periodic state $\mathbb{X} = \{x(t)\}$ for f on \mathcal{G} gives rise to a 2-torus \mathbb{T}^2 for (f, f) on $2\mathcal{G}$, foliated by periodic orbits

$$\mathbb{X}_\theta = \{(x(t), x(t + \theta T)) : t \in \mathbb{R}\} \quad (11.4)$$

where $\theta \in \mathbf{S}^1$. We call \mathbb{X}_θ a θ -*sheared* periodic orbit.

The following result converts rigid phase relations on \mathcal{G} into rigid synchronies on $2\mathcal{G}$; see [34, Section 4]. The proof is immediate.

Lemma 11.4. *Nodes c, d are rigidly phase-related by θ in \mathcal{G} , with c corresponding to c_1 in \mathcal{G}_1 and d corresponding to d_2 in \mathcal{G}_2 , if and only if c_1, d_2 are rigidly synchronous on \mathbb{X}_θ .* ■

The method of proof assumes that \mathcal{G} has the Rigid Synchrony Property for strongly hyperbolic \mathbb{X} , and uses properties of \mathbb{X}_θ , where θ is the assumed rigid phase relation. The idea is to deduce that \mathbb{X}_θ has the Rigid Synchrony Property for the periodic orbit on $2\mathcal{G}$. Then Lemma 11.4 yields the Rigid Synchrony Property for \mathbb{X} on \mathcal{G} . However, the Rigid Synchrony Property for \mathbb{X}_θ is not immediate because the

foliation by tori is an obstacle to hyperbolicity. To avoid this obstacle we establish a modified version of the Rigid Synchrony Property for \mathbb{X}_θ on $2\mathcal{G}$. This is achieved as follows.

Proof of Theorem 11.1. (a) Assume a rigid phase relation

$$x_c(t) = x_d(t + \theta T), \quad t \in J$$

where either $c \neq d$ or $c = d$ and $\theta \not\equiv 0 \pmod{T}$. Assume for a contradiction that this relation does not extend to the corresponding input tuples.

(b) Assume that \mathbb{X} is strongly hyperbolic (or is locally Kupka–Smale or has the strong isolation property). This implies that \mathcal{G} has the Rigid Synchrony Property for \mathbb{X} but in principle is a stronger condition.

(c) Consider the corresponding admissible ODE (11.3) for $2\mathcal{G}$. Let \mathbb{X}_θ be defined by (11.4), and consider an admissible perturbation of the form $(\varepsilon p, \varepsilon p)$ where p is admissible for \mathcal{G} and C^1 -bounded (which follows if p is compactly supported). The perturbed ODE has the form

$$\dot{x} = f(x) + \varepsilon p(x), \quad \dot{y} = f(y) + \varepsilon p(y). \quad (11.5)$$

If θ is rigid, the unique perturbed periodic orbit $\tilde{\mathbb{X}} = \{(\tilde{x}^\varepsilon(t), \tilde{x}^\varepsilon(t + \theta \tilde{T}))\}$ satisfies

$$\tilde{x}_c(t) = \tilde{x}_d(t + \theta \tilde{T}), \quad t \in J.$$

By Lemma 11.4, this corresponds to a rigid synchrony relation for $2\mathcal{G}$.

Theorem 9.2 does not apply directly, as noted above, but we can use the same method of proof with extra conditions. The proof has three key ingredients:

- (a) There exists a generic point $x(t)$ on the periodic orbit with $t \in J$.
- (b) There is a conflicting component of the ODE; that is, one that is formally inconsistent with the equation

$$f_c(x_c, x_{T(c)}) = f_d(x_d(t + \theta T), x_{T(d)}(t + \theta T))$$

in the sense that either $d \not\sim c$ or $x_{T(c)}$ and $x_d(t + \theta T)$ do not lie in the same $B(c, c)$ -orbit.

- (c) $\mathbb{X}^{\mathcal{R}}$ is hyperbolic, if necessary after a pre-perturbation of f .

If we can arrange analogous statements for \mathbb{X}_θ on $P \times P$, the proof goes through and the resulting contradiction establishes Theorem 11.1.

A useful simplifying step is to form the quotient \mathcal{G}/\bowtie where \bowtie is the relation of (local or global) rigid synchrony, which we know is balanced. Now nodes of \mathcal{G}/\bowtie are synchronous if and only if they are identical. Replacing \mathcal{G} by this quotient (and renaming this \mathcal{G}), we may assume that the only synchrony relations for \mathbb{X}_θ on $2\mathcal{G}$ are those of the form $x_c(t) \equiv x_d(t + \theta)$; that is, between $P \times \{0\}$ and $\{0\} \times P$.

Condition (a) is straightforward, except perhaps when $x_1(t) \equiv x_1(t + \theta)$, so the orbit on node 1 has period $\theta < T$. (This could, for example, occur for a multirhythm on node 1.) This possibility corresponds to case (c) of the proof of Lemma 9.4, and is dealt with in the same manner. We are assuming that node 1 does not satisfy the Rigid Phase Property, so $x_{[T(1)]}(t_0)$ and $x_{[T(1)]}(t_0 + \theta)$ lie in distinct $B(1, 1)$ -orbits. Therefore the required conditions on p_1 can be satisfied. Condition (b) is immediate because we are assuming, for a contradiction, that a formal inconsistency occurs. Condition (c) can be dealt with by working only with sheared periodic orbits for $2\mathcal{G}$. We need:

Definition 11.5. A periodic orbit \mathbb{X}_θ is *quasi-hyperbolic* if all of its Floquet multipliers are off the unit circle except for two that equal 1. Of these, one is associated with a phase shift along \mathbb{X}_θ , while the other is associated with a change from \mathbb{X}_θ to \mathbb{X}_ϕ where $\phi \in \mathbf{S}^1$ with ϕ near θ and $\phi \neq \theta$.

Definition 11.6. Let \mathcal{P} be one of the properties ‘strongly hyperbolic’, ‘stably isolated’, ‘locally Kupka–Smale’ for \mathbb{X} . Then the property *quasi- \mathcal{P}* for \mathbb{X}_θ is defined for ‘strongly hyperbolic’ and ‘locally Kupka–Smale’ by replacing ‘hyperbolic’ by ‘quasi-hyperbolic’ in the definition of \mathcal{P} . For ‘stably isolated’ it is defined by being stably isolated except for nearby periodic orbits \mathbb{X}_ϕ .

Now condition (c) follows from:

Lemma 11.7. *Let \mathcal{P} be one of the properties ‘strongly hyperbolic’, ‘stably isolated’, or ‘locally Kupka–Smale’. If \mathbb{X} has property \mathcal{P} for (6.3), then \mathbb{X}_θ has the property quasi- \mathcal{P} for (11.3).*

Proof. First, observe that $2\mathcal{G}$ -admissible perturbed families have the form (11.5).

Strongly hyperbolic: There is a perturbation p such that $\{\tilde{x}^{\mathcal{R}}(t)\}$ is hyperbolic on P . Since $\tilde{x}(t + \theta T)$ is a phase-shifted copy of $\tilde{x}(t)$, the orbit $\{\tilde{x}^{\mathcal{R}}(t + \theta)\}$ is hyperbolic on P for the same perturbation p . Now all Floquet multipliers of \mathbb{X}^θ lie off the unit circle except for two multipliers that are equal to 1: one for the first component $\{\tilde{x}(t)\}$ and the other for the second component $\{\tilde{x}(t + \theta T)\}$. Restricting to any sheared periodic orbit (which includes $\tilde{\mathbb{X}}_\theta$) reduces these to one Floquet multiplier equal to 1 and the rest off the unit circle. Therefore $\mathbb{X}_\theta^{\mathcal{R}}$ is quasi-hyperbolic.

Stably isolated: If $\{\tilde{x}(t)\}$ is stably isolated on P for a small perturbation p , then since $\tilde{x}(t + \theta T)$ is a phase-shifted copy of $\tilde{x}(t)$, the orbit $\{\tilde{x}^{\mathcal{R}}(t + \theta)\}$ is stably isolated on P for the same perturbation p . Therefore $\tilde{\mathbb{X}}_\theta$ is quasi-stably-isolated on $P \times P$.

Locally Kupka–Smale: The proof is similar to case (a) of Lemma 9.4. ■

This completes the proof of Theorem 11.1. ■

Remark 11.8. The example of Section 3 sheds light on the requirement of rigidity. Among the admissible maps (3.1), consider the ‘non-generic’ case when $g = f$. The ODE then has \mathbb{Z}_3 symmetry, and therefore supports rotating waves with $\frac{1}{3}$ -period phase shifts; see for example [42, Chapter XVIII, Section 0], and [36, Section 4.8]. These phase shifts would be rigid if the dashed arrow in Figure 1 had the same arrow-type as the solid ones. However, since this is not the case, a perturbation of the form (p, p, q) can (and, as we proved above, does) change the phase shifts, so that they are no longer one third of the period.

12. Full Oscillation Property

We now deduce the Full Oscillation Property for strongly hyperbolic \mathbb{X} . Following Gandhi et al. [28], say that node d is *upstream* from node c if there is a directed path in \mathcal{G} from d to c . We prove a stronger local version of the property: if some node of a network is rigidly steady for some non-empty open interval of time J , then all upstream nodes are also steady for $t \in J$. The global version follows immediately since ‘oscillate’ is local in time.

A node is always synchronous with itself. In the proof of the Local Rigid Synchrony Property, we do not assume that \mathbb{X} is oscillating at any particular node; only that the overall orbit is periodic (not steady) and nodes 1 and $m + 1$ are distinct.

Phase relations are different. A node can be phase-related to *itself* in a nontrivial manner; indeed, this is precisely what happens in multirhythms. A multirhythm occurs when some nodes oscillate at frequencies rationally related to the overall period, because the phase pattern requires those nodes to oscillate as nontrivially phase-shifted copies of themselves. This phenomenon goes back to [42, Chapter XVIII, Section 0], and is discussed in [36, Section 3.6].

The proof of the Rigid Phase Property allows multirhythms, because the two copies of such a node are distinct in $2\mathcal{G}$. In fact, if node c experiences a multirhythm, rigidly, then Theorem 11.1 implies that $x_{T(c)}$ is invariant under the same phase shift, up to some input automorphism in $\beta \in B(c, c)$. That is,

$$x_c(t) \equiv x_c(t + \theta) \implies x_{T(c)}(t) \equiv \beta^* x_{T(c)}(t + \theta). \quad (12.1)$$

The input automorphism β is essential here. Indeed, without some such automorphism, the phase shift θ would propagate back through the entire (transitive) network and imply that all nodes oscillate with the same *minimal* period. This is false for multirhythms, and is why they are interesting.

Definition 12.1. Let $\mathbb{X} = \{x(t)\}$ be a hyperbolic periodic orbit of a network ODE. A node c is *rigidly steady* on a non-empty open set $J \subseteq \mathbb{R}$ if \tilde{x}_c is steady (that is, $\tilde{x}_c(t)$ is constant) for all sufficiently small perturbed periodic orbits $\tilde{\mathbb{X}} = \{\tilde{x}(t)\}$.

The proof of the Local Rigid Phase Property makes no extra assumptions about the phase shift θ , except that it is rigid. In particular, other phase relations are also permitted (and occur in multirhythms). The connection with the Full Oscillation Property arises because equilibria can be viewed as extreme cases of multirhythms:

Lemma 12.2. *Let $\mathbb{X} = \{x(t)\}$ be a periodic state with period T , and let J be a non-empty open subset of \mathbb{R} . Then $x_c(t)$ is an equilibrium for $t \in J$ if and only if*

$$x_c(t) = x_c(t + \theta T) \quad \forall \theta \in \mathbf{S}^1 = \mathbb{R}/\mathbb{Z}, t \in J. \quad \blacksquare$$

As observed in [34, Section 2], in order to prove the Local Full Oscillation Property it is enough to prove that a small enough perturbation makes at least one additional node oscillate. When \mathcal{G} is transitive, iterating with smaller and smaller perturbations, as in Remark 10.1, makes all nodes oscillate, because ‘node c oscillates on J ’ is an open property.

We prove a stronger result. The key observation is:

Lemma 12.3. *Suppose that \mathcal{G} has the Local Rigid Phase Property and node c is rigidly steady for $t \in J$. Then every node d upstream from c is also rigidly steady for $t \in J$.*

Proof. We prove that any input node $d \in T(c)$ is steady for $t \in J$. That it is rigidly steady then follows by Remark 10.1.

Since node c is rigidly steady, any sufficiently small perturbation creates a unique perturbed periodic orbit $\tilde{\mathbb{X}}$ that is also steady at c . Therefore by Lemma 12.2,

$$\tilde{x}_c(t) = \tilde{x}_c(t + \theta \tilde{T}) \quad \forall \theta \in \mathbf{S}^1$$

where \tilde{T} is the period of $\tilde{\mathbb{X}}$. That is, the phase shift θ at node c is rigid for all $\theta \in \mathbf{S}^1$. (All phase shifts in \mathbf{S}^1 occur, but any particular one is preserved by perturbation, and is distinguished by having that value of θ – which is what the proof of Theorem 11.1 requires.) Theorem 11.1 therefore implies that there exists $\beta \in B(c, c)$ such that

$$x_{T(c)}(t) \equiv \beta^* x_{T(c)}(t + \theta T) \quad \forall \theta \in \mathbf{S}^1.$$

Since $B(c, c)$ is a finite group, $\beta^k = \text{id}$ for some k . Then

$$x_{T(c)}(t) \equiv (\beta^*)^k x_{T(c)}(t + k\theta T) \equiv x_{T(c)}(t + k\theta T) \quad \forall \theta \in \mathbf{S}^1.$$

Now $k\theta$ ranges over the whole of \mathbf{S}^1 since θ does, so by Lemma 12.2, $x_{T(c)}$ is in equilibrium. In particular, any node $d \in T(c)$ is in equilibrium. As noted at the start

of the proof, node d is rigidly steady, so we can iterate. Since \mathcal{G} is finite, after finitely many steps we deduce, using Remark 10.1, that any given upstream node must be rigidly steady. ■

Theorem 12.4. *The Full Oscillation Property holds for all strongly hyperbolic periodic orbits.*

Proof. Suppose not. Then some transitive network has a hyperbolic periodic orbit that is rigidly steady at some node c . But in a transitive network, every node is upstream from c . Therefore, every node is steady, so the state is not periodic – contradiction. ■

13. Cyclic groups of automorphisms and the H/K Theorem

It is known that if conjectures (a, b, c, d) are valid for a network \mathcal{G} , which we have proved is the case for strongly hyperbolic periodic orbits, then there are important consequences for the combinatorial structure of \mathcal{G} . In Golubitsky et al. [34] and [73] it is proved that, on the assumption that these conjectures are valid for a given network \mathcal{G} , there is a natural network analogue of the H/K Theorem of Buono and Golubitsky [14]; see also Golubitsky and Stewart [36] and Golubitsky et al. [31].

In Section 1.1 we stated that for equivariant dynamics with symmetry group Γ , the H/K Theorem characterises, for each Γ , the possible spatiotemporal patterns of periodic states \mathbb{X} that can occur for suitable Γ -equivariant ODEs. This characterisation is stated in terms of the the *spatiotemporal symmetry group* H , which fixes \mathbb{X} setwise, and the *spatial symmetry group* $K \subseteq H$, which fixes \mathbb{X} pointwise. It is easy to prove that $K \triangleleft H$ and (when Γ is finite) the quotient group H/K is cyclic and corresponds to phase shifts through certain rational multiples of the period. The synchrony and phase patterns determined by such subgroups H and K are always rigid [36, Corollary 3.7].

For a network, the natural analogue of K is a balanced colouring \bowtie , determined by the synchrony pattern, and the natural analogue of H is the phase pattern. The crucial feature in common with the equivariant H/K Theorem is proved in [73] under the assumption that \mathcal{G} has the Rigid Phase Property. Namely, when \mathcal{G} is transitive, the existence of a rigid phase pattern implies that the quotient network \mathcal{G}/\bowtie of \mathcal{G} by the synchrony colouring \bowtie has a cyclic symmetry group. Moreover, this symmetry group implies the existence of a discrete rotating wave with the corresponding phase pattern. So rigid phase relations on an arbitrary network occur if and only if they come from a cyclic group symmetry on \mathcal{G}/\bowtie .

The same proof works if we assume only the Local Rigid Phase Property on some interval J . We can also pass from a specific rigid phase relation to the entire phase pattern in the sense of [72], to establish:

Theorem 13.1. *Let \mathcal{G} be a transitive network, and assume that \mathcal{G} has the (global) Rigid Phase Property. Suppose that there is a rigid phase pattern corresponding to a T -periodic state $x(t)$ of an admissible ODE. Then there is a balanced colouring \bowtie of \mathcal{G} and a symmetry γ of the quotient network \mathcal{G}_{\bowtie} , generating a cyclic group $\Gamma \cong \mathbb{Z}_k$, such that*

- (a) $x_c(t) \equiv x_d(t)$ if and only if $c \bowtie d$.
- (b) For each pair (c, d) of nodes that are rigidly phase-related

$$x_c(t) \equiv x_d(t + \theta_{cd}T)$$

we have $\theta_{cd} = \frac{m}{k}$ for some integer m . Moreover, $\gamma\bar{c} = \bar{d}$ where \bar{c}, \bar{d} are the quotient nodes corresponding to c, d . ■

The converse is true with node spaces \mathbb{R}^m for which $m \geq 2$, by results of Josić and Török [52]. The precise characterisation of H, K pairs for network ODEs remains open when $m = 1$ or node spaces are \mathbf{S}^1 (phase oscillators).

Corollary 13.2. *The above theorem holds if \mathbb{X} is strongly hyperbolic.* ■

Transitivity of \mathcal{G} is required in Theorem 13.1 because nodes can be removed from a feedforward network without affecting the phase pattern or its rigidity, but destroying the cyclic symmetry. This issue is raised in Stewart and Parker [73, Section 3.1]. Golubitsky et al. [34, Theorem 1.4] prove a similar result. Finally, Golubitsky et al. [34, Section 7] provide a detailed and thorough discussion of rigid phase patterns in non-transitive networks. Here it may be necessary to *complete* the network \mathcal{G} by adding further downstream nodes and arrows to obtain a larger network \mathcal{G}^* . This extension does not affect the dynamics on \mathcal{G} because nodes in \mathcal{G} force those in $\mathcal{G}^* \setminus \mathcal{G}$, but it restores cyclic group symmetry.

14. Local Rigidity Properties for all 1- and 2-colourings

We now remove the hypothesis of strong hyperbolicity in some special cases. These results are new and add evidence in support of the Rigidity Conjectures.

A colouring with k colours is called a k -colouring. We prove the first three Rigidity Conjectures (RIC, RSC, RPC) for 1- and 2-colourings by proving that strong hyperbolicity is generic. For the Rigid Synchrony Conjecture, the number of colours refers to the number of synchrony classes. For the Rigid Phase Conjecture, the number of colours refers to the number of node waveforms that are the same up to a phase shift. That is, $c \bowtie d$ whenever $x_c(t) \equiv x_d(t + \theta T)$ in the above notation. If the Rigid Phase Property holds, this colouring is balanced. We exclude the Full

Oscillation Conjecture (FOC) because this is not associated with a specific number of colours.

We will prove, for any network \mathcal{G} :

- Theorem 14.1.** (a) *The Rigid Input, Synchrony, and Phase Properties hold for any 1-colouring of any network.*
- (b) *The Rigid Input, Synchrony, and Phase Properties hold for any 2-colouring of any network.*

The proof is deferred to Sections 14.2 and 14.4.

Corollary 14.2. *The Rigid Input, Synchrony, and Phase Properties as well as the Full Oscillation Property hold for any network with 1, 2, or 3 nodes.* ■

Here we include the Full Oscillation Property because this can be defined for a specific number of nodes.

These results are new. They add evidence supporting the Rigidity Conjectures. In particular, if a counterexample exists, it must have at least 4 nodes and the synchrony pattern must involve at least 3 colours. Moreover, the periodic orbit \mathbb{X} must fail to be strongly hyperbolic. Indeed, for any unbalanced synchrony pattern \bowtie^0 , after any small perturbation, each $\mathbb{X}^{\mathcal{R}}$ is the limit of a continuum of periodic orbits of the induced ODE for \mathcal{R} , containing periodic orbits distinct from $\mathbb{X}^{\mathcal{R}}$ that meet any neighbourhood of the point $x^0(t_0)$.

14.1. ODE-equivalence

The proof of Theorem 14.1 depends on the concept of ODE-equivalence [22, 57]. We summarise the definition and basic properties.

A fundamental feature of network dynamics is that for a given choice of node spaces P_c and total state space $P = \prod P_c$, each network \mathcal{G} determines a class of differential equations on P defined by the admissible maps. There is a bijection between network diagrams and these ‘admissible classes’. Despite this, it was pointed out in [57] that networks with different diagrams can define the same *space* of admissible maps. Such networks are said to be ODE-equivalent, because any admissible ODE for one of them can be interpreted as an admissible ODE for the other, and the same goes for the solutions of the ODEs.

Nontrivial ODE-equivalence can occur because of the technical but vital distinction between a component $f_c : P \rightarrow P_c$ and the associated $\hat{f}_c : P_c \times P_{T(c)} \rightarrow P_c$; see Definition 4.5 (b). Although \hat{f}_c determines f_c uniquely, different choices of \hat{f}_c can determine the same f_c . Admissible classes use specific *presentations* \hat{f}_c of the components of the map, whose domain lists the tail nodes of arrows. The space of

admissible maps $P \rightarrow P$ is determined by the f_c , and this difference creates the ambiguity.

Remark 14.3. Properties of specific orbits (such as periodicity, hyperbolicity, and synchrony and phase patterns) are preserved when passing to an ODE-equivalent network.

Example 14.4. Figure 6 shows a simple example of ODE-equivalence, discussed briefly in [43]. In \mathcal{G}_1 both nodes have the same node-type, and similarly for \mathcal{G}_2 . Suppose that the state space for all four nodes is \mathbb{R}^k . The admissible maps for \mathcal{G}_1 have the form

$$F(x_1, x_2) = (f(x_1, x_2), f(x_2, x_1))$$

where $f : \mathbb{R}^k \times \mathbb{R}^k \rightarrow \mathbb{R}^k$ is any smooth map. The admissible maps for \mathcal{G}_2 have the form

$$G(x_1, x_2) = (g(x_1, x_1, x_2), g(x_2, x_2, x_1))$$

where $g : \mathbb{R}^k \times \mathbb{R}^k \times \mathbb{R}^k \rightarrow \mathbb{R}^k$ is any smooth map.

It is now easy to see that the set $\{G\}$ of all G is the same as the set $\{F\}$ of all F . Namely, given f we can define $g(u, v, w) = f(u, w)$ so that $\{G\} \subseteq \{F\}$. Conversely, given g we can define $f(u, v) = g(u, v, v)$, so that $\{F\} \subseteq \{G\}$. Therefore \mathcal{G}_1 and \mathcal{G}_2 are ODE-equivalent.



Figure 6. Two 2-node networks with different network topologies that define the same space of admissible maps. *Left:* \mathcal{G}_1 . *Right:* \mathcal{G}_2 .

Definition 14.5. Networks \mathcal{G}_1 and \mathcal{G}_2 with the same sets of nodes (up to the numbering of the nodes) are *ODE-equivalent*, written $\mathcal{G}_1 \overset{\text{ODE}}{\sim} \mathcal{G}_2$, if, for the same choices of node spaces, \mathcal{G}_1 and \mathcal{G}_2 have the same spaces of admissible maps.

They are *linearly equivalent* if they have the same spaces of linear admissible maps.

Recall that the adjacency matrix A for a given arrow-type is the matrix (a_{ij}) for which a_{ij} is the number of arrows e of that type such that $\mathcal{H}(e) = i$ and $\mathcal{T}(e) = j$. When arrow-types are irredundant (see Section 4.2) and node spaces are 1-dimensional, it is easy to see that the space of linear admissible maps is spanned by the adjacency matrices for the separate arrow types, including internal node ‘arrows’ distinguished by node type. Linear equivalence then becomes ‘the adjacency matrices span the same space’.

The key result on ODE-equivalence is [22, Theorem 7.1]:

Theorem 14.6. *Two networks are ODE-equivalent if and only if they are linearly equivalent. Moreover, when verifying linear equivalence we may assume that all node spaces are 1-dimensional.* ■

This theorem reduces ODE-equivalence to routine linear algebra, applied to the space spanned by the adjacency matrices.

Example 14.7. Consider the networks \mathcal{G}_1 and \mathcal{G}_2 in Figure 6. The linear admissible maps for \mathcal{G}_1 are spanned by

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

(adjacency matrices for internal node arrows, solid arrows, respectively). Those for \mathcal{G}_2 are spanned by

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

(internal node arrows, solid arrows, dashed arrows, respectively). Clearly these spaces are the same, so $\mathcal{G}_1 \stackrel{\text{ODE}}{\sim} \mathcal{G}_2$.

14.2. 1-colour synchrony

The proof of the first three Rigidity Conjectures for 1-colourings (that is, when the orbit is fully synchronous for the RSC and RIC, and has a single waveform up to phase for the RPC) is now straightforward:

Proof of Theorem 14.1 (a). The result is true for equilibria by [3, 43, 69], so we may assume \mathbb{X} is not an equilibrium. Assume for a contradiction that \mathcal{G} is not homogeneous. The periodic orbit \mathbb{X} is fully synchronous; that is, $x_c(t) \equiv x_d(t)$ for all nodes c, d , so the colouring $\bowtie = \{1, 2, \dots, n\}$. Let $\mathcal{R} = \{1\}$.

The quasi-quotient $\mathcal{G}^{\mathcal{R}}$ is a 1-node network, on which the induced ODE is $\dot{u} = f_1(u, \dots, u)$. We claim that $\mathcal{G}^{\mathcal{R}}$ is ODE-equivalent to a 1-node network with no arrows (other than the internal node ‘arrow’). This follows from Theorem 14.6, because all adjacency matrices are integer multiple of the identity. Admissible maps are therefore arbitrary smooth functions of u , so the standard Kupka–Smale Theorem implies that $\mathcal{G}^{\mathcal{R}}$ is a Kupka–Smale network, which implies the Rigidity Properties. ■

The proof of the Rigidity Conjectures for 2-colourings requires further preparation, done in the next subsection. We complete the proof in Section 14.4.

14.3. Classification of 2-node networks up to ODE-equivalence

To deal with 2-colourings we first classify all possible 2-node networks up to ODE-equivalence. This is straightforward, but seems not to be in the literature, so we give details.

Theorem 14.8. *Every 2-node network is ODE-equivalent to precisely one of the networks illustrated in Figure 7.*

Proof. By Theorem 14.6 it is enough to classify 2-node networks up to linear equivalence, assuming all node spaces are 1-dimensional. The classification therefore reduces to considering the vector spaces spanned by the adjacency matrices for each arrow-type (assuming irredundancy). There are three main cases:

- (a) The network \mathcal{G} is disconnected.
- (b) The network \mathcal{G} is feedforward; that is, it is connected and has more than one transitive component.
- (c) The network \mathcal{G} is transitive.

For a 2-node network, transitivity is equivalent to being all-to-all connected. Throughout, let I be the 2×2 identity matrix.

Case (a): This is trivial. Either \mathcal{G} is homogeneous, or not. If it is homogeneous then all adjacency matrices are multiples of I , so \mathcal{G} is ODE-equivalent to network (1) in the figure. If not, all adjacency matrices are positive integer multiples of either of the following two:

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

Now \mathcal{G} is ODE-equivalent to network (2) in the figure.

Case (b): Clearly any network ODE-equivalent to a feedforward network is also feedforward (all adjacency matrices have a common block-triangular structure). Re-




Figure 7. Classification of 2-node networks up to ODE-equivalence. In case (8) we require $p, q > 0$ and $\gcd(p, q) = 1$.

numbering if necessary we can assume that node 1 is the upstream node. Either \mathcal{G} is homogeneous, or not. If it is homogeneous then the internal arrow-type has adjacency matrix a multiple of I . Every other arrow-type also has as adjacency matrix a multiple of I , or it has a adjacency matrix of the form

$$\begin{bmatrix} a+b & 0 \\ a & b \end{bmatrix}$$

with $a \geq 1$, by homogeneity. By connectedness, at least one such arrow-type has $a > 0$. Subtracting bI we obtain

$$a \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}.$$

So the span of the adjacency matrices is that of

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix},$$

however many arrow-types there may be. (Different types may have different entries a, b but the same statement holds.) Therefore \mathcal{G} is ODE-equivalent to network (6) in the figure.

If \mathcal{G} is not homogeneous then the internal arrow-types have adjacency matrices

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

Any other arrow-type either has as adjacency matrix a multiple of one of these, or its adjacency matrix has the form

$$\begin{bmatrix} a & 0 \\ c & d \end{bmatrix}$$

with at least one arrow-type for which $c > 0$. We can subtract the diagonal terms without changing the span of the adjacency matrices, and then divide by c . Therefore the span of the adjacency matrices is that of

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

and \mathcal{G} is ODE-equivalent to network (5) in the figure.

Case (c): Either \mathcal{G} is homogeneous or not. If it is homogeneous then the internal arrow-type has adjacency matrix I . Any other arrow-type is either a self-loop whose adjacency matrix is a multiple of I , or its adjacency matrix has the form

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

where $a + b = c + d$. Swapping nodes if necessary we may assume that $a \leq d$. Subtracting aI , the above matrix becomes

$$\begin{bmatrix} 0 & b \\ c & b - c \end{bmatrix}$$

and by scaling, we may assume that $\gcd(b, c) = 1$.

The largest possible span is the set of all matrices

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

where $a + b = c + d$, which has dimension 3. Therefore there are two possibilities: the span is 3-dimensional or it is 2-dimensional. If it is 3-dimensional it is spanned by

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$$

which is case (7) in the figure. If it is 2-dimensional, then every adjacency matrix is either a multiple of the identity or a multiple of some

$$\begin{bmatrix} 0 & b \\ c & b - c \end{bmatrix}$$

with $\gcd(b, c) = 1$. Let $b = p + q, c = q$, so $b - c = p$.

If $p = 0$ we can scale q to equal 1, giving case (4) in the figure. If $q = 0$ we can scale p to equal 1, giving case (6) in the figure (again). Otherwise $p, q > 0$ and $\gcd(p, q) = 1$. Two such matrices are linearly dependent modulo the identity if and only if they are the same. So now we get a 2-parameter family of networks, with parameters p, q such that $\gcd(p, q) = 1$, which is case (8) in the figure.

Finally, suppose \mathcal{G} is not homogeneous. The nodes are not input equivalent, so by irredundancy they have different types. There must be at least one arrow in each direction, and these have different arrow-types. The node-types give adjacency matrices

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

Any self-loops give adjacency matrices that are multiples of these, and can be deleted. Arrows from node 1 to node 2 and the reverse give adjacency matrices of the form

$$\begin{bmatrix} 0 & 0 \\ a & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & b \\ 0 & 0 \end{bmatrix}$$

respectively, where $a, b \neq 0$ for at least one arrow-type each way. Thus the adjacency matrices span the 4-dimensional space of all 2×2 matrices. The simplest such network is number (3), and all others are ODE-equivalent to it since their adjacency matrices span the same 4-dimensional space. ■

14.4. Rigidity Conjectures for 2-colourings

Some (non-standard) terminology is useful:

Definition 14.9. A periodic orbit \mathbb{Y} of a dynamical system $\dot{x} = f(x)$ is *locally Kupka–Smale* if there exist $\delta_1 > \delta_2 > 0$ and tubular neighbourhoods $\mathbb{Y}_{\delta_1} \supseteq \mathbb{Y}_{\delta_2} \supseteq \mathbb{Y}$ such that

- (a) the flow of f maps \mathbb{Y}_{δ_2} into \mathbb{Y}_{δ_1} ,
- (b) there exists an arbitrarily small perturbation p of f $\delta_3, \delta_4 > 0$ with $\delta_3 \leq \delta_1$ and $\delta_4 \leq \delta_2$ such that the flow of $f + p$ maps \mathbb{Y}_{δ_4} into \mathbb{Y}_{δ_3} ,
- (c) every periodic orbit for $f + p$ that is contained in \mathbb{Y}_{δ_4} is hyperbolic.

Lemma 6.6 states that every periodic orbit \mathbb{X} for a general dynamical system is locally Kupka–Smale. Below we use this to deduce that $\mathbb{X}^{\mathcal{R}}$ is locally Kupka–Smale for certain networks \mathcal{G} . Before embarking on the proof of Theorem 14.1 (b), we explain why it reduces to a case-by-case analysis showing that all networks (3)–(7) in Figure 7 are Kupka–Smale or locally Kupka–Smale.

If \bowtie contains just two colour classes, all induced ODEs are defined by sets of representatives \mathcal{R} of cardinality 2. The corresponding quasi-quotients are 2-node networks. The Rigidity Conjectures then follow from the previous analysis, provided we can prove that every 2-node quasi-quotient network $\mathcal{G}^{\mathcal{R}}$ occurring in the proof has the stable isolation property. In particular this follows if $\mathcal{G}^{\mathcal{R}}$ is Kupka–Smale or locally Kupka–Smale. By Remark 14.3 the stable isolation property and the (local and global) Kupka–Smale properties are preserved by ODE-equivalence. It therefore suffices to prove that all networks in Figure 7 except the disconnected networks (1) and (2) are Kupka–Smale or locally Kupka–Smale. (By Section 6.5, networks (1) and (2) lack the stable isolation property, but we can deal with them by a different, trivial, argument.)

Proof of Theorem 14.1 (b). *Case (1):* This network is disconnected. It arises only if the original network \mathcal{G} has no arrows whose head and tail have different colours. This implies that \mathcal{G} is the disjoint union of two networks, each having all nodes of the same colour. Such a pattern is automatically balanced.

Case (2): This network is also disconnected, and the same argument applies.

Case (3): Admissible ODEs for this network are arbitrary dynamical systems on $P_1 \times P_2$. The standard Kupka–Smale Theorem [55, 62, 66] therefore applies.

Case (4): Admissible ODEs for this network are arbitrary \mathbb{Z}_2 -equivariant dynamical systems on $P_1 \times P_1$. The equivariant Kupka–Smale Theorem of Field [27] for compact group actions therefore applies.

Case (5): This is a general ‘forced’ dynamical system on $P_1 \times P_2$. We prove it is locally Kupka–Smale, which implies the stable isolation property. Admissible ODEs have the ‘forced’ or ‘skew product’ form

$$\dot{x}_1 = f(x_1), \quad \dot{x}_2 = g(x_2, x_1). \tag{14.1}$$

If x_1 is steady, with equilibrium $x_1 = \alpha$, we can perturb f to make α a hyperbolic fixed point. Then the second component is $\dot{x}_2 = g(x_2, \alpha)$, which is an arbitrary ODE in x_2 . By the Kupka–Smale Theorem we can perturb g to make x_2 hyperbolic. Now \mathbb{X} is hyperbolic.

If x_1 oscillates, let the minimal period of \mathbb{X} be T . Let $\Sigma_1 \subseteq P_1$ be a Poincaré section transverse to $\{x_1(t)\}$ at $t = 0$, and define $\Sigma_2 = \{x_1(t_0)\} \times \Sigma_1$ as in Figure 8. The Poincaré map $\pi : \Sigma_2 \rightarrow \Sigma_2$ has the form $\pi(u, v) = (\pi_1(u), \pi_2(u, v))$ since the ODE is feedforward. The derivative is

$$D\pi = \begin{bmatrix} D_1\pi_1 & 0 \\ D_1\pi_2 & D_2\pi_2 \end{bmatrix}$$

and we want to make $D\pi|_{\mathbb{X}^*}$ hyperbolic, where $\mathbb{X}^* = (x_1(t_0), x_2(t_0))$ is the fixed point of π .

Perturbing f and applying the Kupka–Smale Theorem we can make $D_1\pi_1$ hyperbolic on P_1 . The flow of (14.1) leaves the cylinder $\mathbb{X}^1 \times P_2 \cong \mathbf{S}^1 \times P_2$ invariant. The component $\pi_2(x_1(t_0), x_2)$ is a local diffeomorphism near $(x_1(t_0), x_2(t_0)) \in \mathbb{X}$.




Figure 8. Poincaré section for (14.1).

A suitable perturbation of g perturbs π_2 to any nearby diffeomorphism (consider a suspension). By the Kupka–Smale Theorem for diffeomorphisms (discrete dynamics) we can make $D_2\pi_2$ hyperbolic at \mathbb{X}^* . Therefore \mathbb{X} is hyperbolic.

Case (6): This is similar, but the network is homogeneous, so minor modifications are required. Admissible ODEs have the form

$$\dot{x}_1 = f(x_1, x_1), \quad \dot{x}_2 = f(x_2, x_1).$$

The subspace $\Delta = \{(u, u)\}$ is flow-invariant. If \mathbb{X} meets Δ then it remains in Δ for all time, the ODE reduces to $\dot{u} = f(u, u)$, and \mathbb{X} can be made hyperbolic by the Kupka–Smale Theorem applied to f . If \mathbb{X} does not meet Δ then on a small tubular neighbourhood of \mathbb{X} , admissibility imposes no constraints: $f(x_1, x_1)$ and $f(x_1, x_2)$ are independent functions. We can now argue as in case (5), perturbing near \mathbb{X} .

Case (7): Admissible ODEs have the form

$$\dot{x}_1 = f(x_1, x_2, x_2), \quad \dot{x}_2 = f(x_2, x_1, x_2)$$

and $\Delta = \{(u, u)\}$ is flow-invariant. By similar reasoning, either $\mathbb{X} \subseteq \Delta$ and the result follows, or we can reduce this case to a general dynamical system in some tubular neighbourhood of \mathbb{X} . To do so we must show that $\mathbb{X}^1 \cap \mathbb{X}^2 = \emptyset$, where

$$\begin{aligned} \mathbb{X}^1 &= \{(x_1(t), x_2(t), x_2(t)) : t \in \mathbb{R}\}, \\ \mathbb{X}^2 &= \{(x_2(t), x_1(t), x_2(t)) : t \in \mathbb{R}\}. \end{aligned}$$

Suppose that $\mathbb{X}^1 \cap \mathbb{X}^2 \neq \emptyset$. Then there exist times s, t such that

$$(x_1(t), x_2(t), x_2(t)) = (x_2(s), x_1(s), x_2(s)).$$

Therefore $x_1(t) = x_2(s) = x_2(t)$, so $\mathbb{X} \cap \Delta \neq \emptyset$, and $\mathbb{X} \subseteq \Delta$. Now tubular neighbourhoods of \mathbb{X}^1 and \mathbb{X}^2 are disjoint, so we can choose f independently on these neighbourhoods without destroying admissibility. The analysis then reduces to case (3).

Case (8): This is similar to case (7). Admissible ODEs have the form

$$\begin{aligned} \dot{x}_1 &= f(x_1, \underbrace{x_2, \dots, x_2}_{p+q}), \\ \dot{x}_2 &= f(x_2, \underbrace{x_1, \dots, x_1}_p, \underbrace{x_2, \dots, x_2}_q). \end{aligned}$$

Define

$$\begin{aligned} \mathbb{X}^1 &= \{(x_1(t), \underbrace{x_2(t), \dots, x_2(t)}_{p+q}) : t \in \mathbb{R}\}, \\ \mathbb{X}^2 &= \{(x_2, \underbrace{x_1(t), \dots, x_1(t)}_p, \underbrace{x_2(t), \dots, x_2(t)}_q) : t \in \mathbb{R}\}. \end{aligned}$$

Again, we can prove that $\mathbb{X}^1 \cap \mathbb{X}^2 = \emptyset$ by comparing suitable entries of $\mathbb{X}^1(t)$ and $\mathbb{X}^2(s)$ and deducing that $x_1(t) = x_2(s) = x_2(t)$. The rest is as before. ■

14.5. Final remarks

The proof in this paper of the Rigidity Conjectures for strongly hyperbolic periodic orbits uses three unorthodox methods, which together reveal a link with network analogues of the Kupka–Smale Theorem. These methods are:

- (a) The use of overdetermined ODEs in which some components are formally inconsistent with others.
- (b) A construction analogous to the usual quotient network by a balanced colouring, applied to a colouring that is not balanced.
- (c) Construction of an admissible perturbation p that leaves the perturbed periodic orbit $\tilde{\mathbb{X}}$ unchanged. This avoids the main obstacle to proving the Rigidity Conjectures: keeping track of $\tilde{\mathbb{X}}$. Here this is not a problem because $\tilde{\mathbb{X}} = \mathbb{X}$.

This combination works because the rigidity assumption causes enough structure to be preserved for the formal inconsistency (a) to contradict rigidity of the local synchrony pattern for \mathbb{X} . In the current state of knowledge, this contradiction relies on strong hyperbolicity, which is closely related to the Kupka–Smale Theorem and possible network analogues.

A more specific version of strong hyperbolicity is logically equivalent to the Rigid Synchrony Property (local or global). Definition 6.5 is stated for all colourings, but we use it only for colourings determined by a local rigid synchrony pattern. We can therefore weaken the definition of strong hyperbolicity by considering only these colourings. This weaker version still implies the Rigid Synchrony Property, with the same proof. Conversely, the Rigid Synchrony Property implies that any local rigid synchrony pattern is balanced, so its synchrony space is flow-invariant. This implies that the Floquet multipliers of the induced periodic orbit $\mathbb{X}^{\mathcal{R}}$ are a subset of those of \mathbb{X} . Since \mathbb{X} is assumed hyperbolic, so is $\mathbb{X}^{\mathcal{R}}$.

We have shown that if a counterexample to the Rigidity Conjectures exists, its dynamics must be remarkably degenerate from the viewpoint of general dynamical systems theory. Specifically, for some unbalanced colouring and all small admissible perturbations, and for any set of representatives \mathcal{R} of that colouring, the solution $\mathbb{X}^{\mathcal{R}}$ of the induced ODE is non-isolated in an extreme manner: it is included in a continuum of distinct periodic orbits. It is difficult to see how the constraints imposed by network topology could create such degeneracy rigidly. Be that as it may, the focus for proving the Rigidity Conjectures now shifts towards network analogues of the Kupka–Smale Theorem – at least until some alternative method is found.

References

- [1] R. Abraham and J. E. Marsden, *Foundations of mechanics*. Benjamin/Cummings Publishing Co., Inc., Advanced Book Program, Reading, Mass., 1978 Zbl [0393.70001](#) MR [515141](#)
- [2] R. Abraham, J. E. Marsden, and T. S. Raşiu, *Manifolds, tensor analysis, and applications*. Global Analysis Pure and Applied: Series B 2, Addison-Wesley Publishing Co., Reading, Mass., 1983 MR [697563](#)
- [3] J. W. Aldis. *On Balance*, PhD Thesis, University of Warwick, 2010
- [4] F. Antoneli and I. Stewart, Symmetry and synchrony in coupled cell networks. I. Fixed-point spaces. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.* **16** (2006), no. 3, 559–577 Zbl [1145.37019](#) MR [2228833](#)
- [5] F. Antoneli and I. Stewart, Symmetry and synchrony in coupled cell networks. II. Group networks. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.* **17** (2007), no. 3, 935–951 Zbl [1141.37392](#) MR [2324988](#)
- [6] F. Antoneli and I. Stewart, Symmetry and synchrony in coupled cell networks. III. Exotic patterns. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.* **18** (2008), no. 2, 363–373 Zbl [1143.37302](#) MR [2404307](#)
- [7] V. I. Arnol'd, *Mathematische Methoden der klassischen Mechanik*. VEB Deutscher Verlag der Wissenschaften, Berlin, 1988 MR [1007828](#)
- [8] D. K. Arrowsmith and C. M. Place, *An introduction to dynamical systems*. Cambridge University Press, Cambridge, 1990 Zbl [0702.58002](#) MR [1069752](#)
- [9] I. Belykh, V. Belykh, K. Nevidin, and M. Hasler, Persistent clusters in lattices of coupled nonidentical chaotic systems. *Chaos* **13** (2003), no. 1, 165–178 Zbl [1080.37525](#) MR [1964970](#)
- [10] I. Belykh and M. Hasler, Mesoscale and clusters of synchrony in networks of bursting neurons. *Chaos* **21** (2011), no. 1, 016106 Zbl [1345.92035](#) MR [2808181](#)
- [11] V. N. Belykh, I. V. Belykh, and M. Hasler, Hierarchy and stability of partially synchronous oscillations of diffusively coupled dynamical systems. *Phys. Rev. E (3)* **62** (2000), no. 5, part A, 6332–6345 MR [1796442](#)
- [12] S. Boccaletti, L. M. Pecora, and A. Pelaez, A unifying framework for synchronization of coupled dynamical systems. *Phys. Rev. E* **63** (2001), no. 6, article no. 066219
- [13] R. Brown, From groups to groupoids: a brief survey. *Bull. London Math. Soc.* **19** (1987), no. 2, 113–134 Zbl [0612.20032](#) MR [872125](#)
- [14] P.-L. Buono and M. Golubitsky, Models of central pattern generators for quadruped locomotion. I. Primary gaits. *J. Math. Biol.* **42** (2001), no. 4, 291–326 Zbl [1039.92007](#) MR [1834105](#)
- [15] R. J. Butera Jr., J. Rinzal, and J. C. Smith, Models of respiratory rhythm generation in the pre-Bötzing complex I: Bursting pacemaker neurons. *J. Neurophysiol.* **82** (1999), 382–397
- [16] R. J. Butera Jr., J. Rinzal, and J. C. Smith, Models of respiratory rhythm generation in the pre-Bötzing complex II: Populations of coupled pacemaker neurons. *J. Neurophysiol.* **82** (1999), 398–415

- [17] R. Campos, V. Matos, and C. Santos, Hexapod locomotion: a nonlinear dynamical systems approach. *IECON 2010 - 36th Annual Conference on IEEE Industrial Electronics Society* (2010), 1546–1551
- [18] J. D. Chambers, E. A. Thomas, and C. Bornstein, Mathematical modelling of enteric neural motor patterns. *Proc. Austral. Physiol. Soc.* **44** (2013), 75–84
- [19] R. Curtu, Singular Hopf bifurcations and mixed-mode oscillations in a two-cell inhibitory neural network. *Phys. D* **239** (2010), no. 9, 504–514 Zbl [1196.37088](#) MR [2595817](#)
- [20] B. A. Davey and H. A. Priestley, *Introduction to lattices and order*. Camb. Math. Textb., Cambridge University Press, Cambridge, 1990 Zbl [0701.06001](#) MR [1058437](#)
- [21] A. P. S. Dias and I. Stewart, Symmetry groupoids and admissible vector fields for coupled cell networks. *J. London Math. Soc. (2)* **69** (2004), no. 3, 707–736 Zbl [1049.37009](#) MR [2050042](#)
- [22] A. P. S. Dias and I. Stewart, Linear equivalence and ODE-equivalence for coupled cell networks. *Nonlinearity* **18** (2005), no. 3, 1003–1020 Zbl [1079.37010](#) MR [2134081](#)
- [23] C. Diekman, M. Golubitsky, T. McMillen, and Y. Wang, Reduction and dynamics of a generalized rivalry network with two learned patterns. *SIAM J. Appl. Dyn. Syst.* **11** (2012), no. 4, 1270–1309 Zbl [1263.37068](#) MR [3022067](#)
- [24] C. O. Diekman and M. Golubitsky, Network symmetry and binocular rivalry experiments. *J. Math. Neurosci.* **4** (2014), Art. 12 Zbl [1321.92044](#) MR [3202310](#)
- [25] C. O. Diekman, M. Golubitsky, and Y. Wang, Derived patterns in binocular rivalry networks. *J. Math. Neurosci.* **3** (2013), Art. 6 Zbl [1291.92032](#) MR [3083268](#)
- [26] M. Field, Combinatorial dynamics. *Dyn. Syst.* **19** (2004), no. 3, 217–243 Zbl [1058.37008](#) MR [2084549](#)
- [27] M. J. Field, Equivariant dynamical systems. *Trans. Amer. Math. Soc.* **259** (1980), no. 1, 185–205 Zbl [0447.58029](#) MR [561832](#)
- [28] P. Gandhi, M. Golubitsky, C. Postlethwaite, I. Stewart, and Y. Wang, Bifurcations on fully inhomogeneous networks. *SIAM J. Appl. Dyn. Syst.* **19** (2020), no. 1, 366–411 Zbl [1443.34039](#) MR [4059374](#)
- [29] J. Gjorgjieva, J. Berni, J. F. Evers, and S. J. Egle, Neural circuits for peristaltic wave propagation in crawling *Drosophila* larvae: analysis and modeling. *Front. Comput. Neurosci.* **7** (2013), DOI [10.3389/fncom.2013.00024](#)
- [30] A. F. Glova, Phase locking of optically coupled lasers. *Quantum Electronics* **33** (2003), 283–306
- [31] M. Golubitsky, L. Matamba Messi, and L. E. Spardy, Symmetry types and phase-shift synchrony in networks. *Phys. D* **320** (2016), 9–18 Zbl [1364.34045](#) MR [3470657](#)
- [32] M. Golubitsky, M. Nicol, and I. Stewart, Some curious phenomena in coupled cell networks. *J. Nonlinear Sci.* **14** (2004), no. 2, 207–236 Zbl [1136.37359](#) MR [2041431](#)
- [33] M. Golubitsky, D. Romano, and Y. Wang, Network periodic solutions: full oscillation and rigid synchrony. *Nonlinearity* **23** (2010), no. 12, 3227–3243 Zbl [1216.34029](#) MR [2739423](#)
- [34] M. Golubitsky, D. Romano, and Y. Wang, Network periodic solutions: patterns of phase-shift synchrony. *Nonlinearity* **25** (2012), no. 4, 1045–1074 Zbl [1245.34043](#) MR [2904269](#)

- [35] M. Golubitsky and D. G. Schaeffer, *Singularities and groups in bifurcation theory. Vol. I.* Appl. Math. Sci. 51, Springer, New York, 1985 Zbl [0607.35004](#) MR [771477](#)
- [36] M. Golubitsky and I. Stewart, *The symmetry perspective.* Prog. Math. 200, Birkhäuser, Basel, 2002 Zbl [1031.37001](#) MR [1891106](#)
- [37] M. Golubitsky and I. Stewart, Nonlinear dynamics of networks: the groupoid formalism. *Bull. Amer. Math. Soc. (N.S.)* **43** (2006), no. 3, 305–364 Zbl [1119.37036](#) MR [2223010](#)
- [38] M. Golubitsky and I. Stewart, Coordinate changes for network dynamics. *Dyn. Syst.* **32** (2017), no. 1, 80–116 Zbl [1362.37051](#) MR [3606590](#)
- [39] M. Golubitsky and I. Stewart, *Dynamics and Bifurcation in Networks*, SIAM, Philadelphia, to appear
- [40] M. Golubitsky, I. Stewart, P.-L. Buono, and J. J. Collins, A modular network for legged locomotion. *Phys. D* **115** (1998), no. 1-2, 56–72 Zbl [1039.92009](#) MR [1616780](#)
- [41] M. Golubitsky, I. Stewart, J. J. Collins, and P.-L. Buono, Symmetry in locomotor central pattern generators and animal gaits. *Nature* **401** (1999), 693–695
- [42] M. Golubitsky, I. Stewart, and D. G. Schaeffer, *Singularities and groups in bifurcation theory. Vol. II.* Appl. Math. Sci. 69, Springer, New York, 1988 Zbl [0691.58003](#) MR [950168](#)
- [43] M. Golubitsky, I. Stewart, and A. Török, Patterns of synchrony in coupled cell networks with multiple arrows. *SIAM J. Appl. Dyn. Syst.* **4** (2005), no. 1, 78–100 Zbl [1090.34030](#) MR [2136519](#)
- [44] J. Guckenheimer and P. Holmes, *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields.* Appl. Math. Sci. 42, Springer, New York, 1983 Zbl [0515.34001](#) MR [709768](#)
- [45] B. D. Hassard, N. D. Kazarinoff, and Y. H. Wan, *Theory and applications of Hopf bifurcation.* London Math. Soc. Lecture Note Ser. 41, Cambridge University Press, Cambridge, New York, 1981 Zbl [0474.34002](#) MR [603442](#)
- [46] P. J. Higgins, *Notes on categories and groupoids.* Van Nostrand Reinhold Mathematical Studies, No. 32, Van Nostrand Reinhold Co., London, New York, Melbourne, 1971 Zbl [0226.20054](#) MR [0327946](#)
- [47] M. W. Hirsch, C. C. Pugh, and M. Shub, *Invariant manifolds.* Lecture Notes in Math. 583, Springer, Berlin, New York, 1977 Zbl [0355.58009](#) MR [0501173](#)
- [48] M. W. Hirsch and S. Smale, *Differential equations, dynamical systems, and linear algebra.* Pure and Applied Mathematics, Vol. 60, Academic Press [Harcourt Brace Jovanovich, Publishers], New York, London, 1974 Zbl [0309.34001](#) MR [0486784](#)
- [49] K. Itô (ed.), *Encyclopaedic Dictionary of Mathematics*, vol. 1 (2nd ed.), MIT Press, Cambridge, MA, 1993
- [50] E. M. Izhikevich and B. Ermentrout, Phase model. *Scholarpedia* **3** (2008), no. 10, article no. 1487
- [51] R. Joly, Observation and inverse problems in coupled cell networks. *Nonlinearity* **25** (2012), no. 3, 657–676 Zbl [1237.93032](#) MR [2887987](#)
- [52] K. Josić and A. Török, Network architecture and spatio-temporally symmetric dynamics. *Phys. D* **224** (2006), no. 1-2, 52–68 Zbl [1118.34032](#) MR [2301509](#)

- [53] A. Katok and B. Hasselblatt, *Introduction to the modern theory of dynamical systems*. Encyclopedia Math. Appl. 54, Cambridge University Press, Cambridge, 1995
Zbl [0878.58020](#) MR [1326374](#)
- [54] N. Kopell and G. LeMasson, Rhythmogenesis, amplitude modulation, and multiplexing in a cortical architecture. *Proc. Natl. Acad. Sci. USA* (1994) **91**, 10586–10590
- [55] I. Kupka, Contribution à la théorie des champs génériques. *Contrib. Differ. Equations 2* (1963), 457–484 Zbl [0149.41002](#) MR [165536](#)
- [56] P. Lancaster and M. Tismenetsky, *The theory of matrices*. Second edn., Computer Science and Applied Mathematics, Academic Press, Inc., Orlando, FL, 1985 Zbl [0558.15001](#) MR [792300](#)
- [57] M. d. C. A. Leite and M. Golubitsky, Homogeneous three-cell networks. *Nonlinearity* **19** (2006), no. 10, 2313–2363 Zbl [1114.34035](#) MR [2260266](#)
- [58] C. Liu, Q. Chen, and J. Zhang, Coupled van der Pol oscillators utilised as central pattern generators for quadruped locomotion. *IEEE Chinese Control and Decision Conference* (2009), 3677–3682
- [59] S. C. Manrubia, A. S. Mikhailov, and D. H. Zanette, *Emergence of Dynamical Order*. World Scientific, Singapore, 2004 Zbl [1119.34001](#)
- [60] E. Mosekilde, Y. Maistrenko, and D. Postnov, *Chaotic synchronization*. World Scientific Series on Nonlinear Science. Series A: Monographs and Treatises 42, World Scientific Publishing Co., Inc., River Edge, NJ, 2002 MR [1939912](#)
- [61] L. M. Pecora, F. Sorrentino, A. M. Hagerstrom, T. E. Murphy, and R. Roy, Symmetries, cluster synchronization, and isolated desynchronization in complex networks. *Nature Communications* **5** (2014), article no. 4079
- [62] M. M. Peixoto, On an approximation theorem of Kupka and Smale. *J. Differential Equations* **3** (1967), 214–227 Zbl [0153.40901](#) MR [209602](#)
- [63] A. Pogromsky, G. Santoboni, and H. Nijmeijer, Partial synchronization: from symmetry towards stability. *Phys. D* **172** (2002), no. 1–4, 65–87 Zbl [1008.37012](#) MR [1942999](#)
- [64] A. Y. Pogromsky, A partial synchronization theorem. *Chaos* **18** (2008), no. 3, article no. 037107 Zbl [1309.34061](#) MR [2464318](#)
- [65] W. Singer, Neuronal synchrony: a versatile code for the definition of relations. *Neuron* **24** (1999), 49–65
- [66] S. Smale, Stable manifolds for differential equations and diffeomorphisms. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (3)* **17** (1963), 97–116 Zbl [0113.29702](#) MR [165537](#)
- [67] S. Smale, Differentiable dynamical systems. *Bull. Amer. Math. Soc.* **73** (1967), 747–817 Zbl [0202.55202](#) MR [228014](#)
- [68] I. Stewart, The lattice of balanced equivalence relations of a coupled cell network. *Math. Proc. Cambridge Philos. Soc.* **143** (2007), no. 1, 165–183 Zbl [1120.37003](#) MR [2340982](#)
- [69] I. Stewart, Overdetermined constraints and rigid synchrony patterns for network equilibria. *Port. Math.* **77** (2020), no. 2, 163–196 Zbl [1468.34049](#) MR [4163974](#)
- [70] I. Stewart and M. Golubitsky, Symmetric networks with geometric constraints as models of visual illusions. *Symmetry* **11** (2019), article no. 799

- [71] I. Stewart, M. Golubitsky, and M. Pivato, Symmetry groupoids and patterns of synchrony in coupled cell networks. *SIAM J. Appl. Dyn. Syst.* **2** (2003), no. 4, 609–646
Zbl [1089.34032](#) MR [2050244](#)
- [72] I. Stewart and M. Parker, Periodic dynamics of coupled cell networks. I. Rigid patterns of synchrony and phase relations. *Dyn. Syst.* **22** (2007), no. 4, 389–450 Zbl [1162.37312](#)
MR [2377209](#)
- [73] I. Stewart and M. Parker, Periodic dynamics of coupled cell networks. II. Cyclic symmetry. *Dyn. Syst.* **23** (2008), no. 1, 17–41 Zbl [1162.37010](#) MR [2406978](#)
- [74] P. J. Uhlhaas, G. Pipa, B. Lima, L. Melloni, S. Neuenschwander, D. Nikolić, and W. Singer, Neural synchrony in cortical networks: history, concept and current status. *Front. Integr. Neurosci.* **30** (2009), DOI [10.3389/neuro.07.017.2009](#)
- [75] C. van Vreeswijk and D. Hansel, Patterns of synchrony in neural networks with spike adaptation. *Neural Computation* **13** (2001), 959–992 Zbl [1004.92011](#)
- [76] X. F. Wang, Complex networks: topology, dynamics and synchronization. pp. 885–916, 12, 2002 Zbl [1044.37561](#) MR [1913980](#)
- [77] L. Zhang, W. B. Pan, L. Yan, B. Luo, X. Zou, and M. Xu, Cluster synchronization of coupled semiconductor lasers network with complex topology. *IEEE J. Selected Topics in Quantum Electronics* **25** (2019), DOI [10.1109/JSTQE.2019.2913010](#)

Received 24 January 2022.

Ian Stewart

Mathematics Department, University of Warwick, Warwick, United Kingdom;
i.n.stewart@warwick.ac.uk