On Balance

John William Aldis

A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy

Mathematics Institute, University of Warwick

September 2009
Amended September 2010

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgements

I'd like to thank everyone who has contributed to making this thesis possible. Particular thanks go to the staff and students of the University of Warwick Mathematics Institute, especially Carol Wright and Carole Fisher for administrative support. I have profited from very useful conversations with other postgraduate students, especially Tim Sullivan, for which I am very grateful. Thanks also to David Mond for helping me to clarify some of my thoughts during the final months of writing.

I wish to thank my supervision students, for letting me learn mathematical concepts all over again through their own learning experience, and to the supervision system in the Mathematics Institute for giving me the opportunity to teach, and paying me to do so!

Thanks above all to Ian Stewart for his guidance: our meetings and conversations have been illuminating, inspirational, motivational, and above all enjoyable. His guidance has been invaluable, as always. Thanks to Jack Cohen for recommending me to Ian, and for persuading me that I could do research after all.

Innumerable thanks to my parents for their support — emotional and financial — over the last three decades, but especially during the difficult final stages of producing this thesis.

Thanks — and apologies — also to all my friends who have had to put up with me talking about my research to them over the years, especially Sophie, who puts up with more than anyone else.

I would like to extend my thanks also to my examiners, whose constructive criticism has helped to define the final form of this thesis.

This Ph.D. course was funded by the Engineering and Physical Sciences Research Council. The thesis was typeset using LaTeX and BibTeX. Figures were drawn by the author using Adobe Illustrator and Tweakersoft VectorDesigner for Apple Macintosh. The sanity of the author, such as it is, was maintained through judicious use of Twinings Earl Grey tea, for which he is also very grateful.

Declaration

I declare that, to the best of my knowledge, the material contained in this thesis is original and my own work except where otherwise indicated, cited, or commonly known. This work has been carried out under the supervision of Professor Ian Stewart.

The material in this thesis is submitted to the University of Warwick for the degree of Doctor of Philosophy, and has not been submitted to any other university or for any other degree.

1 Introduction

In recent years, there has been some interest in systems which are modelled by networks; in other words, where the structure of the mathematical system is related to the structure of the network.

In terms of systems of differential equations, the relevant systems are called 'admissible' over a network, as described in Golubitsky, Stewart, and Török (2005); Stewart and Parker (December 2007, 2008), for example. In particular, for any system of differential equations, there is a network over which the system is admissible; for any network, there is a family of systems of differential equations which are admissible over this network.

On networks, including networks which model systems of differential equations, there has been interest in 'balanced equivalence relations', which are those equivalence relations on the cells of a network which respect the structure of the network in a well-defined way. Previous work (including Golubitsky et al. (2005); Stewart and Parker (December 2007, 2008)) has investigated these relations; the current author's MSc thesis, Aldis (2005) continued this investigation and showed them to form a partially-ordered set, characterising the maximal element and describing an algorithm for determining this element for arbitrary networks.

This thesis examines some of the outstanding questions about the symmetry properties of admissible systems of differential equations over networks. These questions can be summarised as follows. Given a system of differential equations which is admissible over a network, and a solution of such a system, we may consider the 'pattern of synchrony' of such a solution: that is, the equivalence relation on the cells of the network determined by permanent equality of the quantities represented by those cells. A conjecture of Stewart and others is that all 'realistic' solutions of systems of differential equations have patterns of synchrony which form balanced equivalence relations on the underlying network. Specifically, any pattern of synchrony which is 'rigid', in the sense that small perturbations of the differential equations do not change the pattern of synchrony, forms a balanced equivalence relation. When is this "rigid synchrony conjecture" true?

It has been proved by Golubitsky and coworkers that the possible 'rigid' patterns of synchrony of hyperbolic equilibria over a network are determined by the balanced equivalence relations of the network. Golubitsky's theorem is as follows:

Theorem (Golubitsky et al. (2005) Theorem 7.6)**.** *The equivalence relation $\equiv_{x_0}$ determined by the hyperbolic equilibrium x_0 is rigid if and only if $\equiv_{x_0}$ is balanced.*

We give an alternative proof of this theorem in chapter 7. In fact, the statement we prove is slightly stronger: it only requires the equilibrium to be transverse, not necessarily hyperbolic, and in the case of an equilibrium with a non-balanced pattern of synchrony, it determines the maximal rigid pattern of synchrony of that equilibrium (in this case, Golubitsky's theorem only states that the maximal pattern of synchrony is not rigid).

Rigid Equilibrium Theorem (Theorem 7.2.3)**.** *Let F be an admissible system over a network $\mathcal{N}$, with some transverse equilibrium x^*. The rigid pattern of synchrony of x^* is the maximal balanced equivalence relation refining the pattern of synchrony of x^*.*

It has been conjectured that a similar result holds for hyperbolic periodic trajectories of admissible systems of differential equations. This is a conjecture from Stewart and Parker (December 2007):

Rigid Synchrony Conjecture (Stewart and Parker (December 2007) Conjecture 6.1)**.** *Let G be any coupled cell network, and suppose that $\mathbf{X}$ is a periodic orbit of some G-admissible vector field f. Assume that $\mathbf{X}$ is rigid. Then its pattern of synchrony $\approx_{\mathbf{X}}$ is balanced.*

The notation used in this thesis differs slightly from that paper, which uses G, $\mathbf{X}$, f and $\approx_{\mathbf{X}}$ where this thesis would use $\mathcal{N}$, F, x and $\equiv_x$ respectively. In addition, we would call the pattern of synchrony of $\mathbf{X}$ rigid, where that paper sometimes applies the adjective to $\mathbf{X}$ itself. Thus our statement of the assumption would be that the pattern of synchrony of x is rigid.

Stewart and Parker (December 2007) prove a limited version of this conjecture. Here $\sim_{\mathbf{X}}$ denotes the coarsest balanced equivalence relation refining $\approx_{\mathbf{X}}$. (When we use the same concept in this thesis, we denote it $\bowtie_x$.) A 'tame' trajectory is one which satisfies a couple of technical properties to do with symmetries. We shall briefly revisit this concept in chapter 9, but the reader is referred to the original paper for the full details.

Tame Synchrony Theorem (Stewart and Parker (December 2007) Theorem 10.9)**.** *Let G be a coupled cell network, and let f be a G-admissible vector field. Let $\mathbf{X}$ be a rigid periodic orbit of f. Suppose $G/{\sim_{\mathbf{X}}}$ is an all-to-all coupled cell network and $\mathbf{X}/{\sim_{\mathbf{X}}}$ is tame, then $\sim_{\mathbf{X}} = \approx_{\mathbf{X}}$ and $\approx_{\mathbf{X}}$ is balanced.*

While this thesis was being written, Golubitsky, Romano, and Wang (Preprint) announced a proof of the Rigid Synchrony Theorem using methods derived from singularity theory. We discuss their work briefly in chapter 9. In their notation, their result was as follows:

Rigid Synchrony Theorem (Golubitsky et al. (Preprint) Theorem 6.1)**.** *Suppose $Z_0(t)$ is a hyperbolic periodic solution of $\dot{Z} = F(Z)$. Then the coloring associated to $\Delta(Z_0)$ is rigid if and only if it is balanced.*

We will examine their method briefly in chapter 9.

We give our own limited version of the Rigid Synchrony Theorem, which we prove in chapter 9. This is an extension of the Tame Synchrony Theorem, in that we require a looser property than tameness; we also weaken the hyperbolicity requirement, in the following way. It is well-known that if a periodic trajectory x of a system of differential equations F is hyperbolic, then it satisfies a property we here call the 'hyperbolic property': if F is perturbed by a sufficiently small amount, then the new system $\hat{F}$ will have a trajectory $\hat{x}$ near x. Our theorem only relies on the hyperbolic property, not the full statement of hyperbolicity.

Semi–Tame Synchrony Theorem (Theorem 9.2.3)**.** *If x is a semi-tame Θ–periodic pseudo-hyperbolic trajectory of a system F of differential equations over a network $\mathcal{N}$, and $\equiv_x$ is unbalanced, then the synchrony of x is not rigid.*

We then go on to show that a technical conjecture (introduced in chapter 9) implies the Rigid Synchrony Theorem. Due to the reliance on the conjecture, we call this the 'limited rigid synchrony theorem'.

Limited Rigid Synchrony Theorem (Theorem 9.3.3)**.** *Let F be an admissible system over a network $\mathcal{N}$, with some hyperbolic periodic orbit x. Assume conjecture 9.3.1. Then the rigid pattern of synchrony of x is the maximal balanced equivalence relation refining $\equiv_x$.*

The Rigid Synchrony Conjecture deals with symmetries in the co-ordinates of a system of differential equations. An obvious extension of this conjecture concerns symmetries in both co-ordinates and time. We call these the 'phase relations' of a trajectory of such a system: the relation of θ–shift relates cells with values that are the same, phase-shifted by θ. They are not, in general, equivalence relations, but we define a new kind of relation, which we call a forward relation, which encapsulates the properties of these phase relations. Stewart and Parker (Preprint) speculated about such relations and their structure — in this thesis, we determine a result about this structure. This result rests upon a slightly different technical conjecture to that used in the limited rigid synchrony theorem. In particular, we must deal with a still-weaker version of hyperbolicity. Instead of making a perturbation $\hat{x}$ of a trajectory x, setting up a perturbation $\hat{F}$ of the system F, and then appealing to hyperbolicity (or the hyperbolic property) to ensure that $\hat{x}$ is 'the' unique nearby trajectory, we set up an alternative definition of hyperbolicity that takes into account that the disjoint pieces of the network can be phase-shifted against each other independently.

Note that the technical conjecture used for this limited version of the rigid phase theorem is quite similar to conjecture 9.3.1: although it is technically stronger, the expectation is that a similar method of proof would prove both conjectures — the details of this argument are given in section 10.2.

Limited Rigid Phase Theorem (Theorem 10.3.3)**.** *Assume conjecture 10.2.3 holds. Then the Rigid Phase Conjecture is true: we spell out what this means.*

Let $\mathcal{N}$ be a network and F an admissible system of differential equations on $\mathcal{N}$ with a hyperbolic trajectory x of period Θ. Let $\theta \in [0, \Theta)$, then the rigid θ–shift relation $\dot{\Lleftarrow}_x^\theta$ of x must be the maximal balanced refinement of $\Lleftarrow_x^\theta$ (as a forward relation over $\equiv_x$ on $\mathcal{N}$).

The essence of our proof of theorem 10.3.3 is to duplicate the given network, to make a network with two disconnected pieces, each of which is isomorphic to the given network and then set up a trajectory on this 'double' network which has one piece phase-shifted in relation to the other: copies of phase-related cells are then synchronous: if the pertinent relations are unbalanced, an appropriate perturbation can be used to break this synchrony, in exactly the same way as for the rigid synchrony theorem: provided the trajectory is unchanged at some times, the perturbed trajectory will project back onto the original network. This projection then has a broken pattern of θ–shift, as required.

The remainder of this thesis is structured as follows.

Chapter 2 of this thesis sets up the general underlying concepts and pieces of notation which we use for the remainder of the work; chapter 3 assembles the concepts specific to this area, in a formulation based on existing work but using slightly different algebraic apparatus.

Chapter 4 contains the proof that the balanced equivalence relations on any given network form a complete lattice, in the sense of a partially ordered set of which every subset has a well-defined greatest lower bound and least upper bound. This proof was given in Aldis (2008) but not previously submitted for assessment. We then go on to consider a new form of balanced relation, related to the balance equivalence relation, which we call a 'balanced forward relation'.

Chapter 5 describes the concept of an 'admissible' function, as mentioned above, and gives some basic, but important, results about these functions. This is the basis for the majority of the remaining material: chapter 6 brings together admissibility and systems of differential equations, and asks the outstanding questions in the field, which are summarised above.

Chapters 7,9,10 give proofs of the rigid synchrony conjecture in various cases, as discussed above: firstly, in chapter 7, we prove the result in any equilibrium solution of a system of differential equations (this result was proved by a different method in Stewart and Parker

(December 2007)); in chapter 9 we prove the same result in oscillating solutions (given certain technical constraints, which are different to those in Stewart and Parker (2008), where a similar result was proved); we also briefly discuss a proof of this result from Golubitsky et al. (Preprint).

Chapter 8 is used to set up some of our hyperbolicity terms, and proves our stronger version of the full oscillation conjecture. This is an interesting result in its own right, but its main purpose here is as a waypoint in the proof of a particular case of the Rigid Synchrony Conjecture — and presumably a required waypoint in any proof of the entire conjecture.

We then consider 'phase relations': that is, where the quantity represented by a cell at each time is equal to that represented by another cell at a related time. Chapter 10 builds on the proofs in the previous chapters to show (assuming a technical conjecture similar to that in chapter 9) that rigid relations of this form are also derived from balanced relations: in this case, from the balanced forward relations introduced in chapter 4.

Directions for further study which arise directly from material in the thesis are presented with the material to which they relate: we conclude with a brief look in chapter 11 at some other directions which might be worth pursuing further.

2 Fundamentals

This chapter introduces some fundamental notions, which are in general use to various degrees, as well as some non-standard, but trivial, notation and concepts, which we use for convenience.

To begin, we make some completely non-standard, but entirely trivial, definitions; we then cover some material on lattices from Davey and Priestley (2002), and note that the natural order on equivalence relations, given by refinement, makes the set of equivalence relations on a given set into a lattice. In chapter 4, we shall expand on this, by showing that the 'balanced equivalence relations' we define there also form a lattice.

In section 2.4, we cover the concept of 'multisets', which are described in Girish and John (2009), amongst other places. Although they are not widely used in mathematics, we shall see in chapter 5 that multisets provide a very natural structure to describe the networks and functions which we later define and which form the basis for this thesis. This justifies and makes rigorous the comment in Golubitsky et al. (2005, remark 2.4(b)) to this effect. A more traditional alternative would be to work with n–tuples in which an element may appear several times. The disadvantage of n–tuples is that they impose an ordering on their entries; a symmetry property then has to be satisfied to remove the complications introduced in this way.

Section 2.5 covers the well-known concept of differential equations: the material given here is very standard, as given for example in Anosov and Arnold (1988) — this section exists largely to ensure that the reader understands our use of notation, for which different standards exist.

In section 2.6, some more involved material appears — we deal with 'bump functions', which are functions, smooth in all derivatives, but supported on sets of finite measure. We use these functions to form a function which is equal to one given function in a certain ball, equal to another given function at points which are at least a given distance from that ball, and smooth in all derivatives everywhere. This construction will later be useful in our perturbations of vector fields.

2.1 Sets and Functions

We begin with a few completely non-standard, but trivial, definitions, which will be useful throughout this thesis.

Definition 2.1.1. In this thesis, we often use the symbol ∞ to denote limiting cases and other relevant constructions; these constructions will be introduced individually. However, since this often gives a piece of notation with a position which can feature either a natural number or this symbol, it is useful to use $\mathbb{N}_\infty$ to denote the set $\mathbb{N} \cup \{\infty\}$. Note that this use of the symbol ∞ is completely formal, and the set $\mathbb{N}_\infty$ is not ascribed any algebraic structure other than as an ordered set ($n < \infty$ for all $n \in \mathbb{N}$). In particular, constructions such as $n + \infty$ are never used.

Definition 2.1.2. Given four sets, A, B, C, D where $A \cap C = \varnothing$, a *combined function*

$$f : \begin{cases} A \to B \\ \\ C \to D \end{cases}$$

is a function from $A \cup C$ to $B \cup D$ such that $f(a) \in B,\ f(c) \in D$ for all $a \in A,\ c \in C$.

Given a function $f : X \to Y$, where f is a bijection, we denote the inverse of f by f^{-1}, as is usual. When f is not a bijection, it is common to use the same notation $f^{-1}(y)$ to denote the preimage of y under f, which is the set $\{\ x \in X \mid f(x) = y\ \}$. We eschew this notational collision, by making the following definition.

Definition 2.1.3. Given any function $f : X \to Y$, let $f^{\{-1\}}$ denote the set $\{\ x \in X \mid f(x) = y\ \}$.

Remark 2.1.4. This definition allows us to refer to the preimage (as a set) even in the case of bijections, and to be sure that $f^{-1}(y) \in X$, provided f^{-1} exists.

2.2 Relations and Lattices

The term *lattice* is used here as in Davey and Priestley (1990, 2002), to mean a partially ordered set X in which any two elements $x, y \in X$ have a unique *join*, denoted $x \vee y$, and *meet*, denoted $x \wedge y$.

A lattice X is *complete* if every subset $Y \subseteq X$ has a unique greatest lower bound, or meet, and a unique least upper bound, or join. We denote the meet of Y by $\bigwedge Y$ and its join by $\bigvee Y$. In particular, any finite lattice is complete, and a complete lattice has a maximal and minimal element.

The dual of Theorem 2.16 from Davey and Priestley (1990) will be useful; we give this dual here:

Theorem 2.2.1 (Davey and Priestley, 1990, Theorem 2.16)**.** *Let X be a non-empty partially ordered set. Then the following properties are equivalent:*

1. *X is a complete lattice.*
2. *$\bigvee Y$ exists for all $Y \subseteq X$.*
3. *X has a minimal element $\bigwedge X$, and $\bigvee Y$ exists for all $Y \subset X$ with $Y \neq \varnothing$.*

Definition 2.2.2. Given two equivalence relations $\sim$ and $\approx$ on a set S, $\sim$ is a *refinement* of $\approx$ if $x \sim y \implies x \approx y$ for all $x, y \in S$.

We define a partial order $\leqslant$ on the equivalence relations on a given set S by $\sim \leqslant \approx$ if $\sim$ is a refinement of $\approx$. Naturally, we use the symbol $\geqslant$ for the opposite order.

It should be clear that $=$ is the minimal equivalence relation on any set S, and $\top$ is the maximal relation, where the equivalence relation $\top$ is defined by letting $x \top y \; \forall x, y \in S$. In fact, the set of equivalence relations on S is a complete lattice. The meet $\bigwedge Y$ of a set Y of equivalence relations on S is their intersection when considered as subsets of $S \times S$: that is, $x \bigwedge Y \, y$ if $x \sim y$ for all $\sim \in Y$. The union of a set Y of equivalence relations may not be a transitive relation, but the join $\bigvee Y$ of Y is given by the transitive closure of this union: $x \bigvee Y \, y$ if there is a chain $x = x_0 \sim_1 x_1 \sim_2 \cdots \sim_n x_n = y$ with $x_i \in S$ and $\sim_i \in Y$ for $1 \leqslant i \leqslant n$.

Given any equivalence relation $\sim$ on a set S, a set R is called a *transversal* of $\sim$ if $r_1 \nsim r_2$ for all $r_1, r_2 \in R$, and for all $s \in S$ there is some $r \in R$ such that $r \sim s$.

2.3 Groupoids

As in Brown (1987); Higgins (1971), and used in Golubitsky et al. (2005); Stewart, Golubitsky, and Pivato (2003), we define an algebraic structure similar to a group, called a groupoid. In contrast to groups, two elements of a groupoid need not have a product, although if sufficient products are defined, group-like properties hold.

Definition 2.3.1. Let G be a set with a unary operation $\cdot^{-1} : G \to G$ and a partial function $* : G \times G \to G$: that is, $*$ is similar to a binary operation on G except that it is not necessarily defined for all pairs of elements.

Suppose the following properties hold:

1. If $a * b$ and $b * c$ are defined, then $a * b * c$ is unambiguously defined: that is, $(a * b) * c = a * (b * c)$, and these expressions are defined.

   Also, if either of $(a * b) * c$ or $a * (b * c)$ is defined, then so is the other, and these expressions are equal.

2. $a * a^{-1}$ is always defined, as is $a^{-1} * a$.

3. If $a * b$ is defined, then $a * b * b^{-1} = a$ and $a^{-1} * a * b = b$.

Then we call $(G, *, \cdot^{-1})$ a *groupoid*.

Groupoids naturally arise in a number of situations: for example, consider G as the set of paths through some space, identified by homotopy; let $*$ be the operation of concatenation. Then $a * b$ is only defined for $a, b \in G$ such that a finishes at the same point where b starts. In this case a^{-1} is a path tracing the same route as a but in the other direction. More relevant to the use of groupoids here is the example of a set of bijective functions between the sets of a collection: let S be some collection of sets, and $f_i : s_i \to t_i$, where $s_i, t_i \in S$. Then

$f = \{ f_i \}$ forms a groupoid under composition of functions, with usual function inverse. Note that the 'product' (composition) $f_i f_j$ is only defined where $t_j = s_i$.

2.4 Multisets

We define 'multisets' as in Wildberger (2003) and Girish and John (2009). Although infrequently used in mathematics as a whole, multisets will be a useful and natural way of describing the networks we later define. This approach was suggested in Golubitsky et al. (2005): we will act on this suggestion to define the whole formalism in this thesis in terms of multisets.

A *multiset* is, informally, a set-like collection of objects S where each object in S has a natural-valued *multiplicity*, which can be considered to be the 'number of times' it appears in the multiset. As suggested in Golubitsky et al. (2005, Remark 2.4(b)), multisets are very natural objects to use in the theory of coupled-cell networks. We now give the important definitions to formalise multisets — more details can be found in Wildberger (2003); Girish and John (2009).

We use the notation $[\![x, x, y, z]\!]$ to denote the multiset with elements x, y, z, where x has multiplicity 2 and y, z each have multiplicity 1. In this way, although the order of representation of a multiset S is ignored, the multiplicity of distinct elements in preserved: $[\![x, x, y, z]\!] = [\![x, y, z, x]\!] \neq [\![x, y, z, z]\!]$. Effectively, a multiset S is a function ϵ_S from some *index set*, $\operatorname{index} S$, to the 'content set', or *support* of S, $\operatorname{supp} S$; without loss of generality, we choose $\operatorname{supp} S$ such that ϵ_S is a surjection. Two multisets are considered equal if there is a bijection $\varphi : \operatorname{index} A \to \operatorname{index} B$ such that $\epsilon_B \varphi = \epsilon_A$. We use the notation $x \in X$ to mean $x \in \operatorname{supp} X$, and define the cardinality $|X|$ of a multiset X to be equal to $|\operatorname{index} X|$.

Let $\epsilon^S(x)$ denote the multiplicity of x in S: that is, $\epsilon^S(x) = |\{ i \in \operatorname{index} S \mid \epsilon_S(i) = x \}|$.

When defining a multiset, we may enumerate the elements as above — or, given sets X, Y and a function $f : X \to Y$, we may use the notation $S = \llbracket f(x) \mid x \in X \rrbracket$, which defines $\operatorname{index} S = X$, $\operatorname{supp} S = f(X)$, and $\epsilon_S(x) = f(x)$; that is, $\epsilon_S = f$. Crucially, this gives a multiset with the same cardinality as X, even if f is not injective.

Definition 2.4.1. If S, T are multisets such that $\operatorname{supp} S \subseteq \operatorname{supp} T$, and $\epsilon^S(s) \leqslant \epsilon^T(s)$ for all $s \in S$, we write $S \Subset T$; if, in addition, $|S| = \lambda$, we may write $S \overset{\lambda}{\Subset} T$. Notice that with this definition, $S \Subset T$ means that S contains 'some of' T, just as $X \subseteq Y$ means X contains 'some of' Y. If $\operatorname{supp} S = \{s\}$, and $\epsilon^S(s) = \epsilon^T(s)$, we may write $S \Cap T$; here we say that S is a *multiplet* in T.

Definition 2.4.2. Let S be a multiset. Then the power set $\mathbb{P}(S)$ of S is defined in the obvious way: $\mathbb{P}(S) = \{ A \mid A \Subset S \}$.

We now define the notion of a function of multisets: our definition differs slightly from that given in, for example, Girish and John (2009); in fact, the definition given here is more general. We shall require this generality in this thesis.

Definition 2.4.3. Let S, T be two multisets, and $f : \mathbb{P}(S) \to \mathbb{P}(T)$ a function such that if $A \Subset B \Subset S$ then $f(A) \Subset f(B) \Subset T$. Then we consider f to be a function between A and B, and write $f : A \to B$. We can define a function $f : S \to T$ by defining $f(X)$ for all $X \Cap S$, and letting $f(Y) = \bigcup_{X \Cap Y} f(X)$: this sets $f(Y) = \varnothing$ for many $Y \Subset S$, but gives a well-defined value $f(Y)$ for all $Y \Subset S$.

Suppose $f : S \to T$ is a function such that $f(S) = T$. Then we call f a multiset surjection, by obvious extension from the set construction. Alternatively, suppose $f : S \to T$ such that

$|f(A)| = |A|$ for all $A \Subset S$. Then we call f a multiset injection. If f is both a surjection and an injection, we call it a bijection.

Remark 2.4.4. Suppose we take two multisets A, B with all multiplicities 1: that is, A and B are sets, considered as multisets. Then each of the possible functions $f : A \to B$ must assign, to each $[\![a]\!] \Subset A$, some $f([\![a]\!]) \subseteq B$. If each of these images also has multiplicity 1, then f is a function $A \to B$ as sets. Conversely, any function of sets $f : A \to B$ must be a function $A \to B$, where A, B are considered as multisets. So set functions are multiset functions with all multiplicities 1.

Girish and John (2009) have another definition of a multiset function: they treat it as a function of multiplets, $\{ a \Subset A \} \to \{ b \Subset B \}$. Our definition is more general, in that each of these multiplet functions can be represented as a multiset function as defined above, but although we allow a function which takes $[\![1]\!] \mapsto [\![10]\!]$ and $[\![1,1]\!] \mapsto [\![10,11]\!]$, it is not a multiset function in the sense of Girish and John (2009).

These multisets are intuitively somewhere between sets and tuples: like tuples, multiplicity is preserved; like sets, order is disregarded. It is completely standard to represent vector spaces by sets of tuples. In chapter 5, we shall see that this concept of vector space can be extended to make what we call a 'multispace', which is a set of multisets which behaves in a somewhat similar way to a vector space.

2.5 Differential Equations

As is standard (for example, see Anosov and Arnold (1988, section 1.2)), given a vector field $F : U \to V$, the *differential equation* corresponding to F is the equation $\dot{x} = F(x)$, where

$\dot{x}$ denotes the image of the standard tangent vector to the t–axis under the derivative of the mapping $t \mapsto x$. Due to the large number of perturbations we make, which we denote by a mark above a variable name, in the manner of $\tilde{x}$, we shall generally use the alternative notation $\frac{\mathrm{d}x}{\mathrm{d}t}$ in place of $\dot{x}$. In this thesis, we only consider autonomous equations; that is, the function F is independent of the time parameter t.

For our purposes, a solution of $\frac{\mathrm{d}x}{\mathrm{d}t} = F(x)$ is a differentiable map $x : \mathbb{R} \to U$ such that $\frac{\mathrm{d}x}{\mathrm{d}t}(t) = F(x(t))$ for all $t \in \mathbb{R}$, where $\frac{\mathrm{d}x}{\mathrm{d}t}(t) = \left[\frac{\mathrm{d}}{\mathrm{d}t}x\right](t)$ denotes the derivative of x at t, which is sometimes (e.g., Anosov and Arnold (1988)) denoted $\frac{\mathrm{d}}{\mathrm{d}t}x(\tau)\big|_{\tau=t}$ We may call a solution of $\frac{\mathrm{d}x}{\mathrm{d}t} = F$ a *trajectory* of the vector field F: the same term is often used (as in Anosov and Arnold (1988)) for the image of such a solution. If x is stationary (that is, $x(t) = x^* \in V$ for all $t \in \mathbb{R}$), we call x an *equilibrium* of F.

Given a vector field F, its Jacobian DF is the map (or matrix) $\frac{\partial f}{\partial x}$: that is, $DF = (a_{i,j})$, where $a_{i,j} = \frac{\partial f_i}{\partial x_j}$, for some co-ordinate mapping $x = (x_1, \ldots, x_n)\,,\; y = (y_1, \ldots, y_n)\,,\; F = (f_1, \ldots, f_n)$ such that $f_k = y_k F$. We use the notation $DF(x)$ to denote the evaluation of this Jacobian at x, which is sometimes written $DF\,|_x$.

2.5.1 Isolation, Transversality, and Hyperbolicity

We now introduce three properties of trajectories of differential equations which we shall examine in more detail later.

Definition 2.5.1 (Isolation)**.** Let F be any system of differential equations. An equilibrium x^* of F is *isolated* if there is some $\delta > 0$ such that $\{\, x \in \mathbb{B}(x^*, \delta) \mid F(x) = 0 \,\} = \{x^*\}$. We may specify the δ by using the term δ–isolated.

Definition 2.5.2 (Transversality)**.** Let F be a system of differential equations with an equilibrium x^*. Then x^* is *transverse* in F if $\det DF(x^*) \neq 0$.

Definition 2.5.3 (Hyperbolicity)**.** Let F be a system of differential equations with an equilibrium x^*. Then x^* is *hyperbolic* in F if $DF(x^*)$ has no purely imaginary eigenvalues. (Here, 0 is considered purely imaginary, as it has zero real part.)

The property of hyperbolicity extends to periodic solutions of differential equations as follows.

Definition 2.5.4 (Hyperbolicity)**.** Let F be a system of differential equations with a periodic solution x. Consider a region of a hyperplane Π through $x(0)$ normal to $\dot{x}(0)$ such that this region contains $x(0)$. Define a map $f : \Pi \to \Pi$ which takes each point $y_0 \in \Pi$ to the next point of intersection between Π and a trajectory y with $y(0) = y_0$. Then x is a *hyperbolic* trajectory of F if y is a hyperbolic equilibrium of f.

2.6 Bump Functions, Norms, and Smoothing

As mentioned above, much of this thesis rests on the concept of perturbing a trajectory of a vector field by a small amount. This motivates several definitions, which we give in this section. Firstly, it is useful to have a simple notation which we introduce for common norms on vector spaces of functions. We also introduce the idea of an ε–perturbation of a function, which has distance less than ε from F, and a δ–local perturbation, which is equal to F outside some ball of radius δ. Secondly, we wish to be able to construct a smooth function which is non-zero on a set of finite measure; we give such a construction. Thirdly, we state and prove some simple facts which result in being able to 'patch together' two smooth functions f, g to create a function h which is equal to f in some places and equal to g in others, but smooth everywhere.

We now define our notation for common norms on vector spaces of functions.

Given two normed vector spaces X, Y with the norm of Y denoted by $\|\cdot\|$, and a function $f : X \to Y$, let $|f|$ denote the supremum norm of f, $|f| = \sup_{x \in X} \|f(x)\|$. Let $\|f\| = \max\{|f|, |Df|\}$, the C^1–norm of f.

If F is a set of functions, let $|F| = \sup_{f \in F} |f|$ and $\|F\| = \sup_{f \in F} \|f\|$.

Definition 2.6.1. Given a function $f : X \to Y$ and $\varepsilon > 0$:

- $g : X \to Y$ is an ε*–perturbation* of f if $\|F - G\| < \varepsilon$.
- $g : X \to Y$ is a perturbation of f which is *supported* on $Z \subseteq X$ if $g(x) = f(x)$ for all $x \notin Z$.
- $g : X \to Y$ is a perturbation of f which is δ*–local* to $x \in X$ if it is supported on $\mathbb{B}(x, \delta)$ as a perturbation of f.

Remark 2.6.2. A trajectory of a differential equation is a function $x : \mathbb{R} \to U$, and so we may consider perturbations of trajectories.

2.6.1 Bump Functions and Smoothing

Definition 2.6.3. Let $\mathrm{Bump} : \mathbb{R} \to \mathbb{R}$ be defined by $\mathrm{Bump}(x) = \exp\left(\frac{x^2}{|x|-1}\right)$ for $|x| < 1$ and 0 elsewhere. This function is C^∞–smooth (although clearly not analytic), is 1 at 0, and is 0 in all derivatives for $|x| \geqslant 1$.

Let $M_{\mathrm{Bump}} = \sup\left\{ \frac{\mathrm{d}}{\mathrm{d}x} \mathrm{Bump}(x) \;\middle|\; x \in \mathbb{R} \right\} = 2.156\ldots < 2.5$.

Given an interval $T = [a, b]$ and $t \in T$, define $\mathrm{Bump}_{t \in T} \in C^\infty(\mathbb{R} \to [0, 1])$ by $\mathrm{Bump}_{t \in T}(x) = \mathrm{Bump}\left(\frac{x-t}{d}\right)$, where $d = \min\{t - a, b - t\}$. This function is also C^∞–smooth (but not analytic), and has the property that $\mathrm{Bump}_{t \in T}(t) = 1$, but $\mathrm{Bump}_{t \in T}(\partial T)$ is 0 in all derivatives, and $\mathrm{Bump}_{t\in T}(u) = 0$ for all $u \notin T$.

Theorem 2.6.4 (Smoothing Theorem)**.** *Given two spaces X, Y, a point $x^* \in X$, two radii $R > r > 0$, and any $\varepsilon > 0$, there is some $\delta > 0$ such that for any two smooth functions $f, g : X \to Y$ such that $\|f - g\| < \delta$, there exists some function $h : X \to Y$ such that:*

1. *$h(x) = g(x)$ for $x \in \mathbb{B}(x^*, r)$;*
2. *$h(x) = f(x)$ for all $x \notin \mathbb{B}(x^*, R)$;*
3. *h is a smooth function;*
4. *$\|h - f\|, \|h - g\| < \varepsilon$.*

We call this δ a function bound *for x^*, R, r, ε.*

Proof. Let $\mathrm{d}_{r,R}(x, y) = \frac{\max(\mathrm{d}(x,y)-r,0)}{R-r}$: this function ranges from 0 when x and y are closer than r to 1 when x and y are distance R apart. Now let $H : X \to [0, 1]$ be defined by $H(x) = \mathrm{Bump}(\mathrm{d}_{r,R}(x^*, x))$. Then H is a positive bump function on X which is equal to 1 on $\mathbb{B}(x^*, r)$ and supported on $\mathbb{B}(x^*, R)$, with $|\det DH| \leqslant M = \frac{M_{\mathrm{Bump}}}{R-r}$. Let $K = M + 1$, and define $\delta = \frac{\varepsilon}{K}$.

For smooth functions $f, g : X \to Y$ with $\|f - g\| < \delta$, let $h = f + (g - f)H$. Clearly, property 1 and property 2 hold. Also, h is smooth by the algebra of smooth functions, so property 3 holds.

Now $|h - f| = H|g - f| < |g - f| < \delta$, and:

$$\begin{aligned}
|\det D[h - f]| &= |\det(Dh - Df)| \\
&= |\det(Df + DH \times (g - f) + H \times (Dg - Df) - Df)| \\
&= |\det(DH \times (g - f) + H \times (Dg - Df))| \\
&\leqslant |\det DH| \times |g - f| + H \times |\det(Dg - Df)| \\
&< \delta |\det DH| + \delta H < \delta (M + 1) = \delta K = \varepsilon
\end{aligned}$$

So property 4 holds for f. It holds for g by a symmetrical argument. □

Theorem 2.6.5 (Bumping Theorem)**.** *Given a function $f : X \to Y$, some point $x^* \in X$, and any $R, \varepsilon > 0$, there exists a function $\hat{f} : X \to Y$ such that:*

1. *$\hat{f}(x^*) \neq f(x^*)$.*
2. *$\hat{f}(x) = f(x)$ for all $x \notin \mathbb{B}(x^*, R)$;*
3. *$\hat{f} - f$ is a smooth function;*
4. *$\left\| \hat{f} - f \right\| < \varepsilon$;*

Proof. Take $r = R/2$, and let $y \in Y$ be some non-zero vector with $\|y\| < \delta$ for δ as in theorem 2.6.4. Let g be defined by $g(x) = f(x) + y$ for all $x \in X$, then $\|f - g\| < \delta$. Hence we may apply theorem 2.6.4 to get a function h which has the properties desired in the function $\hat{f}$. □

Remark 2.6.6. If F is some system of differential equations with an equilibrium x^*, and $\varepsilon, r > 0$ are given, we may apply theorem 2.6.5 to F to obtain an ε–perturbation $\tilde{F}$ of F which is δ–local to x^* such that x^* is not an equilibrium of $\tilde{F}$.

3 Coupled–Cell Networks

This thesis brings together two major concepts: differential equations, and 'balanced equivalence relations.' As we shall see in chapter 4, this kind of relation is defined on 'coupled-cell networks'; we define these networks here. As such this chapter is fundamental to the work in this thesis. The definitions in this chapter are largely derived from Stewart et al. (2003); Golubitsky et al. (2005); Dias and Stewart (2004); Stewart and Parker (December 2007, 2008, Preprint); Wilson (1985), although we depart slightly from previous treatments of these networks by using multisets to describe edges of networks. However, we introduce these 'coupled cell networks' in a manner that is compatible with the one that is becoming standard, for example, in Aldis (2008); Golubitsky and Stewart (2006).

In section 3.2, we also provide results on 'trees,' which are connected networks with no undirected cycles: as we shall see in chapter 5, certain trees are the natural network-theoretic way to determine which values in a differential equation can affect which others. We also describe 'restrictions' and 'limits' of trees. In section 3.3, we define 'bunching', which is a natural operation on trees, which helps us to see which trees are equivalent (that is, have equivalent differential equations). This sense of 'equivalence' motivates us to turn the class of trees into a category in a particular way: most of the details are omitted here, but enough information is provided to demonstrate that this would be an interesting direction for further study.

3.1 Network Terminology

We now define the concept of a 'coupled-cell network', on which all of the material in this thesis is based. We define some simple properties of networks, and define what it means for a relation to be an equivalence relation on a network, or across two networks; we also give analogues of these definitions for more general relations, which will be useful in chapter 10, since the 'phase relations' we deal with there are, in general, not equivalence relations.

A (coupled cell) network $\mathcal{N} = (C, E, \sim)$ is a set of cells C, a multiset of edges E, and an equivalence relation $\sim$, of a certain kind described below. Each edge (or 'arrow') is of the form $e = (c, d, t)$, where c and d are cells, called the *tail* and *head* of e respectively, and t is some *type marker* (of an arbitrary form). Given an edge e, let $\mathcal{H}(e)$ denote its head, $\mathcal{T}(e)$ its tail, and $[\![e]\!]$ its type marker. We may denote the edge (c, d, t) by $d \underset{t}{\rightarrow} c$, and let $d \overset{w}{\underset{t}{\rightarrow}} c$ denote a multiset of identical edges: $\left| d \overset{w}{\underset{t}{\rightarrow}} c \right| = w$, $\operatorname{supp}\left(d \overset{w}{\underset{t}{\rightarrow}} c \right) = \left\{ d \underset{t}{\rightarrow} c \right\}$.

In previous work (including Aldis (2005)), multiplicities have sometimes been taken from arbitrary groups. In general, we do not allow this in this thesis, which simplifies the nature of the edge sets considerably. We revisit this simplification in chapter 11.

The relation $\sim\, = \,\sim_C \cup \sim_E$, where $\sim_C$ is a (chosen) equivalence relation on C, and $\sim_E$ is the equivalence relation on E given by $e \sim_E f$ if $[\![e]\!] = [\![f]\!]$. For a cell c or edge e, let $[c]$ or $[e]$ denote its equivalence class under $\sim$; call this class the *cell type* of c or *edge type* of e.

For any network $\mathcal{N}$, let $\mathcal{C}(\mathcal{N})$ and $\mathcal{E}(\mathcal{N})$ denote the cell and edge sets of $\mathcal{N}$, respectively; let $\sim_{\mathcal{N}}$ denote the network's equivalence relation. If $c, d \in \mathcal{C}(\mathcal{N})$ then let $\mathcal{E}(c, d)$ denote the multiset of arrows from c to d; that is, $\mathcal{E}(c,d) = [\![\, e \in \mathcal{E}(\mathcal{N}) \mid \mathcal{H}(e) = d, \mathcal{T}(e) = c \,]\!]$.

When we draw networks, we use the same symbol (circle, square, etc.) for cells c, d where $c \sim_C d$, and the same style of arrow (dotted, dashed, etc.) for edges e, f such that $e \sim_E f$. An example is shown in figure 3.1.

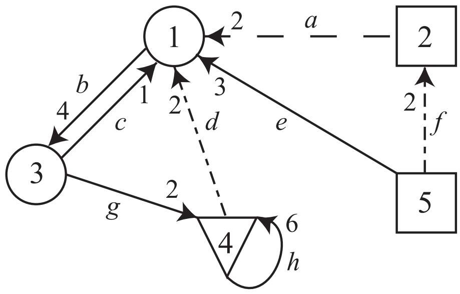

Figure 3.1: An example network, with 5 cells, $C = \{1,2,3,4,5\}$ and 8 edges, $E = \{a,b,c,d,e,f,g,h\}$. As shown by the shapes of cells and edges, $1 \sim_C 3,\ 2 \sim_C 5$; $b \sim_E c \sim_E e \sim_E g \sim_E h$, also $d \sim_E f$. Small numbers show edge multiplicities.

Definition 3.1.1. A network with only one cell type and one arrow type, as in figure 3.2, is called a *homogeneous* network.

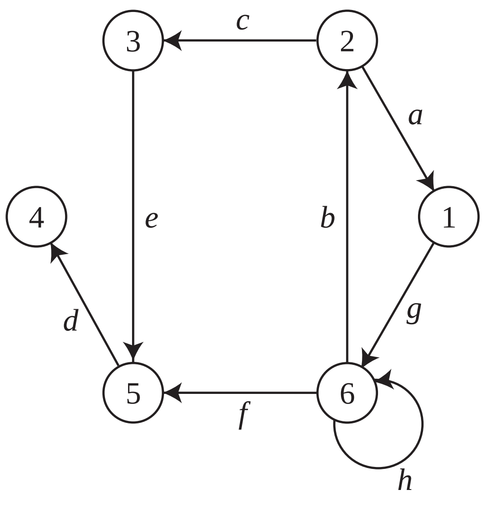

Figure 3.2: A homogenous network with 6 cells and 8 edges.

Definition 3.1.2. Networks with a countable or finite number of cells and edges are called *countable* or *finite* networks, respectively. Networks are called *locally countable* or *locally* *finite* if they have only a countable or finite number of edges meeting any given cell, that is, for all cells $c \in \mathcal{C}(\mathcal{N})$, the set $\{ e \in \mathcal{E}(\mathcal{N}) \mid \mathcal{H}(e) = c$ or $\mathcal{T}(e) = c \}$ is countable or finite, respectively. (We take the term "countable" to include "finite"; that is, all finite numbers are countable).

Definition 3.1.3. Given two networks $\mathcal{N} = (C, E, \sim_{\mathcal{N}})$, $\mathcal{M} = (D, F, \sim_{\mathcal{M}})$, a *network homomorphism* $\varphi : \mathcal{N} \to \mathcal{M}$ is a combined function $\varphi : \begin{cases} C \to D \\ E \to F \end{cases}$, such that:

1. φ preserves the cell and edge equivalence relations: if $x \sim_{\mathcal{N}} y$ for $x, y \in C \cup E$, then $\varphi(x) \sim_{\mathcal{M}} \varphi(y)$.
2. φ preserves network structure: if e is an edge in $\mathcal{N}$, then $\mathcal{H}(\varphi(e)) = \varphi(\mathcal{H}(e))$ and $\mathcal{T}(\varphi(e)) = \varphi(\mathcal{T}(e))$.

A *network isomorphism* is a bijective network homomorphism. It should be clear that the inverse of a network isomorphism is also a network isomorphism: in particular, the isomorphisms are the homomorphisms with homomorphism inverses.

Definition 3.1.4. Given a network $\mathcal{N}$, let $\sim$ denote a pair of relations, one on the cells of $\mathcal{N}$, the other on its edges. We call this kind of relation pair a *relation on* $\mathcal{N}$. For example, $\sim_{\mathcal{N}}$ is a relation on $\mathcal{N}$. If, in addition, $\sim$ is an equivalence relation which refines $\sim_{\mathcal{N}}$, then we call $\sim$ an *equivalence relation on* $\mathcal{N}$.

Definition 3.1.5. Given two networks $\mathcal{N}, \mathcal{M}$, let $\sim$ denote a pair of relations, one on the disjoint union of the cell sets of $\mathcal{N}$ and $\mathcal{M}$, the other on the disjoint union of their edge sets. We call this kind of relation pair a *relation across* $\mathcal{N}$ *and* $\mathcal{M}$.

If $\sim$ is an equivalence relation such that the restriction of $\sim$ to $\mathcal{N}$ is an equivalence relation on $\mathcal{N}$, and the restriction of $\sim$ to $\mathcal{M}$ is an equivalence relation on $\mathcal{M}$, then we call $\sim$ an *equivalence relation across $\mathcal{M}$ and $\mathcal{N}$.*

Definition 3.1.6. Let $\sim$ be an equivalence relation across two networks $\mathcal{M}, \mathcal{N}$. We say $\mathcal{N}$ is *isomorphic to $\mathcal{M}$ with respect to* $\sim$ if there is an isomorphism $F : \mathcal{M} \xrightarrow{\cong} \mathcal{N}$ which preserves equivalence classes under $\sim$ — that is, $x \sim F(x)$ for cells and edges x in $\mathcal{M}$. We denote this by $\mathcal{M} \cong_{\sim} \mathcal{N}$, or just $\mathcal{M} \sim \mathcal{N}$.

Similarly, if $\rightsquigarrow$ is any relation across two networks $\mathcal{M}, \mathcal{N}$, we say $\mathcal{M}$ *relates to* $\mathcal{N}$, and write $\mathcal{M} \rightsquigarrow \mathcal{N}$, if there is an isomorphism $F : \mathcal{M} \to \mathcal{N}$ such that $x \rightsquigarrow F(x)$ for all $x \in \mathcal{C}(\mathcal{M}) \cup \mathcal{E}(\mathcal{M})$.

3.2 Trees

A particular kind of coupled-cell network is the 'tree', a term which here means a rooted directed tree in the usual sense: such a tree has some unique root cell and a unique directed path from any cell to this root.

Trees will be very important in this thesis, as they are the natural way to track which values in a differential equation can affect which others. In particular, although the networks we define from the differential equation application will be finite, in the general case the trees we deal with may be infinite. In order to make intelligible statements about these infinite trees, we show that an infinite tree is the limit, in a well-defined way, of a sequence of finite trees. We may then reason about the finite trees, and show that our results carry 'upwards' to the infinite limiting tree of any convergent sequence.

We start by defining some basic relations on the cells of a network.

Definition 3.2.1. We define the relation $\to$ on the cells of a network by $c \to d$ if there is an edge $\left(c \underset{t}{\rightarrow} d\right) \in \mathcal{E}(\mathcal{N})$ for some edge type t. We write $c \rightharpoonup d$ (or $d \leftharpoondown c$) if there is a sequence of cells $c = c_0, c_1, \ldots, c_n = d$ such that for each i, we have that $c_i \rightarrow c_{i+1}$. Such a sequence is called a path from c to d.

Definition 3.2.2. Two cells c, d in a network are *path-connected* if $c \rightharpoonup d$ and $c \leftharpoondown d$: in this case we write $c \rightleftharpoons d$. A *path-connected network* is a network whose cells are pairwise path-connected.

Definition 3.2.3. Two cells c, d in a network are *connected* if there is a sequence of cells $c = c_0, c_1, \ldots, c_n = d$ such that for each i there is an edge e_i with c_i, c_{i+1} as its ends. In this case, we write $c \leftrightsquigarrow d$. Note that, unlike in the definition of path-connectedness, there is no requirement for the edges used in connectedness to 'point in the same direction': the edge e_i can be $c_i \rightarrow c_{i+1}$ or $c_i \leftarrow c_{i+1}$, independently of whether the edge e_j is $c_j \rightarrow c_{j+1}$ or $c_j \leftarrow c_{j+1}$. In this way, connectedness in our network sense is a weaker property than path-connectedness.

A *connected network* is a network whose cells are pairwise connected.

Definition 3.2.4. A *component* of a network $\mathcal{N}$ is the network $\mathcal{M}$ with cell set equal to some subset C of $\mathcal{C}(\mathcal{N})$ and edge set $\bigcup_{c,d \in C} \mathcal{E}(c, d)$.

A component $\mathcal{M}$ of $\mathcal{N}$ is called a *connected component* if $\mathcal{C}(\mathcal{M})$ is an equivalence class of $\leftrightsquigarrow$ on $\mathcal{C}(\mathcal{N})$. Similarly, $\mathcal{M}$ is a *path-connected component* if $\mathcal{C}(\mathcal{M})$ is a $\rightleftharpoons$–equivalence class of $\mathcal{C}(\mathcal{N})$. If $\mathcal{M}$ is a path-connected component of $\mathcal{N}$ such that there are no cells $c \in \mathcal{C}(\mathcal{M}),\ d \in \mathcal{C}(\mathcal{N}) \backslash \mathcal{C}(\mathcal{M})$ where $d \rightarrow c$, then we call $\mathcal{M}$ an *upstream component* of $\mathcal{N}$.

Definition 3.2.5. We define a metric d on the cells of a network. If c, d are connected, then $\mathrm{d}(c, d)$ is the smallest n such that $c = c_0, c_1, \ldots, c_n = d$ is a sequence of cells connecting c and d. If c, d are not connected, $\mathrm{d}(c, d) = \infty$.

Definition 3.2.6. A *tree* is a connected network which has some cell c_0, called the *root*, such that:

1. There is no edge with c_0 as its tail.

2. Given any other cell $c \neq c_0$, there is a unique edge with c as its tail.

A *leaf* of a tree is a cell c such that there is no edge with c as its head. The *generation* of a cell c in a tree is the distance $\mathrm{d}(c, c_0)$. The *depth* of a tree T, $\mathrm{depth}(T)$, is the maximal generation of cells in the tree. If there are cells in T at generation n for all $n \in \mathbb{N}$, then we define $\mathrm{depth}(T) = \infty$. See figure 3.3.

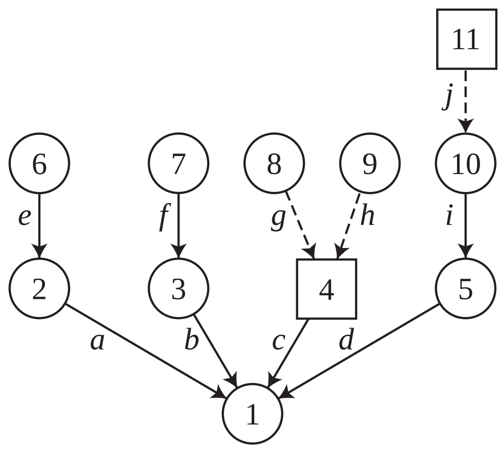

Figure 3.3: A tree with root 1 and leaves $6, 7, 8, 9, 11$. Cells $2, 3, 4, 5$ are at generation 1, cells $6, 7, 8, 9, 10$ at generation 2 and 11 is at generation 3; the depth of this tree is 3.

Definition 3.2.7. Given a tree T containing some cell c, the *subtree rooted at* c, $T_{(c)}$, is the tree containing all cells d of T which have a directed path from d to c in T, and all the edges used in such paths. See figure 3.4.

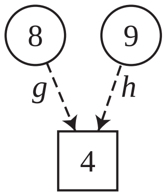

Figure 3.4: The subtree of figure 3.3 rooted at cell 4.

3.2.1 Operations on Trees

Here we give two straightforward definitions. Restriction of a tree to a given depth is the procedure by which we make infinite trees into finite ones. The second definition is the 'join' of two trees, which involves 'gluing' the root of one tree to a leaf of another.

Definition 3.2.8. The *restriction* of a tree T to depth n, denoted $T|_n$, is the tree with cells c from T where the generation of c is at most n, and all edges (c, d, t) from T where both c and d are of generation at most n. See figure 3.5. The cell and edge equivalence relations on $T|_n$ are precisely those on T, restricted to the cells and edges of $T|_n$. Obviously, the depth of $T|_n$ is at most n: specifically, it is $\min\{n, \operatorname{depth}(T)\}$, and where the depth of T is less than or equal to n, we have $T|_n = T$.

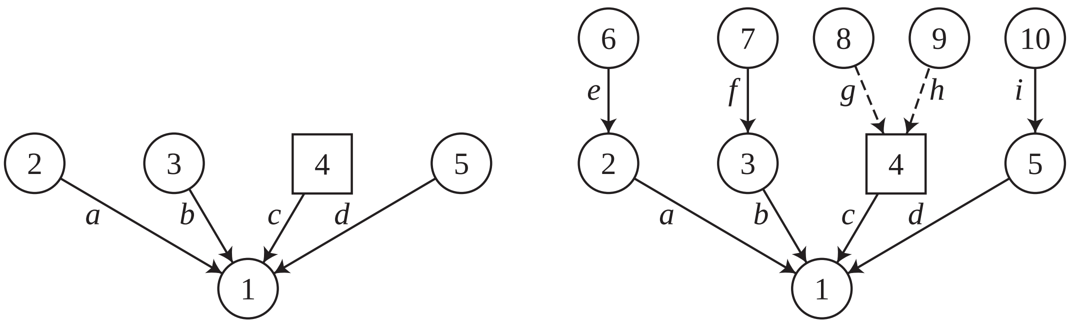

Figure 3.5: The restriction of figure 3.3 to depths 1 (left) and 2 (right).

Lemma 3.2.9. *For locally finite trees T, the tree $T|_n$ is finite for all n.*

Proof. The proof is routine, and omitted. □

We now define the 'join' of two trees: informally, we shall ensure that the cell sets of the trees are disjoint, and then identify the root of one tree with a leaf in the other, while keeping all edges 'the same'. The following definition makes this precise.

Definition 3.2.10. Given two trees, $\mathcal{T} = (C, E)$ with a leaf c and $\mathcal{U} = (D, F)$ with root d, the *join* of $\mathcal{U}$ to $\mathcal{T}$ at c is a tree with cell set given by the union $C_* \cup D_*$, where $C_* = C \times \{1\} = \{\ (c', 1) \mid c' \in C\ \}$, $D_* = (D \backslash \{d\}) \times \{2\} = \{\ (d', 2) \mid d' \in D \backslash \{d\}\ \}$, and edge set $E_* \cup F_*$, where

$$E_* = \left[\!\left[\ (c_1, 1) \underset{t}{\rightarrow} (c_2, 1)\ \middle|\ c_1 \underset{t}{\rightarrow} c_2 \in E\ \right]\!\right]$$

and

$$F_* = \left[\!\left[\ (d_1, 2) \underset{t}{\rightarrow} (d_2, 2)\ \middle|\ d_1 \underset{t}{\rightarrow} d_2 \in F \text{ and } d_1 \neq d\ \right]\!\right] \cup \left[\!\left[\ (c, 1) \underset{t}{\rightarrow} (d_2, 2)\ \middle|\ d \underset{t}{\rightarrow} d_2 \in F\ \right]\!\right]$$

(Here, as is usual, notation such as $(c_2, 1)$ represents the Cartesian pair taken from $(C \cup D) \times \{1, 2\}$.)

See figure 3.6 for an example. Cell and edge equivalences are, by default, defined as the union of those of $\mathcal{T}$ and $\mathcal{U}$, with the additional condition that all cells which are equivalent to d in $\mathcal{U}$ become equivalent to all cells which are equivalent to c in $\mathcal{T}$. The following remark describes cases where a more particular equivalence relation makes sense.

Remark 3.2.11. With the 'default equivalences' described in the definition of the join of trees $\mathcal{T}, \mathcal{U}$ at c, d respectively, no cell in $\mathcal{T}$ which is not equivalent to c becomes equivalent to a cell in $\mathcal{U}$, and no cell in $\mathcal{U}$ which is not equivalent to d becomes equivalent to any cell in $\mathcal{T}$. However, there are cases when the cell and edge equivalence relations on $\mathcal{T}$ and $\mathcal{U}$

are naturally considered as the restrictions of a single pair of equivalence relations $\sim_C, \sim_E$ on the unions of the cells and edges of these trees; for example, where the trees are both derived from some underlying network, as we shall see in chapter 4. In that case, we may take the restrictions of the relations $\sim_C, \sim_E$ to the join of $\mathcal{T}$ and $\mathcal{U}$ as the equivalences of this new tree; again, we define $(c', 1) \sim_C (d', 2)$ whenever $c' \sim_C c$ and $d \sim_C d'$. Usually, we shall join cells c, d where $c \sim_C d$, making this step unnecessary.

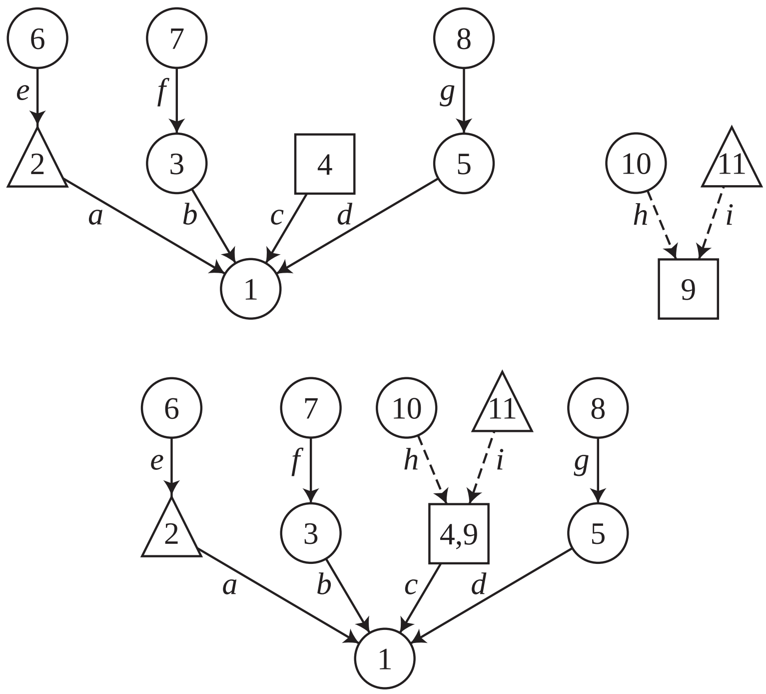

Figure 3.6: Trees $\mathcal{T}$ (top left) and $\mathcal{U}$ (top right) and the join of $\mathcal{T}$ to $\mathcal{U}$ at 4. Notice that some cells in $\mathcal{T}$ are equivalent to some in $\mathcal{U}$, as shown by shapes here (for example, $2 \sim_C 11$), and in particular that $4 \sim_C 9$.

Given two pairs of trees, with homomorphisms between them, we may define a homomorphism from the join of one pair of trees to the other in the obvious way: refer to Aldis (2008) for the details.

3.2.2 Sequences of Trees

The previous section dealt with the restriction of trees to given depths; this procedure made finite trees from infinite ones. We now consider sequences of trees, and their limits: this process in general takes finite trees into infinite ones. We note that this procedure is in a sense the inverse of restriction: given a tree $\mathcal{T}$, its restriction to depth n tends to $\mathcal{T}$ as n tends to infinity. Proofs of the results presented here are given in Aldis (2008).

Definition 3.2.12. Given a sequence of trees $(\mathcal{T}_i)$ with a sequence of isomorphisms f_i : $\mathcal{T}_i|_{n_i} \stackrel{\cong}{\to} \mathcal{T}_{i+1}|_{n_i}$ for some sequence $n_i \to \infty$, we call $\mathcal{T}_i$ a *convergent* sequence. We call (f_i) the *convergence isomorphism* of $(\mathcal{T}_i)$.

Given a convergent sequence of trees $(\mathcal{T}_i)$, the direct limit $\mathcal{T}_\infty$ of $(\mathcal{T}_i)$ is constructed by taking disjoint unions of cell and edge sets of the $\mathcal{T}_i$, and then identifying all cells and edges mapped to each other by any of the individual isomorphisms f_i from the convergence isomorphism.

Remark 3.2.13. This identification gives a well-defined injection $f_i^\infty : \mathcal{T}_i \hookrightarrow \mathcal{T}_\infty$ for all i; further, these injections cause this diagram to commute for all i:

$$\begin{array}{ccc} \mathcal{T}_i & \xrightarrow{\quad f_i \quad} & \mathcal{T}_{i+1} \\ & {\scriptstyle f_i^\infty} \searrow \quad \swarrow {\scriptstyle f_{i+1}^\infty} & \\ & \mathcal{T}_\infty & \end{array}$$

Definition 3.2.14. The *direct limit* of a sequence of functions $(F_i : \mathcal{T}_i \to \mathcal{U}_i)$ is defined if, for all i greater than some sufficiently large $n \in \mathbb{N}$, F_i commutes with the convergence isomorphisms $(f_i), (g_i)$ of $(\mathcal{T}_i)$ and $(\mathcal{U}_i)$, as in the following diagram:

$$\begin{array}{ccc} \mathcal{T}_i & \xrightarrow{F_i} & \mathcal{U}_i \\ {\scriptstyle f_i}\downarrow & & \downarrow{\scriptstyle g_i} \\ \mathcal{T}_{i+1} & \xrightarrow{F_{i+1}} & \mathcal{U}_{i+1} \end{array}$$

This direct limit F_∞ is then defined by:

$$\begin{array}{ccc} \mathcal{T}_i & \xrightarrow{F_i} & \mathcal{U}_i \\ {\scriptstyle f_i^\infty}\downarrow & & \downarrow{\scriptstyle g_i^\infty} \\ \mathcal{T}_\infty & \xrightarrow{F_\infty} & \mathcal{U}_\infty \end{array}$$

where $i > n$.

Lemma 3.2.15. *This limit is well-defined.*

Proof. Consider this diagram:

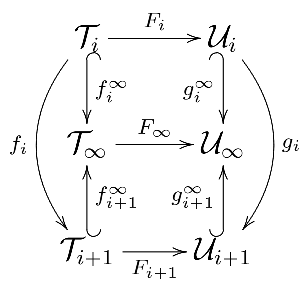

Since the outer arrows commute, the F_∞ defined by the upper square (from F_i) is the same as that defined by the lower square (from F_{i+1}). Induction shows that the same F_∞ is defined for any choice of i. □

Lemma 3.2.16. *If $F : \mathcal{T} \to \mathcal{U}$ is a homomorphism of infinite trees, then each restriction $F\big|_{\mathcal{T}|_n} : \mathcal{T}|_n \to \mathcal{U}|_n$ is also a homomorphism. If F is an isomorphism, so is each $F\big|_{\mathcal{T}|_n}$.*

Further, if $\sim$ is an equivalence relation across $\mathcal{T}$ and $\mathcal{U}$ which is preserved by F, then each $F\big|_{\mathcal{T}|_n}$ preserves the restriction of $\sim$ to $\mathcal{T}|_n$, $\mathcal{U}|_n$.

Conversely, if $(F_i : \mathcal{T}_i \to \mathcal{U}_i)$ is a sequence of tree homomorphisms for convergent sequences of trees, then F_∞ is also a homomorphism, $F_\infty : \mathcal{T}_\infty \to \mathcal{U}_\infty$. If the F_i are isomorphisms, then so is F_∞. If $\sim$ is an equivalence relation across $\mathcal{T}_\infty$ and $\mathcal{U}_\infty$ such that F_i respects $\sim$ for all i, then so does F_∞.

Proof. Given in Aldis (2008). □

Lemma 3.2.17. *For any tree $\mathcal{T}$, $\mathcal{T} \cong \lim_{n\to\infty} \mathcal{T}|_n$.*

Proof. Given in Aldis (2008). □

3.3 Bunching

In systems modelled by the networks we are defining here, we consider any set of arrows of the same type between the same cells with the same total multiplicity to be equivalent. The use of multisets for edge sets enforces this choice. More interestingly, identical cells with identical inputs will behave identically, and so they may be considered as 'one cell' (or, multiple copies of the same cell). In effect, any output arrows from one of the identically-behaving cells can be transferred to the other without affecting the behaviour of the network; transferring all output arrows from some cell removes the effect of that cell on the network: this cell can then be removed. To inspect the value of the removed cell, we may take the value of an equivalent cell that is left in the network. Aldis (2008) formalises this process on trees, resulting in the definition of a 'bunched tree', $\mathcal{T}/\!\!/\!\sim$.

For the next set of definitions, we shall let $\mathcal{T}$ be a finite tree, and $\sim$ an equivalence relation on $\mathcal{T}$. Notice that bunching a tree will not change its depth (we will prove this statement as we proceed), and in particular, trees of depth 0 (consisting of a single cell) are unchanged by bunching.

Definition 3.3.1. With $\mathcal{T}, \sim$ as above, define a relation $\dot{\sim}$ on $\mathcal{E}(\mathcal{T})$ by $e \dot{\sim} f$ where either $e = f$, or all of the following: $e \sim f$, $c = \mathcal{T}(e) \sim \mathcal{T}(f) = d$, and the subtrees rooted at c, d are isomorphic with respect to $\sim$ after being bunched by $\sim$.

Lemma 3.3.2. *This relation $\dot{\sim}$ is well-defined.*

Proof. The apparent self-reliance is dealt with by descending induction. Let $\mathcal{T}$ be a tree of depth n, and let all bunchings of trees of depth less than n by equivalence relations be well-defined. Then for any two edges e, f, their tails are two non-root nodes c, d. The depth of the subtrees rooted at c and d must be less than n (recall that $\mathcal{T}$ is finite), and so the bunched trees rooted at c and d are well-defined by the inductive hypothesis. Thus we can determine whether these bunched trees are isomorphic with respect to $\sim$. □

Definition 3.3.3. Let $\mathcal{T}, \sim, \dot{\sim}$ be as above. Define a third relation, $\hat{\sim}$, between cells and edges: $c \hat{\sim} d$ if $c = d$ or c and d are at the tail of edges e and f respectively such that $e \hat{\sim} f$; $e \hat{\sim} f$ if $e \dot{\sim} f$ and the heads of e and f are equivalent under $\hat{\sim}$.

Lemma 3.3.4. *This relation $\hat{\sim}$ is well-defined.*

Proof. This proof will use ascending induction on the generation of cells, and hence only work for finite trees. (Recall that $\mathcal{T}$ is assumed to be finite, so this isn't a problem.) We will further prove inductively that if $c \hat{\sim} d$ then c and d are at the same generation.

Firstly, notice that the root of $\mathcal{T}$ cannot be equivalent to any other cell under $\hat{\sim}$, since it is not at the tail of any edge. Now suppose that $\hat{\sim}$ is defined for cells and edges of generation at most m, and no two cells of different generations at most m are equivalent.

Let c, d be non-root cells at generation at most $m + 1$. Then they are at the tails of some unique edges e, f respectively. The relation $\dot{\sim}$ is already defined (on edges), and so it only remains to determine whether the heads of e, f are equivalent. But these heads are cells of generation at most m; thus the relation $\hat{\sim}$ is sufficiently defined to decide this question.

By induction, $\hat{\sim}$ is well-defined on all cells (and edges) of $\mathcal{T}$. □

Definition 3.3.5. With $\mathcal{T}, \sim$ as above, we can now define the bunched tree $\mathcal{T} /\!\!/ \sim$ to be the usual tree quotient $\mathcal{T}/\hat{\sim} = (\mathcal{C}\,(\mathcal{T})/\hat{\sim}, \mathcal{E}\,(\mathcal{T})/\hat{\sim})$.

We also extend this definition to infinite trees, by letting $\mathcal{T} /\!\!/ \sim\, = \lim_{n\to\infty} (\mathcal{T}|_n /\!\!/ \sim)$. Aldis (2005) shows that this is well-defined.

Note that for trees $\mathcal{T}, \mathcal{U}$ with integer edge multiplicities, as allowed in this thesis, $\mathcal{T} /\!\!/ \sim\, \cong_\sim \mathcal{U} /\!\!/ \sim \iff \mathcal{T} \cong_\sim \mathcal{U}$. We continue to refer to bunched trees where the results are straightforward, as they are useful in the more general setting, but we will not examine this use in detail here.

The following results are proved in Aldis (2005).

Lemma 3.3.6. *Let $\mathcal{T}$ be a tree and $\sim$ an equivalence relation on $\mathcal{T}$. Let c be a cell of $\mathcal{T}$, corresponding to a cell bunch $c_\sim$ in $\mathcal{T} /\!\!/ \sim$. Then $\mathcal{T}_{(c)} /\!\!/ \sim\, \cong_\sim (\mathcal{T} /\!\!/ \sim)_{(c_\sim)}$.*

Proposition 3.3.7. *If $\mathcal{T}$ and $\mathcal{U}$ are two trees, and $\sim\, \leqslant\, \approx$ are equivalence relations across $\mathcal{T}$ and $\mathcal{U}$ such that $\mathcal{T}$ is equivalent to $\mathcal{U}$ with respect to $\sim$, then $\mathcal{T}$ is also equivalent to $\mathcal{U}$ with respect to $\approx$. That is, $\mathcal{T} /\!\!/ \sim\, \cong_\sim \mathcal{U} /\!\!/ \sim \implies \mathcal{T} /\!\!/ \approx\, \cong_\approx \mathcal{U} /\!\!/ \approx$.*

Lemma 3.3.8. *If $\sim$ is an equivalence relation on a tree $\mathcal{T}$, then $(\mathcal{T}/\!\!/\sim)|_n = \mathcal{T}|_n /\!\!/\sim$.*

Lemma 3.3.9. *Let $\mathcal{T}$,$\mathcal{U}$ be two trees, with some equivalence relation $\sim$ across them. Let $c_\sim$ be a cell bunch which is a leaf in $\mathcal{T}/\!\!/\sim$ — that is, all $c \in c_\sim$ are leaves. Let $\mathcal{J}$ be the join of $\mathcal{U}$ to $\mathcal{T}$ at each $c \in c_\sim$. Then $\mathcal{J}/\!\!/\sim$ is (isomorphic to) the join of $\mathcal{U}/\!\!/\sim$ to $\mathcal{T}/\!\!/\sim$ at $c_\sim$, considering $\mathcal{T}/\!\!/\sim$ as a tree in its own right, as usual.*

3.4 A Category–Theoretic Treatment

As a direction for future research, consider a category of trees based on the 'bunched isomorphisms' we have seen in this chapter. There is already a formalism of 'the category of trees', **Tree**: see, for example Snydal (1999). In this treatment, the objects of **Tree** are trees, and there is a morphism from tree t to u when u is a 'refinement' of t: that is, we can pass from t to u by a succession of moves, each of which either shrinks an edge or replaces a vertex by an edge. These moves are not relevant to our use of trees: our treatment leads us to consider instead the category **Bunched Tree**, which again has trees as objects, but has a morphism from t to u if u is obtained from t by the bunching procedure.

Further study should be able to link this category to the 'natural behaviour' we examine in the remainder of this thesis.

4 The Lattice of Balanced Equivalence Relations

In the previous chapter, we described coupled-cell networks, which are a prerequisite for the explanation of balanced equivalence relations. In this chapter, we make the definition of these relations. The later chapters of this thesis will show that balanced equivalence relations are the relations of 'rigid patterns of synchrony', which is a useful property of solutions of differential equations.

In order to define balanced equivalence relations, we use 'input trees' of cells of networks, which we also define. In contrast to previous work, but in common with the previous chapter, we shall make extensive use of multisets in our definitions. Additionally, several results make use of the infinite limit constructions of the previous chapter.

Worthy of particular interest is section 4.1.3, which sets out constructions for determining 'when and where' relations are not balanced — that is, under what conditions and at which cells of the network. This work expands in a specific and useful way on a similar construction presented in Stewart and Parker (Preprint).

Section 4.2 is more interesting still: in it we prove that the balanced equivalence relations on a given network form a lattice under the partial order given by refinement. This result is more general, and its proof more elementary, than the similar result given by Stewart (2007); it has previously appeared in Aldis (2008), but was not included for examination in Aldis (2005). After this, we touch very briefly on an idea for further study, which would aim to determine more fully the structure of the lattice of balanced equivalence relations on a network.

Finally, section 4.3.1 introduces a new kind of relation, the 'forward relation', which is the natural set-theoretic framework in which to study the 'phase relations' of chapter 10. These 'forward relations' may also be balanced, in their own way: we define this, and see that its definition is compatible with that of balanced equivalence relations; we note that the similarity in definitions directly allows us to show that there is some maximal balanced forward relation, and to give its structure.

4.1 Input Trees and Balanced Equivalence Relations

We now define "balanced equivalence relations" on a network $\mathcal{N}$ in terms of certain trees, derived from the structure of $\mathcal{N}$, which we call "input trees". We then consider the sets of cells which summarise the places where a given equivalence relation 'fails to be balanced': this is based on Stewart and Parker (Preprint), but introduces the new idea of the 'duplicate set'.

In contrast to previous work, we consider relations other than equivalence relations and define what it means for these to be 'balanced'.

4.1.1 Input Sets and Input Trees

In this section, we define the input tree of a cell, in much the same way as Aldis (2005). However, in contrast to the approach there, we use multisets throughout: we show that this new definition is equivalent to the old one. We then define n–th input trees for any n, as before, and the infinite input tree of any cell. We define a relation on cells, n–th input tree equivalence, for any $n \in \mathbb{N}_\infty$. This material will allow the definition of balanced equivalence relations in the next section.

Definition 4.1.1. The *input set* of a cell c in a network $\mathcal{N}$ is the multiset of edges of the form $d \to c$. The *input tree* of c is a network $\mathcal{I}(c) = (C_c, E_c)$, with cell set

$$C_c = \{c\} \cup \left\{ (d,e) \,\middle|\, e = \left(d \overset{w}{\underset{t}{\to}} c\right) \Subset \mathcal{E}(\mathcal{N}) \right\}$$

and edge set

$$E_c = \left[\!\left[(d,e) \underset{t}{\to} c \,\middle|\, e = \left(d \underset{t}{\to} c\right) \in \mathcal{E}(\mathcal{N}) \right]\!\right]$$

This $\mathcal{I}(c)$ is clearly a tree of depth 1 with root c (unless c has no inputs, in which case it is a tree of depth 0). Cell and edge types are preserved: we extend $\sim_C$ and $\sim_E$ such that $(d,e) \sim_C d$ and $\left((d,e) \underset{t}{\to} c\right) \sim_E \left(d \underset{t}{\to} c\right) = e$, and use the appropriate restrictions of these relations as cell and edge equivalences on the input set of c.

This definition is slightly different from the one given in Aldis (2008) and other places, which does not use multisets explicitly. However, the two definitions are compatible in the sense that the identifying marks in Aldis (2008) exist to create a multiset: let $\tilde{E}_c$ denote the set E_c as defined in previous work, and define a function f from $\tilde{E}_c$ to the multiset E_c defined above: $f((d,e), c, t, i) = (d,e) \underset{t}{\to} c$. Then this f is a multiset bijection: the image of f is

E_c as above, and the number of edges in $\tilde{E}_c$ mapped to the same item $e' = (d, e) \underset{t}{\to} c$ of the multiset E_c is the same as the multiplicity $\epsilon^{E_c}(e')$.

Definition 4.1.2. Given a cell c in a network $\mathcal{N}$, we inductively define the *n–th input tree of c*, denoted $\mathcal{I}^n(c)$ so that each $\mathcal{I}^n(c)$ has root c. The first input tree of c, $\mathcal{I}^1(c) = \mathcal{I}(c)$, is the input tree of c. The $(n-1)$–th input tree of c has certain leaves at depth $n-1$. Each such leaf d is associated with a cell $d_{\mathcal{N}}$ in $\mathcal{N}$. Join the input set of $d_{\mathcal{N}}$ to $\mathcal{I}^{n-1}(c)$ at d. Having joined all of these input sets, the result is the n–th input tree $\mathcal{I}^n(c)$. We define $\mathcal{I}^0(c)$ to be the tree consisting only of the cell c, with no edges.

The *infinite input tree*, $\mathcal{I}^\infty(c)$, is the direct limit of the input trees $\mathcal{I}^n(c)$. This may or may not be an infinite network, even for finite networks $\mathcal{N}$ — although if $\mathcal{N}$ is finite with no directed cycle, $\mathcal{I}^\infty$ is finite for all cells c of $\mathcal{N}$. Also, if $\mathcal{N}$ is locally finite, $\mathcal{I}^\infty$ is also locally finite for all cells $c \in \mathcal{N}$.

Lemma 4.1.3. *For cells c from a locally finite network, the n–th input tree of c is finite for any finite n. Further, the infinite input tree of such a cell c is countable: that is, it has a countable number of cells and edges. (Recall that our use of 'countable' includes 'finite'.)*

Proof. Routine, and omitted. □

Definition 4.1.4. Given an equivalence relation $\sim$ on a network $\mathcal{N}$, we define the relation of *n–th input tree equivalence under* $\sim$ between cells of $\mathcal{N}$, denoted $c \cong^n_\sim d$ (for $n \in \mathbb{N} \cup \{\infty\}$), by $c \cong^n_\sim d$ whenever the n–th input trees of c and d are equivalent under $\sim$: $\mathcal{I}^n(c) /\!\!/ {\sim} \cong_\sim \mathcal{I}^n(d) /\!\!/ {\sim}$. (Recall that $\mathcal{I}^n(c) /\!\!/ {\sim} \cong_\sim \mathcal{I}^n(d) /\!\!/ {\sim} \iff \mathcal{I}^n(c) \cong \mathcal{I}^n(d)$ for the natural-valued multiplicities we consider here.)

For brevity and clarity, we denote $\cong^n_{\sim_\mathcal{N}}$ by $\cong^n_\mathcal{N}$.

Remark 4.1.5. Since adding extra generations to two trees can only make them more different, it is clear that if $m \geqslant n$ then $\cong_{\sim}^{m} \leqslant \cong_{\sim}^{n}$.

Definition 4.1.6. Given a sequence of input tree homomorphisms between the input trees of cells c and d, we now define its limit. Let $\left(F_{c,d}^{(n)}\right)$ be a sequence of homomorphisms $F_{c,d}^{(n)} : \mathcal{I}^{n}(c) \to \mathcal{I}^{n}(d)$. We define $F_{c,d}^{(\infty)} = \lim_{n\to\infty} F_{c,d}^{(n)}$.

Count, that is, uniquely enumerate from 1, the cells in $\mathcal{I}^{\infty}(c)$ in increasing order of generation, with cells at the same generation enumerated in arbitrary order. (This is possible because the input tree is countable, by lemma 4.1.3.) Let c_j denote the cell enumerated by $j \in \mathbb{N}$. For each j, we define a strictly increasing sequence $n_i^{(j)}$ of natural numbers such that the subsequence $(f_{i,j}) = \left(F_{c,d}^{\left(n_i^{(j)}\right)}\right)$ of $\left(F_{c,d}^{(n)}\right)$ is constant on cells c_k for $k \leqslant j$. We then let $d_{i,j} = F_{c,d}^{\left(n_i^{(j)}\right)}(c_j)$. This is constant for all i, so we can define $F_{c,d}^{(\infty)}(c_j) = d_{i,j}$.

The details of this process are given in Aldis (2005, 2008), and not repeated here.

4.1.2 Balanced Equivalence Relations

We now define what it means for a relation to be balanced, in the standard way, as introduced in Golubitsky and Stewart (2006) and used throughout the field: for example, Golubitsky, Stewart, and Nicol (2004); Golubitsky et al. (2005); Stewart et al. (2003); Stewart and Parker (December 2007). These uses of balance have also described the groupoid of input tree isomorphisms: given cells c, d of a network, with $c \cong_{\mathcal{N}}^{1} d$, an input tree isomorphism from c to d is an isomorphism of the input trees $\mathcal{I}(c)$ and $\mathcal{I}(d)$ which respects $\sim_{\mathcal{N}}$. The groupoid of all such input tree isomorphisms for c, d has previously been denoted $B(c,d)$: here we introduce the new notation $\mathrm{Iso}^{1}(c,d)$; taking bunching into account, we define the similar structure $\mathrm{Iso}_{\mathcal{N}}^{1}(c,d)$.

Definition 4.1.7. Let $\sim$ be an equivalence relation on a network $\mathcal{N}$; recall that this implies that $\sim \leqslant \sim_{\mathcal{N}}$. Then we call $\sim$ a *balanced equivalence relation* if:

1. $c \sim d$ for cells c and d only where the bunched input sets of c and d are equivalent under $\sim$. That is, $\sim \leqslant \cong^1_\sim$ on cells.

2. $e \sim f$ for edges e and f only where their tails are equivalent: $e \sim f \implies \mathcal{T}(e) \sim \mathcal{T}(f)$. In general, we will define $\sim$ on cells, and let $e \sim f$ precisely when $e \sim_E f$ and $\mathcal{T}(e) \sim \mathcal{T}(f)$. Clearly, this produces the maximal $\sim$ on edges of a network for a given $\sim$ on its cells.

Figure 4.1 shows an example of a balanced equivalence relation.

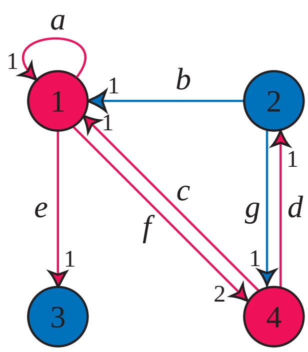

Figure 4.1: A balanced equivalence relation on a network: equivalence classes under the balanced relation are represented by colours.

Definition 4.1.8. Let $\mathcal{N}$ be a network, and $c, d \in \mathcal{C}(\mathcal{N})$. Then $\mathrm{Iso}(\mathcal{I}(c), \mathcal{I}(d))$ is the groupoid of isomorphisms $\mathcal{I}(c) \overset{\cong_{\mathcal{N}}}{\to} \mathcal{I}(d)$ of the input trees of those cells. Define $\mathrm{TreeIso}(c, d)$ to be this group of isomorphisms. For $n \in \mathbb{N}_\infty$, let $\mathrm{Iso}^n(c, d)$ denote the group of n–th input tree isomorphisms, $\mathrm{Iso}(\mathcal{I}^n(c), \mathcal{I}^n(d))$.

Where $c = d$, this group $\mathrm{Iso}^n(c, c)$ is denoted $\mathrm{Aut}^n(c)$.

We also give a version of this definition which accounts for bunching: as remarked previously, this will not alter the algebraic structure here, but is the more general formulation.

Definition 4.1.9. Let $\mathcal{N}$ be a network, and $c, d \in \mathcal{C}(\mathcal{N})$; let $\sim \leqslant \sim_{\mathcal{N}}$. Then $\mathrm{Iso}_{\sim}(\mathcal{I}(c), \mathcal{I}(d))$ is the group of isomorphisms of the input trees of those cells which respect $\sim$:

$$\mathrm{Iso}_{\sim}(\mathcal{I}(c), \mathcal{I}(d)) = \left\{ f : \mathcal{I}(c) \overset{\cong}{\underset{\sim}{\to}} \mathcal{I}(d) \right\}$$

For $n \in \mathbb{N}_{\infty}$, let $\mathrm{Iso}^{n}_{\sim}(c, d)$ denote the group of n–th input tree isomorphisms bunched by $\sim$, respecting $\sim$, $\mathrm{Iso}_{\sim}(\mathcal{I}^{n}(c) /\!\!/ \sim, \mathcal{I}^{n}(d) /\!\!/ \sim)$.

Where $c = d$, this group $\mathrm{Iso}^{n}_{\sim}(c, c)$ is denoted $\mathrm{Aut}^{n}_{\sim}(c)$.

Define $\mathrm{Iso}^{n}_{\sim}(\mathcal{N}) = \bigcup_{c,d \in \mathcal{C}(\mathcal{N})} \mathrm{Iso}^{n}_{\sim}(c, d)$. This forms a groupoid structure, in the sense that it has a binary operation (function composition) which is not defined for all pairs $f, g \in \mathrm{Iso}^{n}_{\sim}(\mathcal{N})$, but satisfies the remainder of the group axioms where it is defined.

Where $\sim = \sim_{\mathcal{N}}$ in these definitions, we may abbreviate it: for example, $\mathrm{Iso}^{n}_{\mathcal{N}}(c, d)$.

The following results are proved in Aldis (2005).

Theorem 4.1.10. *Cells of a network $\mathcal{N}$ which are equivalent under a balanced equivalence relation $\sim$ on $\mathcal{N}$ have infinite input trees which are equivalent with respect to $\sim$. That is, $\sim \leqslant \cong^{\infty}_{\sim}$.*

Corollary 4.1.11. *If $\sim$ is a balanced equivalence relation on a network $\mathcal{N}$, then $\cong^{\infty}_{\sim}$ is the same relation as $\sim$.*

Lemma 4.1.12. *Given a network $\mathcal{N}$ and any equivalence relation $\sim$ on $\mathcal{N}$, the relation $\cong^{\infty}_{\sim}$ is a balanced equivalence relation on $\mathcal{N}$.*

Proposition 4.1.13. *If $\sim$ and $\approx$ are two relations on a network $\mathcal{N}$ such that $\sim \leqslant \approx$, then also $\cong^{\infty}_{\sim} \leqslant \cong^{\infty}_{\approx}$.*

4.1.3 Reduced, Duplicate, and Constraint Sets

In later chapters, we will be interested in how relations fail to be balanced. Here we take an equivalence relation $\sim$ on a network $\mathcal{N}$ and work with a kind of quotient of $\mathcal{N}$ called the 'reduced network', which is equal to the balanced quotient $\mathcal{N}/\!\!/\!\sim$ if and only if $\sim$ is balanced; we are mostly interested in what happens when $\sim$ is not balanced. This is based on Stewart and Parker (Preprint) — that paper uses 'reduced' and 'constraint' sets: we make a new definition of the 'duplicate' set. We also show that if a relation is not balanced, then it is not balanced at some cell, c, in the sense that it is not balanced at (c, r), where r is the cell in a given transversal R of $\sim$ such that $r \sim c$. We later use this to find appropriate cells for perturbation. A new concept we introduce here is that of an 'augmented reduced system', which adds constraint cells back into the reduced system in a controlled way.

Definition 4.1.14. Let $\sim$ be an equivalence relation on (the cells of) a network $\mathcal{N}$.

Take some transversal R of $\sim$. For each cell c, let $R(c)$ be defined by $c \sim R(c) \in R$. Partition the remainder of $\mathcal{C}(\mathcal{N})$ into two parts, C and D:

1. If $\mathcal{I}(c) \sim \mathcal{I}(R(c))$, then $c \in D$.
2. Otherwise, $c \in C$.

Then we call R, D and C the reduced, duplicate and constraint sets for this transversal of $\sim$.

Remark 4.1.15. The cells d in the duplicate set are precisely those where the relation $\sim$ is balanced at (r, d) for all $r \in R$. The constraint set $C \neq \varnothing$ precisely if $\sim$ is unbalanced. The following lemma takes this idea a little further.

Lemma 4.1.16. *Let $\sim$ be an equivalence relation on a network $\mathcal{N}$, and R a transversal of*

$\sim$. If $\sim$ is 1–balanced at $(c, R(c))$ for all $c \in \mathcal{C}(\mathcal{N})$, then $\sim$ is balanced on $\mathcal{N}$.

Proof. Suppose c, d are cells of $\mathcal{N}$ such that $\sim$ is unbalanced at (c, d). Then $c \sim d$, so $R(c) = R(d)$: however, $c \ncong^1_\sim d$.

Consider the following chain of relations:

$$c \overset{?}{\cong}{}^1_\sim R(c) = R(d) \overset{?}{\cong}{}^1_\sim d$$

Since $c \ncong^1_\sim d$, at least one of the relations $\overset{?}{\cong}{}^1_\sim$ must be false. Suppose without loss of generality that it is the left-hand relation. Then $c \ncong^1_\sim R(c)$, but $c \sim R(c)$. So the relation $\sim$ is not balanced at $(c, R(c))$. □

Definition 4.1.17. Suppose R is a reduced set for an equivalence relation $\sim$ on $\mathcal{N}$ — that is, R is a transversal of $\sim$. Then we define the *reduced network* $\mathcal{R} = \mathcal{N} /\!\!/_R \sim\ = (R, \mathcal{E}_\mathcal{R})$ as a kind of 'pseudo-quotient' of $\mathcal{N}$. The cell set of this network is the reduced set R itself. For each $r \in R$, the input set of r in this network, $I_\mathcal{R}(r)$, is defined (with reference to its input set $I_\mathcal{N}(r)$ in the network $\mathcal{N}$) by $I_\mathcal{R}(r) = R(I_\mathcal{N}(r))$, where $R(c, d, t) = (R(c), R(d), t)$. Then the edge set $\mathcal{E}_\mathcal{R} = \bigcup_{r \in R} (I_\mathcal{R}(r))$.

Remark 4.1.18. If d is any cell in D, the pseudo-quotient $\mathcal{N} /\!\!/_R \sim$ is isomorphic to the pseudo-quotient $\mathcal{N} /\!\!/_{R'} \sim$, where $R' = (R \backslash \{R(d)\}) \cup \{d\}$: this is why we call D the 'duplicate' set. For $c \in C$, however, $\mathcal{N} /\!\!/_R \sim\ \ncong \mathcal{N} /\!\!/_{R^*} \sim$, where $R^* = (R \backslash \{R(c)\}) \cup \{c\}$.

If $\sim$ is balanced, then this pseudo-quotient is (isomorphic to) the balanced quotient $\mathcal{N} /\!\!/ \sim$, independently of the choice of R.

We now define a new concept, the importance of which will become apparent later.

Definition 4.1.19. Let R, D, C be reduced, duplicate and constraint sets for an equivalence relation $\sim$ on a network $\mathcal{N}$. Let $c \in C$, then we define the *augmented reduced network* $\mathcal{R}_c = \mathcal{N} /\!\!/_R^c \sim$. The cell set of this network is $R \cup \{c\}$; for each $r \in R$, the input set of r in this network is equal to the input set $I_{\mathcal{R}}$ of r in $\mathcal{R} = \mathcal{N} /\!\!/_R \sim$. The input set of c is defined by $I_{R_c}(c) = R(I_{\mathcal{N}}(c))$, in the same way as the input sets of $r \in R$. Note that c features in no input sets, and therefore has no output edges.

4.2 The Lattice of Balanced Equivalence Relations

As remarked earlier, the ordering $\leqslant$ on equivalence relations gives a lattice with minimal element $=$ and maximal element $\top$. We now prove that under the same partial ordering, the set of balanced equivalence relations also forms a lattice with a minimal element (the relation $=$) and a maximal element (which we denote $\bowtie$). A proof of this was given in Stewart (2007): that proof uses a number of sophisticated techniques to obtain the result for a class of not-necessarily-finite networks which includes all locally finite networks. Here, we detail the proof given in Aldis (2008) which is the major result added to that paper after its submission as Aldis (2005).

Note that the paper where this result appeared used a more general notion of multiplicity than we use here. However, the proof is the same in either context.

Theorem 4.2.1. *For a given network $\mathcal{N}$, the set of balanced equivalence relations on $\mathcal{N}$ forms a complete lattice under the partial order given by refinement. In particular, $\mathcal{N}$ has a maximal and a minimal balanced equivalence relation.*

Proof. For any network $\mathcal{N}$, $=$ is (trivially) a balanced equivalence relation, so the set of balanced equivalence relations on $\mathcal{N}$ is non-empty and bounded below.

Take a non-empty set Y of balanced equivalence relations on $\mathcal{N}$. Then let $\approx$ denote the equivalence relation $\bigvee Y$, its join (taken in the lattice of equivalence relations on $\mathcal{N}$). We show that this relation $\approx$ is balanced: if $c \approx d$, then by definition there is some finite chain $c = c_0 \sim_1 c_1 \sim_2 \cdots \sim_n c_n = d$ with $\sim_i \in Y$ for each i. In particular, each $\sim_i$ is balanced. This means that the input sets of c_{i-1} and c_i are equivalent under $\sim_i$, so they are equivalent under $\approx$. Since equivalence under $\approx$ is an equivalence relation, the bunched input set of $c = c_0$ is isomorphic to that of $c_n = d$ with respect to $\approx$. Finally, we show that $\approx = \bigvee Y$ is a refinement of $\sim_{\mathcal{N}}$. Since all the relations in Y are refinements of the cell and edge equivalences $\sim_{\mathcal{N}}$, this $\sim_{\mathcal{N}}$ is an upper bound for Y in the set of equivalence relations on $\mathcal{N}$. Hence $\bigvee Y \leqslant \sim_{\mathcal{N}}$, as required.

We have shown that the partially-ordered set of balanced equivalence relations on $\mathcal{N}$ has a minimal element, and that any non-empty subset Y of this set has a join. By 2.2.1, the set of balanced equivalence relations on $\mathcal{N}$ is therefore a complete lattice, and so, in particular, it has a maximal element. $\square$

We have shown that there are two relations on $\mathcal{N}$ which are maximal and minimal balanced equivalence relations. As remarked above, the minimal balanced equivalence relation is $=$. The structure of the maximal balanced equivalence relation $\bowtie$ was examined at length in Aldis (2005, 2008). We give the result again here, for completeness.

Theorem 4.2.2. *Given a network $\mathcal{N}$, the maximal balanced equivalence relation $\bowtie$ on $\mathcal{N}$ is given by $c \bowtie d$ if and only if $\mathcal{I}^{\infty}(c) /\!\!/ \sim_{\mathcal{N}} \;\cong_{\mathcal{N}}\; \mathcal{I}^{\infty}(d) /\!\!/ \sim_{\mathcal{N}}$. That is, $\bowtie = \cong_{\mathcal{N}}^{\infty}$.*

Corollary 4.2.3. *If $\sim$ is an equivalence relation which is balanced on $\mathcal{N}$, then $\sim$ is ∞–balanced at all pairs of cells (c, d) from $\mathcal{N}$.*

4.2.1 Submaximal Balanced Equivalence Relations

Aldis (2008) determined the existence of a maximal balanced equivalence relation $\bowtie$ on any network, and described its structure. An interesting question which has not yet been answered would be: what balanced equivalence relations $\sim$ are there on $\mathcal{N}$ such that $\sim < \bowtie$, but if $\approx < \bowtie$ then $\sim \nless \approx$? This would lead to a greater understanding of the structure of the lattice of balanced equivalence relations, and might inform further study.

4.3 Forward Relations

This chapter has so far followed the usual treatment in this field by considering only equivalence relations. We now extend the concept of balance to a more general class of relations: these relations will be useful in examining the 'phase relations' of chapter 10. Most of the concepts transfer directly, although a small amount of formal work would be required to describe bunching.

Definition 4.3.1. Let $\asymp$ be an equivalence relation on a set S, and $\prec$ a relation on S such that:

1. $a \prec b, a \prec c \implies b \asymp c$,
2. $a \prec b, c \prec b \implies a \asymp c$,
3. $a \asymp b \prec c \implies a \prec c$,
4. $a \prec b \asymp c \implies a \prec c$.

For example, with $S = \mathbb{N}$, the relation $\asymp$ could be congruence modulo 13, $\equiv_{13}$, and $n \prec m$ precisely if $n + 6 \equiv_{13} m$.

Then we call $\prec$ a *forward relation* over $\asymp$.

Remark 4.3.2. Despite the notation, forward relations are not, in general, order relations. On the contrary, they describe the relation given by some 'step forward': in addition to the numerical example given above, consider the relation on points on the Earth's surface given by 'is 3km due east of'. Notice that if points x, y are both 3km due east of some other point z, then x and y are the same place; in the other direction, if z is 3km due east of both x and y, then again x and y are equal. Thus this relation is a forward relation over equality.

As an alternative geographical example, let $\asymp$ denote the equivalence relation which relates points of equal longitude. Then let $\prec$ denote 'has a longitude of thirty degrees more than' — this relation would be a forward relation over $\asymp$.

Remark 4.3.3. Any equivalence relation $\sim$ is a forward relation over itself.

Lemma 4.3.4. *Consider the set R of forward relations on a set S over a given equivalence relation $\asymp$ as a partially-ordered set, with refinement order $\leqslant$. While this set is not a lattice, every non-empty subset X of R has a well-defined meet, $\bigwedge X = \bigcap X$, when relations in X are considered as subsets of $S \times S$.*

Proof. Let $\lll \neq \varnothing$ be a set of forward relations on a set S over a given equivalence relation $\asymp$. Let $\curlywedge = \bigwedge \lll$. Then we show that $\curlywedge$ is a forward relation on S over $\asymp$.

Let $\prec$ be any relation in $\lll$. Take $a, b, c \in S$ such that $a \curlywedge b$ and $a \curlywedge c$. Then $a \prec b$ and $a \prec c$, so $b \asymp c$. Similarly, if $a \curlywedge b$ and $c \curlywedge b$, then $a \prec b$ and $c \prec b$, so $a \asymp c$.

Suppose $a \asymp b \curlywedge c$. Then $b \prec c$ for all $\prec \in \lll$, so $a \prec c$ for all these $\prec$. Hence $a \curlywedge c$. Similarly, if $a \curlywedge b \asymp c$, then $a \prec b$ and so $a \prec c$ for all $\prec \in \lll$, hence $a \curlywedge c$.

This shows that $\curlywedge$ is a forward relation, as required.

To see that R is not a lattice (for non-trivial choices of S and $\asymp$), consider $a, b \in S$. Then there is a forward relation $\prec$ such that $a \prec b$ — for example, the relation such that $a' \prec b'$ for $(a', b') \asymp (a, b)$ and $c \nprec d$ for $(c, d) \not\asymp (a, b)$. In addition, $=$ is a forward relation over $\asymp$. If $\curlyvee \leqslant =, \prec$, then $a \curlyvee b$ (by $\prec$) and $a \curlyvee a$ (by $=$), so $a \asymp b$, by property 1 of section 4.3.1. Thus $a \asymp b$ for all $a, b \in S$. This shows that if $\asymp \neq \top$, the set of forward relations over $\asymp$ is not a lattice. □

4.3.1 Balanced Forward Relations

Having defined forward relations, we can now say what it means for such a relation to be balanced. The definitions and results are exactly analogous to those for balanced equivalence relations given earlier in this chapter.

Definition 4.3.5. Let $\prec$ be a forward relation over $\asymp$, where $\prec, \asymp$ are relations on some network $\mathcal{N}$. Then define a relation $\prec_{\asymp}$ on $\mathcal{C}(\mathcal{N})$ by $c \prec_{\asymp} d$ if $\mathcal{I}(c) /\!\!/ \!\asymp\, \prec \mathcal{I}(d) /\!\!/ \!\asymp$.

Further, let $n \in \mathbb{N}_{\infty}$. Define $\prec^n_{\asymp}$ by $c \prec^n_{\asymp} d$ if $\mathcal{I}^n(c) /\!\!/ \!\asymp\, \prec \mathcal{I}^n(d) /\!\!/ \!\asymp$.

Definition 4.3.6. Let $\prec$ be a forward relation over $\asymp$ on a network $\mathcal{N}$. Then we call $\prec$ a *balanced forward relation* (over $\asymp$) if $\prec \leqslant \prec^1_{\asymp}$.

Remark 4.3.7. Recall that an equivalence relation $\sim$ on $\mathcal{N}$ is a forward relation over itself. Notice that $\sim^n_{\sim} = \cong^n_{\sim}$, and so $\sim$ is a balanced equivalence relation exactly when it is a balanced forward relation over itself.

The results from this chapter carry over directly to the class of balanced forward relations.

Theorem 4.3.8. *Cells of a network $\mathcal{N}$ which are equivalent under a balanced forward relation $\prec$ over $\asymp$ on $\mathcal{N}$ have infinite input trees which are equivalent with respect to $\prec$. That is, $\prec \leqslant \prec^{\infty}_{\asymp}$.*

Proof. Since $\prec$ is balanced, there is a $\prec$–preserving isomorphism between the input sets of any pair of $\prec$–related cells. Take two cells c, d of $\mathcal{N}$, and as in Aldis (2005, Theorem 4.2), inductively define an isomorphism $\varphi^{(\infty)}$ of bunched input sets by 'joining' copies of appropriate isomorphisms of input sets at each generation. □

Corollary 4.3.9. *If $\prec$ is a balanced forward relation over $\asymp$ on a network $\mathcal{N}$, then $\prec_{\asymp}^{\infty}$ is the same relation as $\prec$.*

Proof. For cells $c, d \in \mathcal{N}$, theorem 4.3.8 ensures that $c \prec d \implies c \prec_{\asymp}^{\infty}$, so we need only show $c \nprec d \implies c \nprec_{\asymp}^{\infty}$. This is trivial, as c, d are the roots of the input trees $\mathcal{I}^{\infty}(c)$ and $\mathcal{I}^{\infty}(d)$; if $c \nprec d$ then these trees cannot be $\prec$–equivalent. □

Lemma 4.3.10. *Given an equivalence relation $\asymp$ on a network $\mathcal{N}$, and any forward relation $\prec$ over $\asymp$ on $\mathcal{N}$, the relation $\prec_{\asymp}^{\infty}$ is a balanced forward relation on $\mathcal{N}$.*

Proof. The relation $\prec_{\asymp}^{\infty}$ can trivially be seen to refine $\prec$. □

Theorem 4.3.11. *Given a network $\mathcal{N}$ with a forward relation $\prec$ over $\asymp$ on $\mathcal{N}$, the maximal balanced refinement of $\prec$ is $\prec_{\asymp}^{\infty}$.*

Proof. Exactly as theorem 4.2.2. Let $\lhd$ denote the maximal refinement of $\prec$ which is balanced over $\asymp$. Since it is balanced over $\asymp$, we have that $\lhd_{\asymp}^{\infty} = \lhd$ by corollary 4.3.9. Meanwhile, since $\lhd$ refines $\prec$, $\lhd_{\asymp}^{\infty}$ refines $\prec_{\asymp}^{\infty}$. Also, $\prec_{\asymp}^{\infty}$ is balanced, by lemma 4.3.10, so $\prec_{\asymp}^{\infty} \leqslant \lhd$. Thus:

$$\lhd = \lhd_{\asymp}^{\infty} \leqslant \prec_{\asymp}^{\infty} \leqslant \lhd$$

Hence all of these relations are equal, and in particular $\lhd = \prec_{\asymp}^{\infty}$, as required. □

5 Admissible Functions

As mentioned previously, this thesis deals with the concepts of balanced equivalence relations and differential equations. This chapter begins to bring together these two strands, defining the link between functions and networks: 'admissible functions'. In order to do this in a clear and concise way, we introduce a new kind of algebraic structure in section 5.1. Derived from the multiset, and sharing properties with vector spaces, we call this new structure a 'multispace'.

In section 5.2, we use this new 'multispace' concept to define cell and input spaces, over which our 'admissible functions' will be defined. The motivation for this definition is to ensure that $c \cong^1_{\mathcal{N}} d$ precisely when the input space of c is equal to that of d. This result ensures that admissible functions 'fit with' the structure of the given network in a well-defined way. In section 5.2.2, we define these admissible functions and examine what it means for a function to be admissible, in terms of the induced functions on the underlying vector spaces of cell input spaces. Section 5.3 expands upon this by providing some basic 'machinery' to work with admissible functions, starting with lifts and quotients, and moving on to basic extension theorems. The section — and this chapter — finishes with an example to demonstrate these theorems.

5.1 Multispaces

As remarked earlier, a multiset lies conceptually between a set and a tuple. Considering it as an unordered tuple, we now consider the derived analogue of a vector space. Motivated by the representation of vector spaces as sets of tuples, we define a new kind of space, the 'multispace', as a set of multisets. We shall see that this is not a vector space — although, like a normed vector space, it is a metric space, and therefore a topological space. We shall assume from now on that all our vector spaces are normed.

As is well understood, a vector space over a field K, with basis $e_1, \ldots, e_n$, may be represented faithfully by (i.e., is isomorphic to) the vector space $V = \{ (x_1, \ldots, x_n) \mid x_i \in K \; \forall i \}$, with operations of addition and scalar multiplication defined in the obvious manner. Given a norm on a vector space, $\|\cdot\| : V \to \mathbb{R}^+$, we can define a metric d on V by $\mathrm{d}(v, w) = \|w - v\|$.

Given two (or more) vector spaces, V, W, we can define the Cartesian product of these spaces, $V \times W$ in the obvious way; as is well known, this product is also a vector space, and the operation $\times$ is commutative up to isomorphism: that is, $V \times W \cong W \times V$, although $V \times W$ is not generally equal to $W \times V$. Multispaces are also a kind of product of vector spaces, the difference being that our product is commutative up to equality.

Definition 5.1.1. Given a (finite) set of distinct vector spaces $\{V_1, \ldots, V_n\}$, and a corresponding set of multiplicities $\{\lambda_1, \ldots, \lambda_n\}$, we define the *multispace over* $V_1, \ldots, V_n$ *with multiplicities* $\lambda_1, \ldots, \lambda_n$ to be the set $V_1^{[\lambda_1]} \circledast \cdots \circledast V_n^{[\lambda_n]} = \left\{ M_1 \cup \cdots \cup M_n \;\middle|\; M_i \overset{\lambda_i}{\Subset} V_i \right\}$.

We refer to λ_i as the multiplicity of V_i in $V_1^{[\lambda_1]} \circledast \cdots \circledast V_n^{[\lambda_n]}$; if W is a vector space which does not feature in the set $\{V_1, \ldots, V_n\}$, we say that W has multiplicity 0 in $V_1^{[\lambda_1]} \circledast \cdots \circledast V_n^{[\lambda_n]}$. Let $(V : A)$ denote the multiplicity of V in A.

For brevity and clarity, we may also write the multispace over $\{ V_i \}$ with multiplicities $\{ \lambda_i \}$

respectively as $\circledast_i V_i^{[\lambda_i]}$.

Note that if two or more vector spaces are equal, but we wish them to be distinct for use in forming a multispace, we may 'mark' them by forming the usual vector space Cartesian product of each with the singleton set containing a distinct natural number. This trivial fact is important, since it is essential to be clear which spaces are equal and which are merely isomorphic, as we now see as we define what it means for multispaces to be isomorphic.

Definition 5.1.2. Two multispaces $A = \circledast_{1 \leqslant i \leqslant m} V_i^{[\lambda_i]}$, $B = \circledast_{1 \leqslant i \leqslant n} W_i^{[\mu_i]}$ are called *isomorphic* if $n = m$, and there is some permutation $\pi \in S_n$ such that $V_i \cong W_{\pi(i)}$ and $\lambda_i = \mu_{\pi(i)}$ for all i.

For comparison, note that A and B are equal if $n = m$ and there is some permutation $\pi \in S_n$ such that $V_i = W_{\pi(i)}$ and $\lambda_i = \mu_{\pi(i)}$ for all i.

Definition 5.1.3. Given a pair of multispaces A, B over $V_1, \ldots, V_n$ and $W_1, \ldots, W_m$ with multiplicities $\lambda_1, \ldots, \lambda_n$ and $\mu_1, \ldots, \mu_n$ respectively, we define the 'star product' $A * B$ to be a multispace over the union $U = \{ V_i \mid 1 \leqslant i \leqslant n \} \cup \{ W_j \mid 1 \leqslant j \leqslant m \}$. Notice that this union may not be disjoint, in which case it will contain fewer than $m + n$ vector spaces. The multiplicity of $U_k \in U$ in $A * B$ is the sum of the multiplicities of U_k in A and B.

Thus $A * B = \circledast_{u \in U} u^{[(u:A)+(u:B)]}$.

Remark 5.1.4. Recall that if a vector space U_k only appears in one of these multispaces — say, A — then its multiplicity in the other (B) will be 0, and so its multiplicity in $A * B$ will equal that in A; thus if the union $\{ V_i \mid 1 \leqslant i \leqslant n \} \cup \{ W_j \mid 1 \leqslant j \leqslant m \}$ is disjoint, then the multiplicity of each vector space in the star product will simply equal its multiplicity in whichever of A and B it appears.

Whether this union is disjoint or not, the sum of multiplicities of all U_k in $A * B$ will equal the sum of all λ_i and μ_j: that is, $\sum_{u \in U} (u : A * B) = \sum_{i=1}^{n} \lambda_i + \sum_{j=1}^{m} \mu_j$.

Every multispace has a 'reduced form' $M = V_1^{[\lambda_1]} * \cdots * V_n^{[\lambda_n]}$, unique up to order of terms, where $V_i \neq V_j$ for $i \neq j$ and $\lambda_i \neq 0$ for all i. Where we wish to emphasise (or ensure) that this is the case, we may write $V_1^{[\lambda_1]} \circledast \cdots \circledast V_n^{[\lambda_n]}$ in place of $V_1^{[\lambda_1]} * \cdots * V_n^{[\lambda_n]}$.

We may extend the operation $*$ to an accumulative operation on a set of multispaces: given such a set $\mathfrak{M}$, where for each $M \in \mathfrak{M}$, $M = V_{M,1}^{[\lambda_{M,1}]} * \ldots * V_{M,n_M}^{[\lambda_{M,n_M}]}$, let $\mathfrak{V} = \{ V_{M,i} \mid M \in \mathfrak{M},\ 1 \leqslant i \leqslant n_M \}$, and enumerate the distinct spaces of $\mathfrak{V}$ as $\mathfrak{V} = \{V_1, \ldots, V_n\}$. For $1 \leqslant j \leqslant n$, let $\lambda_j = \sum_{M \in \mathfrak{M}} (V_j : M)$. Then $\ast_{\mathfrak{M}} M$ is the multispace $\circledast_i V_i^{[\lambda_i]}$.

Definition 5.1.5. Let A be a multispace, $A = V_1^{[\lambda_1]} * \cdots * V_n^{[\lambda_n]}$. The *underlying space* of A is a vector space $\underset{\sim}{A} = V_1^{\lambda_1} \times \cdots \times V_n^{\lambda_n}$.

Remark 5.1.6. Notice that 'the' underlying space of A is in fact dependent on the expression of A. Different ways of expressing A as a star product give different, but isomorphic, underlying vector spaces. Thus when we consider the underlying space of some multispace, we must make a choice; having chosen an underlying space for a particular multispace, we must stick to our choice. To facilitate this, let all the vector spaces in use be put into an arbitrary, but fixed, order, $\mathcal{V}_1, \mathcal{V}_2 \ldots$, and take the canonical form of a multispace A to be the reduced form defined above, with the terms listed in this order.

Definition 5.1.7. Let $v \in V$ be a vector in the underlying space of a multispace A. Write $v = (v_1, \ldots, v_n)$, where $v_i \in V_i^{\lambda_i}$, and each $v_i = \left(v_i^1, \ldots, v_i^{\lambda_i}\right)$. We define the projection of v into A by $\underset{\sim A}{v} = \underset{\sim A}{\cdot}(v) = \llbracket\ v_i^j \mid 1 \leqslant i \leqslant n,\ 1 \leqslant j \leqslant \lambda_i\ \rrbracket$. Notice that the preimage

$S_{A}^{\{-1\}} \subseteq V$ of a multiset $S \in A$ is a finite set of vectors: in particular, we see that

$$\left|S_{A}^{\{-1\}}\right| = \frac{\prod_{1 \leqslant i \leqslant n} (\lambda_i!)}{\prod_{s \in S} (\epsilon^{A}(s)!)}$$

We also extend the concept of underlying space to Cartesian products of multispaces: $\underset{\smile}{A \times B} = \underset{\sim}{A} \times \underset{\sim}{B}$.

We use this underlying space to define a natural metric on multispaces.

Definition 5.1.8. Let A be a multispace, $A = V_1^{[\lambda_1]} \circledast \cdots \circledast V_n^{[\lambda_n]}$. Let V be the underlying space of A, with metric d_V. Let $S, T \in A$ be two multisets. Define the distance $\mathrm{d}_A(S,T) = \mathrm{d}_V\left(\underset{\sim}{S}_A^{\{-1\}}, \underset{\sim}{T}_A^{\{-1\}}\right) = \inf\left\{ \mathrm{d}_V(v,w) \;\middle|\; v \in \underset{\sim}{\cdot}_A^{\{-1\}}(S),\ w \in \underset{\sim}{\cdot}_A^{\{-1\}}(T) \right\}$.

Lemma 5.1.9. *This d_A is a metric on A.*

Proof. Straightforward: we prove this from first principles. Let $S, T, U \in A$.

1. $\mathrm{d}_A(S,S) = \inf\left\{ \mathrm{d}_V(v,w) \;\middle|\; v, w \in \underset{\sim}{S}_A^{\{-1\}} \right\} = 0.$
2. $\mathrm{d}_A(S,T) = \inf\left\{ \mathrm{d}_V(v,w) \;\middle|\; v \in \underset{\sim}{S}_A^{\{-1\}},\ w \in \underset{\sim}{T}_A^{\{-1\}} \right\}$
   $= \inf\left\{ \mathrm{d}_V(w,v) \;\middle|\; v \in \underset{\sim}{S}_A^{\{-1\}},\ w \in \underset{\sim}{T}_A^{\{-1\}} \right\} = \mathrm{d}_A(T,S).$
3. Since $\underset{\sim}{S}_A^{\{-1\}}$ is finite, for all $s \in \underset{\sim}{S}_A^{\{-1\}}$ there is some $t \in \underset{\sim}{T}_A^{\{-1\}}$ such that $\mathrm{d}_V(s,t) = \mathrm{d}_A(S,T)$. Similarly, let $u \in \underset{\sim}{U}_A^{\{-1\}}$ such that $\mathrm{d}_V(t,u) = \mathrm{d}_A(T,U)$. Then $\mathrm{d}_A(S,U) \leqslant \mathrm{d}_V(s,u) \leqslant \mathrm{d}_V(s,t) + \mathrm{d}_V(t,u) = \mathrm{d}_A(S,T) + \mathrm{d}_A(T,U)$.

□

Remark 5.1.10. Although (A, d) is a metric space, it is not a vector space, since we cannot define a reasonable definition of 'addition'. For example, let $A = \mathbb{R}^{[3]}$. Take $a = [\![1,1,2]\!] \in A$ and $b = [\![1,0,0]\!] \in A$. Then $a + b$ could be defined as $[\![2,1,2]\!]$ or $[\![1,1,3]\!]$, and there is no way to choose between these alternatives. Forming identifications of points to make this addition well defined would always result in a trivial one-point space, which would be useless.

5.2 Functions on Networks

We now define the fundamental link between functions (and therefore systems of differential equations) and networks: the notation of an 'admissible function' on a network, which is a function which 'fits with' the network in a sensible way. While the structure and outcome are similar to those given, for example, in Golubitsky and Stewart (2006); Stewart et al. (2003), we use a slightly different underlying definition, which uses multispaces to make the formulation more obvious. The next section will prove some basic results about these functions.

5.2.1 Cell and Input Spaces

In this section, we define the spaces over which our 'admissible functions' will be defined. These are multispaces, as defined above.

Definition 5.2.1. Given a network $\mathcal{N}$, a *choice of phase space* for $\mathcal{N}$ is a choice of a vector space for each cell type of $\mathcal{N}$. We denote the space associated with the type $[\![c]\!]$ of cell $c \in \mathcal{N}$ by $\mathcal{P}([\![c]\!])$, or just $\mathcal{P}(c)$. The *phase space* of the network $\mathcal{N}$ is then the product space $\mathcal{P}(\mathcal{N}) = \prod_{c \in \mathcal{C}(\mathcal{N})} \mathcal{P}(c)$.

A point $x \in \mathcal{P}(\mathcal{N}) = \mathcal{P}\mathcal{C}(\mathcal{N})$ is called a *value* on the network: the component x_c is called the *value on the cell* c.

We require that the phase spaces of distinct cell types are themselves distinct, although they may be isomorphic. Notice that this is easy to ensure by 'marking' a space with the cell type with which it is associated. Conversely, notice that phase spaces of cells of the same type are actually equal, not merely isomorphic as vector spaces.

We now define some phase spaces in an appropriate way to work well with input trees and bunching as defined earlier.

Definition 5.2.2. Given an edge $e = c \underset{t}{\rightarrow} d$, the phase space of e is the phase space of its tail c, marked with the type t of the edge: $\mathcal{P}\left(c \underset{t}{\rightarrow} d\right) = \mathcal{P}(c) \times \{t\}$.

Given a multiset of edges S, its phase space is the star product of the phase spaces of the edges: $\mathcal{P}(S) = \circledast_{e \in S} \mathcal{P}(e)$.

Definition 5.2.3. Given a tree $\mathcal{T}$ of depth 1 with root r, and a choice of phase spaces on $\mathcal{T}$, we define the *input space* of $\mathcal{T}$ to be $\mathcal{QT} = \mathcal{P}(r) \times \mathcal{P}(\mathcal{E}(\mathcal{T}))$.

Given a cell c in a network $\mathcal{N}$, we define its *input space* to be the input space of its first input tree, $\mathcal{QI}(c) = \mathcal{P}(c) \times \mathcal{P}(\mathcal{E}(\mathcal{I}(c)))$.

Remark 5.2.4. This definition ensures that, if $\mathcal{I}(c) \cong_{\mathcal{N}} \mathcal{I}(d)$, then $\mathcal{QI}(c) = \mathcal{QI}(d)$. In fact, this implication is true in both directions.

5.2.2 Admissible Functions

We now define the central notion of an 'admissible function'. We make this initial definition in terms of the multispaces $\mathcal{QI}(c)$ — however, we then consider these functions as functions on the underlying spaces of the $\mathcal{QI}(c)$, which returns us to the same definitions as Golubitsky and Stewart (2006).

Definition 5.2.5. A function f is said to be *admissible* at a cell c of a network $\mathcal{N}$ if $f : \mathcal{QI}(c) \to \mathcal{P}(c)$. A function $F : \bigcup_c \mathcal{QI}(c) \to \bigcup_c \mathcal{P}(c)$ is called an *admissible function* on $\mathcal{N}$ if it can be expressed as the union of admissible functions at the cells, $F = \bigcup_c f_c : \mathcal{QI}(c) \to \mathcal{P}(c)$. Note that because F must be a well-defined function, this ensures that $c \sim_{\mathcal{N}} d \implies f_c \equiv f_d$.

When defining a function f on $\mathcal{QI}(c)$, it is natural to consider it as a function $\underset{\sim}{f}$ on the underlying space $\underset{\sim}{\mathcal{QI}}(c)$ of $\mathcal{QI}(c)$. This space is $\mathcal{PCI}(c)$, which is the product space of the phase spaces of the cells of the input tree of c. Note that this space includes both the internal phase space $\mathcal{P}(c)$ and the 'coupling phase space' $\mathcal{TI}(c)$ of Golubitsky et al. (2005). The following diagram shows the relationship between f and $\underset{\sim}{f}$:

$$\begin{array}{ccc} \mathcal{QI}(c) & \xrightarrow{\overset{\cdot}{\sim}_{\mathcal{QI}(c)}} & \mathcal{PCI}(c) \\ & {\scriptstyle f} \searrow \quad \swarrow {\scriptstyle \underset{\sim}{f}} & \\ & \mathcal{P}(c) & \end{array}$$

Motivated by this method of defining functions, we now characterise the admissible functions on $\mathcal{N}$ in terms of functions on the underlying spaces $\mathcal{PCI}(c)$.

Definition 5.2.6. A function f is said to satisfy the *domain condition at a cell* c of a network $\mathcal{N}$ if

$$f : \mathcal{PCI}(c) \to \mathcal{P}(c)$$

We may describe a function $F_c : \mathcal{P}(\mathcal{N}) \to \mathcal{P}(c)$ as satisfying the domain condition at c if F_c depends only on the values on cells in $\mathcal{CI}(c)$: in other words, there is some f_c such that this diagram commutes:

$$\begin{array}{ccc} \mathcal{P}(\mathcal{N}) & \xrightarrow{\pi} & \mathcal{PCI}(c) \\ & {\scriptstyle F_c}\searrow \quad \swarrow{\scriptstyle f_c} & \\ & \mathcal{P}(c) & \end{array}$$

where π is the natural projection $\mathcal{P}(\mathcal{N}) \to \mathcal{PCI}(c)$. (Note that to appear in this diagram in this way, f_c must satisfy the domain condition as formulated above.)

Definition 5.2.7. A function f is said to be *admissible at a cell* c of a network $\mathcal{N}$ if it satisfies the domain condition at c, and:

$$f\left(x_{\mathcal{C}\beta\mathcal{I}(c)}\right) = f\left(x_{\mathcal{CI}(c)}\right) \quad \forall\beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)\,,\ x \in \mathcal{P}(\mathcal{C}(\mathcal{N}))$$

A pair of functions (f, g) are said to be *admissible at a pair of cells* (c, d) if f, g satisfy the domain condition at c, d respectively, and:

$$g\left(x_{\mathcal{C}\beta\mathcal{I}(c)}\right) = f\left(x_{\mathcal{CI}(c)}\right) \quad \forall\beta \in \mathrm{Iso}^1_{\mathcal{N}}(c, d)\,,\ x \in \mathcal{P}(\mathcal{C}(\mathcal{N})) \quad (*)$$

Remark 5.2.8. The condition $(*)$ is called the *pushforward condition* — it is exactly equivalent to the standard 'pullback condition' — as used, for example, in Golubitsky and Stewart (2006); Stewart et al. (2003); Stewart and Parker (December 2007, 2008), and Golubitsky et al. (2005, equation 2.1). It is easy to see that f is admissible at c precisely when (f, f) is admissible at (c, c).

Definition 5.2.9. A set of functions $F = \{\ f_c : \mathcal{PCI}(c) \to \mathcal{P}(c) \mid c \in \mathcal{C}(\mathcal{N})\ \}$ is said to be *admissible over* $\mathcal{N}$ if (f_c, f_d) is admissible at (c, d) for all $c, d \in \mathcal{C}(\mathcal{N})$.

We describe $F : \mathcal{PC}(\mathcal{N}) \to \mathcal{PC}(\mathcal{N})$ as admissible if each component $F_c : \mathcal{PC}(\mathcal{N}) \to \mathcal{P}(c)$ is admissible.

Definition 5.2.10. For any $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(c,d)$ which is a function $\beta : \mathcal{I}(c) \to \mathcal{I}(d)$, let β also denote a function $\mathcal{PCI}(c) \to \mathcal{PCI}(d)$, defined by $(\beta x)_{\mathcal{CI}(c)} = x_{\mathcal{C}\beta\mathcal{I}(c)}$. Then this diagram commutes:

$$\begin{array}{ccc} \mathcal{CI}(c) & \xrightarrow{\beta} & \mathcal{CI}(d) \\ {\scriptstyle \mathcal{P}}\big\downarrow & & {\scriptstyle \mathcal{P}}\big\downarrow \\ \mathcal{PCI}(c) & \xrightarrow{\beta} & \mathcal{PCI}(d) \end{array}$$

We call $\beta : \mathcal{PCI}(c) \to \mathcal{PCI}(d)$ the *pushforward map* of $\beta : \mathcal{CI}(c) \to \mathcal{CI}(d)$.

Remark 5.2.11. With this definition, we can characterise the pushforward condition by saying that the following diagram commutes for all $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(c,d)$:

$$\begin{array}{ccc} \mathcal{PCI}(c) & \xrightarrow{\beta} & \mathcal{PCI}(d) \\ {\scriptstyle f_c}\big\downarrow & & {\scriptstyle f_d}\big\downarrow \\ \mathcal{P}(c) & = & \mathcal{P}(d) \end{array}$$

Lemma 5.2.12. *Suppose $\sim$ is a balanced equivalence relation on a network $\mathcal{N}$, and $x \in \mathcal{PC}(\mathcal{N})$ such that $c \sim d \implies x_c = x_d$. Let F be an admissible function on the network. Then $F(x)$ has the same symmetry property: $c \sim d \implies [F(x)]_c = [F(x)]_d$.*

Proof. Immediate from the definition of admissible functions. □

Lemma 5.2.13 (Vector Fields by Input Class)**.** *Let $C_{\mathcal{I}}$ be a transversal of the input classes of $\mathcal{N}$, and for each $c \in C_{\mathcal{I}}$, let f_c be a function which is admissible at c.*

Then there is an admissible vector field F over $\mathcal{N}$ such that $F_c = f_c\pi_c$ (where π_c is the natural projection $\mathcal{PC}(\mathcal{N}) \to \mathcal{PCI}(c)$).

Proof. The vector field F defined by the property $F_c = f_c\pi_c$ is admissible by definition. □

5.3 Operations with Admissible Functions

Lemma 5.2.13 shows that we can construct an admissible function piecewise by defining admissible components for each input class and then putting them together: in other words, unions of admissible functions are admissible. In this section, we give examples of other operations which preserve admissibility; we then go on to give a method of extending a function which is 'admissible in a small region of $\mathcal{PCI}(c)$' (for a given cell c) to an admissible function on the whole of $\mathcal{N}$.

5.3.1 Lifts and Quotients

We now give some results which show that admissible functions remain admissible under pseudo-quotients and lifts; we also provide extensions of these results to augmented quotient systems.

Lemma 5.3.1 (Admissible Quotient Lemma)**.** *Suppose R is a reduced set for an equivalence relation $\sim$ on $\mathcal{N}$; further, suppose F is an admissible function on $\mathcal{N}$. Let the* natural quotient system *F_R be a function on $\mathcal{R} = \mathcal{N} /\!\!/_R \sim$ given by $F_R(x) = F\left(x^R\right)$ where for every $x \in \mathcal{P}(\mathcal{R})$, $x^R \in \mathcal{P}(\mathcal{N})$ is defined by $x^R_c = x_{R(c)}$. Then F_R is an admissible function on $\mathcal{R}$.*

Proof. This proof will require us to use both input networks and functions in $\mathcal{N}$ and those on $\mathcal{R}$, so we introduce some temporary notation to distinguish between them for now. Let $\mathcal{I}_{\mathcal{N}}(c)$ denote the input network of c in the network $\mathcal{N}$, and similarly $\mathcal{I}_{\mathcal{R}}(c)$ for $\mathcal{R}$. Let $[f_{\mathcal{N}}]_c$ here denote the function we usually write $f_c : \mathcal{PCI}_{\mathcal{N}}(c) \to \mathcal{P}(c)$: the corresponding function in $\mathcal{R}$ will be denoted $[f_{\mathcal{R}}]_c : \mathcal{PCI}_{\mathcal{R}}(c) \to \mathcal{P}(c)$. The first stage of the proof is to ensure that this function $[f_{\mathcal{R}}]_c$ exists.

The domain condition requires that the component of $F_R(x)$ associated with the cell r, denoted $[F_R(x)]_r$, must only depend upon $x_{\mathcal{CI}_{\mathcal{R}}(r)}$, so it may reasonably be described by a function $[f_{\mathcal{R}}]_c$ as above. By definition $[F_R(x)]_r = F\left(x^R\right)$; by the domain condition for F, this only depends upon $x^R_{\mathcal{CI}_{\mathcal{N}}(r)}$. By construction of x^R, $x_c = x_{R(c)}$ for each cell $c \in \mathcal{CI}_{\mathcal{N}}(r)$. But $R(c) \in \mathcal{CI}_{\mathcal{R}}(r)$. Hence $x^R_{\mathcal{CI}_{\mathcal{N}}(r)} = x_{\mathcal{CI}_{\mathcal{R}}(r)}$, proving the domain condition. Taking $\pi_r, \pi_{\mathcal{CI}_{\mathcal{N}}(r)}, \pi_{\mathcal{CI}_{\mathcal{R}}(r)}$ as the natural projections, and letting $\pi^R(x) = x^R$ for $x \in \mathcal{PC}(\mathcal{R})$, this can be summarised using this diagram:

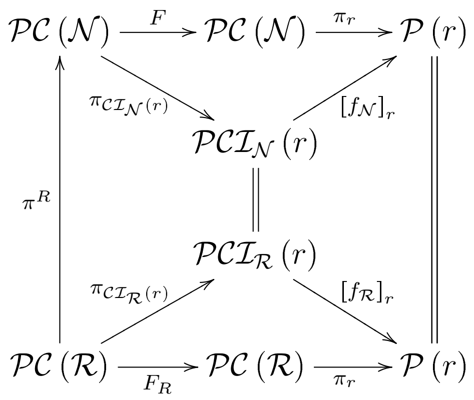

The definition of F_R is precisely that the outermost arrows commute. Note that $\mathcal{CI}_{\mathcal{R}}(r) \sim_{\mathcal{C}} \mathcal{CI}_{\mathcal{N}}(r)$, so $\mathcal{PCI}_{\mathcal{R}}(r) = \mathcal{PCI}_{\mathcal{N}}(r)$; the left section commutes by trivial properties of projections. The domain condition in $\mathcal{N}$ guarantees the existence of $[f_{\mathcal{N}}]_r$ such that the top section commutes — the right section then gives the existence of $[f_{\mathcal{R}}]_r$ such that the diagram commutes: the commutativity of the bottom section proves the domain condition in $\mathcal{R}$.

We prove the pushforward condition entirely by diagram chasing. As above, we recall that $\mathcal{CI}_{\mathcal{R}}(c) \sim_{\mathcal{C}} \mathcal{CI}_{\mathcal{N}}(c)$, so $\mathcal{PCI}_{\mathcal{R}}(c) = \mathcal{PCI}_{\mathcal{N}}(c)$. Now notice that for all $c \in R$, the construction of $[f_{\mathcal{R}}]_c$ is such that this diagram commutes:

$$\begin{array}{ccc} \mathcal{PCI}_{\mathcal{R}}(c) & = & \mathcal{PCI}_{\mathcal{N}}(c) \\ & {\scriptstyle [f_{\mathcal{R}}]_c} \searrow \quad \swarrow {\scriptstyle [f_{\mathcal{N}}]_c} & \\ & \mathcal{P}(c) & \end{array}$$

Suppose that there is some $\beta \in \mathrm{Iso}^1_{\mathcal{R}}(c,d)$ for $c, d \in R$. Define β^R such that this diagram commutes:

$$\begin{array}{ccc} \mathcal{CI}_{\mathcal{N}}(c) & \xrightarrow{\beta^R} & \mathcal{CI}_{\mathcal{N}}(d) \\ {\scriptstyle \cong}\uparrow & & \uparrow{\scriptstyle \cong} \\ \mathcal{CI}_{\mathcal{R}}(c) & \xrightarrow{\beta} & \mathcal{CI}_{\mathcal{R}}(d) \end{array}$$

This $\beta^R \in B_{\mathcal{N}}(c,d)$: it is an isomorphism by definition. By the pushforward condition in $\mathcal{N}$, the following diagram commutes:

$$\begin{array}{ccc} \mathcal{PCI}_{\mathcal{N}}(c) & \xrightarrow{\beta^R} & \mathcal{PCI}_{\mathcal{N}}(d) \\ {\scriptstyle [f_{\mathcal{N}}]_c}\downarrow & & \downarrow{\scriptstyle [f_{\mathcal{N}}]_d} \\ \mathcal{P}(c) & = & \mathcal{P}(d) \end{array}$$

Then put these diagrams together to see that this diagram commutes:

$$\begin{array}{ccccc} \mathcal{PCI}_{\mathcal{N}}(c) & & \xrightarrow{\beta^R} & & \mathcal{PCI}_{\mathcal{N}}(d) \\ \| & {\scriptstyle [f_{\mathcal{N}}]_c}\searrow & & \swarrow{\scriptstyle [f_{\mathcal{N}}]_d} & \| \\ \| & & \mathcal{P}(c) = \mathcal{P}(d) & & \| \\ \| & {\scriptstyle [f_{\mathcal{R}}]_c}\nearrow & & \nwarrow{\scriptstyle [f_{\mathcal{R}}]_d} & \| \\ \mathcal{PCI}_{\mathcal{R}}(c) & & \xrightarrow[\beta]{} & & \mathcal{PCI}_{\mathcal{R}}(d) \end{array}$$

The commutativity of the bottom section of this diagram is the required pushforward condition in $\mathcal{R}$. □

Lemma 5.3.2 (Admissible Lift Lemma). *Let $\mathcal{N}$ be a network, and $\sim$ an equivalence relation on $\mathcal{N}$ which respects cell type. Let R, D, C be some reduced, duplicate and constraint systems for $\sim$. Let F_R be an admissible function over the network $\mathcal{R} = \mathcal{N} /\!\!/_R \sim$.*

Then the lift $F = F_R^{\sim}$, defined for each $c \in \mathcal{C}(\mathcal{N})$ as $[f_R^{\sim}]_c = [f_R]_{R(c)}$, is an admissible function on $\mathcal{N}$, with $F /\!\!/_R \sim = F_R$.

Further, if F_1, F_2 are two admissible functions over $\mathcal{R}$, with $\|F_1 - F_2\| < \varepsilon$, then $\|F_1^{\sim} - F_2^{\sim}\| < \varepsilon$.

Proof. The domain condition holds trivially by the method of definition.

The proof of the pushforward condition is almost exactly the reverse of that in the proof of lemma 5.3.1: we continue to use the notation $[f_{\mathcal{R}}]_c$ and $[f_{\mathcal{N}}]_c$ as before. Take $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(c, d)$ for $c, d \in \mathcal{C}(\mathcal{N})$. Define β_R such that the following diagram commutes:

$$\begin{array}{ccc} \mathcal{CI}_{\mathcal{N}}(c) & \xrightarrow{\beta} & \mathcal{CI}_{\mathcal{N}}(d) \\ \cong \uparrow & & \uparrow \cong \\ \mathcal{CI}_{\mathcal{R}}(c) & \xrightarrow{\beta_R} & \mathcal{CI}_{\mathcal{R}}(d) \end{array}$$

This $\beta_R \in B_{\mathcal{R}}(c, d)$: it is an isomorphism by definition. By the pushforward condition in $\mathcal{R}$, the following diagram commutes:

$$\begin{array}{ccc} \mathcal{PCI}_{\mathcal{R}}(c) & \xrightarrow{\beta_R} & \mathcal{PCI}_{\mathcal{R}}(d) \\ {\scriptstyle [f_{\mathcal{R}}]_c} \downarrow & & \downarrow {\scriptstyle [f_{\mathcal{R}}]_d} \\ \mathcal{P}(c) & = & \mathcal{P}(d) \end{array}$$

Then put these diagrams together to see that this diagram commutes:

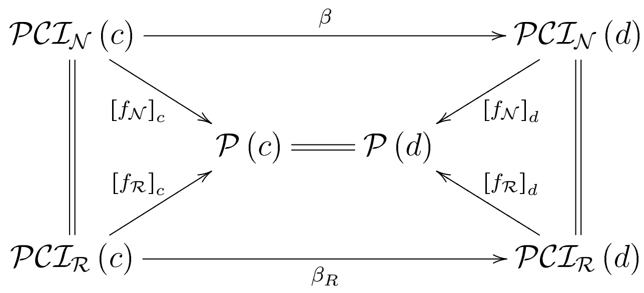

The commutativity of the top section of this diagram is the required pushforward condition in $\mathcal{N}$.

Finally, we note that for functions F_1, F_2 as specified,

$$\|F_1^{\sim} - F_2^{\sim}\| = \sup_{c \in \mathcal{C}(\mathcal{N})} \|[f_1^{\sim}]_c - [f_2^{\sim}]_c\| = \sup_{c \in \mathcal{C}(\mathcal{N})} \left\| [f_1^{\sim}]_{R(c)} - [f_2^{\sim}]_{R(c)} \right\|$$

$$= \sup_{r \in R} \|[f_1^{\sim}]_r - [f_2^{\sim}]_r\| = \|F_1 - F_2\| < \varepsilon$$

as required. □

Lemma 5.3.3. *Let $\sim$ be an equivalence relation on a network $\mathcal{N}$ which respects cell type. Let R, D, C be some reduced, duplicate and constraint systems for $\sim$, and $c \in C$ some constraint cell. Let F be an admissible function on $\mathcal{N}$.*

Then the following equalities hold:

1. $(\mathcal{N} /\!\!/_R^c \sim) /\!\!/_R \sim = \mathcal{N} /\!\!/_R \sim$
2. $[F_{R^+}]_R = F_R$

Proof. Trivial, and omitted. □

Lemma 5.3.4. *Suppose R is a reduced set for an equivalence relation $\sim$ on $\mathcal{N}$, with* *constraint set C. Let $c \in C$, and F be an admissible function on $\mathcal{N}$. Let $R^+ = R \cup \{c\}$. Let the natural augmented quotient system F_R^c be given by $F_R^c(x)|_R = F_R(x_R)$. Then F_R^c is an admissible function on $\mathcal{R}^+ = \mathcal{N} /\!\!/_R^c \sim$.*

Proof. Consider the natural quotient system F_R on $\mathcal{R}$. Then by lemma 5.3.1, this is an admissible function over $\mathcal{R} = \mathcal{N} /\!\!/_R \sim = (\mathcal{N} /\!\!/_R^c \sim) /\!\!/_R \sim = \mathcal{R}^+ /\!\!/_R \sim$. By lemma 5.3.2, we have the admissble function F_R^c. □

Lemma 5.3.5. *Let $\mathcal{N}$ be a network, and $\sim$ an equivalence relation on $\mathcal{N}$ which respects cell type. Let R, D, C be some reduced, duplicate and constraint systems for $\sim$. Let $c \in C$, and $R^+ = R \cup \{c\}$. Let F_{R^+} be an admissible function over the network $\mathcal{R} = \mathcal{N} /\!\!/_R^c \sim$.*

Then the lift $F_{R^+}^{\sim}$, defined as in lemma 5.3.2, is an admissible function.

Proof. The proof is exactly identical to that of lemma 5.3.2. □

5.3.2 Symmetrisation Extensions

This section deals with making admissible functions which are extensions of other, appropriate, functions.

The nature of admissible functions means that constructing a desired function 'all at once' is often difficult. To simplify matters, we use this section to show that we can define admissible functions piecewise, in a number of ways. Firstly, the 'pushforward proposition' extends a function defined at one cell of a network $\mathcal{N}$ into a function defined on all cells of that network. Then we show (the 'symmetrisation proposition') that a suitable function defined in a small region of $\mathcal{PCI}(c)$ can be extended to an admissible function on the whole of $\mathcal{PCI}(c)$. Putting these two results together, we see that a suitable function defined on a small region of $\mathcal{PCI}(c)$ can be extended to an admissible function on $\mathcal{N}$.

We conclude the section with an example which demonstrates this process.

Proposition 5.3.6 (Pushforward Proposition)**.** *Let $\mathcal{N}$ be a network with a cell c, and $f : \mathcal{PCI}(c) \to \mathcal{P}(c)$ some function which is admissible at c in $\mathcal{N}$.*

Then there exists a function $F : \mathcal{P}(\mathcal{C}(\mathcal{N})) \to \mathcal{P}(\mathcal{C}(\mathcal{N}))$ such that:

1. *The function F is admissible over $\mathcal{N}$.*

2. *The component of F corresponding to the cell c, $f_c = f$.*

Proof. For each $d \in \mathcal{C}(\mathcal{N})$ such that $d \sim_I c$, take some $\beta_{d,c} \in \mathrm{Iso}^1_{\mathcal{N}}(d, c)$. Then we let $f_d = f_c\beta_{d,c}$, so this diagram commutes:

$$\begin{array}{ccc} \mathcal{PCI}(d) & \xrightarrow{\beta_{d,c}} & \mathcal{PCI}(d) \\ {\scriptstyle f_d}\downarrow & & \downarrow{\scriptstyle f} \\ \mathcal{P}(d) & = & \mathcal{P}(c) \end{array}$$

We wish to show that the pair $(f_d, f_{d'})$ is admissible at (d, d') for all $d, d' \sim_I c$; in other words, $f_{d'}(\mathcal{C}\beta\mathcal{I}(d)) = f_d(\mathcal{CI}(d))$ for all $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(d, d')$.

Take $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(d, d')$. Then set $\beta_{c,c} = \beta_{d',c}\beta\beta^{-1}_{d,c}$, so $\beta = \beta^{-1}_{d',c}\beta_{c,c}\beta_{d,c}$. In other words, this diagram commutes:

$$\mathcal{PCI}(d) \xrightarrow[\beta_{d,c}]{} \mathcal{PCI}(c) \xrightarrow[\beta_{c,c}]{} \mathcal{PCI}(c) \xrightarrow[\beta^{-1}_{d',c}]{} \mathcal{PCI}(d'), \qquad \mathcal{PCI}(d) \xrightarrow{\beta} \mathcal{PCI}(d')$$

Now $\beta_{c,c} \in \mathrm{Aut}^1_{\mathcal{N}}(c)$, so $f\beta_{c,c} = f$ since f is admissible at c; this diagram commutes:

$$\begin{array}{ccc} \mathcal{PCI}(c) & \xrightarrow{\beta_{c,c}} & \mathcal{PCI}(c) \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f} \\ \mathcal{P}(c) & = & \mathcal{P}(c) \end{array}$$

Finally,

$$f_{d'}\beta = f_c\beta_{d',c}\beta = f_c\beta_{d',c}\beta^{-1}_{d',c}\beta\beta_{d,c} = f_c\beta\beta_{d,c} = f_c\beta_{d,c} = f_d$$

as required. This is illustrated in the following commutative diagram:

$$\begin{array}{ccccccc} & & & \overset{\beta}{\curvearrowright} & & & \\ \mathcal{PCI}(d) & \xrightarrow[\beta_{d,c}]{} & \mathcal{PCI}(c) & \xrightarrow[\beta_{c,c}]{} & \mathcal{PCI}(c) & \xrightarrow[\beta^{-1}_{d',c}]{} & \mathcal{PCI}(d') \\ {\scriptstyle f_d}\downarrow & & \downarrow{\scriptstyle f} & & {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f_{d'}} \\ \mathcal{P}(d) & = & \mathcal{P}(c) & = & \mathcal{P}(c) & = & \mathcal{P}(d') \end{array}$$

The required property is shown by the outside set of arrows. □

Definition 5.3.7. Let $\mathcal{N}$ be a network with a cell c, and $x^* \in \mathcal{P}(\mathcal{C}(\mathcal{N}))$.

For each $\beta \in \mathrm{Iso}^n_{\sim}(\mathcal{N})$, let $r_\beta = \frac{1}{2}\,\mathrm{d}\left(x^*_{\mathcal{CI}(c)}, x^*_{\mathcal{C}\beta\mathcal{I}(c)}\right)$; set $r = \min\{\, r_\beta \mid \beta \in \mathrm{Iso}^n_{\sim}(\mathcal{N}), r_\beta \neq 0 \,\}$.

Then $r > 0$ is called the *radius of disjoint symmetric neighbourhoods* of x^* at c.

We now come to the main result of this section, which shows that a sufficiently symmetrical function defined on a sufficiently small region in $\mathcal{PCI}(c)$ can be extended to an admissible function on $\mathcal{N}$. The proof is somewhat technical; we follow it with an illustrative example.

Proposition 5.3.8 (Symmetrisation Proposition)**.** *Let $\mathcal{N}$ be a network with a cell c, and $X \subseteq \mathcal{PCI}(c)$.*

Let $\tilde{X} = \bigcup_{\beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)} \beta X$.

Let $f^* : X \to \mathcal{P}(c)$ *be some function such that* $f^*(x) = f^*\beta(x)$ *for all* $x \in X$ *where* $\beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)$ *such that* $\beta x \in X$*; we call a function with this property* admissible in X at c*.*
Let $f^+ : \mathcal{PCI}(c) \to \mathcal{P}(c)$ *be some function which is admissible at* c*.*

Then there is some $f : \mathcal{PCI}(c) \to \mathcal{P}(c)$ *which is admissible at* c*, such that* $f|_X = f^*$*, and* $f\Big|_{\mathcal{PCI}(c)\setminus\tilde{X}} = f^+\Big|_{\mathcal{P}(c)\setminus\tilde{X}}$*.*

Thus by proposition 5.3.6, there is an admissible function F *on* $\mathcal{N}$ *such that* $f_c|_X = f^*$*, where* f_c *is the component of* F *corresponding to the cell* c*.*

Proof. For $\beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)$, and $x \in X$, we let $f(\beta x) = f(x)$, which defines $f(x)$ for all $x \in \tilde{X}$ exactly as required.

The crucial observation is that for $\beta, \beta' \in \mathrm{Aut}^1_{\mathcal{N}}(c)$, if the points βx and $\beta' x'$ are equal for some $x, x' \in X$ and $\beta, \beta' \in \mathrm{Aut}^1_{\mathcal{N}}(c)$, then $\beta^* = \beta'\beta^{-1} \in \mathrm{Aut}^1_{\mathcal{N}}(c)$ such that $\beta^* x = x'$. Thus $f^*\beta^* = f^*$ in X; this means $f^*\beta'\beta^{-1} = f^*$, so $f^*\beta = f^*\beta'$, and f is well-defined, as required.

Now we may let f be any admissible function outside $\tilde{X}$ (here, f^+), and we have an admissible f. □

Remark 5.3.9. The functions f^* and f^+ interact in the definition of f only at the boundary $\partial\tilde{X}$, where the value of f^* is derived from the value at ∂X, and the value of f^+ must satisfy the pushforward condition. Thus, if f^* and f^+ are C^n functions which agree in value and all n derivatives at the boundary ∂X, f will also be C^n.

We now give the promised illustrative example of the Symmetrisation Proposition.

Example 5.3.10. Consider the network pictured in figure 5.1. Let $\mathcal{P}(\bigcirc) = \mathbb{R}$, and $\mathcal{P}(\square) = \mathbb{C}$. Let the underlying space of the input space of cell 1 be considered in the

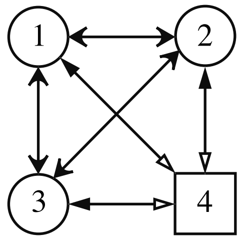

Figure 5.1: An example network, with 4 cells. In order to simplify the diagram, edge types are shown by arrowhead styles.

order $(x_1; x_2, x_3, x_4)$. By symmetry, $(x_1; x_3, x_2, x_4)$ would be an equally valid representation: this information is represented by the cell input isomorphism which transposes cells 2 and 3 in the input tree of cell 1. The pushforward of this isomorphism takes $(w; x, y, z)$ to $(w; y, x, z)$. As is common in this area, we use a line over the permuted coordinates to express that a function has the appropriate symmetry for this condition: writing a function as $f\left(w; \overline{x, y}, z\right)$ means that $f\left(w; x, y, z\right) = f\left(w; y, x, z\right)$.

Let $c = 1$, and let $f^{+}\left(w; \overline{x, y}, z\right) \equiv 0$, so f^{+} trivially satisfies the conditions for admissibility at c.

Let C be the cube $\{\ (w, x, y, z) \mid |w|, |x|, |y|, |z| < 1\ \}$, as shown in figure 5.2, and set $X = C + (1, 2, 3, 4)$. Let $b\left(w, x, y, z\right) = \operatorname{Bump} \frac{\left|\left(\frac{w}{2}, x, y, \frac{z}{3}\right)\right|}{10}$.

Now let

$$f^{*}\left(w; x, y, z\right) = b\left((w, x, y, z) - (1, 2, 3, 4)\right) + \operatorname{Bump}\left(\operatorname{d}\left((w, x, y, z), (10, 9, 8, 7 + 6i)\right)\right)$$

Note that f^{*} is not admissible at cell 1 due to the second term in this definition: for example, let $\boldsymbol{x} = (10, 9, 8, 7 + 6i)$ and $\beta\left(w, x, y, z\right) = (w, y, x, z)$, then $f^{*}\left(\boldsymbol{x}\right) = 1 \neq 0 = f^{*}\left(10, 8, 9, 7 + 6i\right) = f^{*}\beta\boldsymbol{x}$. Inside X, however, the function f^{*} is symmetric, since the second term is identically zero there, and therefore f^{*} is admissible at cell 1 inside X.

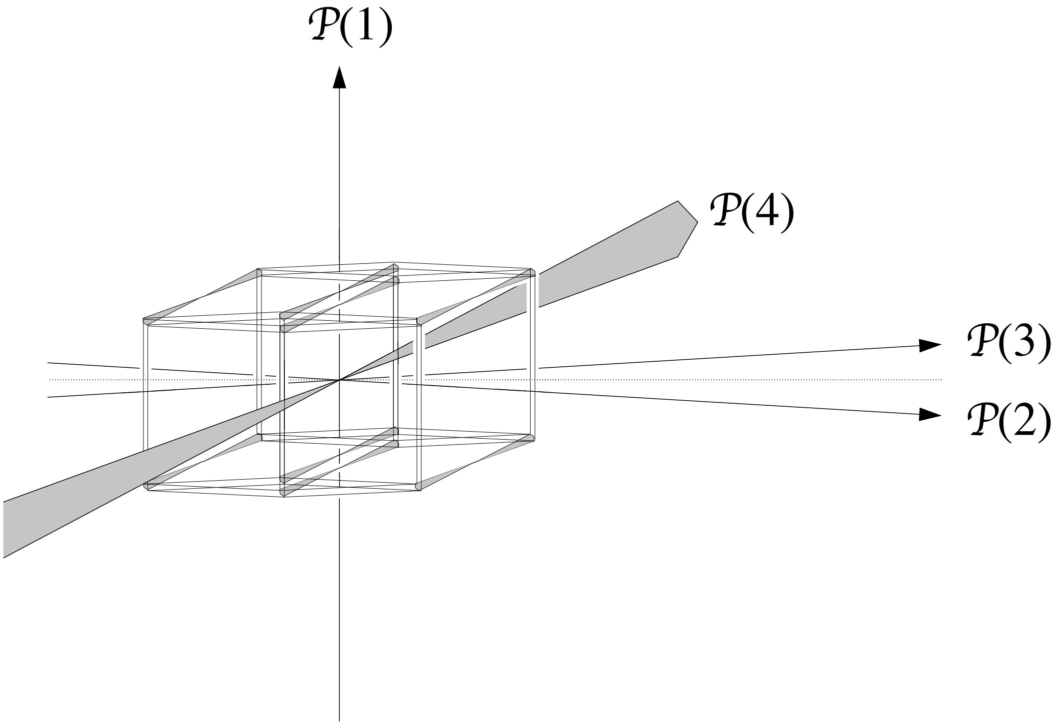

Figure 5.2: The (5–dimensional) cube C.

As in proposition 5.3.8, we let $\tilde{X}$ denote the symmetric images of X, so $\tilde{X} = (C + (1,2,3,4)) \cup (C + (1,3,2,4))$. Let $f = f^*$ inside $\tilde{X}$, and $f = f^+$ elsewhere. The situation is shown in figure 5.3.

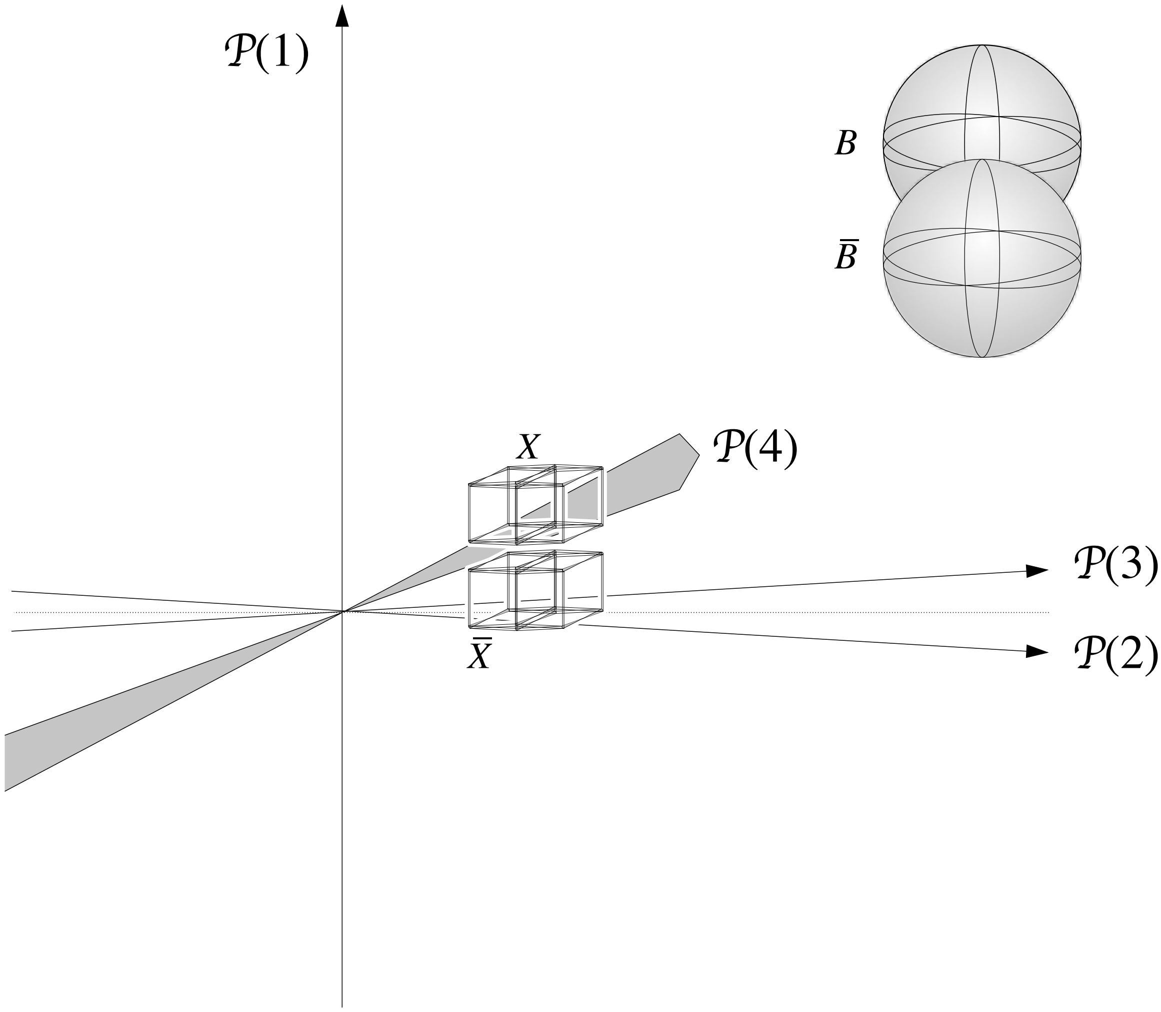

Figure 5.3: The construction of f.

Let $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{PCI}(c)$, and take $\beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)$ such that $\boldsymbol{x} = \beta\boldsymbol{y}$. We show that $\boldsymbol{x} \in \tilde{X} \iff \boldsymbol{y} \in \tilde{X}$.

As remarked above, $\mathrm{Aut}^1_{\mathcal{N}}(c) = \{\mathrm{id}, (2\ 3)\}$. If $\beta = \mathrm{id}$, the result is obvious. Let $\beta = (2\ 3)$. Since $\beta = \beta^{-1}$, we may assume without loss of generality that $\boldsymbol{x} = (w, x, y, z) \in \tilde{X}$, and show that $\boldsymbol{y} = (w, y, x, z) \in \tilde{X}$. If $\boldsymbol{x} = (w, x, y, z) \in \tilde{X}$, then either $\boldsymbol{x} \in X$, so $0 < w < 2,\ 1 < x < 3,\ 2 < y < 4,\ 3 < z < 5$, or $\boldsymbol{x} \in \beta X$, so $0 < w < 2,\ 1 < y < 3,\ 2 < x < 4,\ 3 < z < 5$. In the first case, $\boldsymbol{y} \in \beta X$; in the second, $\boldsymbol{y} \in X$. In either case, this shows that $\boldsymbol{y} \in \tilde{X}$.

This shows that $\boldsymbol{x} \in \tilde{X} \iff \boldsymbol{y} \in \tilde{X}$, so we need only consider two cases: either $\boldsymbol{x}, \boldsymbol{y} \in \tilde{X}$ or $\boldsymbol{x}, \boldsymbol{y} \notin \tilde{X}$.

Suppose $\boldsymbol{x}, \boldsymbol{y} \notin \tilde{X}$. Then $f(\boldsymbol{x}) = f^+(\boldsymbol{x}) = 0 = f^+(\boldsymbol{y}) = f(\boldsymbol{y})$.

Now suppose $\boldsymbol{x}, \boldsymbol{y} \in \tilde{X}$. Take points $\boldsymbol{x}', \boldsymbol{y}' \in X$ such that $\boldsymbol{x}' = \beta_1 \boldsymbol{x}$ and $\boldsymbol{y}' = \beta_2 \boldsymbol{y}$. Then since f^* is admissible inside X, $f^*(\boldsymbol{x}') = f^*(\boldsymbol{y}')$. By the symmetrisation procedure in proposition 5.3.8, this value is the value we choose for both $f(\boldsymbol{x})$ and $f(\boldsymbol{y})$. This shows that the new function f is admissible.

In fact, this new function is precisely given by:

$$f(w; x, y, z) = b((w, x, y, z) - (1, 2, 3, 4)) + b((w, y, x, z) - (1, 2, 3, 4))$$

.

6 Coupled–Cell Networks and Differential Equations

Chapter 5 began the process of linking the two strands of this thesis, by defining our admissible functions. We now complete this process, by defining an admissible vector field, in the obvious way.

The main results of this thesis concern the 'patterns of synchrony and phase shift' of trajectories. Unlike the specific trajectory curves themselves, these are symmetry properties of the trajectories, and therefore applicable to many different trajectories — they may also be subject to examination without obtaining a full 'solution' to the differential equation, in the normal sense. We introduce these 'patterns of synchrony and phase shift' formally in section 6.2, and then define 'rigidity', which is a property of patterns of synchrony designed to disregard 'physically impossible' solutions, where the precise mathematical setup of the equations gives a possible solution which disappears completely under small perturbations. Note that our notation differs slightly from Stewart and Parker (2008), although our definitions are analogous: the notation here is chosen to accord well with the other notation used in this thesis.

In section 6.2.2, we extend the results from section 5.3.1 concerning lifts and quotients of admissible functions to results on equilibria and periodic trajectories of the differential equations defined by these functions. It is section 6.2.3, however, that contains in some sense the most important results of this chapter, showing that there are unique maximal rigid patterns of synchrony and phase for a given trajectory. Without these results, it would be meaningless to try to consider the structure of these patterns.

The chapter concludes with the presentation of four previously conjectured statements which have been outstanding questions in this field: see, for example Stewart and Parker (2008) for a discussion of the importance of these conjectures. We prove these statements, in various special cases, in chapters 7,8,9,10.

6.1 Admissible Differential Equations

Recall that chapter 5 defined an admissible function on a network, and chapter 2 defined a system of differential equations. As is usual in this field, for example, in Golubitsky et al. (2005); Stewart and Parker (December 2007, 2008, Preprint), we now put these two concepts together, linking systems of differential equations and networks in the following (obvious) definition:

Definition 6.1.1. An admissible system of differential equations on a network $\mathcal{N}$ is a differential equation $\dot{x} = F(x)$ where $F : \mathcal{P}(\mathcal{N}) \to \mathcal{P}(\mathcal{N})$ forms an admissible function on $\mathcal{N}$.

6.2 Patterns of Synchrony and Phase Shift

The main results of this thesis — in chapters 7,8,9,10 — concern properties of periodic trajectories which are preserved under small admissible perturbations. We introduce these properties here: we define the 'pattern of synchrony' of an equilibrium or trajectory of a differential equation, and what it means for such a pattern to be 'robust' and 'rigid'.

Let $\mathcal{N}$ be a network, and F an admissible system of differential equations on $\mathcal{N}$.

Definition 6.2.1. Let $x : \mathbb{R} \to \mathcal{P}(\mathcal{N})$: for example, x could be a trajectory of F. Define the *pattern of synchrony* of x to be an equivalence relation $\equiv_x$ on the cells of $\mathcal{N}$ defined by $c \equiv_x d$ if $x_c \equiv x_d$; that is, $x_c(t) = x_d(t)$ for all $t \in \mathbb{R}$. Call any refinement of $\equiv_x$ a *synchrony relation* of x.

Note that x could be an equilibrium trajectory of F: that is, $x(t) = x^* \in \mathcal{P}(\mathcal{N})$ for all $t \in \mathbb{R}$; the pattern of synchrony of x^* would then be defined by $c \equiv_{x^*} d$ if $x_c = x_d$. To simplify the notation in this case, let $\equiv_*$ denote $\equiv_{x^*}$.

Definition 6.2.2. Let x be a periodic function $x : [0, \Theta) \to \mathcal{P}(\mathcal{N})$: for example, x could be a periodic trajectory of F. Let $\theta \in [0, \Theta)$. Define the *pattern of θ–shift* of x to be a forward relation $\Leftarrow_x^\theta$ over $\equiv_x$ on the cells of $\mathcal{N}$, defined by $c \Leftarrow_x^\theta d$ if $x_c(t) = x_d(t + \theta)$ for all t. In this way, $\equiv_x = \Leftarrow_x^0$.

Call any refinement of $\Leftarrow_x^\theta$ a *θ–shift relation* of x: for arbitrary θ, call these relations *phase relations* of x.

Recall that for an equivalence relation $\sim$ on $\mathcal{N}$, we denote n–th input tree equivalence by $\cong_\sim^n$. In the case where $\sim = \equiv_*$, we denote $\cong_\sim^n$ by $\cong_*^n$; where $\sim = \sim_\mathcal{N}$, we denote $\cong_\sim^n$ by $\cong_\mathcal{N}^n$. Similarly, where $\sim = \equiv_x$, denote $\cong_\sim^n$ by $\cong_x^n$. Where $\sim = \Leftarrow_x^\theta$, denote $\cong_\sim^n$ by $\cong_{x,\theta}^n$.

As mentioned before, our notation differs slightly from that of Stewart and Parker (2008): where that paper uses $\approx_{\mathbf{X}}$ to denote the maximal pattern of synchrony of a trajectory x, we use $\equiv_x$. Also, the previous work uses $\dot{\approx}_{\mathbf{X}}$ to denote the general 'phase relation' of cells: this relation is equal to $\bigvee_\theta \Lleftarrow_x^\theta$ in our notation, where $\bigvee$ denotes the join in the lattice of arbitrary relations. We prefer the new notation described here because, as shown in Stewart and Parker (2008, section 4.1), 'the phase shift' between two cells is not necessarily uniquely determined. So-called 'multirhythms', where a given cell has a shorter period than that of the whole system, can give multiple phase shifts between two cells. For this reason, our notation incorporates an explicit phase shift θ, and we consider whether two cells have this specific phase relation.

To illustrate this issue, we now give an example.

Example 6.2.3. Let x be a Θ–periodic trajectory of an admissible system of differential equations F over a network $\mathcal{N}$, such that cell c has period Θ and cell d has period $\Theta/6$. Further, suppose c' is a cell such that $x_{c'}(t) = x_c(t + \Theta/2)$, and d' a cell such that $x_{d'}(t) = x_d(t + \Theta/2)$. We say that this network exhibits 'multirhythms', as in Stewart and Parker (2008), because some cells have a lesser period than that of the whole trajectory. As remarked there, we find it difficult to consider 'the' phase shift between two cells. The phase shift $\theta_{c,c'}$ between c and c' must be equal to $\left(n + \frac{1}{2}\right)\Theta$ for $n \in \mathbb{Z}$: it is reasonable to choose $n = 0$ so that $0 \leqslant \theta_{c,c'} = \frac{1}{2}\Theta < \Theta$. However, the phase shift $\theta_{d,d'}$ between d and d' is $n\frac{\Theta}{6} + \frac{1}{2}\Theta$. A reasonable choice for this might be to make $\theta_{d,d'}$ as small as possible: $\theta_{d,d'} = \frac{1}{3}$. However, this does not signify clearly that $d \Lleftarrow_x^{\Theta/2} d'$, which is important if we are to compare 'the' phase shifts of the pairs (c, c') and (d, d').

These 'multirhythms' will not cause our notation or proofs problems: we shall examine how our method works with an example later.

6.2.1 Rigidity

In this thesis, we aim to show that certain patterns of synchrony are 'generic', whereas some others are 'rare', and therefore do not occur except in systems specially constructed to have them. We now give a precise definition of this 'rarity' by defining what it means for a pattern of synchrony to be 'rigid': the other patterns (which we call 'fragile') are the rarities.

Definition 6.2.4 (Stewart et al. (2003, Definition 6.2))**.** Let $\sim$ be an equivalence relation on $\mathcal{N}$. Then $\sim$ is *robustly polysynchronous* if, for all admissible systems of differential equations F, and all trajectories $x(t)$ of F with initial conditions $x(0)$ with synchrony relation $\sim$, then $\sim$ is a synchrony relation of $x(t)$ for all $t \in \mathbb{R}$.

Definition 6.2.5. Let x^* be an equilibrium of F. Suppose $\sim$ is some synchrony relation of x^*, and there exist $\varepsilon, \delta > 0$ such that for $\hat{F}$ any ε–perturbation of F, and $\hat{x}^*$ an equilibrium of $\hat{F}$ in $\mathbb{B}(x^*, \delta)$, $\sim$ is a synchrony relation of $\hat{x}^*$, then we call $\sim$ a *rigid* synchrony relation for x^*. Otherwise, we call it *fragile*.

Definition 6.2.6. Let S be a metric space; suppose $x : \mathbb{R} \to S$ is a Θ–periodic function and $\hat{x} : \mathbb{R} \to S$ is a $\hat{\Theta}$–periodic function. Then we say that x and $\hat{x}$ are (δ, η)*–close* if x and $\hat{x}$ are δ–close between 0 and the larger of $\Theta, \hat{\Theta}$, with $\left|\hat{\Theta} - \Theta\right| < \eta$.

If, in addition, $\hat{x}|_R = x|_R$ for some set $R \subseteq \mathbb{R}$, then we say that x and $\hat{x}$ are (δ, η)*–close respecting* R.

Definition 6.2.7. Let x be a periodic orbit of F. Suppose $\sim$ is some synchrony relation of x, and there exist $\varepsilon, \delta, \eta > 0$ such that for $\hat{F}$ any ε–perturbation of F, and $\hat{x}$ a periodic orbit of $\hat{F}$ which is (δ, η)–close to x, the same relation $\sim$ is a synchrony relation of $\hat{x}$, then we call $\sim$ a *rigid* synchrony relation for x. Otherwise, we call it *fragile*.

6.2.2 Lifts and Quotients

Recall (section 5.3.1) that given an admissible function F on a network $\mathcal{N}$, and a balanced equivalence relation $\sim$ on $\mathcal{N}$, there is a function $F/\!\!/\!\sim$ on $\mathcal{N}/\!\!/\!\sim$, which is admissible over this network. We now show that equilibria and periodic trajectories of the system $\frac{\mathrm{d}x}{\mathrm{d}t} = F(x)$ are preserved when passing to this $F/\!\!/\!\sim$, subject to certain conditions on their pattern of synchrony.

Lemma 6.2.8 (Equilibrium Quotient). *Let R be a reduced set for an equivalence relation $\sim$ on $\mathcal{N}$. Let F be an admissible system of differential equations on $\mathcal{N}$, and x^* an equilibrium of F with pattern of synchrony $\equiv_* \geqslant \sim$. Then x^*_R is an equilibrium of $F_{\mathcal{R}}$.*

Proof. For any cell $c \in R$, $\mathcal{I}_{\mathcal{R}}(c) = \{\ R(d) \mid d \in \mathcal{I}_{\mathcal{N}}(c)\ \}$. Since $\equiv_* \geqslant \sim$, $R(d) \equiv_* d$, so $x_{\mathcal{I}_{\mathcal{R}}(c)} = x_{R(\mathcal{I}_{\mathcal{N}}(c))} = x_{\mathcal{I}_{\mathcal{N}}(c)}$.

Hence $f_c\left(x_{\mathcal{I}_{\mathcal{R}}(c)}\right) = f_c\left(x_{\mathcal{I}_{\mathcal{N}}(c)}\right) = 0.$ □

Lemma 6.2.9 (Trajectory Quotient). *Let R be a reduced set for an equivalence relation $\sim$ on $\mathcal{N}$. Let F be an admissible system of differential equations on $\mathcal{N}$, and x a trajectory of F with pattern of synchrony $\equiv_x \geqslant \sim$. Then x_R is a trajectory of $F_{\mathcal{R}}$.*

Proof. For any cell $c \in R$, $\mathcal{I}_{\mathcal{R}}(c) = \{\ R(d) \mid d \in \mathcal{I}_{\mathcal{N}}(c)\ \}$. Since $\equiv_x \geqslant \sim$, $R(d) \equiv_x d$, so $x_{\mathcal{I}_{\mathcal{R}}(c)} \equiv x_{R(\mathcal{I}_{\mathcal{N}}(c))} \equiv x_{\mathcal{I}_{\mathcal{N}}(c)}$.

Hence $f_c\left(x_{\mathcal{I}_{\mathcal{R}}(c)}\right) \equiv f_c\left(x_{\mathcal{I}_{\mathcal{N}}(c)}\right) \equiv 0.$ □

Proposition 6.2.10 (Augmented Quotient). *If $\sim$ is a rigid synchrony relation of x on $\mathcal{N}$, and R is a reduced set for $\sim$ with some constraint cell c, then $\sim$ is also a rigid synchrony relation of x^+ on $\mathcal{R}^+ = \mathcal{N}/\!\!/^c_R \sim$.*

Proof. On R_+, the only non-trivial equivalence is $c \sim R(c)$.

Suppose $\hat{F}^+$ is an admissible, smooth ε–perturbation of F^+. Then this gives functions $\hat{f}_r : \mathcal{PCI}(r) \to \mathcal{P}(r)$ for all $r \in R^+$. For $d \in \mathcal{N}$, let

$$\hat{f}_d = \begin{cases} \hat{f}_r & d \sim_{\mathcal{I}} r \in R^+ \\ f_d & \text{otherwise} \end{cases}$$

Then this gives a small perturbation $\hat{F}$ of F such that $\left[\hat{F}\right]^+ = \hat{F}^+$.

Since $\sim$ is a rigid synchrony relation of x on $\mathcal{N}$, then $\sim$ is also a synchrony relation of $\hat{x}$. Therefore $\hat{x}_c \equiv \hat{x}_{R(c)}$.

Now $\hat{x}^+ = \hat{x}_{R^+}$ is an orbit of $\hat{F}^+$. So $\hat{x}_c^+ \equiv \hat{x}_c \equiv \hat{x}_{R(c)} \equiv \hat{x}_{R(c)}^+$, as required. □

Lemma 6.2.11 (Trajectory Lift)**.** *Let R be a reduced set for an equivalence relation $\sim$ on a network $\mathcal{N}$. Let F be an admissible system of differential equations on $\mathcal{N}$, and x_R a trajectory of $F_{\mathcal{R}}$. Then x, defined by $x_c = [x_R]_{R(c)}$ is a trajectory of F with some pattern of synchrony $\equiv_x \geqslant \sim$.*

Naturally, if x_R is an equilibrium of $F_{\mathcal{R}}$, then x is an equilibrium of F.

Proof. Again trivial, and exactly complementary to those of lemma 6.2.8 and lemma 6.2.9.

□

6.2.3 Maximal Patterns of Synchrony and Phase

The results in chapters 7, 9 and 10 go some way towards determining the value of the maximal rigid pattern of synchrony (or phase shift) — in particular, towards proving the conjecture that this pattern is balanced. In order for this line of enquiry to make any sense, we must show that there *is* a (unique) maximal rigid pattern of synchrony (phase shift) for a given trajectory. This section proves this in the three cases — synchrony for equilibria, synchrony for periodic trajectories, and phase shift for periodic trajectories — lending evidence to the important questions which we give in the next section, and then prove, subject to some limitations, in the later chapters.

Again, our notation differs from Stewart and Parker (2008): where that paper uses a superscript of 'rig' to denote the maximal rigid patterns of synchrony and phase, we use a dot over the relation: thus, their $\approx_{\mathbf{X}}^{\mathrm{rig}}$ is our $\dot{\equiv}_x$. Our notation makes sense here, in the following intuitive way. We usually denote perturbations of trajectories by overmarks: for example, $\hat{x}$ or $\tilde{x}$ for perturbations of x. In common with our earlier notational simplification, we might write $\hat{\equiv}_x$ for $\equiv_{\hat{x}}$, and $\tilde{\equiv}_x$ for $\equiv_{\tilde{x}}$. Then the relation $\dot{\equiv}_x$ is inspired by the idea of making 'any' perturbation to x, in effect, putting any overmark in place of the dot, just as the notation $f(t+\cdot)$ uses the dot as a placeholder.

Lemma 6.2.12 (Rigid Pattern of Synchrony of Equilibrium)**.** *There is some maximal rigid synchrony relation of x^*, which we denote by $\dot{\equiv}_*$: that is, $\dot{\equiv}_*$ is a rigid synchrony relation for x^*, such that, if $\sim$ is any rigid synchrony relation of x^*, we have $\sim \leqslant \dot{\equiv}_*$.*

We call this $\dot{\equiv}_*$ the rigid pattern of synchrony *of x^*.*

Proof. For a point $x \in \mathcal{P}(\mathcal{N})$, let $\mathcal{S}_x$ denote the set of synchrony relations for x. Let $\mathcal{R}_*$ denote the set of rigid synchrony relations for x^*.

The set $\mathcal{S}_x$ is a complete lattice, with respect to the refinement order $\leqslant$: take $\approx \subseteq \mathcal{S}_x$. Then $c\left[\bigvee \approx\right] d \implies c = c_0 \sim_1 c_1 \sim_2 c_2 \sim_3 \cdots \sim_n c_n = d$ for some $\{\sim_i\} \subseteq \approx \implies x_c = x_{c_0} = x_{c_1} = \cdots = x_{c_n} = x_d \implies x_c = x_d$, so $\bigvee \approx \in \mathcal{S}_x$. Similarly, take any $\sim \in \approx$, then

$c\,[\bigwedge \approx]\,d \implies c \sim d \implies x_c = x_d$, so $\bigwedge \approx \in \mathcal{S}_x$.

We show that the set $\mathcal{R}_*$ is also a complete lattice. That is, given a set of rigid synchrony relations, $\approx \subseteq \mathcal{R}_*$, its join $\bigvee \approx$ and meet $\bigwedge \approx$ are also rigid. Take any perturbation $\hat{F}$ of F, inducing a perturbed equilibrium $\hat{x}^*$, and let $\hat{\mathcal{S}}_*$ denote the set of synchrony relations of $\hat{x}^*$. Then, by rigidity, $\approx \subseteq \hat{\mathcal{S}}_*$, so $\bigvee \approx,\ \bigwedge \approx \in \hat{\mathcal{S}}_*$, as required.

Thus $\mathcal{R}_*$, as a complete lattice, has a maximal element $\doteqdot_* = \bigvee \mathcal{R}_*$. □

Lemma 6.2.13 (Rigid Pattern of Synchrony of Trajectory)**.** *Let x be a hyperbolic periodic orbit of F. Then there is some maximal rigid synchrony relation of x, which we denote by $\doteqdot_x$: that is, $\doteqdot_x$ is a rigid synchrony relation for x, such that, if $\sim$ is any rigid synchrony relation of x, we have $\sim \leqslant \doteqdot_x$.*

We call this $\doteqdot_x$ the rigid pattern of synchrony *of x.*

Proof. For any periodic orbit x of F, let $\mathcal{S}_x$ denote the set of synchrony relations for x. Let $\mathcal{R}_x$ denote the set of rigid synchrony relations for x.

The set $\mathcal{S}_x$ is a complete lattice, by exactly the same argument as before (with $\equiv$ in place of $=$).

We show that the set $\mathcal{R}_x$ is also a complete lattice. That is, given a set of rigid synchrony relations $\approx \subseteq \mathcal{R}_x$, its join $\bigvee \approx$ and meet $\bigwedge \approx$ are also rigid. Take any perturbation $\hat{F}$ of F, inducing a perturbed trajectory $\hat{x}$, and let $\hat{\mathcal{S}}_x$ denote the set of synchrony relations of $\hat{x}$. Then, by rigidity, $\approx \subseteq \hat{\mathcal{S}}_x$, so $\bigvee \approx,\ \bigwedge \approx \in \hat{\mathcal{S}}_x$, as required.

Thus $\mathcal{R}_x$, as a complete lattice, has a maximal element $\doteqdot_x = \bigvee \mathcal{R}_x$. □

Lemma 6.2.14 (Rigid Pattern of θ–shift)**.** *Let x be a hyperbolic periodic orbit of F, with period Θ. Let $0 \leqslant \theta < \Theta$. Then there is some maximal rigid θ–shift relation of x, which we* *denote by $\dot{\Lleftarrow}^{\theta}_{x}$: that is, $\dot{\Lleftarrow}^{\theta}_{x}$ is a rigid θ–shift relation for x such that, if $\prec$ is any rigid θ–shift relation of x, we have $\prec \leqslant \dot{\Lleftarrow}^{\theta}_{x}$.*

In fact, the set of rigid θ–shift relations for x is a complete lattice.

Proof. The proof of this lemma is again similar to those above: the main difference lies in the use of the forward relation meet operator. Note that in contrast to the lattice of equivalence relations on a set, the set of all forward relations on a set is not a lattice. However, this will not harm our proof, as we now demonstrate.

For any periodic orbit x of F and phase shift $\theta \in \mathbb{R}$, let $\mathcal{S}_{x,\theta}$ denote the set of θ–shift relations for x. Let $\mathcal{R}_{x,\theta}$ denote the set of rigid θ–shift relations for x.

The set $\mathcal{S}_{x,\theta}$ is a complete lattice, with respect to the refinement order $\leqslant$: take $\lll \subseteq \mathcal{S}_{x,\theta}$. Then the meet $\bigwedge\lll$ is defined as in lemma 4.3.4 by $c \left[\bigwedge\lll\right] d$ if and only if $c \prec d$ for all $\prec \in \lll$. We show that this is also a θ–shift relation for x. Take $\prec \in \lll$: if $c \left[\bigwedge\lll\right] d$ then $c \prec d$, so $x_c(t) = x_d(t+\theta)$ for all t. Thus $\bigwedge\lll$ is a θ–shift relation. As remarked earlier, arbitrary sets of forward relations need not have well-defined joins. However, in the order of this set $\mathcal{S}_{x,\theta}$, we set the relation $\bigwedge\varnothing$ to be defined by $c \left[\bigwedge\varnothing\right] d$ if $x_c(t) = x_d(t+\theta)$ for all t. This is clearly a θ–shift relation for x; by definition, any θ–shift relation for x must be a refinement of this relation. Then by theorem 2.2.1, $\mathcal{S}_{x,\theta}$ is a complete lattice under the refinement order.

We also show that the set $\mathcal{R}_{x,\theta}$ is also a complete lattice. That is, given a set of rigid θ–shift relations $\lll$, its join $\bigwedge\lll$ and meet $\bigvee\lll$ are also rigid. Take any perturbation $\hat{F}$ of F, inducing a perturbed trajectory $\hat{x}$, then by rigidity $\lll \subseteq \mathcal{S}_{\hat{x},\theta}$. So $\bigvee\lll$, $\bigwedge\lll \in \hat{\mathcal{S}}_{x}$, as required.

Thus $\mathcal{R}_{x,\theta}$ is a complete lattice; in particular, it has a maximal element $\dot{\Lleftarrow}^{\theta}_{x} = \bigvee \mathcal{R}_{x,\theta}$. □

6.3 Important Questions

We now come to the discussion of the questions in the field which are particularly interesting: how does balance interact with existence of equilibria or trajectories? Recall from chapter 1 that any balanced equivalence relation is robustly polysynchronous; further, that any robustly polysynchronous equivalence relation is balanced. Since robust polysynchrony of $\sim$ implies rigidity of $\sim$ as a pattern of synchrony for any trajectory (or equilibrium) with synchrony relation $\sim$, all balanced equivalence relations make rigid patterns of synchrony. In the following chapters, we attempt to prove the opposite implication: given a trajectory with a rigid pattern of synchrony $\sim$, the relation $\sim$ is balanced.

Recall the Rigid Synchrony Conjecture of Stewart and Parker (December 2007), as stated in chapter 1:

Rigid Synchrony Conjecture (Stewart and Parker (December 2007) Conjecture 6.1)**.** *Let G be any coupled cell network, and suppose that $\mathbf{X}$ is a periodic orbit of some G-admissible vector field f. Assume that $\mathbf{X}$ is rigid. Then its pattern of synchrony $\approx_{\mathbf{X}}$ is balanced.*

As mentioned in chapter 1, we prove some specific cases (and a generalisation) of this conjecture in chapters 7,9,10. We give these cases below.

Rigid Equilibrium Theorem (Theorem 7.2.3)**.** *Let F be an admissible system of differential equations over a network $\mathcal{N}$, with some transverse equilbrium x^*. The rigid pattern of synchrony of x^* is the maximal balanced equivalence relation refining $\equiv_*$.*

The following two theorems rely on technical conjectures 9.3.1 and 10.2.3. Justification for these conjectures appears with their statements, in chapters 9 and 10.

Limited Rigid Synchrony Theorem (Theorem 9.3.3)**.** *Let F be an admissible system over* *a network $\mathcal{N}$, with some hyperbolic periodic orbit x. Assume conjecture 9.3.1. Then the rigid pattern of synchrony of x is the maximal balanced equivalence relation refining $\equiv_x$.*

Limited Rigid Phase Theorem (Theorem 10.3.3)**.** *Assume conjecture 10.2.3 holds.*

Let $\mathcal{N}$ be a network and F an admissible system of differential equations on $\mathcal{N}$ with a hyperbolic trajectory x of period Θ. Let $\theta \in [0, \Theta)$, then the rigid θ–shift relation $\dot{\Leftarrow}_x^\theta$ of x must be the maximal balanced refinement of $\Leftarrow_x^\theta$ (as a forward relation over $\equiv_x$ on $\mathcal{N}$).

In addition, the very recent work of Golubitsky et al. (Preprint) has proved the Rigid Synchrony Theorem, in the following form:

Rigid Synchrony Theorem (Golubitsky et al. (Preprint) Theorem 6.1)**.** *Suppose $Z_0(t)$ is a hyperbolic periodic solution of $\dot{Z} = F(Z)$. Then the coloring associated to $\Delta(Z_0)$ is rigid if and only if it is balanced.*

We discuss their method briefly in chapter 9.

As has been suggested by Stewart and Parker (2008), the results in periodic systems are easiest to prove if a property called 'full oscillation' holds: a trajectory has this property if every cell is oscillating. It has been conjectured that this property holds for almost all admissible systems on all networks — in fact, the conjecture states that the property is generic: given any such system, it may be perturbed (admissibly) by an arbitrarily small amount to get a system where the property holds.

Stewart–Parker Conjecture 3.3 (Full Oscillation Conjecture)**.** *In a path-connected network, whenever x is a hyperbolic periodic orbit which is not an equilibrium, there exists a small admissible perturbation of the vector field having a perturbed periodic orbit $\hat{x}(t)$ in which no cell is in equilibrium.*

Golubitsky et al. (Preprint) have proved this result recently; we also prove this result in chapter 8: in fact, we prove a strictly stronger result. We call our result 'strong oscillation'; it is stated here for motivation.

Strong Oscillation Property (Property 8.2.1)**.** *Let F be an admissible system of differential equations over a network $\mathcal{N}$, and x a hyperbolic periodic trajectory of F. Then x has the* strong oscillation property *if given a cell c, the set of times $\left\{ t \in \mathbb{R} \mid \frac{\mathrm{d}x_c}{\mathrm{d}t}(t) = 0 \right\}$ has zero measure.*

Strong Oscillation Theorem (Theorem 8.2.2)**.** *For every hyperbolic periodic trajectory x of an admissible system of differential equations F over a path-connected network $\mathcal{N}$, and every $\varepsilon > 0$, there is some ε–perturbation $\hat{F}$ of F such that the trajectory $\hat{x}$ near x satisfies the Strong Oscillation Property.*

7 The Rigid Equilibrium Theorem

The first of the three 'rigid-balance' theorems we prove in this thesis is the rigid equilibrium theorem; in contrast to the limited proofs of the other versions of the rigid synchrony and phase conjectures, given in the following chapters, the proof of the rigid equilibrium conjecture given in this chapter is complete, and does not rely on an additional 'niceness' property of the trajectory (we will see these properties, called 'tameness' and 'semi-tameness', in chapter 9). This is because in an equilibrium, every 'trajectory' of a cell is a single point in its phase space. Any two of these points are either identical or disjoint, which makes it considerably easier to prove the rigid equilibrium theorem than to answer the other questions of section 6.3.

The rigid equilibrium theorem has already been proved in Golubitsky et al. (2005), however the proof given here is new and very different to the previous proof. It is usual in perturbation problems to make use of a transversality argument, but this method is not used in the previous proof. We show in this chapter that it is possible, and indeed natural, to use a transversality argument. This argument extends, with some modifications, to the results in the following chapters.

The first section of this chapter deals with basic results on transversality. Most can be found in Hirsch and Smale (1974); the only one of these that is at all original shows that transverse equilibria remain transverse after taking balanced quotients. However, the transversality theorem of section 7.1.1 is both novel and important: it shows that an admissible perturbation may be made to an admissible system of differential equations to turn any equilibrium into a transverse equilibrium. This is very important for the proof of the rigid equilibrium theorem.

The proof of the rigid equilibrium theorem follows in section 7.2, and uses multiple perturbations: we first perturb the given vector field to give it certain desirable properties, and then we use these properties to show the existence of a perturbation which is synchrony-breaking, as required. A similar method of proof will also be used to prove results relating to the rigid synchrony and rigid phase conjectures in the following chapters.

7.1 Isolation and Transversality

This section starts with some definitions of well-known properties and simple results about equilibria of systems of differential equations. We then show that one of these properties, transversality, may be attained by making a small admissible perturbation to any admissible system of differential equations on a network. This property will form the basis of our proof of the rigid equilibrium theorem in the next section.

In this section, let F be any system of differential equations.

Definition 7.1.1. An equilibrium x^* of F is *isolated* if there is some $\delta > 0$ such that $\{\ x \in \mathbb{B}(x^*, \delta) \mid F(x) = 0\ \} = \{x^*\}$. We may specify the value of δ by using the term δ–isolated.

Definition 7.1.2 (Transversality)**.** Let F be a system of differential equations with an equilibrium x^*. Then x^* is *transverse* in F if $\det DF(x^*) \neq 0$.

The following property is well-understood in the context of the lemma following it, although this definition is original to this thesis.

Definition 7.1.3 (Transversality Property)**.** Suppose x^* is an equilibrium of F, and there are some $\varepsilon, \delta > 0$ such that for all ε–perturbations $\hat{F}$ of F, $\hat{F}$ has a unique equilibrium in $\mathbb{B}(x^*, \delta)$. Then we say that x^* satisfies the *transverse property*. We call this ε the *transversality value* of F.

Lemma 7.1.4. *If x^* is a transverse equilibrium of F, then x^* satisfies the transversality property.*

Proof. Let x^* be a transverse equilibrium of a system F of differential equations which is admissible over some network $\mathcal{N}$.

We state Theorem 1 of Chapter 16 of Hirsch and Smale (1974) using the notation from that source. In this notation, E is a vector space, and the space of linear maps from E to itself is denoted $L(E)$. Further, W is an open set in E, and $\mathfrak{U}(W)$ denotes the set of all C^1 vector fields on W; that is, C^1 functions $f : W \to E$. This set is equipped with the C^1 norm and becomes a normed vector space. A 'neighbourhood' of a point v in a vector space V (for example, $x \in E$ or $f \in \mathfrak{U}(W)$) is any subset $N \subset V$ that contains an open ball around v (under the vector space norm). A final difference is that Hirsch and Smale (1974) uses x' where Anosov and Arnold (1988) uses $\dot{x}$, and we prefer $\frac{\mathrm{d}}{\mathrm{d}t}x$.

Hirsch–Smale Theorem 16.1. *Let $f : W \to E$ be a C^1 vector field and $\bar{x} \in W$ an equilibrium of $x' = f(x)$ such that $Df(\bar{x}) \in L(E)$ is invertible. Then there exists a neighbourhood $U \subset W$ of $\bar{x}$ and a neighbourhood $\mathfrak{N} \subset \mathfrak{U}(W)$ of f such that for any $g \in \mathfrak{N}$ there is a unique equilibrium $\bar{y} \in U$ of $y' = g(y)$. Moreover, if E is normed, for any $\varepsilon > 0$ we can choose $\mathfrak{N}$ so that $|\bar{y} - \bar{x}| < \varepsilon$.*

In this proof, we shall refer to this theorem simply as 'the Theorem'.

Let $W = \mathbb{B}(x^*, \delta) \subseteq \mathcal{PC}(\mathcal{N})$ and $E = \mathcal{PC}(\mathcal{N})$. Then let f denote the restriction of F to W. Thus $f : W \to E$ is a C^1 vector field. Let $\bar{x} = x^*$.

Note that transversality is a local property, so since x^* is a transverse equilibrium of F it is likewise transverse in f — that is, $D(\bar{x})$ is invertible. So by the main part of the Theorem, we can find U and $\mathfrak{N}$ which satisfy the given properties.

Since $\mathfrak{N}$ is a neighbourhood of f, it contains some open ball around f. Let ε be the radius of this ball. Thus if $\hat{F}$ is any (admissible) ε–perturbation of F, its restriction $g = \hat{F}\big|_W \in \mathfrak{N}$. Thus by the Theorem, g has a unique equilibrium $\bar{y} \in U$. To make the notation consistent with the rest of this thesis, define $\hat{x}^* = \bar{y}$.

We let δ be the radius of a ball around $x^* = \bar{x}$ which is contained in U. Then using the second part of the Theorem ('Moreover...'), we can ensure that $\hat{x}^*$ is inside $\mathbb{B}(x^*, \delta)$. Since this ball is entirely contained in U, and the equilibrium $\hat{x}^*$ is unique in U, we have that $\hat{x}^*$ is the unique equilibrium of g (and therefore $\hat{F}$) in $\mathbb{B}(x^*, \delta)$.

It remains to ensure that $\det D\hat{F}(\hat{x}^*) \neq 0$. Notice that although we have chosen a particular ε and δ above, the result will also hold with any smaller values, as long as their calculation is only based upon f and x^*. We therefore revisit our choice of ε and δ to ensure that the ball $\mathbb{B}(x^*, \delta)$ has no points x where $|\det Df(x)| < \varepsilon$ (this is possible since Df is continuous and $\det Df(x^*) \neq 0$). Thus when $\hat{F}$ and therefore g are chosen to be ε–close to F and f, we have that $\det Dg(x) \neq 0$ for all $x \in \mathbb{B}(x^*, \delta)$ — in particular, $\det D\hat{F}(\hat{x}^*) \neq 0$. □

Corollary 7.1.5. *Let F be a system of differential equations with a transverse equilibrium x^*. Then x^* is an isolated equilibrium of F.*

Proof. By lemma 7.1.4, there are some $\delta, \varepsilon > 0$ such that for all ε–perturbations $\hat{F}$ of F, $\hat{F}$ has a unique equilibrium in $\mathbb{B}(x^*, \delta)$. $\hat{F} = F$ is such an ε–perturbation, so it has a unique equilibrium in $\mathbb{B}(x^*, \delta)$, which must be x^*. Hence x^* is a δ–isolated equilibrium of F. $\square$

In general, lifts and quotients with respect to some transversal of an equivalence relation do not preserve isolation or transversality of equilibria. However, where the relation is balanced, we may take the balanced quotient without disrupting these properties.

Lemma 7.1.6. *Let $\sim$ be a balanced equivalence relation on a network $\mathcal{N}$; let F be an admissible system of differential equations on $\mathcal{N}$ and x^* an ε–transverse equilibrium of F, with $\sim$ a synchrony relation of x^*.*

Then $x^* /\!\!/ \sim$ is an ε–transverse equilibrium of $F /\!\!/ \sim$ on $\mathcal{N} /\!\!/ \sim$.

Proof. Let $v \in \mathcal{P}(\mathcal{C}(\mathcal{N} /\!\!/ \sim))$ be an eigenvector of $D[F /\!\!/ \sim](x^* /\!\!/ \sim)$ with eigenvalue 0. Define $v^\sim$ by $[v^\sim]_c = v_c$. Then $v^\sim$ is an eigenvector of $DF(x^*)$, with eigenvalue 0. So if $F /\!\!/ \sim$ is not transverse, neither is F. $\square$

7.1.1 The Transversality Theorem

The theorem in this section shows that any admissible system of differential equations may be admissibly perturbed to make a given equilibrium into a transverse equilibrium. Importantly, it does this without moving the equilibrium itself. This level of control will become important in the following proof of the rigid equilibrium theorem.

Theorem 7.1.7 (Transversality Theorem)**.** *Let $\mathcal{N}$ be a network and F a system of differential equations which is admissible over $\mathcal{N}$. Let $x^* \in \mathcal{P}(\mathcal{N})$ be an equilibrium of F. Suppose $\varepsilon > 0$ is given. Then there is an admissible ε–perturbation $\tilde{F}$ of F such that x^* is a transverse equilibrium of $\tilde{F}$.*

Proof. Without the restriction to admissible perturbations, this result is a well-known and direct consequence of Sard's Theorem. However, we restrict here to admissible perturbations, which a priori could cause problems.

For each $c \in \mathcal{C}(\mathcal{N})$, let R_c be the radius of disjoint symmetric neighbourhoods of x^* at c. Let $r = \min(\{ R_c \mid c \in \mathcal{C}(\mathcal{N}) \} \cup \{1\}) > 0$.

Now use theorem 2.6.4, with $X, Y, x^*, R, r, \varepsilon$ in that theorem respectively equal to $\mathcal{P}(\mathcal{N}), \mathcal{P}(\mathcal{N}), x^*, r, r/2, \varepsilon$ here, to obtain a function bound $\Delta > 0$.

Let $J = DF(x^*)$. Assume that $\det J = 0$ (otherwise F itself is a sufficient example of $\tilde{F}$). Now J has a number of eigenvalues λ_i, some of which are equal to 0. Take $\delta = \min(\{\Delta\} \cup \{ |\lambda_i| \neq 0 \})/2 > 0$. Then the matrix $\tilde{J} = J + \delta I$ has no zero eigenvalues.

For each $c \in \mathcal{C}(\mathcal{N})$, define $\tilde{f}_c^*(x_c; x_1, x_2, \ldots) = f_c(x_c; x_1, x_2, \ldots) + \delta(x_c - x_c^*)$ in the region $\mathbb{B}\left(x^*_{\mathcal{CI}(c)}, r/2\right)$. Note that $\tilde{f}_c^*\left(x^*_{\mathcal{CI}(c)}\right) = f_c\left(x^*_{\mathcal{CI}(c)}\right) = 0$, and $\left|\tilde{f}_c^*(x) - f_c(x)\right| = \delta|(x_c - x_c^*)| < \delta$ for all $x \in \mathbb{B}\left(x^*_{\mathcal{CI}(c)}, r/2\right)$. Also, $\frac{\partial \tilde{f}_c^*}{\partial x_d}(x^*) = \frac{\partial f_c^*}{\partial x_d}(x^*)$ for all $d \neq c$, but $\frac{\partial \tilde{f}_c^*}{\partial x_c}(x^*) = \frac{\partial f_c^*}{\partial x_c}(x^*) + \delta$. This being the case, $\left\|\tilde{f}_c^*(x) - f_c(x)\right\| \leqslant \delta < \Delta$ for all $x \in \mathbb{B}\left(x^*_{\mathcal{CI}(c)}, r/2\right)$, and we may use theorem 2.6.4 to define $\tilde{f}_c^*$ in $\mathbb{B}\left(x^*_{\mathcal{CI}(c)}, r\right) \backslash \mathbb{B}\left(x^*_{\mathcal{CI}(c)}, r/2\right)$ which is smooth and equal to f_c at the outer boundary of this region such that $\left\|\tilde{f}_c^*(x) - f_c(x)\right\| < \varepsilon$ for all $x \in \mathbb{B}\left(x^*_{\mathcal{CI}(c)}, r\right)$.

We show that $\tilde{f}_c^*(x) = \tilde{f}_c^*(\beta x)$ for all $\beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)$ such that $x^* = \beta x^*$ and $x \in \mathbb{B}\left(x^*_{\mathcal{CI}(c)}, r\right)$. If $\beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)$ with $x^* = \beta x^*$, then $\beta : \mathcal{I}(c) \to \mathcal{I}(c)$ must preserve the root of $\mathcal{I}(c)$. Hence $\beta(c) = c$. Therefore, for $x \in \mathcal{PCI}(c)$, $(\beta x)_c = x_c$. So $\tilde{f}_c^*(\beta x) = f_c(\beta x) + \delta((\beta x)_c - x_c^*) = f_c(\beta x) + \delta(x_c - x_c^*)$. Since f_c is admissible at c, $f_c(\beta x) = f_c(x)$, so $\hat{f}_c^*(\beta x) = f_c(x) + \delta(x_c - x_c^*) = \tilde{f}_c^*(x)$, as required.

Hence, by proposition 5.3.8, we can define $\tilde{F}$ to be an admissible function on $\mathcal{N}$, with components $\tilde{f}_c^*$ for each $c \in \mathcal{C}(\mathcal{N})$ inside the regions $\mathbb{B}\left(x^*_{\mathcal{C}\beta\mathcal{I}(c)}, r\right)$, and equal to f_c for points outside the union of these regions.

We immediately see that $\left\|\tilde{f}_c^*(x) - f_c(x)\right\| < \varepsilon$ for all $x \in \mathcal{PCI}(c)$, thus $\tilde{F}$ is an ε–perturbation of F. The Jacobian of $\tilde{F}$ at x^* is $D\tilde{F}(x^*) = D[F + \delta(\mathrm{id} - x^*)](x^*) = DF(x^*) + \delta D[\mathrm{id} - x^*](x^*) = J + \delta I = \tilde{J}$, which has no zero eigenvalues by construction. So x^* is a transverse equilibrium of $\tilde{F}$. $\square$

7.2 The Rigid Equilibrium Theorem

This section contains the main result of this chapter, in two halves: first, we prove that balanced synchrony relations of transverse equilibria are rigid. Following that, we come to the more technically involved direction: an unbalanced pattern of synchrony of an (arbitrary, not necessarily transverse) equilibrium is fragile.

Our line of attack is as follows. Firstly, perturb to ensure that the given equilibrium x^* is transverse. Then take some reduced set R for the pattern of synchrony $\equiv_*$ of x^*, and consider the quotient system $F /\!\!/_R \equiv_*$: perturb this to ensure that the quotient of x^* is also transverse in the quotient system. Further perturb this perturbed $F /\!\!/_R \equiv_*$ in the component for a cell where $\equiv_*$ is not balanced, and reassemble a perturbed F using the results of chapter 5; the equilibrium near x^* in this perturbed system has broken symmetry, by the method of construction.

Proposition 7.2.1. *Suppose $\sim$ is balanced on $\mathcal{N}$. Let F be an admissible system of differential equations on $\mathcal{N}$ and x^* a transverse equilibrium of F with $\sim$ a synchrony relation of x^*. Then $\sim$ is a rigid synchrony relation for x^*.*

Proof. Let R be a reduced set for $\sim$, and ε the transversality value of x^*. Then, since $\sim$ is balanced, x^*_R is an ε–transverse equilibrium of F_R.

Suppose $\hat{F}$ is an ε–perturbation of F. Since F is transverse, $\hat{F}$ has a unique equilibrium $\hat{x}^*$ near x^*. Also, $\hat{F}_R$ is an ε–perturbation of F_R, and has some unique equilibrium $\hat{x}^*_R$ near x^*_R. This induces an equilibrium $[\hat{x}^*_R]^\sim$ of $\hat{F}$; $\sim$ is a synchrony relation of $[\hat{x}^*_R]^\sim$. By uniqueness, $\hat{x}^* = [\hat{x}^*_R]^\sim$, and so $\sim$ is a synchrony relation of $\hat{x}^*$, as required. □

Theorem 7.2.2. *Let F be an admissible system over a network $\mathcal{N}$, with some equilibrium x^*. Let $\equiv_*$ denote the pattern of synchrony of x^*. Suppose $\equiv_*$ is unbalanced.*

Let $\varepsilon > 0$ be given. Then there is some ε–perturbation $\hat{F}$ of F such that, if $\hat{x}^*$ denotes the equilibrium of $\hat{F}$ near x^*, the pattern of synchrony of $\hat{x}^*$ is a strict refinement of $\equiv_*$, or the two patterns are incomparable. That is, the synchrony of x^* is not rigid.

Hence, for equilibria, unbalanced patterns of synchrony are not rigid, or, in the contrapositive, rigid patterns of synchrony are balanced.

Proof. By theorem 7.1.7, we may take an $\varepsilon/2$ perturbation F' of F such that x^* is a transverse equilibrium of F'. Let ε' be the smaller of $\varepsilon/2$ and the transversality value of x^* in F'. Then all ε'–perturbations of F' are ε–perturbations of F, and have a unique, transverse, equilibrium δ'–near x^* for some δ', by lemma 7.1.4. $(*)$

Now let R, D, C be some reduced, duplicate and constraint sets for $\equiv_*$. Let $\mathcal{R} = \mathcal{N} /\!\!/_R \equiv_*$. Since this symmetry is not balanced, $C \neq \varnothing$. Take $c \in C$, so $\equiv_*$ is not balanced at $(c, R(c))$.

Consider $F'_R = F' /\!\!/_R \equiv_*$ as a system of differential equations over $\mathcal{R}$. Since $F'_R(x) = F'_R(x^*)$, the point $x^*_R = \pi_R x^* \in \mathcal{P}(\mathcal{R})$ is an equilibrium of F'_R. Using theorem 7.1.7 again, we may make an admissible $\varepsilon'/2$–perturbation taking this F'_R to $\tilde{F}_R$ such that x^*_R is a transverse equilibrium of $\tilde{F}_R$.

Let $\tilde{F}$ denote the lift $\tilde{F}_R^{\equiv *}$; this is an $\varepsilon'/2$–perturbation of F'. Hence x^* is transverse in the system $\tilde{F}$, by $(*)$ above. Let $\tilde{\varepsilon}$ be the smaller of the transversality value of x^* in $\tilde{F}$ and $\varepsilon'/2$. Then all $\tilde{\varepsilon}$–perturbations of $\tilde{F}$ are ε–perturbations of F which have a unique, transverse, equilibrium δ'–near x^* and a unique equilibrium $\tilde{\delta}$–near x^* for some $\tilde{\delta}$; these must be the same equilibrium.

Let r_1 be the radius of disjoint symmetric neighbourhoods of x^* at c.

Then $r_1 < \frac{1}{2}\,\mathrm{d}\left(x_{\mathcal{C}\mathcal{I}(c)}, x_{\mathcal{C}\beta\mathcal{I}(R(c))}\right)\ \forall \beta \in \mathrm{Aut}^1_{\mathcal{N}}(c)$.

Also, let $r_2 = \frac{1}{2}\min\left\{\ \mathrm{d}\left(x_{\mathcal{C}\mathcal{I}(c)}, x_{\mathcal{C}\mathcal{I}(d)}\right)\ \middle|\ c \sim_{\mathcal{I}} d \wedge c \not\equiv_* d\ \right\}$.

Now let $r = \min\{r_1, r_2\}$. Note that this means that
$\mathbb{B}\left(x^*_{\mathcal{C}\beta\mathcal{I}(c)}, r\right) \cap \mathbb{B}\left(x^*_{\mathcal{C}\mathcal{I}(d)}, r\right) = \varnothing$ for all $d \in R$ and $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(c,d)$. $(+)$

Now, as in remark 2.6.6, we form $\hat{f}_c$, an admissible $\tilde{\varepsilon}$–perturbation of $\tilde{f}_c$ in $\mathbb{B}\left(x^*_{\mathcal{C}\mathcal{I}(c)}, r\right)$, in which the equilibrium local to x^*_c is some $\hat{x}^*_c \neq x^*_c$. We know this equilibrium $\hat{x}^*_c \in \mathbb{B}\left(x^*_{\mathcal{C}\mathcal{I}(c)}, r\right)$.

Then we may symmetrise $\hat{f}_c$ into the system F, by proposition 5.3.8: let $\hat{F}$ denote this symmetrisation. This $\hat{F}$ is equal to F outside the regions $\mathbb{B}\left(x^*_{\mathcal{C}\beta\mathcal{I}(c)}, r\right)$, and so F and $\hat{F}$ are equal inside the regions $\mathbb{B}\left(x^*_{\mathcal{C}\mathcal{I}(d)}, r\right)$ for $d \in R$, by $(+)$. Hence $\hat{F}_{\mathcal{R}}\Big|_{\mathbb{B}\left(x^*_R, r\right)} = F_{\mathcal{R}}\Big|_{\mathbb{B}\left(x^*_R, r\right)}$: in particular, x^*_R is an equilibrium of $\hat{F}_{\mathcal{R}}$.

Since this $\hat{F}$ is a $\tilde{\varepsilon}$–perturbation of $\tilde{F}$, it has a unique equilibrium δ–near the equilibrium x^* of F'. Let $\hat{x}^*$ denote this equilibrium, and let $\hat{\equiv}_*$ denote its pattern of synchrony. The component of $\hat{x}^*$ associated to the cell c must be the $\hat{x}^*_c$ defined above, which is not equal to x^*_c.

Suppose the synchrony is unbroken by this perturbation: $\hat{\equiv}_* \geqslant \equiv_*$. Then, by lemma 6.2.8, $\hat{x}^*_R$ is an equilibrium of $F_{\mathcal{R}}$; also, $\mathrm{d}\left(\hat{x}^*_R, x^*_R\right) \leqslant \mathrm{d}\left(\hat{x}, x\right) < \delta$. But $\hat{x}^*_R$ is a δ–isolated equilibrium of $\hat{F}_{\mathcal{R}}$, so $\hat{x}^*_R = x^*_R$. This creates a contradiction: $x^*_R = \hat{x}^*_R$, but $[x^*_R]_{R(c)} = x^*_{R(c)} = x^*_c \neq \hat{x}^*_c = \hat{x}^*_{R(c)} = [\hat{x}^*_R]_{R(c)}$.

Therefore, the pattern of synchrony $\equiv_*$ must be broken by this perturbation. □

Theorem 7.2.3 (Rigid Equilibrium Theorem)**.** *Let F be an admissible system over a network $\mathcal{N}$, with some transverse equilibrium x^*. The rigid pattern of synchrony of x^* is the maximal balanced equivalence relation refining $\equiv_*$.*

Proof. Let $\bowtie_*$ denote the maximal balanced refinement of $\equiv_*$. This relation $\bowtie_*$ is a pattern of synchrony for x^*, and is balanced. Hence $\bowtie_*$ is a rigid synchrony relation for x^*, by proposition 7.2.1, so $\bowtie_* \leqslant \dot{\equiv}_*$.

All rigid patterns of synchrony for x^* are balanced, by theorem 7.2.2, and must refine $\equiv_*$. Thus $\bowtie_* \geqslant \sim$ for all rigid patterns of synchrony, by maximality. In particular, $\bowtie_* \geqslant \dot{\equiv}_*$.

This shows that $\bowtie_* = \dot{\equiv}_*$ — that is, the maximal balanced equivalence relation refining $\equiv_*$ is the rigid pattern of synchrony of x^* — and completes the proof of the Rigid Equilibrium Theorem. □

8 The Strong Oscillation Theorem

Stewart and Parker (2008) consider the property of 'full oscillation'. A (periodic) trajectory of a system of differential equations over a network $\mathcal{N}$ is said to be 'fully oscillatory' when no cell is in equilibrium. The results given by Stewart and Parker (2008) assume full oscillation as a hypothesis.

The results in this chapter prove a stronger version of full oscillation than that conjectured in this previous work: we call this property 'strong oscillation'.

The full oscillation theorem has recently been proved in Golubitsky et al. (Preprint) by a rather technical method, but the proof we give here, in addition to proving a slightly stronger property, is a more straightforward geometrical argument.

We start with some results about hyperbolicity, and then proceed to strong oscillation.

8.1 Hyperbolicity

Recall that a trajectory x of an admissible system F of differential equations over a network $\mathcal{N}$ is hyperbolic if $[DF](x)$ has no purely imaginary eigenvalues. We now state a property which is a well-known consequence of hyperbolicity; we will in fact use this property, and not hyperbolicity itself, for all our proofs. This is interesting, since a slightly modified version of this property also holds for some non-hyperbolic trajectories, as we shall see in chapter 10.

Property 8.1.1 (Hyperbolic Property)**.** Let F be a system of differential equations, and x be a Θ–periodic trajectory of F. Then x has the *hyperbolic property* if there is some $\varepsilon_0 > 0$ (which we call a *hyperbolic bound*) such that for all $\varepsilon \in (0, \varepsilon_0)$, there exist $\delta, \eta > 0$ such that for all ε–perturbations $\hat{F}$ of F, there exists a unique trajectory $\hat{x}$ of $\hat{F}$ which is δ–close to x, and this $\hat{x}$ has period $\hat{\Theta} \in (\Theta - \eta, \Theta + \eta)$.

Definition 8.1.2. A trajectory which satisfies the hyperbolic property is called a *pseudo-hyperbolic trajectory*.

Lemma 8.1.3. *Every hyperbolic trajectory of a system of differential equations satisfies the hyperbolic property.*

Proof. Standard result, given for example in Anosov and Arnold (1988). □

Note that when we talk of 'unique' trajectories here, we mean unique up to phase shift. That is, if we say that x is the unique trajectory that satisfies some property, we mean that if y is a trajectory which satisfies the same property, then there is some $\theta \in \mathbb{R}$ such that $y(t) = x(\theta + t)$ for all $t \in \mathbb{R}$. Since the differential equation $\dot{x} = F(x)$ is autonomous, this is as unique as it can be: given a trajectory x of a system of differential equations F, all trajectories $y = x(\theta + \cdot)$ are also trajectories of F. We will revisit this idea in chapter 10.

Using this lemma, we may restate the property of rigidity.

Definition 8.1.4. Let x be a hyperbolic periodic orbit of F with hyperbolic bound ε_0. Suppose $\sim$ is some synchrony relation of x such that for all ε_0–perturbations $\hat{F}$ of F, the same relation $\sim$ is a synchrony relation of the nearby trajectory $\hat{x}$ of $\hat{F}$.

Then we call $\sim$ a *rigid synchrony relation* of x.

This definition easily extends to pseudo-hyperbolic trajectories.

Definition 8.1.5. Let x be a pseudo-hyperbolic periodic orbit of F with hyperbolic bound ε_0. Suppose $\sim$ is some synchrony relation of x such that for all ε_0–perturbations $\hat{F}$ of F, the same relation $\sim$ is a synchrony relation of the nearby trajectory $\hat{x}$ of $\hat{F}$.

Then we call $\sim$ a *rigid synchrony relation* of x.

When we consider hyperbolic trajectories in this chapter from now on, we in fact only use the hyperbolic property. While this is immaterial to the current proofs, this will become critical in the next chapter. Additionally, many of the results in this chapter will not depend upon hyperbolicity at all, and although it seems integral to the statement of the rigid synchrony theorem, in fact we only use the hyperbolic property at the final stage of the proof. This fact will be very useful in the next chapter.

8.2 The Strong Oscillation Theorem

It is usual in this area to consider only those trajectories which are 'fully oscillating': that is, given a cell c, the function $x_c : \mathbb{R} \to \mathcal{P}(c)$ is not constant. This assumption is spelled out, for example, in Stewart and Parker (2008, conjecture 3.3):

Stewart–Parker Conjecture 3.3 (Full Oscillation Conjecture)**.** *In a path-connected network, whenever x is a hyperbolic periodic orbit which is not an equilibrium, there exists a small admissible perturbation of the vector field having a perturbed periodic orbit $\hat{x}(t)$ in which no cell is in equilibrium.*

Stewart and Parker (2008, example 3.6) shows that the full oscillation conjecture, or a similar statement, is necessary to derive the rigid synchrony theorem for general networks and trajectories. Here, we shall prove a stronger statement than the given full oscillation conjecture: we call it the strong oscillation theorem.

Property 8.2.1 (Strong Oscillation Property)**.** Let F be an admissible system of differential equations over a network $\mathcal{N}$, and x a periodic trajectory of F. Then x has the *strong oscillation property* if given a cell c, the set of times $\left\{ t \in \mathbb{R} \;\middle|\; \frac{\mathrm{d}x_c}{\mathrm{d}t}(t) = 0 \right\}$ has zero measure.

The strong oscillation property effectively says that a trajectory is almost always fully oscillating.

Recall that a network $\mathcal{N}$ is path-connected if for any pair of cells c, d in $\mathcal{N}$ there is a path $c = c_0 \xrightarrow{e_1} c_1 \xrightarrow{e_2} \cdots \xrightarrow{e_{n-1}} c_{n-1} \xrightarrow{e_n} c_n = d$ of edges $e_i \in \mathcal{E}(\mathcal{N})$.

Theorem 8.2.2 (Strong Oscillation Theorem)**.** *For every pseudo-hyperbolic periodic trajectory x of an admissible system of differential equations F over a path-connected network $\mathcal{N}$, and every $\varepsilon > 0$, there is some ε–perturbation $\hat{F}$ of F such that the trajectory $\hat{x}$ near x satisfies the Strong Oscillation Property. In other words, the Strong Oscillation Property is generic.*

Proof. In this proof, we shall use the term 'interval' to mean an interval of non-zero Lebesgue measure in $\mathbb{R}$.

Let x be a pseudo-hyperbolic periodic trajectory of an admissible system of differential equations F over a path-connected network $\mathcal{N}$, and $\varepsilon > 0$ be given. Suppose (for contradiction) that x does not satisfy the Strong Oscillation Property. Then there is some non-empty set of cells $C \subset \mathcal{C}(\mathcal{N})$ such that for all $c \in C$, the set of times $\left\{ t \in \mathbb{R} \;\middle|\; \frac{\mathrm{d}x_c}{\mathrm{d}t}(t) = 0 \right\}$ has non-zero measure.

We now show that we have some cell c and interval $T \subsetneq \mathbb{R}$ where c has constant value but never constant inputs. It is clear that there cannot be an interval T such that x is constant in T (that is, x_c is constant in T for all c), as the behaviour of the trajectory is entirely determined by its current state, so if it is in equilibrium over an interval, then it is always in equilibrium. Thus, given any interval T, there is some cell which does not have constant value on T. Let $d \in C$, and T be some interval over which x_d is constant. Then there is some cell d' such that $x_{d'}$ is not constant on T. Since $\mathcal{N}$ is path-connected, there is some path $d' = d_0 \to d_1 \to \cdots \to d_n = d$. Let i be the smallest integer such that d_i has constant value on T. Then let $c = d_i$.

Now c is a cell with non-constant inputs such that $\frac{\mathrm{d}x_c}{\mathrm{d}t} = 0$ on some interval T. By definition, $f_c\left(x_{\mathcal{CI}(c)}(t)\right)$ is identically 0 for $t \in T$. This defines a curve in $\mathcal{PCI}(c)$ where f_c is identically zero. Let ε_1 be a hyperbolic bound for F, and $\varepsilon_2 = \varepsilon_1/2$. Then there some $\delta > 0$ such that for all ε_2–perturbations $\hat{F}$ of F, there exists a unique trajectory $\hat{x}$ of $\hat{F}$ which is δ–close to x.

Take $\varepsilon_3 = \min\left\{ \left|\frac{\mathrm{d}x_d}{\mathrm{d}t}\right| \;\middle|\; d \to c \right\}$. Then if $\hat{F}$ is an ε_3–perturbation of F, $\hat{x}_d$ is not constant for all cells $d \to c$. Let $\varepsilon = \min\{\varepsilon_2, \varepsilon_3\}$. Then all ε–perturbations of F are ε_1–perturbations, ε_2–perturbations and ε_3–perturbations of F.

Because all functions under consideration must be smooth, we know that for any $\hat{F}$, there must be some curve of points $y : T \to \mathcal{PCI}(c)$ near $x_{\mathcal{CI}(c)}(T)$ on which f_c is identically zero. However, we can use an appropriate bump function to ensure that $\left\{ t \in T \;\middle|\; \frac{\mathrm{d}y_c}{\mathrm{d}t}(t) = 0 \right\}$ has measure zero. Now, let $\hat{x}$ be the nearby trajectory of $\hat{F}$. If $\hat{x}_{\mathcal{CI}(c)}(t) = y(t')$, then consider some interval τ around t. We show that $\hat{x}_c$ cannot be constant on τ. If $x(\tau) = \{x^*\}$, then the function $f_c\left(x_{\mathcal{CI}(c)}(\cdot)\right)$ is identically zero on τ. But the places where this is the case locally form the curve y. Hence $x_{\mathcal{CI}(c)}(\tau) \subseteq y(T)$. Since the inputs of c are non-constant, this means $x_{\mathcal{CI}(c)}(\cdot)$ follows the curve y, and also cannot be constant. This contradiction completes the proof. □

Corollary 8.2.3. *Let x be a periodic trajectory of an admissible system of differential equations F over a network $\mathcal{N}$, such that for every path-connected component $\mathcal{M}$ of $\mathcal{N}$, the restricted trajectory $x|_{\mathcal{M}}$ is oscillating and hyperbolic over $\mathcal{M}$, then for every $\varepsilon > 0$, there is some ε–perturbation $\hat{F}$ of F such that the trajectory $\hat{x}$ near x satisfies the Strong Oscillation Property.*

Proof. Given such a trajectory and $\varepsilon > 0$, we may apply the Strong Oscillation Theorem to each path-connected component in sequence, each time obtaining an ε/N perturbation, where N is the number of path-connected components in $\mathcal{N}$. We consider the path-connected components in some order such that if C is upstream of D, the component C is considered before D. The resulting perturbation satisfies the Strong Oscillation Property. □

Corollary 8.2.4. *If we can show that the rigid synchrony theorem holds for all trajectories of systems of ODEs on networks which satisfy the strong oscillation property, then the rigid synchrony theorem holds for all trajectories of systems of ODEs on path-connected networks.*

The strong oscillation theorem only holds for networks and trajectories where each upstream component has at least one oscillating cell. This condition is required, as we now see.

Lemma 8.2.5. *Let F be an admissible system of differential equations over a network $\mathcal{N}$, and x a hyperbolic trajectory of F. Suppose C is an upstream component of $\mathcal{N}$ such that x_c is stationary for all $c \in C$. Then there is some ε such that, for all ε–perturbations $\hat{F}$ of F, the trajectory $\hat{x}$ of $\hat{F}$ near x is stationary on all $c \in C$.*

Proof. Since C is an upstream component, it contains all cells in the infinite input trees of its own cells: thus it is dynamically self-contained. So we may consider the network $\mathcal{C}$ with cells

C and all edges from $\mathcal{N}$ whose ends are in C: F induces an admissible system of differential equations F_C on this network; restricting x to the cells of C gives an equilibrium x_C of this system. This equilibrium is hyperbolic, and therefore transverse, so there is some ε such that for all ε–perturbations $\hat{F}_C$ of F_C, the trajectory $\hat{x}_C$ of $\hat{F}_C$ near x_C is an equilibrium. □

9 The Rigid Synchrony Theorem

Recall from chapter 1 the statement of the Rigid Synchrony Conjecture originally proposed in Stewart and Parker (December 2007).

Rigid Synchrony Conjecture (Stewart and Parker (December 2007) Conjecture 6.1). *Let G be any coupled cell network, and suppose that $\mathbf{X}$ is a periodic orbit of some G-admissible vector field f. Assume that $\mathbf{X}$ is rigid. Then its pattern of synchrony $\approx_{\mathbf{X}}$ is balanced.*

As described in the introduction, the notation in Stewart and Parker (December 2007) differs from the notation in this thesis. Consider the following statement, in our notation:

Theorem 9.0.6 (Rigid Synchrony Theorem). *Let F be an admissible system over a network $\mathcal{N}$, with some pseudo-hyperbolic periodic orbit x. The rigid pattern of synchrony of x is the maximal balanced equivalence relation refining $\equiv_x$.*

This thesis examines the rigid synchrony theorem, and proves it from a conjectured result, but does not contain a direct proof of the theorem. Rather, during its preparation, a proof of an equivalent theorem was given by Golubitsky et al. (Preprint) — the chapter finishes with a brief examination of this proof. In its original notation, this theorem is as follows.

Rigid Synchrony Theorem (Golubitsky et al. (Preprint) Theorem 6.1). *Suppose $Z_0(t)$ is* *a hyperbolic periodic solution of* $\dot{Z} = F(Z)$*. Then the coloring associated to* $\Delta(Z_0)$ *is rigid if and only if it is balanced.*

It is worth considering the apparent difference between these three statements. The conjecture given by Stewart and Parker and theorem of Golubitsky et al. merely state that all rigid patterns of synchrony are balanced, whereas theorem 9.0.6 precisely categorises the maximal rigid pattern of synchrony. These two conclusions are equivalent, by applying exactly the argument of theorem 7.2.3 to patterns of synchrony of periodic trajectories (orbits) rather than those of equilibria.

In this chapter, we examine the Rigid Synchrony Theorem, theorem 9.0.6 (RST). Firstly, we consider the Tame Synchrony Theorem of Stewart and Parker (December 2007), which proves the RST for a certain class of trajectories, called 'tame' trajectories in that paper. We then expand the Tame Synchrony Theorem to a more general result with a more specific conclusion; we call our new result the Semi-Tame Synchrony Theorem. It is noteworthy that our result covers all cases where the Tame Synchrony Theorem holds, in addition to some other cases. Our result hinges on the idea of assuming rigidity of a non-balanced synchrony, and then using this rigidity to ensure that certain operations on the network do not break the synchrony of the trajectory. Finally, we arrive at a situation where we can perturb to break the synchrony, giving a contradiction. This same overall method will be useful in chapter 10.

Notably, the proof of the RST given by Golubitsky et al. (Preprint) does not immediately guarantee the existence of a contradictory perturbation of the kind discussed above, a necessary condition for the arguments of the next chapter. It is possible that the method of Golubitsky et al. can be adapted to give a contradictory synchrony-breaking perturbation for an assumed rigid synchrony; however, it is not immediately apparent how — we leave this direction to further study. Such a contradictory perturbation would be useful in proving the

Rigid Phase Conjecture, as we will see in chapter 10.

9.1 The Tame Synchrony Theorem

Stewart and Parker (December 2007) define a property of 'tameness' as follows:

Definition 9.1.1 (Stewart and Parker (December 2007) Definition 9.6)**.** The periodic orbit $\mathbf{X}$ is *tame* if it contains a point x^0 that has trivial isotropy and is not an orbital overlap point. We call x^0 a *tame point*.

Here, the property that x^0 "has trivial isotropy and is not an orbital overlap point" means precisely that $x^0 = x(t^0)$ such that there are no cells $c, d \in \mathcal{C}(\mathcal{N})$ time $u \in [0, \Theta)$ and network isomorphism $\beta \in \mathrm{Iso}^1(c, d)$ such that:

$$x_{\mathcal{C}\mathcal{I}(c)}\left(t^0\right) = x_{\mathcal{C}\beta\mathcal{I}(c)}(u)$$

Using this property, Stewart and Parker (December 2007) go on to prove the following theorem.

Theorem 9.1.2 (Tame Synchrony Theorem, Stewart and Parker (December 2007) Theorem 10.2)**.** *Let G be an all-to-all coupled network, and let f be a G–admissible vector field. Let $\mathbf{X}$ be a periodic orbit of f. If $\mathbf{X}$ is tame and rigid, then the pattern of synchrony $\approx_{\mathbf{X}}$ is balanced.*

As noted before, this theorem uses slightly different notation to this thesis. Here a G–admissible vector field means a vector field admissible over the network G.

In overview, their method is as follows. The reader is directed to the original paper for the details.

Pass to an appropriate quotient network such that the (quotient) trajectory does not intersect any proper balanced synchrony subspace of the network phase space (here a 'synchrony subspace' is a subspace of the phase space constructed by considering those points where some synchrony relation applies). The (quotient) point x^0 remains tame in this quotient network.

Now there is a diffeomorphism of the phase space which moves x^0 out of the synchrony subspace associated with the pattern of synchrony of x — that is, which breaks the synchrony of x at x^0. Then construct a vector field $\hat{f}$ by conjugating the existing vector field f by this diffeomorphism. This $\hat{f}$ will have a trajectory with appropriately broken synchrony.

The following idea of 'semi-tameness' is closely related to that of tameness — the results presented here should be understood as direct generalisations of the Tame Synchrony Theorem.

9.2 The Semi–Tame Synchrony Theorem

The Tame Synchrony Theorem proves the Rigid Synchrony Theorem in the case of a tame trajectory; as discussed previously, this is a trajectory with 'tame point' x^0, which is effectively determined by a time t^0, in that $x^0 = x\left(t^0\right)$.

We define a looser property than tameness: we call it 'semi-tameness'.

Definition 9.2.1. Let x be a Θ–periodic trajectory of a system F of differential equations over a network $\mathcal{N}$. Let R, D, C be some reduced, duplicate and constraint sets of $\equiv_x$.

Then a *semi-tame point* of x, R, D, C is a pair $(c,t) \in C \times [0,\Theta)$ such that there are no $r \in R,\ u \in [0,\Theta)$ and $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(c,r)$ such that $x_{\mathcal{C}\mathcal{I}(c)}(t) = x_{\mathcal{C}\beta\mathcal{I}(c)}(u)$.

If x has a semi-tame point, then call x a *semi-tame trajectory*. (If C is empty, allow x to be semi-tame by convention.)

Remark 9.2.2. Semi-tameness is a looser property than tameness, at least intuitively. This is because, in order to find a tame point, one must find a time t^0 such that $x(t^0)$ has no orbital overlaps or (non-trivial) isotropy. We discussed the technical implications of these requirements in the previous section. Assuming such a point is found, taking any reduced, duplicate and constraint sets R, D, C and any constraint $c \in C$ gives a semi-tame point (c, t^0). (And if C is empty, x is semi-tame by convention.)

For, if there are r, β, u as in section 9.2.1, then $r \in R \subseteq \mathcal{C}(\mathcal{N})$, so $x(t^0)$ has an orbital overlap given by $x_{\mathcal{C}\mathcal{I}(c)}(t^0) = x_{\mathcal{C}\beta\mathcal{I}(c)}(u)$. (This is a smilar observation to the remark in Stewart and Parker (December 2007, proof of Theorem 10.2) that the only tameness condition which is actually required is tameness in an appropriate quotient network.)

Theorem 9.2.3 (Semi–Tame Synchrony Theorem)**.** *If x is a semi-tame Θ–periodic pseudo-hyperbolic trajectory of a system F of differential equations over a network $\mathcal{N}$, and $\equiv_x$ is unbalanced, then the synchrony of x is not rigid.*

Notice that in addition to requiring the looser semi-tameness in place of tameness, the semi-tame synchrony theorem does not require the network to be all-to-all coupled. Thus it is a more general theorem than the tame synchrony theorem.

Proof. Assume the synchrony of x is rigid; we argue for contradiction.

Consider the system of differential equations F^c_R on $\mathcal{R}^+ = \mathcal{N} /\!\!/^c_R \equiv_x$. Then by proposition 6.2.10, x^+ is a trajectory of this system.

We may make a perturbation $\hat{x}^+$ of x^+ such that $\hat{x}^+_c \neq x^+_c$ using a perturbation $\hat{F}^c_R$ of F^c_R supported on a small neighbourhood of $x^+_{\mathcal{CI}(c)}(t)$ and its symmetric images. This neighbourhood can be chosen small enough not to intersect $x_{\mathcal{CI}(r)}$ for all $r \in R$, since (c, t) is a semi-tame point. Thus $\hat{x}^+_r \equiv x^+_r$ for all $r \in R$, since R is self-contained and no cell in R has its trajectory perturbed. By rigidity, this gives a lifted trajectory $\hat{x}$ on $\mathcal{N}$; this lift would give $\hat{x}_c = \hat{x}^+_{R(c)} = x^+_{R(c)} = x_{R(c)}$. However, c would then have the same inputs in both $\mathcal{R}^+$ and $\mathcal{N}$, giving an alternative trajectory $\hat{x}_c = \hat{x}^+_c$ lifted directly from the cell c in $\mathcal{R}^+$. These two trajectories are not equal by construction; if both trajectories were possible, this would contradict pseudo-hyperbolicity. Reaching this contradiction completes the proof. □

9.3 A Programme for Proof

We suggest that a similar method to the Tame and Semi–Tame Synchrony Theorems (assuming rigidity, perturbing a constraint cell's trajectory, and deriving a contradiction) applies for *any* constraint cell $c \in C$ and time t such that $x_{\mathcal{CI}(c)}(t) \neq x_{\mathcal{C}\beta\mathcal{I}(c)}(t)$ for all $\beta \in \mathrm{Iso}^1_{\mathcal{N}}(c, R(c))$. (Notice that such a time t must exist for any given $c \in C$, or $c \equiv_x R(c)$, and c would not be a constraint cell.) This method would prove the Rigid Synchrony Theorem.

To prove the RST in this way, we suggest the following programme.

Follow the same method as above, except that for any $r \in R$ such that $x_{\mathcal{CI}(c)}(t) = x_{\mathcal{C}\beta\mathcal{I}(c)}(u)$ we accept that $\hat{x}^+_r$ near u will equal $\hat{x}^+_c$ near t. Notice that by choice of t, we can be sure even here that $(r, u) \neq (R(c), t)$.

However, technical difficulties may arise when $r \to r_1$, $R(c) \to r_2$, and $x_{\mathcal{CI}(r_1)}(u) = x_{\mathcal{C}\beta\mathcal{I}(r_2)}(t)$. Then if x_r is perturbed by a small amount and $R(c)$ is not, we have suggested input trajectories for r_1 and r_2 which are close together (since $\hat{F}$ must be a small perturbation of F, and so $\hat{x}$ must be close to x) but have the same gradient $\frac{\mathrm{d}x_{r_1}}{\mathrm{d}t}(u) = \frac{\mathrm{d}x_{r_2}}{\mathrm{d}t}(t)$ — thus $f_{r_1} = f_{r_2}$ must have the same value at those points. By the mean value theorem, this requires any curve γ between these points to go through a point where f_{r_1} has zero gradient along γ, which places an extra constraint on (the gradient of) $\hat{F}$, which may cause $\hat{F} - F$ to have 'large' C^1 norm.

A possible solution to this problem might be found by perturbing $R(c)$ in the opposite direction to r. The required points of zero gradient in $\hat{F}$ would then be found on the diagonal where $r = R(c)$; any admissible function (for example, F) must have zero gradient across this diagonal. This might allow an appropriate admissible perturbation to be constructed with small C^1 norm.

It is hoped that further research should be able to flesh out this argument and thus prove the RST in the way suggested by the following conjecture. Notice that hyperbolicity is not assumed in the first paragraph.

Conjecture 9.3.1. *Given a periodic trajectory x of a system of differential equations F which is admissible over some network $\mathcal{N}$, suppose the pattern of synchrony $\equiv_x$ is unbalanced but rigid. Then we can find an admissible perturbation $\hat{F}$ of F with a trajectory $\hat{x}$ near x such that $\equiv_{\hat{x}} \not\geq \equiv_x$.*

Therefore, assuming x is pseudo-hyperbolic so that this $\hat{x}$ is the unique trajectory of F near x gives a contradiction.

Given previous results such as the Tame and Semi-Tame Synchrony Theorems, it seems reasonable to expect that a proof of this conjecture might run along the following lines. Firstly, take some balanced quotient network to acquire trivial balanced synchrony — thus any remaining synchrony is unbalanced, and should be fragile. Choose a point where some non-balanced synchrony is exemplified, and perturb F in a small region around this point. By an appropriate choice of perturbation, this can be made to break the synchrony of the trajectory in a neighbourhood of the chosen point (while leaving the remainder of the trajectory essentially unchanged, at least in terms of synchrony). As we shall see in the next chapter, a proof of the RST that follows these lines would go a long way towards also proving the Rigid Phase Conjecture.

Proposition 9.3.2. *Suppose $\sim$ is balanced on $\mathcal{N}$. Let F be an admissible system of differential equations on $\mathcal{N}$ and x a hyperbolic periodic orbit of F with $\sim$ a synchrony relation of x. Then $\sim$ is a rigid synchrony relation for x.*

Proof. Let R be a reduced set for $\sim$, and ε the hyperbolicity value of x. Then, since $\sim$ is balanced, x_R is an ε–hyperbolic trajectory of F_R.

Suppose $\hat{F}$ is an ε–perturbation of F. Since F is hyperbolic, $\hat{F}$ has a unique periodic trajectory $\hat{x}$ near x. Also, $\hat{F}_R$ is an ε–perturbation of F_R, and has some unique periodic trajectory $\hat{x}_R$ near x_R. This induces a periodic trajectory $[\hat{x}_R]^\sim$ of $\hat{F}$; $\sim$ is a synchrony relation of $[\hat{x}_R]^\sim$. By uniqueness, $\hat{x} = [\hat{x}_R]^\sim$, and so $\sim$ is a synchrony relation of $\hat{x}$, as required. □

Theorem 9.3.3 (Limited Rigid Synchrony Theorem)**.** *Let F be an admissible system over a network $\mathcal{N}$, with some hyperbolic periodic orbit x. Assume conjecture 9.3.1. Then the rigid pattern of synchrony of x is the maximal balanced equivalence relation refining $\equiv_x$.*

Proof. We follow the same line of argument as theorem 7.2.3. As before, $\dot{\equiv}_x$ denotes the maximal rigid balanced pattern of synchrony.

Let $\bowtie_x$ denote the maximal balanced refinement of $\equiv_x$. This relation $\bowtie_x$ is a pattern of synchrony for x, and is balanced. Hence $\bowtie_x$ is a rigid synchrony relation for x, by proposition 9.3.2, so $\bowtie_x \leqslant \doteqdot_x$, by maximality of $\equiv_x$ as a rigid synchrony relation of x. Notice that this direction of implication does not use the conjecture.

All rigid patterns of synchrony for x are balanced, by conjecture 9.3.1, and must refine $\equiv_x$. Thus $\bowtie_x \geqslant \sim$ for all rigid patterns of synchrony $\sim$, by maximality of $\bowtie_x$ as a balanced equivalence relation refining $\equiv_x$. In particular, $\bowtie_x \geqslant \doteqdot_x$.

This shows that $\bowtie_x = \doteqdot_x$ — that is, the maximal balanced equivalence relation refining $\equiv_x$ is the rigid pattern of synchrony of x — and completes the proof of the Limited Rigid Synchrony Theorem. □

9.4 The Rigid Synchrony Theorem: a proof from Golubitsky et al.

During the preparation of this thesis, Golubitsky et al. (Preprint) proved the Rigid Synchrony Theorem by a completely different method.

For completeness, we restate their result here in their notation, which is straightforward, although different from that used elsewhere in this thesis.

Theorem 9.4.1 (Rigid Synchrony Theorem)**.** *Suppose $Z_0(t)$ is a hyperbolic periodic solution of $\dot{Z} = F(Z)$. Then the coloring associated to $\Delta(Z_0)$ is rigid if and only if it is balanced.*

In brief, their method of proof involved proving that the space of admissible perturbations of a system of periodic hyperbolic ODEs on a network was infinite-dimensional, whereas the synchrony-preserving perturbations of non-balanced patterns of synchrony were finite-dimensional (they lie within a particular finite-dimensional subspace) — showing that there must be some infinite-dimensional 'remaining' space of available perturbations which are synchrony-breaking.

It would be interesting (and useful!) to see if this method could be adapted to prove a conjecture such as conjecture 9.3.1, since this is the approach to the (unproved) Rigid Phase Conjecture which we examine in the next chapter.

10 The Rigid Phase Conjecture

The previous chapter dealt with the rigid synchrony theorem, proving that it followed from what appears to be a reasonable conjecture in a fundamentally different way to the existing proof by Golubitsky et al. (Preprint). We now proceed to apply the same approach to the rigid phase conjecture: in fact, the conjecture we use in this chapter is very close to the conjecture in the previous chapter. Our proof proceeds in exactly the same way as the one given in that chapter. This scheme has two benefits. Firstly, it seems likely that a proof of conjecture 9.3.1 from the previous chapter — for example, using some of the techniques of Golubitsky et al. (Preprint) — would generalise to a proof of the conjecture given in this chapter, which would mean that the results presented in this chapter prove the Rigid Phase Conjecture with no further work. Secondly, even if the method of proof of the Rigid Synchrony Theorem does not generalise, the rigid phase theorem will still hold, by the same proof given here, provided another proof of the conjecture in this chapter can be found.

Our technique, in overview, is to take a trajectory x of a system of differential equations F on a network $\mathcal{N}$ with an unbalanced pattern of phase, and duplicate that network, putting a trajectory on the duplicate network equal to the original trajectory, phase-shifted so that cells in the duplicate network are synchronous with cells in the original. We would like to then use the rigid synchrony theorem discussed in the previous chapter to complete the proof. The complication arises as the trajectory on the union of the network and its duplicate will not be hyperbolic. However, it will still satisfy a modified version of the hyperbolic property, which we examine in detail here.

Recall the statement of the rigid phase conjecture.

Conjecture 10.0.2 (Rigid Phase Conjecture)**.** *Given a network $\mathcal{N}$ and an admissible system of differential equations F on $\mathcal{N}$ with a hyperbolic trajectory x with period Θ, and $\theta \in [0, \Theta)$, then the rigid θ–shift relation $\Leftarrow\!\!\dot{=}^{\theta}_{x}$ of x must be a balanced forward relation over $\equiv_x$ on $\mathcal{N}$.*

Section 4 of Golubitsky et al. (2005) gives an example of 'multirhythms' (A Three-Cell Ring Coupled to a Two-Cell Ring), which seem at first to cause problems with the rigid phase conjecture. We shall allay these fears by restating the example here, and using it as an illustration of our method. See figure 10.1.

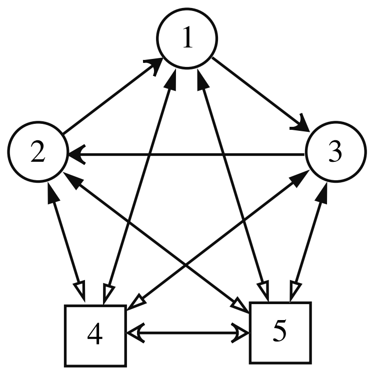

Figure 10.1: An example network, with 5 cells, as given in section 3 of Golubitsky et al. (2005). As before, edge types are shown by arrowhead styles.

The cells in the three-cell ring $(1, 2, 3)$ are assumed to be a different (cell) type from those in the two-cell ring $(4, 5)$. The general system of differential equations which is admissible over this network is:

$$\dot{x}_1 = f\left(x_1; x_2, \overline{y_4, y_5}\right)$$

$$\dot{x}_2 = f\left(x_2; x_3, \overline{y_4, y_5}\right)$$

$$\dot{x}_3 = f\left(x_3; x_1, \overline{y_4, y_5}\right)$$

$$\dot{y}_4 = g\left(y_4; y_5, \overline{x_1, x_2, x_3}\right)$$

$$\dot{y}_5 = g\left(y_5; y_4, \overline{x_1, x_2, x_3}\right)$$

The cell trajectories are labelled x and y to emphasise the two rings that make up this network.

Golubitsky et al. (2005) goes on to show that this network can exhibit a multirhythm: $x(t) = \left(x_1(t), x_1\left(t + \frac{1}{3}\Theta\right), x_1\left(t + \frac{2}{3}\Theta\right), y_4(t), y_4\left(t + \frac{1}{2}\Theta\right)\right)$, where $x_1(t) \equiv x_1\left(t + \frac{1}{2}\Theta\right)$ and $y_4(t) \equiv y_4\left(t + \frac{1}{3}\Theta\right)$. We shall continue to work with this solution throughout this chapter.

10.1 Network Unions

We introduce a simple, but powerful, operation on networks, which we will use in our results on the rigid phase conjecture.

Definition 10.1.1. Suppose $\mathcal{N}$ is any network, and a is any symbol — for example, $a \in \mathbb{N}$. Then let $\mathcal{N}_a$ be a network with cell set $\mathcal{C}(\mathcal{N}) \times \{a\}$, and edge set $\left[\!\left[(c, a) \underset{t}{\rightarrow} (d, a) \;\middle|\; \left(c \underset{t}{\rightarrow} d\right) \in \mathcal{E}(\mathcal{N}) \right]\!\right]$; let $(c, a) \sim_{\mathcal{N}_a} (d, a)$ precisely when $c \sim_{\mathcal{N}} d$. In this way, $\mathcal{N}_a$ is an isomorphic, but disjoint, copy of $\mathcal{N}$.

Definition 10.1.2. Suppose $\mathcal{M}, \mathcal{N}$ are disjoint networks. Then their union $\mathcal{M} \cup \mathcal{N}$ is defined in the obvious way to be a network with cell and edge sets formed from the union of the cell and edge sets of $\mathcal{M}$ and $\mathcal{N}$. In this general case, the types of edges may be marked, and the cell equivalence relation defined, so that no cell or edge of $\mathcal{M}$ is equivalent to one of $\mathcal{N}$.

To extend upon this definition, suppose $\mathcal{N}$ is a network. For $a \neq b$, we let $\mathcal{N}_{a\cup b}$ be the network $\mathcal{N}_a \cup \mathcal{N}_b$, with the additional cell equivalences given by $(c,a) \sim_C (c,b)$ and transitivity, and edge types left unmarked so that $\left((c,a) \underset{t}{\rightarrow} (d,a)\right) \sim_E \left((c,b) \underset{t}{\rightarrow} (d,b)\right)$.

In the case of our example network, figure 10.1, $\mathcal{N}_{1\cup 2}$ is the network pictured in figure 10.2.

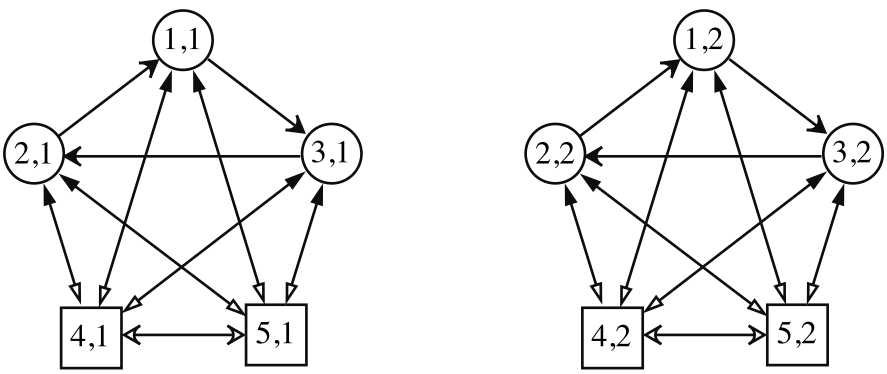

Figure 10.2: The network $\mathcal{N}_{1\cup 2}$, for $\mathcal{N}$ as shown in figure 10.1. $\mathcal{N}_1$ is shown on the left, $\mathcal{N}_2$ on the right.

Given a choice of phase spaces, an admissible system of differential equations, or a trajectory, on $\mathcal{N}$, we may immediately form a choice of phase spaces, admissible system of differential equations or trajectory on $\mathcal{N}_a$ by the trivial isomorphism. We can also form a choice of phase spaces, admissible system of differential equations or trajectory on $\mathcal{N}_{a\cup b}$ in the same way, such that the restrictions of these constructions to $\mathcal{N}_a$ or $\mathcal{N}_b$ give the same objects as those constructed by the isomorphism.

Remark 10.1.3. Let F be an admissible system of differential equations over a (connected) network $\mathcal{N}$, and x a trajectory of F. Let $F_{1\cup 2}$ be the admissible system of differential equations on $\mathcal{N}_{1\cup 2}$ as defined above. Define $x^{\theta}(t) = x(t+\theta)$ for all $t \in \mathbb{R}$.

Then x^{θ} is also a trajectory of F on $\mathcal{N}$. Since $\mathcal{N}_{1\cup 2}$ has two connected components, each isomorphic to $\mathcal{N}$, then the function $x^{0,\theta} : \mathbb{R} \to \mathcal{P}\left(\mathcal{N}_{1\cup 2}\right)$ given by $x^{0,\theta}_{(c,1)} = x_c$, $x^{0,\theta}_{(c,2)} = x^{\theta}_c$ is a trajectory of $F_{1\cup 2}$. For example, let $\theta = \Theta/2$ in our standard example of figure 10.1. Then $x^{0,\theta}$ is a trajectory on the network $\mathcal{N}_{1\cup 2}$ in figure 10.2, defined as follows: $x_{\cdot,1} = \left(x_1\left(t\right), x_1\left(t+\frac{1}{3}\Theta\right), x_1\left(t+\frac{2}{3}\Theta\right), y_4\left(t\right), y_4\left(t+\frac{1}{2}\Theta\right)\right)$, as above; $x_{\cdot,2}\left(t\right) = \left(x_1\left(t+\frac{1}{2}\Theta\right), x_1\left(t+\frac{5}{6}\Theta\right), x_1\left(t+\frac{7}{6}\Theta\right), y_4\left(t+\frac{1}{2}\Theta\right), y_4\left(t+\Theta\right)\right)$. Recall that $x_1\left(t\right) = x_1\left(t+\frac{1}{2}\Theta\right)$ and $y_4\left(t\right) = y_4\left(t+\frac{1}{3}\Theta\right)$: this means that the input sets of cells 1 to 3 are phase shifted by $\frac{1}{3}\Theta$ and $\frac{2}{3}\Theta$, and similarly the input sets of cells 4 and 5 are shifted by $\frac{1}{2}\Theta$, even though the phase relations written here do not suggest that.

Remark 10.1.4. As remarked above, given a trajectory x of an admissible system of differential equations F on $\mathcal{N}$, we may form a trajectory $x^{0,\theta}$ of $F_{1\cup 2}$. However, even if x is a hyperbolic trajectory, $x^{0,\theta}$ is certainly not hyperbolic: duplicating the network and trajectory in this way duplicates the eigenvalues of the Poincaré map of the trajectory, and so the single eigenvalue on the unit circle becomes two eigenvalues on the unit circle. In terms of the hyperbolic property, notice that even the unperturbed $F_{1\cup 2}$ has other trajectories $\hat{x}^{0,\theta}$ near $x^{0,\theta}$: let $\hat{x}^{0,\theta} = x^{0,\theta+\eta}$ for sufficiently small η. In the next section, we will outline another approach to hyperbolicity and rigidity that allows the use of an extended version of the rigid synchrony theorem — really, an extended version of conjecture 9.3.1.

10.2 Hyperbolicity in Disjoint Networks

Recall that if a given trajectory of a system of ODEs on a network satisfies the hyperbolic property, property 8.1.1, then a perturbation of the system perturbs the trajectory 'uniquely', at least when considering trajectories of the new system near the original trajectory. Note that in the absence of specified initial conditions, *any* solution may be phase-shifted by *any* amount to create another solution. Assuming the solution x to be uniformly continuous, for every $\varepsilon > 0$ there is some $T > 0$ such that for $\tau < T$ we have $\mathrm{d}\left(x\left(\tau+t\right), x\left(t\right)\right) < \varepsilon$ for all t. Hence we may always find 'extra' solutions in any neighbourhood of a given solution x.

In order to say that a solution in 'unique' in this case, we form an equivalence relation between phase-shifted solutions: $x \sim y$ if there is some $\theta \in \mathbb{R}$ such that $x\left(t\right) = y\left(t+\theta\right)$; we understand that the solution of the perturbed equation is unique (in some neighbourhood of the original solution) *up to equivalence* under $\sim$.

In the case of a network with (at least) two connected components $\mathcal{N}_1$ and $\mathcal{N}_2$, this equivalence is not sufficient, since the solution may be phase-shifted independently on the pieces $\mathcal{N}_1$ and $\mathcal{N}_2$: that is, for a solution $x = (x_1; x_2)$ to an admissible system of ODEs, ε, T as before, we have that for all $\tau_1, \tau_2 < T$, $x_\tau = \left(x_1\left(\tau_1 + \cdot\right); x_2\left(\tau_2 + \cdot\right)\right)$ must be a solution of the same system of ODEs. In this case, the equivalence relation must be defined by $x \approx y$ if $\exists \theta_1, \theta_2$ such that $\left(x_1\left(t\right); x_2\left(t\right)\right) \equiv \left(y_1\left(t+\theta_1\right); y_2\left(t+\theta_2\right)\right)$. Then we can consider uniqueness of solutions up to equivalence by $\approx$.

This 'disjoint phase-shift identification' process extends easily to networks with any number of disjoint pieces. With this relation in hand, we can now state a version of the hyperbolic property for networks with disjoint pieces.

Property 10.2.1 (Disjoint Hyperbolic Property)**.** Let F_1, F_2 be two systems of differential equations, and $F = (F_1; F_2)$ the combined system. Let $x = (x_1; x_2)$ be a Θ–periodic trajectory of F. Then x has the *disjoint hyperbolic property* if there is some $\varepsilon_0 > 0$ such that for all $\varepsilon \in (0, \varepsilon_0)$, there exist $\delta, \eta > 0$ such that for all ε–perturbations $\hat{F}$ of F, there exists a trajectory $\hat{x} = (\hat{x}_1; \hat{x}_2)$ of $\hat{F}$ which is δ–close to x, and this trajectory is unique up to disjoint phase-shift identification $\approx$ as above; further, this $\hat{x}$ has period $\hat{\Theta} \in (\Theta - \eta, \Theta + \eta)$.

This property allows us to use an extended version of the first paragraph of conjecture 9.3.1 on networks with more than one connected component.

Definition 10.2.2. For a network with more than one connected component with an admissible system of ODEs F which has a trajectory x with the disjoint hyperbolic property, call a pattern of synchrony $\equiv_x$ of x *component-rigid* if for every admissible perturbation $\hat{F}$ of F the unique nearby $\approx$–class of trajectories of $\hat{F}$ contains some $\hat{x}$ with pattern of synchrony $\equiv_{\hat{x}} \geqslant \equiv_x$.

Conjecture 10.2.3. *Given a periodic trajectory x of a system of differential equations F which is admissible over some network $\mathcal{N}$ of two connected components $\mathcal{N}_1, \mathcal{N}_2$, suppose the pattern of synchrony is unbalanced but component-rigid. Then there is an admissible perturbation $\hat{F}$ of F with an $\approx$–class of trajectories $\hat{X}$ near x such that $\equiv_{\hat{x}} \ngeqslant \equiv_x$ for all $\hat{x}$ in $\hat{X}$.*

Notice that this conjecture is stronger than conjecture 9.3.1, although if the expected structure is followed for proving that conjecture then the proof will involve making a perturbation which breaks synchrony on some time interval, but leaves the synchrony intact elsewhere (compare with the Patch Lemma from Stewart and Parker (December 2007)). Then a phase shift which 'unbroke' the synchrony where it had been broken would break the intact synchrony — and so no phase shift could completely restore the synchrony.

We now examine the case where we duplicate a network with a hyperbolic trajectory; we show that the trajectory given by lifting satisfies the disjoint hyperbolic property.

Proposition 10.2.4. *Let $\mathcal{N}$ be a network with an admissible system of ODEs F; let x be a pseudo-hyperbolic trajectory of F. Define $\mathcal{N}_{1\cup 2}$ and $F_{1\cup 2}$ as before. Then the lifted trajectory $x_{1\cup 2} = (x; x(\theta + \cdot))$ is a trajectory of $F_{1\cup 2}$ that satisfies the disjoint hyperbolic* *property.*

Proof. That $x_{1\cup 2}$ is a trajectory of $F_{1\cup 2}$ is clear. Since x is pseudo-hyperbolic, it has some hyperbolic bound ε. Let θ be given, so that $x_{1\cup 2}$ is uniquely defined. Then for any given $\hat{F}_{1\cup 2}$ which is an admissible ε–perturbation of $F_{1\cup 2}$, admissibility ensures that the projection $\hat{F}$ onto $\mathcal{N}$ is a well-defined ε–perturbation of F. Therefore by the pseudo-hyperbolicity of x, $\hat{F}$ has some unique trajectory $\hat{x}$ near x; lifting to $\mathcal{N}_{1\cup 2}$ yields a family of trajectories $\hat{x}_{1\cup 2} = (\hat{x}(\theta_1 + \cdot)\,; \hat{x}(\theta_2 + \cdot))$. Any other trajectory $\tilde{x}_{1\cup 2}$ of $\hat{F}_{1\cup 2}$ also projects to trajectories $\tilde{x}_1$ and $\tilde{x}_2$ of $\hat{F}$: if $\tilde{x}_{1\cup 2}$ is close to $x_{1\cup 2}$ then both of these trajectories are close to x (possibly after a phase-shift). Hence these trajectories must both be (phase-shifted versions of) the unique trajectory $\hat{x}$ near x. Thus $\tilde{x}_{1\cup 2} = (\hat{x}_1(\theta_1 + \cdot)\,; \hat{x}_2(\theta_2 + \cdot))$, as required. □

10.3 Programme for Proof of the Rigid Phase Conjecture

We now proceed to show that the Rigid Phase Conjecture follows from conjecture 10.2.3. Once again, the proof that balance implies rigidity is trivial, but a technically intricate proof is needed for the other direction.

Lemma 10.3.1. *Balanced phase-shift relations are rigid.*

Proof. Let x be a hyperbolic trajectory of F on $\mathcal{N}$, let $\equiv_x$ be the synchrony relation of x, and $\prec_x$ be some balanced phase-shift relation of x with phase shift θ. Let $\hat{F}$ be an admissible perturbation of F with trajectory $\hat{x}$ near x. Suppose there exist $c, d \in \mathcal{C}(\mathcal{N})$ such that $c \prec_{\hat{x}} d$, but $c \nprec_{\hat{x}} d$. Then let $\tilde{x}$ be such that $\tilde{x}_c(t) = \hat{x}_d(t - \theta)$. This gives a second trajectory of $\hat{F}$ which is close to x, which is impossible since x is hyperbolic. □

Proposition 10.3.2. *Assume conjecture 10.2.3 holds. Then rigid phase-shift relations are balanced.*

Proof. Let F be an admissible system of differential equations on a network $\mathcal{N}$ with Θ–periodic hyperbolic trajectory x. Perturb x to break any fragile synchronies; we now assume that the pattern of synchrony of x is rigid.

Let $\theta \in [0, \Theta)$ be given, and suppose that the phase shift relation of x of phase θ is unbalanced. Let $\prec$ denote this relation. Then there are cells $c, d \in \mathcal{N}$ such that $c \prec d$ but $\mathcal{I}(c) \not\prec_x^1 \mathcal{I}(d)$.

Let $\mathcal{N}_{1\cup2},\ F_{1\cup2},\ x_{1\cup2}$ be as in proposition 10.2.4 — that is, $x_{1\cup2} = (x; x(\theta + \cdot))$: by that proposition, $x_{1\cup2}$ satisfies the disjoint hyperbolic property. Also d_1 is synchronous with c_2, but $\mathcal{I}(d_1) \not\equiv_{1\cup2} \mathcal{I}(c_2)$ (as usual, we use the shortened notation $\equiv_{1\cup2}$ to avoid stacked subscripts). By conjecture 10.2.3, there is an admissible perturbation which breaks the synchrony $\equiv_{1\cup2}$ of $x_{1\cup2}$. Call the perturbed function $\hat{F}_{1\cup2}$ and the trajectory $\hat{x}_{1\cup2}$ — call the broken synchrony $\hat{\equiv}_{1\cup2}$. This perturbation cannot break the synchronies internal to x_1, x_2 since we ensured these were rigid to start with. Hence there are some $c', d' \in \mathcal{N}$ such that $c'_2 \equiv_{1\cup2} d'_1$ but $c'_2 \hat{\not\equiv} d'_1$. (These c', d' could be c, d, but this is not neccesarily the case.)

Now the perturbed trajectories $\hat{x}_1, \hat{x}_2$ on the pieces of the network are perturbed versions of (phase-shifted copies of) x. Undo the phase shift by letting $\hat{x} = \hat{x}_1$ (this copy is not phase-shifted) and $\hat{x}' = \hat{x}_2(\cdot - \theta)$. These trajectories of $\hat{F}$ on $\mathcal{N}$ are both close to x. Therefore they are equal (up to a *small* phase shift η), by the hyperbolicity of x. Hence synchrony of c'_2, d'_1 in $\hat{x}_{1\cup2}$ is equivalent to phase shift of c', d' in $\hat{x}$ by a shift $\hat{\theta} = \theta + \eta$ near θ.

So c' is phase related to d' in x, but not in $\hat{x}$, proving that the phase relation is fragile, as required. □

Theorem 10.3.3 (Limited Rigid Phase Theorem)**.** *Assume conjecture 10.2.3 holds. Then the Rigid Phase Conjecture is true: we spell out what this means.*

Let $\mathcal{N}$ be a network and F an admissible system of differential equations on $\mathcal{N}$ with a hyperbolic trajectory x of period Θ. Let $\theta \in [0, \Theta)$, then the rigid θ–shift relation $\dot{\Lleftarrow}^{\theta}_{x}$ of x must be the maximal balanced refinement of $\Lleftarrow^{\theta}_{x}$ (as a forward relation over $\equiv_x$ on $\mathcal{N}$).

Proof. We follow the same line of argument as theorem 9.3.3. We assume that conjecture 10.2.3 holds throughout this proof.

Let $\lhd_x$ denote the maximal balanced refinement of $\Lleftarrow^{\theta}_{x}$. This relation $\lhd_x$ is a pattern of θ–shift for x, and is balanced. Hence $\lhd_x$ is a rigid θ–shift relation for x, by lemma 10.3.1, so $\lhd_x \leqslant \dot{\Lleftarrow}^{\theta}_{x}$, by maximality of $\dot{\Lleftarrow}^{\theta}_{x}$ as a rigid θ–shift relation.

All rigid patterns of phase shift for x are balanced, by proposition 10.3.2 (we use conjecture 10.2.3 here), and all these patterns must refine $\Lleftarrow^{\theta}_{x}$. Thus $\lhd_x \geqslant \sim$ for all patterns of θ–shift $\prec$, by maximality of $\lhd_x$ as a balanced forward relation: in particular, $\lhd_x \geqslant \dot{\Lleftarrow}^{\theta}_{x}$.

This shows that $\lhd_x = \dot{\Lleftarrow}^{\theta}_{x}$ — that is, the maximal balanced forward relation refining $\Lleftarrow^{\theta}_{x}$ is the rigid pattern of θ–shift of x — and completes the proof of the Rigid Phase Theorem. □

11 Further Directions

Directions for further study have been discussed throughout the thesis, as natural extensions of included material. This chapter concludes the thesis by examining some other directions which do not follow directly from the individual topics examined in detail in the previous chapters.

The first of these topics concerns an application of coupled-cell networks to the models introduced by Nowak (1990); Nowak and May (1992); Lieberman, Hauert, and Nowak (2005). There follow comments on expanding the idea of multiplicity; these ideas are foreshadowed in Aldis (2005), and should not be too complicated to incorporate into the formalism used in this thesis. After a brief reminder of the category $\mathbf{Tree}$, we briefly consider Markov chains as another potential application of coupled-cell networks.

11.1 Network Models for Population Dynamics

Nowak (1990); Nowak and May (1992); Lieberman et al. (2005) describe models for population dynamics based on grids or interconnected networks of cells. These 'Nowak models' should be easily described in terms of our coupled-cell networks. In contrast to the systems of ODEs in this thesis, these Nowak models are discrete-time in nature. The analogue of cell phase space is a finite set S of 'cell states' (which has no necessary structure, unlike our vector space phase spaces): the product of these sets, analogous to network phase space, would be the set of network states — a point in this space $\mathbf{s} : \mathcal{C}(\mathcal{N}) \to S$ could be called a 'network state'. The analogue of the cell functions f_c would then be some rule regarding the updating of the state of cell c. Nowak's models make this function 'choose' from among the states of the inputs of c.

Nowak's results make predictions on the eventual states of these systems. We would expect the symmetries of these states to follow similar patterns to those found in coupled-cell networks representing ODEs, as examined in this thesis: the 'generic' symmetries should be derived from balanced equivalence relations — or, to put it another way, from symmetries of a balanced quotient network.

We also conjecture that in the limit (as time $t \to \infty$), some collections of interconnected cells will (almost surely) attain the same state, and therefore equilibrium: we suggest the name 'homogenised component' for these collections. We would expect these homogenised to be the upstream fully-connected components of the network. The probability distribution of the eventual states of these components could then be considered.

Assuming these suggestions play out well, we present two more possible routes for further study of these Nowak models.

Nonhomogenised Behaviour

We have suggested the classification of distribution of states on homogenised components. Another interesting question is the distribution of states on non-homogenised components. It is clear that for an initial network state with a relation $\sim$ defined by $c \sim d$ where the initial states of cells c and d are equal, cells c, d such that $c \cong_{\sim}^{\infty} d$ will have identical (stochastic) behaviour. We conjecture that small (admissible) changes in edge weights can ensure that other identical behaviour is not preserved. Thus the relation of rigidly identical future behaviour would be the maximal balanced equivalence relation.

A Local Theory of Nowak Models

We now discuss a possible way of reducing the complexity of a Nowak process on a network.

Let $\mathcal{N}$ be a network. Let $r \in \mathbb{N}_\infty$ be given and consider a neighbourhood $\mathrm{B}_r(c)$ of radius r around each cell in $\mathcal{N}$. If $\mathcal{N}$ has some symmetry (possibly a groupoid symmetry, $c \cong_{\mathcal{N}}^{n} d$), then several of these neighbourhoods will be isomorphic (or bunched isomorphic). This gives a concept we call the 'balance metric': for each pair of cells, c, d, let $\mathrm{d}(c, d)$ be 0 if $c = d$, otherwise let it be $1/n$, where n is the smallest integer such that $c \ncong_{\mathcal{N}}^{n} d$.

For a given network state $\mathrm{s} : \mathcal{C}(\mathcal{N}) \to S$, each of these neighbourhoods $\mathrm{B}_r(x)$ will have its own state $\mathrm{s}_r(c)$, which is derived from s by assigning the same states to the same cells. Now, several combinations of network layout and network state will have identical sets $\{ \mathrm{s}_r(c) \mid c \in \mathcal{C}(\mathcal{N}) \}$. As the network's state performs the Markov process dictated by the Nowak model, this neighbourhood state set will have its own random walk. A question for future study is: to what extent does a value of the neighbourhood state set at some time t determine the future (stochastic) behaviour of its random walk? Nowak and May (1992) shows by simulation that in the case $r = 0$, where each neighbourhood consists of only one cell, the answer to this question is "not very much" — this $r = 0$ case is the one which discards all spatial structure, so it is apparent that some spatial structure is important. However, when $r = \infty$, the neighbourhood of a cell contains its entire input tree, and so the behaviour is determined by the current state; the neighbourhood state set follows a Markov chain itself. A question which occurs naturally is: how does this transition occur? Does the accuracy of this neighbourhood model increase faster than its complexity? If so, this would form a good basis for approximation in simulated Nowak models.

11.2 Non-natural Multiplicities

In this thesis, we have considered multisets as the natural objects over which to form a space (the multispace). However, one could go further, and define 'weighted sets' to be sets of objects with multiplicities, like multisets, but where the multiplicities are drawn from some group-like structure, instead of restricting them to be natural numbers. The most obvious extensions are to the rational and negative numbers, and to real numbers. However, it would be possible to take multiplicities from any abelian groupoid.

11.2.1 Bunching

The notion of bunching and bunched tree equivalence was touched on briefly in chapter 4, and then set aside for the remainder of the thesis to simplify the discussion: having natural-valued multiplicities meant that ignoring bunching made no difference to the structures we obtained. However, as our multiplicities become more general, it becomes essential to consider bunching again. Having examined the situation in this simplified 'unbunched' system, it would be useful to consider the technicalities of bunching: this would involve using multispaces themselves more thoroughly, removing most of the underlying spaces which are used in this thesis.

11.2.2 A Category–Theoretic 'Tree Nonsense'

With the use of bunching, we should return to the idea of a category of trees. In chapter 3, we touched on the category $\mathbf{Bunched\ Tree}$, which has trees as its objects and bunching operations as its arrows. It should be possible to frame much of the content of this thesis in terms of this category.

11.3 Balanced Equivalence Relations on Markov Chains

An interesting question arises with reference to Markov chains. Given a network $\mathcal{N}$, a Markov chain is admissible over $\mathcal{N}$ if the states of the chain can be identified with the cells of the network, with edges of the network showing possible state transitions. We conjecture that there is a type of equivalence relation on a network, called a 'Markov-balanced' equivalence relation, which is a balanced equivalence relation with certain extra properties such that if M is a Markov process defined on a network $\mathcal{N}$, and $\sim$ is a 'Markov-balanced' equivalence relation, then there is a Markov process $M_\sim$ on the balanced quotient $\mathcal{N}/\!\!/\!\sim$ such that the behaviour of $M_\sim$ is closely related to that of M.

It seems probable that these relations are those balanced equivalence relations which are also balanced equivalence relations of the network with all edge directions reversed. Further study would be required to ascertain whether this is in fact the case, and to work out the combinatorial issues involved in characterising these relations.

Finding a Markov-balanced equivalence relation would obviously help the computational simulation of Markov chains, as passing to the balanced quotient of a network under such a relation could dramatically reduce the number of states required in many chains.

Bibliography

John W. Aldis. A polynomial time algorithm to determine maximal balanced equivalence relations. Master's thesis, University of Warwick, 2005.

John W. Aldis. A polynomial time algorithm to determine maximal balanced equivalence relations. *Internat. J. Bifur Chaos Appl. Sci. Engrg.*, 18(2):407–428, 2008. doi: 10.1142/S0218127408020367.

D. V. Anosov and V. I. Arnold, editors. *Dynamical Systems I*, volume 1 of *Encyclopedia of Mathematical Sciences*. Springer-Verlag, Berlin Heidelberg New York, 1988.

Ronald Brown. From Groups to Groupoids: a Brief Survey. *Bull. London Math. Soc.*, 19(2): 113–134, 1987. doi: 10.1112/blms/19.2.113. URL http://blms.oxfordjournals.org.

B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, first edition, 1990. ISBN 0521365848.

B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, second edition, 2002. ISBN 0521784514. doi: 10.2277/0521784514.

Ana Paula S. Dias and Ian Stewart. Symmetry groupoids and admissible vector fields for coupled cell networks. *Journal of the London Mathematical Society*, 69(03):707–736, 2004. doi: 10.1112/S0024610704005241. URL http://journals.cambridge.org/action/displayAbstract?fromPage=online&aid=224639&fulltextType=RA&fileId=S0024610704005241.

K.P. Girish and Sunil Jacob John. Relations and functions in multiset context. *Information Sciences*, 179(6):758 – 768, 2009. ISSN 0020-0255. doi: DOI:10.1016/j.ins.2008.11.002. URL http://www.sciencedirect.com/science/article/B6V0C-4TY9MTY-2/2/4a8771229bb84b250e38ce2eb96d542c.

Martin Golubitsky and Ian Stewart. Nonlinear dynamics of networks: the groupoid formalism. *Bull. Amer. Math. Soc*, 43:305–364, 2006.

Martin Golubitsky, Ian Stewart, and Matthew Nicol. Some curious phenomena in coupled cell networks. *J. Nonlin. Sci.*, 14(2):207–236, 2004. doi: 10.1007/s00332-003-0593-6.

Martin Golubitsky, Ian Stewart, and Andrei Török. Patterns of synchrony in coupled cell networks with multiple arrows. *SIAM J. Appl. Dynam. Sys*, 4:78–100, 2005.

Martin Golubitsky, David Romano, and Yunjiao Wang. Transitivity and rigidity of oscillations. Preprint. URL http://people.mbi.ohio-state.edu/mgolubitsky/reprintweb-0.5/output/entireyr.html#2010.

P. J. Higgins. *Notes on categories and groupoids*. Van Nostrand Reinhold, 1971.

M. W. Hirsch and S. Smale. *Differential Equations, Dynamical Systems and Linear Algebra*, volume 60 of *Pure and applied mathematics*. Academic Press, 1974. ISBN 0-12-349550-4.

E. Lieberman, C. Hauert, and M. A. Nowak. Evolutionary dynamics on graphs. *Nature*, 433: 312–316, 2005.

M. A. Nowak. An evolutionarily stable strategy may be inaccessible. *J. Theor. Biol.*, 142: 237–241, 1990.

M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, 1992.

Craig T. Snydal. Equivalence of borcherds g-vertex algebras and axiomatic vertex algebras, 1999. URL http://www.citebase.org/abstract?id=oai:arXiv.org:math/9904104.

Ian Stewart. The lattice of balanced equivalence relations of a coupled cell network. *Mathematical Proceedings of the Cambridge Philosophical Society*, 143(01):165–183, 2007. doi: 10.1017/S0305004107000345. URL http://journals.cambridge.org/action/displayAbstract?fromPage=online&aid=1209128&fulltextType=RA&fileId=S0305004107000345.

Ian Stewart and Martyn Parker. Periodic dynamics of coupled cell networks II: Cyclic symmetry. *Dyn. Syst.*, 23(1):17–41, 2008.

Ian Stewart and Martyn Parker. Periodic dynamics of coupled cell networks I: Rigid patterns of synchrony and phase relations. *Dyn. Syst.*, 22:389–450(62), December 2007. doi: 10.1080/14689360701450410. URL http://www.ingentaconnect.com/content/tandf/cdss/2007/00000022/00000004/art00001.

Ian Stewart and Martyn Parker. Periodic dynamics of coupled cell networks III: Rigid phase patterns. Preprint.

Ian Stewart, Martin Golubitsky, and Marcus Pivato. Symmetry groupoids and patterns of synchrony in coupled cell networks. *SIAM J. Appl. Dynam. Sys*, pages 609–646, 2003.

N. J. Wildberger. A new look at multisets. 2003. URL http://web.maths.unsw.edu.au/~norman/papers/NewMultisets5.pdf.

R. J. Wilson. *Introduction to Graph Theory*. Longman, Harlow, 3 edition, 1985.

Index