Lorentzian Structures on Branching Spacetimes

MSc Thesis *(Afstudeerscriptie)*

written by

David O'Connell
(born 7^{th} August, 1993 in Stevenage, United Kingdom)

under the supervision of **Dr. Alexandru Baltag**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
28^{th} February, 2019	Dr. Floris Roelofsen
	Dr. Nick Bezhanishvili
	Levin Hornischer MSc
	Dr. Klaas Pieter Hart

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis investigates logical models for branching spacetimes, in the tradition started by Belnap. To get a proper generalization of Minkowski spacetime that is closer to physical reality, the thesis adds to Belnap's causal (i.e. order-theoretic) structure appropriate topological, differentiable and metric structures, and investigates the resulting setting, proving a number of interesting original results. Notably, it is shown that the Minkowskian Branching Spacetimes of Belnap's school can be seen as non-Hausdorff, time-oriented Lorentzian manifolds in which each history is an open, Lorentzian submanifold isometric to a fixed Minkowski spacetime. The approach used in this thesis naturally lends itself to the construction of a new class of models, which are named *Lorentzian Branching Spacetimes*. It is shown that these Lorentzian BSTs can be constructed from arbitrary spacetimes, in such a way that many enjoyable causal properties are preserved.

Acknowledgements

Now is the time to express my gratitude to those who contributed to the completion of this thesis. Naturally I'd like to thank Alexandru for all the supervision and engaging coversation, as well as Floris, Nick, Levin and KP for agreeing to be part of the committee. Of course, I also have to thank Benno and Tanja for the support. Special thanks must also go to Thomas Müller and Tomasz Placek for the correspondence and the inspiration, as well as Bob Geroch, Michel Walter, and Sebastian de Haro Olé for responding to my emails and facilitating my learning in various ways. I have to thank Robert, Thijs, and Leo for listening to my mad rambles about my thesis, as well as Sam and Yvette for the much-needed advice when writing it.

I would also like to thank my various friends and family who have supported me – I am a lucky man in that there are too many of you to mention here. Last, and by no means least, I would like to give thanks to my heavenly Father, for all of this.

Introduction

Since at least the 1960's, indeterministic processes have been modelled within a branching framework. The idea, which was historically used for temporal models, is that each branch looks like a conventional timeline, whereas the global structure is tree-like, as pictured in Figure 1. These "branches" are normally referred to as *histories*, and can be seen as the maximally-linear subsets of the model.[1]

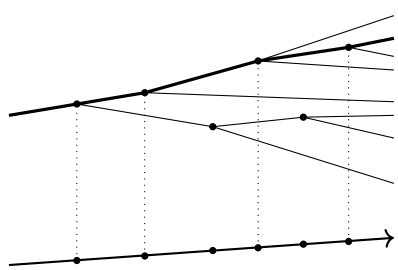

Fig. 1. A branching temporal model with 8 histories, each of which is order-isomorphic to a linear temporal model. There are a certain number of "splitting points" which determine where the histories split from one another.

An important characteristic of branching temporal models is the existence of points that are temporally unrelated. This means that there are points in the model that share a common past, but have disconnected futures. Formally, this manifests as a relaxation of the totality[2] of the temporal ordering, so that the model is only a partial order.

[1] Here maximality is with respect to set-theoretic inclusion, i.e. a history is a linear subset that is not properly contained in any other linear subset.

[2] Recall a binary relation R on a set X is called *total* (or sometimes, *linear*) if for each pair of elements x and y of X, either xRy or yRx.

In 1992, Nuel Belnap generalised this idea by proposing a theory that merged the branching process of an indeterministic model with the causal structure of Minkowski spacetime [1].[3] Belnap's approach was very similar to the description of branching time models – first he provided an axiomatisation of the causal properties of Minkowski spacetime, and then he relaxed the property of directedness[4] in order to facilitate branching. Histories can then be analogously defined as subsets of the model that are maximally-directed.

Belnap's axiomatisation, which we will call BST92,[5] has been widely studied.[6] Notably, there have been discussions of probability theory within BSTs [6], analyses of quantum-mechanical phenomena [7]–[10], and even some applications to the theory of computation [11], [12]. It can be suggested that this broad range of applications has arisen from the ontological simplicity of BST92 – the theory only takes a set W and a binary relation $\leq$ as it's primitives.

The intended class of models of BST92 are the so-called *Minkowskian Branching Spacetimes* (hereafter Minkowskian BSTs, or MBSTs for short). As the name suggests, MBSTs are models of BST92 in which every history is order-isomorphic to a fixed Minkowski spacetime. Figure 2 depicts a simple MBST in which there are two histories, both of which are isomorphic to the 2-dimensional Minkowski spacetime. It should be noted that Minkowskian BSTs were not explicitly constructed until a decade later by Müller [13], and this was later refined by Placek/Wronski [14].

Despite BST92 being a novel generalisation of branching temporal models, Belnap had other ambitions. With regards to his motives for introducing BST92, he writes

> "The aim was to contribute to the problem of uniting relativity with indeterminism in a fully rigorous theory" [1, Pg. 43].

Although Belnap's work was a groundbreaking contribution to this problem, there are some obvious senses in which BST92 and its Minkowskian BSTs do not suffice as a full resolution. In particular, there are two main limitations of Belnap's theory BST92, namely:

---

[3] The original paper can be found as [2], though we will generally refer to the 2003 post-print [1], since it contains more detail.

[4] Recall that a binary order R on a set X is *directed* iff for any pair of elements x and y in X, there exists some z in X such that both xRz and yRz.

[5] BST92 is an abbreviation of "Branching Spacetime 1992", and is also called BST1992 by other authors.

[6] There have been other proposals for BSTs, for instance McCall [3] and Douglas [4] have both proposed topological models that exhibit a different kind of branching known as *individual branching.* Given the arguments of Earman [5] against individual branching, we will not consider such models in this thesis.

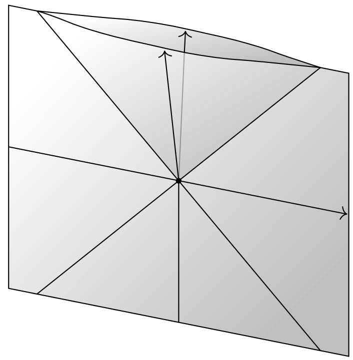

Fig. 2. A 2-dimensional Minkowskian BST in which the origin has two causal futures. These two futures are superimposed on top of one another, and are in some sense "in the same place". For a more-complicated model, see [15, Fig. 12.3(a)]

1. Structurally speaking, models of BST92 are too *coarse* to be interpreted as models of relativity. Specifically, models of general relativity have topological, differentiable and metric structures that need to be accounted for.
2. BST92 and its MBSTs can only deal with special-relativistic branching, and do not treat the models of *general* relativity adequately.

In the last decade, various attempts have been made to resolve these issues. For the first item above, there have been a number of discussions of the topological properties of branching spacetimes – see e.g. [16], and [17] for a more general discussion. A notable contribution came from Placek et. al. [18], who introduced a natural topological extension of BST92 models known as the *Bartha topology*. As of yet there have been no concrete discussions of the potential differentiable and metric structures of the MBSTs of BST92.

As for the second item, Placek has recently provided a logico-mathematical generalisation of BST92 [19]. The models, which are called *genBSTs* are constructed by pasting together MBSTs in a manner analogous to the standard construction of smooth manifolds from local data (as in, say [20, Lem. 1.35]). It is also shown that in certain situations, a topological structure can be defined by pasting together locally-defined Bartha topologies.

In this thesis, we will introduce a novel resolution to these two limitations. As it turns out, this is far easier to do if we adopt a modified version of BST92 recently introduced by Placek in [19], which we will call BST92*. By the end of this thesis, we will have constructed Minkowskian BSTs of BST92* that possess topological, differentiable and metric structures. Moreover, we will argue that these structural extensions are natural under the same criteria used to justify the naturalness of the Bartha topology as in [18, Sec. 6].

The concept that underpins our solution is our choice of construction of the MBSTs of BST92*. In general, there are two ways to take a non-branching object and make it branch at a point. These are

1. to glue on another future at the point, and
2. to take two disjoint copies of the whole space, and then glue them together everywhere *outside* of the future of the point.

Figure 3 depicts the idea behind these two constructions. In the branching spacetime setting, a type-1 construction has been used in [16] and sketched in [21], and a type-2 construction has been used in [13] and [14]. We will opt for a type-2 construction, since this will allow us to easily develop the mathematical machinery needed to complete our task.

Our approach is as follows: first, we construct the Minkowskian BSTs of the modified theory BST92* by gluing together copies of a fixed Minkowski spacetime. The branching processes of a given MBST will be encoded by a collection of interacting subsets which we will call "splitting data". This construction, which will be an adaption of the type-2 construction used by Müller in [13], is performed at the level of the causal structure of Minkowski spacetime. The idea is then to glue together Minkowski spacetimes at the level of their topological, differentiable, and metric structures, in accordance with the same splitting data. The resulting space will be a non-Hausdorff, smooth manifold possessing a metric structure that is isomorphic to the metric structure of the starting Minkowski space, once restricted to its histories.

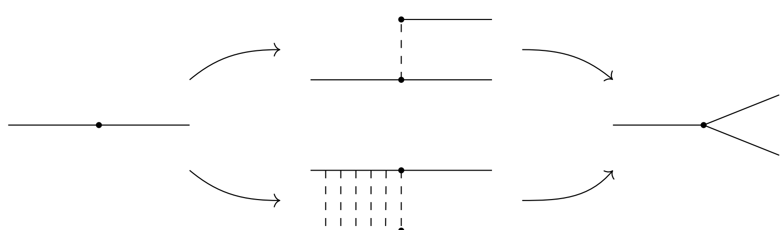

Fig. 3. Two methods for constructing a branching model

This thesis is divided into two parts. The first part develops the metatheory of BST92*. In Chapter 1, we remind the reader of the relevant causal properties of Minkowski spacetime, as well as the basic tenets of Belnap's original theory BST92. In Chapter 2, we introduce BST92* and take the opportunity to prove some yet-to-be-documented basic facts about the system. We will also construct the Minkowskian BSTs of BST92*. Chapter 3 is more of an *intermezzo* in which we will introduce the reader to a general theory of *adjunction spaces*, which are a well-known type of topological space that are formed by gluing together other topological spaces. In Chapter 4, we complete the first part of this thesis by analysing the Bartha topology on BST92* models. This can be seen as a recreation of the results of [18] in the BST92* setting. We will strike a comparison with the BST92 Bartha topology, and draw the conclusion that from a topological perspective, BST92* is actually a *better* theory.

In Part 2 of this thesis, we will extend the theory of adjunction spaces introduced in Chapter 3 to the Lorentzian manifold setting. In Chapter 5, we will show that smooth manifolds can be glued along diffeomorphic, open submanifolds. We will also discuss the nature of vector bundles over adjoined manifolds. In Chapter 6, we complete the thesis by extending the results of the previous chapter one more time, by identifying conditions under which Lorentzian manifolds can be glued to each other. We will then synthesise the results of Part 2 with the developments in Part 1, and express the Minkowskian Branching Spacetimes of BST92* as an adjunction of Minkowski spacetimes at the level of the *Lorentzian* structure. Finally, we will complete the thesis by introducing and discussing a new class of general branching spacetimes, which we will call *Lorentzian Branching Spacetimes.*

Before getting on with things, we should first make a few remarks. First, a note on assumed knowledge. If the reader has not already guessed, we will assume familiarity with the basic notions of topology, order theory, and set theory, as well as the basic tenets of calculus and linear algebra. We will not assume familiarity with differential or Lorentzian geometry – the uninitiated reader is invited to read Appendix A1 before beginning Part 2 of this thesis.[7] Second, we should make explicit the scope of this paper. In this thesis, we will not be pushing any arguments suggesting that the branching spacetime framework is an accurate description of reality. This is a *technical* paper, and we will be intentionally agnostic with regards to any descriptive or ontological claims.

[7] In this appendix we introduce the relevant background material in Lorentzian geometry, all from a non-Hausdorff perspective. This chapter can be seen as a mathematical introduction to spacetime (in the physicists' sense of the word), without assuming the Hausdorff property.

Contents

Part II Lorentzian Branching Spacetimes

Part I

BST92* and its Topological Extension

1

Preliminaries: Minkowski Spacetime and BST92

Before discussing the theory BST92* and its topological development, we will start by introducing the required background material. First, we will introduce the Minkowski spacetime and discuss some of its causal properties. We will then provide a brief overview of Belnap's original theory BST92. It will become very clear that in essence, Belnap is using an axiomatisation of the causal relation on Minkowski spacetime in which a particular property is relaxed in order to facilitate branching. We will also remind the reader of the Bartha topology on BST92 models, since this will be a very useful reference for Chapter 4.

1.1 The Causal Structure of Minkowski Spacetime

A Minkowski spacetime is an n-dimensional Euclidean space $\mathbb{R}^n$ equipped with a *(pseudo)metric* η. We will denote this object by (M^n, η), or where the context is clear, simply M^n. The metric η acts on points $x := (x_0, ..., x_{n-1})$ and $y := (y_0, ..., y_{n-1})$ of M^n by:

$$\eta(x, y) = -x_0 y_0 + x_1 y_1 + ... + x_{n-1} y_{n-1}.$$

The negation sign in front of the first components (which are physically interpreted as the temporal components of x and y) is needed in order to guarantee that the speed of light is constant for all observers. It can also be shown that the metric η is both symmetric and non-degenerate.[1] We also have the following definitions.

Definition 1.1. *Let $x, y \in M^n$. We say that x and y are*

1. *timelike related iff $\eta(x, y) < 0$,*

[1] Recall that η is symmetric if $\eta(x, y) = \eta(y, x)$ for all $x, y \in M^n$, and η is non-degenerate if the only element $x \in M^n$ such that $\eta(x, y) = 0$ for all y is the zero element.

2. *lightlike related iff* $\eta(x, y) = 0$*, and*
3. *spacelike related iff* $\eta(x, y) > 0$*.*

We will call an element x of M^n timelike (lightlike, spacelike) if it is timelike (lightlike, spacelike) related to itself. We can also generalise the above definition to curves[2] in M^n, by saying that a curve γ is timelike (lightlike, spacelike) if its derivative is everywhere timelike (lightlike, spacelike). In particular, we will say that a curve γ is *causal* if its derivative at each point is not spacelike.

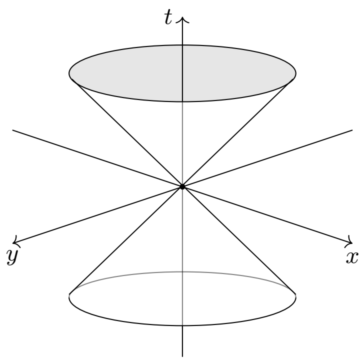

Fig. 1.1: The lightcone at the origin in M^3.

This notion of causal curves can be used to define the *light cone* at a point, which would look something a bit like Figure 1.1. Roughly speaking, the lightcone at a point x in Minkowski spacetime represents the region of M^n that is accessible from x at light speed or less. Since there is no *a priori* notion of the flow of time, we will also assert that time flows "upwards" through M^n.[3] This allows us to divide our lightcones into two parts – the future-facing part (which points upwards), and the past-facing part (which points downwards).

Causal Relations

We can use the concept of future and past lightcones to define some binary relations on Minkowski spacetime that encode its causal structure. Given two points $x, y \in M^n$, we say that x *causally precedes* y, written $x \leq^M y$, if there is a future-directed, causal curve γ connecting x to y, and we will say that x *temporally precedes* y, written $x \ll^M y$, if there is a future-directed, timelike

[2] Recall that a curve is a continuous map $\gamma : \mathbb{R} \to M^n$, and a curve with endpoints is a continuous map $\gamma : [0, 1] \to M^n$.

[3] Such a choice of future/past is called a *time-orientation*. For our current purposes we do not need to formalise this notion, though the interested reader can go to Section A.3.2, where we discuss time-orientations of general-relativistic spacetimes.

curve connecting x to y. Clearly $x \ll^M y$ implies that $x \leq^M y$. Intuitively speaking, x causally precedes y whenever y lies in the future lightcone of x. We also have the following fact about the causal relation $\leq^M$.

We will now introduce some standard notation that will simplify things going forward.

Definition 1.2. *Let x be a point in M^n, and $A \subseteq M^n$.*

- *The causal future of x is the set $J^+(x) := \{y \in M^n \mid x \leq^M y\}$,*
- *The temporal future of x is the set $I^+(x) := \{y \in M^n \mid x \ll^M y\}$, and*
- *The horismotic future of x is the set $E^+(x) := J^+(x) \backslash I^+(x)$.*

We define the causal, temporal and horismotic pasts analogously, and denote these by $J^-(x), I^-(x)$ and $E^-(x)$, respectively.

The relations $\leq^M$ and $\ll^M$ can be rephrased in terms of the sets introduced in the definition above. For a pair of elements $x, y \in M^n$, we have that:

$$x \ll^M y \text{ iff } y \in I^+(x) \text{ iff } \; x \in I^-(y), \text{ and}$$

$$x \leq^M y \text{ iff } y \in J^+(x) \text{ iff } x \in J^-(y).$$

We also have the following result, which follows routinely from the definition of $\leq^M$.

Proposition 1.3. *$(M^n, \leq^M)$ is a dense, directed partial order.*

The directedness of the Minkowski ordering is depicted in Figure 1.2.

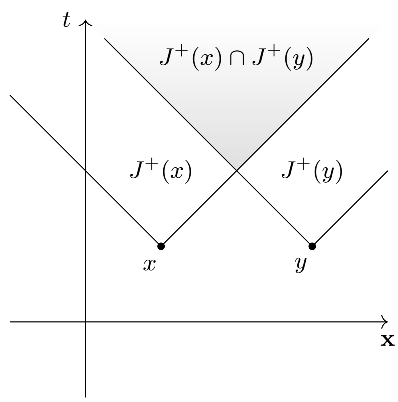

Fig. 1.2: The directedness of the causal relation $\leq^M$.

The causal and temporal sets have the following enjoyable properties.

Proposition 1.4 (Facts about J and I). *Let x, y and z be points in M^n. Then:*

1. *If* $x \in J^+(y)$ *and* $y \in I^+(z)$*, then* $x \in I^+(z)$*.*
2. *If* $x \in I^+(y)$ *and* $y \in J^+(z)$*, then* $x \in I^+(z)$*.*
3. $I^+(x)$ *and* $I^-(x)$ *are open sets in the Euclidean topology*[4] *on* M^n*.*
4. $Cl(I^+(x)) = Cl(J^+(x)) = J^+(x)$*.*

Proof. Items 1 and 2 follow routinely from the definitions of the I and J sets; the proofs can be found in [22, Prop 2.8]. Items 3 and 4 are corollaries of [23, Prop 2.16]. □

It is also possible to generalise the I, J and E formulations to regions of Minkowski spacetime. Given a subset A of M^n, we can define the causal, temporal and horismotic futures of A as

$$J^+(A) := \{y \in M^n \mid \exists x \in A(x \leq^M y)\} = \bigcup_{a \in A} J^+(a),$$

$$I^+(A) := \{y \in M^n \mid \exists x \in A(x \ll^M y)\} = \bigcup_{a \in A} I^+(a), \text{ and}$$

$$E^+(A) := J^+(A) \backslash I^+(A),$$

respectively. The past sets of A are again defined similarly. A corollary of the previous proposition is that the sets $I^+(A)$ and $I^-(A)$ are topologically open.

Causal Chains

We will finish our discussion of the causal properties of Minkowski spacetime by listing a few facts about $\leq^M$-chains that will be needed in subsequent arguments.[5] Our first result should come as no surprise.

Proposition 1.5. *Every causal (timelike) curve* $\gamma : [0,1] \to M^n$ *induces an* $\leq^M$*-chain* C *(an* $\ll^M$*-chain* C*). Moreover,* γ *acts as an order-isomorphism from* $[0,1]$ *to* C*.*

Proof. We claim that the subset $C := \gamma([0,1])$ is the required $\leq^M$-chain. Consider two distinct elements x and y in $\gamma([0,1])$. To see that γ is an order-isomorphism, suppose first that $\gamma^{-1}(x) \leq \gamma^{-1}(y)$. Then we can restrict the curve γ to the interval $[\gamma^{-1}(x), \gamma^{-1}(y)]$, and this will be a future-directed causal curve connecting x to y. Thus $x \leq^M y$. Conversely, if $x \leq^M y$, then since $[0,1]$ is linearly-ordered, either $\gamma^{-1}(x) \leq \gamma^{-1}(y)$ or $\gamma^{-1}(y) \leq \gamma^{-1}(x)$. But the latter case is impossible – if this were true, we could restrict γ to the interval $[\gamma^{-1}(y), \gamma^{-1}(x)]$, and this would imply that $y \leq^M x$. Then we would have that $x \leq^M y \leq^M x$, which contradicts the antisymmetry of $\leq^M$

[4] Recall that the Euclidean topology is the topology generated by the open balls of rational radii, centred at rational coordinates.

[5] Recall a $\leq^M$-chain is a subset of M^n that is linearly-ordered by the relation $\leq^M$.

(since by assumption x and y are distinct). Thus it must be the case that $\gamma^{-1}(x) \leq \gamma^{-1}(y)$.

To see that γ is a bijection from $[0,1]$ to C, it suffices to show that γ is injective. So, suppose towards a contradiction that there are two distinct elements $a, b \in [0,1]$ such that $\gamma(a) = \gamma(b)$. Since $[0,1]$ is linearly ordered, without loss of generality we may assume that $a < b$. Since γ is order-preserving, it follows that $\gamma(a) \leq^M \gamma(b)$. If it were the case that $\gamma(a) = \gamma(b)$, then it would also be the case that $\gamma(b) \leq^M \gamma(a)$, and thus $b \leq a < b$, a contradiction. Hence $\gamma(a) \neq \gamma(b)$, and we may conclude that γ is injective. Since we defined C to be the image of γ, it follows that $\gamma : [0,1] \to C$ is a bijective map that preserves order in both directions, that is, γ is an order-isomorphism. $\square$

Before discussing any more properties of causal chains, we first remind the reader of the following definitions.

Definition 1.6. *Let $(P, \leq)$ be a poset, and C a chain of P. An element p of P is called the supremum of C in P if:*

1. *p upper-bounds C, i.e. $c \leq p$ for every c in C (we will write this as $C \leq p$), and*
2. *p is the minimum element of P with this property, i.e. if there is some other $q \in P$ such that $C \leq q$, then $p \leq q$.*

We will denote the supremum of C by $sup(C)$. Additionally, if A is a subset of P, and $C \subset A$, we denote by $sup_A(C)$ the supremum of C in A.

We define the infimum $inf(C)$ and subset-relative infimum $inf_A(C)$ in an analogous way (i.e. by reversing the ordering $\leq$).

The following result shows that the partial order $(M^n, \leq^M)$ is in some sense *chain-complete*.

Proposition 1.7. *In $(M^n, \leq^M)$ every lower-bounded chain has an infimum and every upper-bounded chain has a supremum.*

Proof. Let C be some $\leq^M$-chain and let $C \leq^M x$ for some $x \in M^n$. Without loss of generality we can extend C to a maximal chain D that passes through x. By the previous result, maximal chains are order-isomorphic to the real line under some mapping $f : D \to \mathbb{R}$. Thus we can consider the subset C as a subset $f(C)$ of $\mathbb{R}$. Since $f(C)$ is upper-bounded by $f(x)$ and $\mathbb{R}$ has the upper-bound property, there is some supremum r of $f(C)$. Order-isomorphisms preserve suprema, so it follows that $f^{-1}(r)$ is the supremum of C in M^n. A similar argument can be made for the existence of an infimum in the case where C is lower-bounded. $\square$

The above result, together with Proposition 1.3 gives us the following corollary.

Corollary 1.8. *Let M^n be a Minkowski spacetime. Then $(M^n, \leq^M)$ is a directed partial order in which every upper-bounded chain has a supremum, and every lower-bounded chain has an infimum.*

We will now finish our discussion of the causal properties of Minkowski spacetime with a few results that will be needed when we discuss the Bartha topology on BST92(*) models.

Proposition 1.9. *If $y \in I^-(x)$ then there exists an $\leq^M$-chain C such that $y \in C$ and $sup(C) = x$.*

Proof. By definition, $y \in I^-(x)$ implies that there is a future-directed, timelike curve connecting y to x. By the previous proposition, the curve γ induces a maximal $\leq^M$-chain C on M^n, and acts as an order-isomorphism from C to the interval $[0, 1]$. We now show that $\gamma(1) = x$ is the supremum of C in M^n. Since γ is an order-isomorphism, and 1 upper-bounds the interval $[0, 1]$, it follows that x upper-bounds C in M^n. Clearly x is the least such element – if there were some $z \in M^n$ such that $C \leq^M z$ and $z <^M x$, then $x \in C$ would imply that $x \leq^M z <^M x$, which is a contradiction. Thus $sup(C) = x$ as required. □

Proposition 1.10. *Let x, y, z be distinct elements of M^n such that $z \in I^+(x) \cap I^-(y)$, and let C be a maximal $\leq^M$-chain of M^n passing through z. Then:*

1. *there is an element $c_1 \in C$ such that $x \ll^M c_1 <^M z$, and*
2. *there is an element $c_2 \in C$ such that $z <^M c_2 \ll^M y$.*

Proof. The hypothesis is displayed in Figure 1.3. We will only show prove the first item, since the other proof is similar. Consider the subchain $D := C \backslash I^+(x)$. Observe that D is upper-bounded by z, and non-empty since C is maximal.[6] Thus by Prop. 1.7 the supremum $d := sup(D)$ exists. We now show that $d \in C$. Let $c \in C$ be arbitrary. If $c \in D$ then $c \leq^M d$, and we are done. So, suppose that $c \notin D$, that is, $c \in I^+(x)$. Let c' be any element of D. If it were the case that $c' \not\leq^M c$, then since C is a chain, $c \leq^M c'$ and thus $c' \in I^+(x)$, contradicting c' as a member of D. Thus $c' \leq^M c$. Since we chose c' arbitrarily, it follows that c upper-bounds D and thus $d \leq^M c$. It follows that d is comparable with every c in C, that is, $C \cup \{d\}$ is a $\leq^M$-chain. Since C is maximal, it follows that $d \in C$ as required.

To obtain the element c_1 as in Fig. 1.3, we can use the density of $\leq^M$ to pick c_1 in C such that $d <^M c_1 <^M z$. Observe that since d is the supremum of D, $c_1 \notin I^+(x)$ would imply that $c_1 \in D$ and thus $c_1 \leq^M d <^M c_1$, a contradiction. Thus $c_1 \in I^+(x)$ as required. □

[6] In the case that $D = \emptyset$, it would follow that $C \subseteq I^+(x)$, and thus we can form a chain $C \cup \{x\}$ that extends C, contradicting its maximality.

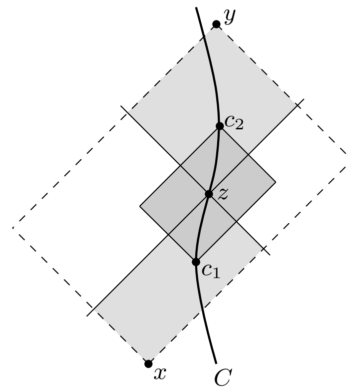

Fig. 1.3: The set-up of Proposition 1.10.

1.2 The Theory BST92

We will now introduce the logical theory BST92 introduced by Belnap [1]. Given the scope of this thesis, we will not delve into the philosophical motivations of BST92,[7] instead we will focus mainly on the technical details. We will first introduce the axioms of BST92, then discuss its models, and finally we will recall the basic facts about the Bartha topology.

As mentioned in the introduction, the theory BST92 has two primitives, namely a set W and a binary relation $\leq$ on W. The idea here is that the set W represents the collection of spacetime points, and the binary relation $\leq$ encodes the causal structure of W. There is an intentional similarity between $(W, \leq)$ and the causal structure $(M^n, \leq^M)$ of a Minkowski spacetime. In fact, most of the axioms of the theory are postulated in order to force the relation $\leq$ to mimic the properties of the causal ordering $\leq^M$. The first axiom of BST92 takes inspiration from Proposition 1.3, and can be stated as follows.

BST1: The tuple $(W, \leq)$ is a dense partial order with no maxima.

As with the formation of branching temporal models, Belnap enables branching by relaxing a property of the ordering $\leq$. Histories (i.e. branches) of a given model are then described as maximal subsets of the model that retain this globally-relaxed property. In BST92, the relaxed property is the directedness of the ordering $\leq$ (which is why the above axiom only assumes that $\leq$ is a partial order, cf. Prop. 1.3). This motivates the following definition of histories.

Definition 1.11. *A subset h of W is a history iff it is maximally-directed. We denote the set of histories of W by $H(W)$.*

[7] For a discussion of the philosphical underpinnings of BST92, see Belnap's paper.

Note that in the above definition, we mean maximal with respect to set-theoretic inclusion.[8] Each history naturally possesses a causal structure inherited from the ordering $\leq$ restricted to h. From this notion of a history and the axiom **BST1**, we have the following results.

Proposition 1.12 (Basic facts about histories).

1. *Histories are downwards-closed.*
2. *Every directed subset of W can be extended to a history.*
3. *Every element x of W belongs to at least one history.*
4. *For every pair of distinct histories h_1 and h_2, the differences $h_1 \backslash h_2$ and $h_2 \backslash h_1$ are non-empty.*
5. *Every finite subset A of a history h has an upper bound in h.*

Proof. 1. Suppose that $h \subseteq W$ is a history with $x \in h$ and $y \leq x$. Pick any $z \in h$. Since h is directed, there is some element $w \in h$ such that $z \leq w$ and $x \leq w$. Since $y \leq x$, it follows from the transitivity of $\leq$ that $y \leq w$ as well. Since z was an arbitrary member of h, we can conclude that $h \cup \{y\}$ is a directed subset of W. The maximality of h then implies that $y \in h$.
2. This result follows as an application of Zorn's lemma.
3. Follows from the previous item, since the singleton $\{x\}$ is directed.
4. If the difference $h_1 \backslash h_2$ was empty, then h_1 would be a subset of h_2. Since we assumed that h_1 and h_2 are distinct, h_1 would be a proper subset of h_2, which contradicts h_1 as maximal.
5. A standard inductive argument. □

The next two axioms of BST92 regard the behaviour of $\leq$-chains. These axioms assert that the infima and suprema of a bounded chain C exist in every *history* containing C, which amounts to postulating that histories behave like $(M^n, \leq^M)$ as in 1.7. The axioms are stated as follows.

BST2: If C is a lower-bounded chain of W then C has an infimum in W, which we denote by $inf(C)$.

BST3: If C is an upper-bounded chain of W then C has a suprema in every history h such that $C \subseteq h$. We denote such a supremum by $sup_h(C)$.

Note that in the statement of **BST2**, we can unambiguously refer to *the* infimum of C, since whenever C lies in the intersection of two histories h_1 and h_2, the downward closure of histories ensures that $inf_{h_1}(C) = inf_{h_2}(C)$. We should also remark that the above axioms are required since it is not guaranteed from **BST1** alone that such suprema and infima exist.[9] Before

[8] That is, h is a history iff h is directed and there is no directed subset h' of W such that $h \subsetneq h'$.

[9] As a counterexample, consider the rational line $(\mathbb{Q}, \leq)$. This satisfies **BST1**, since $\mathbb{Q}$ is a dense linear order with no maxima. However, the rational line is famously *not* complete: a subset such as $C := \mathbb{Q} \cap [0, \pi)$ is an upper-bounded chain that has no suprema in $\mathbb{Q}$.

introducing the final axiom of BST92, we first give some results about chains.

Proposition 1.13 (Basic facts about chains). *Let* $(W, \leq)$ *be a tuple that satisfies axioms* ***BST1-3****, and let* C *be a chain of* W*.*

1. C *can be extended to a maximal chain.*
2. $C \subseteq h$ *for some history* h *of* W*.*
3. *If* $C \subset h_1 \cap h_2$ *and* $sup_{h_1}(C) \neq sup_{h_2}(C)$*, then* $sup_{h_1}(C) \in h_1 \backslash h_2$ *and* $sup_{h_2}(C) \in h_2 \backslash h_1$*.*

Proof. 1. A standard Zorn's lemma argument.
2. Since C is a linear order, in particular it is directed, so we can use Prop. 1.12.2 to extend C to a history.
3. If $sup_{h_1}(C) \in h_2$ or $sup_{h_1}(C) \in h_2$, then this would contradict suprema as unique. □

The final axiom of BST92 regards the splitting of histories. Roughly speaking, the axiom says that if a chain C is in the set-theoretic difference of histories h_1 and h_2, then there was some event prior to C that triggered the splitting of h_1 and h_2. This axiom, known as the *Prior Choice Principle*, is stated as follows.

PCP: (Prior Choice Principle) If C is a chain of W such that $C \subseteq h_1 \backslash h_2$, then there is some element x of W such that $x \leq C$ and x is maximal in $h_1 \cap h_2$.

We will call an element x a *choice point for histories* h_1 *and* h_2 if x is a maximal element of the intersection $h_1 \cap h_2$. The requirement of maximality with respect to $\leq$ ensures that choice points are the last points before two histories split.

Models of BST92

A trivial example of a BST92 model is the Minkowski spacetime M^n, viewed as a set, together with its causal ordering $\leq^M$. Corollary 1.8 shows that axioms **BST1-3** are satified, and the axiom **PCP** is trivially met since $(M^n, \leq^M)$ has a single history.

A less trivial example used by Belnap is depicted in Figure 1.4. Intuitively, Belnap's model is a modification of the standard 2-dimensional Minkowski spacetime $(M^2, \leq^M)$ in which the future lightcone of the origin is replaced with two copies. There are then two histories, depicted as h_1 and h_2 in the figure. The origin is left in the intersection of h_1 and h_2, and acts as the sole choice point for these histories.

Models such as in Figure 1.4 in which every history is order-isomorphic to some Minkowski spacetime have been studied and constructed by a number of authors, e.g. , [13], [14] and [24]. In this thesis we will refer to this class

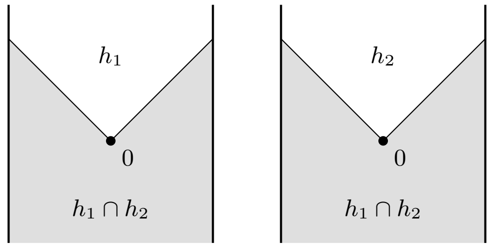

Fig. 1.4: A model of BST92 in which each history is order-isomorphic to M^2, the two-dimensional Minkowski spacetime. Note that the point 0 is included in the intersection $h_1 \cap h_2$, whereas the rest of the border of $h_1 \cap h_2$ is *not* included in the intersection.

of models as the *Minkowskian Branching Spacetimes* (Minkowskian BSTs, or MBSTs for short).

It is also possible to define models in which histories split in multiple places, as in Figure 1.5. However, it has to be the case any choice points x, y for given histories h_1 and h_2 are incomparable under $\leq$, since otherwise this would contradict x and y as maximal elements of $h_1 \cap h_2$. This leads to the possibility of MBSTs in which there are arbitrarily-many histories, all splitting from each other at uncountably-many points. Given this plethora of available models, we will need to restrict our attention to more manageable ones. In this thesis we will work with the following conventions.

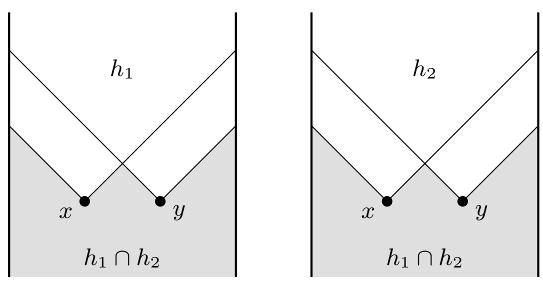

Fig. 1.5: A model of BST92 in which two histories have multiple choice points. Again the points x and y are included in the intersection $h_1 \cap h_2$, whereas the causal futures of these points lie in the differences $h_1 \backslash h_2$ and $h_2 \backslash h_1$.

Remark 1.14 (Assumptions in this thesis).

1. *Each model has at most countably-many histories.* This is necessary for the discussion of non-Hausdorff manifolds in Part 2.

2. *Each model has at most finitely-many choice points.* This is used for two reasons: simplicity, and to exclude models in which splitting occurs at a spacelike hypersurface.

In the remainder of this thesis, unless specified otherwise, we will assume that all models of both BST92 and BST92* satisfy these requirements.

1.2.1 The Bartha Topology on BST92 Models

In this section we introduce a topology on BST92, known as the *Bartha topology*.[10] This topology has been discussed in a number of papers, the most exahustive of which is [18]. As such, we will omit some proofs and instead provide references where necessary. Throughout this section, we fix a model $(W, \leq)$ of BST92.

Before introducing the Bartha topology, we need to define the notion of a *causal diamond*.

Definition 1.15. *Given two elements x and y of W, we define the causal diamond d_{xy} to be the set:*

$$d_{x,y} = \uparrow x \cap \downarrow y = \{z \in W \mid x \leq z \leq y\}.$$

We can now define what it means for a set to be open in the Bartha topology on W.

Definition 1.16. *A subset U of W is open in the Bartha topology iff for every $x \in U$ and every maximal chain C of W passing through x, there are elements c_1 and c_2 in C such that $c_1 < x < c_2$ and $d_{c_1 c_2} \subseteq U$. We denote the collection of all such sets by τ_B^W.*

Where the context is clear, we will sometimes simplify τ_B^W to τ_B. We will also casually refer to the condition listed in the above definition as the *Bartha condition*. Figure 1.6 depicts the intuition behind the Bartha condition.

We now confirm that the Bartha topology is well-defined.

Proposition 1.17. *For every BST92 model $(W, \leq)$, the tuple (W, τ_B^W) is a topological space.*

Proof. Clearly $\emptyset$ and W are open, so it suffices to show that τ_B^W is closed under finite intersections and arbitrary unions.

Let U and V be open sets such that $U \cap V \neq \emptyset$. Let $x \in U \cap V$, and let C be a maximal chain such that $x \in C$. Since $x \in U$, there exist $c_1, c_2 \in C$ such that $c_1 < x < c_2$ and $d_{c_1 c_2} \subset U$. Similarly, x is an element of V, so there exist $d_1, d_2 \in C$ such that $d_1 < x < d_2$ and $d_{d_1 d_2} \subset V$. Since C is a

[10] The name comes from P. Bartha, who Belnaps credits with the suggestion of this topology. See [1, Fn.26].

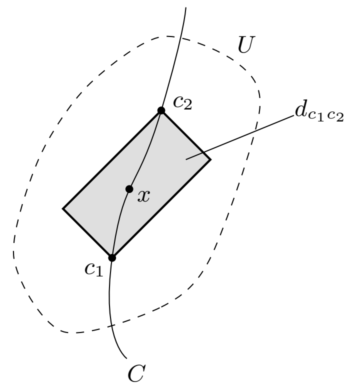

Fig. 1.6: An open set U in the Bartha topology.

chain, the elements x_i and y_i are comparable, so we can always pick the two elements closest to x, and use these for our choice of causal diamond. For instance, suppose that $c_1 \leq d_1 < x < d_2 \leq c_2$. Then $d_{d_1 d_2} \subseteq d_{c_1 c_2} \subset U$ and thus $d_{d_1 d_2} \subset U \cap V$. The other cases are similar. It follows that $U \cap V$ is open.

Suppose now that $\{U_i\}_{i \in I}$ is a collection of open sets, and consider $\bigcup_{i \in I} U_i$. Let $x \in \bigcup_{i \in I} U_i$, and let C be a maximal chain passing through x. Since x lies in a union of sets, there is some $i \in I$ such that $x \in U_i$. By assumption U_i is open, so there are $c_1, c_2 \in C$ such that $c_1 < x < c_2$ and $d_{c_1 c_2} \subset U_i$. Since $U_i \subseteq \bigcup_{i \in I} U_i$, it follows that $d_{c_1 c_2} \subset \bigcup_{i \in I} U_i$ and thus $\bigcup_{i \in I} U_i$ is open. □

We also have the following fact, which will be useful in subsequent arguments.

Lemma 1.18. *Suppose $(W', \leq')$ is a BST92 model and $f : W \to W'$ an order-isomorphism. Then f is also a homeomorphism of Bartha topologies.*

Proof. Let U be open in W, and consider its image $f(U)$. We will show that $f(U)$ is open in W'. Consider an element x of $f(U)$, and a maximal $\leq'$-chain C that passes through x. Since f is an order-isomorphism, the preimage $f^{-1}(C)$ is a maximal $\leq$-chain in W.[11] Since f is a bijection, there is a unique element $f^{-1}(x)$ in U. By assumption U is open in W, so there are elements $c_1, c_2 \in f^{-1}(C)$ such that $f^{-1}(x) \in d_{c_1 c_2} \subset U$. Thus $f(d_{c_1 c_2}) = d_{f(c_1) f(c_2)} \subset f(U)$, from which it follows that $f(U)$ is open in W'. A symmetric argument holds for the converse direction. We may conclude that both f and its inverse f^{-1} are open maps, and consequently f is a homeomorphism. □

[11] Since f is order-preserving, the set $f^{-1}(C)$ is a chain of W. Moreover, if it were the case that there were some $d \in W \backslash f^{-1}(C)$ such that $f^{-1}(C) \cup \{d\}$ were a chain, then $\{f(d)\} \cup C$ would be a chain in W' extending C, which contradicts the maximality of C. Hence no such $dinW$ exists, and thus $f^{-1}(C)$ is a maximal chain of W.

It will also be useful to discuss a history-relative Bartha topology, which we will denote by τ_B^h. This is defined as follows:

Definition 1.19. *Let h be a history of W. A subset U of h is open in the history-relative Bartha topology iff for every $x \in U$ and every maximal chain C of h passing through x, there are c_1 and c_2 in C such that $c_1 < x < c_2$ and $d_{c_1 c_2} \subseteq U$. We denote the collection of all such sets by τ_B^h.*

Strictly speaking we should verify that τ_B^h is well-defined, however the proof is near-identical to that of 1.17. The following proposition is a summary of facts about the Bartha topology on BST92 models. These will be very useful when we discuss the Bartha topology on BST92* models in Chapter 4.

Lemma 1.20 (Basic facts about the Bartha topology on BST92 models). *Let $(W, \leq)$ be a model of BST92 and let h a history of W. Then*

1. *If $U \subseteq h$ is an open set of the topology τ_B^h and U contains a choice point for h and some other history h', then U is not open in τ_B^W.*
2. *If W is a multi-history model, then no histories are open.*
3. *U is open in W iff for every $h \in H(W)$, the set $U \cap h$ is open in (h, τ_B^h).*
4. *Both (W, τ_B^W) and (h, τ_B^h) are connected.*
5. *If W is a multi-history model then W is not Hausdorff.*
6. *If W is a multi-history model then W is not locally Euclidean.*
7. *Given some natural assumptions, each history h is Hausdorff.*

Proof. 1. See [18, Fact 6].
2. An immediate consequence of the previous item.
3. See [18, Fact 7].
4. See [19, Facts 53,54].
5. See [18, Thm. 44].
6. See [19, Lem. 59].
7. See [18, Thm. 35].[12] □

It is also suggested in [18, Fn.10] that the subspace topology

$$\tau_S^h = \{U \cap h \mid U \in \tau_B^W\}$$

induced on h may be strictly coarser than τ_B^h, though no proof is provided.

[12] The "natural assumptions" which are assumed can be found earlier in the same paper as Conditions 15,17,18 and 19. These are imposed in order to restrict the class of models to ones in which the up-sets of events in W, that is the sets of the form $\uparrow x := \{y \in W \mid x \leq y\}$ behave in a manner closer to the light cones of Minkowski spacetime.

Naturality of the Bartha Topology

It has been argued in [18, Sec. 6] that the Bartha topology is a natural extension of BST92's structure. Placek et. al. identify three main criteria for the naturality of the Bartha topology. These can be paraphrased as follows.

1. Given certain circumstances the history-relative Bartha topology coincides with the Euclidean topology,
2. The Bartha topology possesses a certain universal property, and
3. The topological notion of convergences agrees coincides with the order-theoretic notions of suprema and infima.

We can generalise these three criteria to accommodate different structures other than the topological. For any structure S that is proposed as extension of BST92 (or BST92*), we will say that S is a natural extension of the BST92(*) structure whenever the following *naturality criteria* are satisfied.

N1) The structure, once restricted to a single-historied model, should be isomorphic (where relevant) to a Minkowski spacetime.
N2) The structure possesses a certain universal property, in that it can be canonically reconstructed from its history-relative structures.
N3) The structure is compatible with any pre-existing BST92(*) concepts.

2

BST92* and its Minkowskian Branching Spacetimes

In this chapter we will introduce the theory BST92* and its Minkowskian Branching Spacetimes. The key difference between BST92 and BST92* is that in the latter, choice points are replaced by *choice pairs*. At this point, one might ask – why bother to make the leap to BST92* ? Well, we will see in Chapter 4 that the Bartha topologies on BST92* models are far easier to deal with, and that they are truly natural in a way that would be difficult to show in the BST92 case. In fact, these topologies are so well-behaved that they catalyse an extension to the manifold setting, as we will discuss in Part 2 of this thesis. With this in mind, we will work with BST92* .

There is only one resource to date that discusses BST92* , and this is a recent paper by Placek [19]. We will introduce a simplified fragment of this theory, and take the opportunity to prove some basic results. After this, we will construct the Minkowskian BSTs associated to BST92* . Technically speaking these are new models, however, conceptually speaking they are not. We will closely follow the construction of MBSTs used by Müller in [13], and adapt the notation according to our needs.

2.1 The Theory BST92*

For pedagogical reasons, we will start with the intended models of BST92* , and reverse-engineer its axioms. Roughly speaking, we would like the models of BST92* to be near-identical to the models of BST92, except that choice points are replaced by pairs. As such, the causal ordering should still behave similarly to the causal ordering $\leq^M$ of some Minkowski spacetime, and the histories of a BST92* still ought to be isomorphic to some M^n. As such, BST92* should maintain the axioms **BST1-3**.

2.1.1 The Definition of BST92*

We begin by making precise this notion of a choice pair. Throughout this section, we will fix a model $(W, \leq)$ of the axioms **BST1-3**.[1] For simplicities sake, suppose that W has only two histories, say h_1 and h_2. Perhaps a suitable criterion to capture our intuition of choice pairs is the following: two elements x and y of W should form a choice pair for h_1 and h_2 iff once suitably identified, the resultant structure is a BST92 model with a single choice point.

Recall that in a BST92 model $(W', \leq')$, we defined a choice pair for two histories h_1' and h_2' to be an element x of W' that is maximal in the intersection $h_1' \cap h_2'$. So, we could satisfy the criterion above by simply adding in a new point y in the same place as x.[2] That way, identifying y and x would yield a BST92 isomorphic to W'. However the problem with this approach is that in such a model, subsets like $h_1' \cup \{y\}$ would be histories, but are clearly not isomorphic to Minkowski spacetime.[3]

A naive approach to avoiding this problem would be to replace the BST92 choice point x with two points, say y and z, such that y is minimal in $h_1 \backslash h_2$ and z is minimal in $h_2 \backslash h_1$.[4] That way the histories would only contain at most one of the doubled points, which means that they could potentially be isomorphic to some Minkowski spacetime. A naive definition of a choice pair might be something like the following.

Definition 2.1 (Naive Definition of a Choice Pair). *The set $\{y, z\}$ forms a choice pair for histories h_1 and h_2 iff y is minimal in $h_1 \backslash h_2$ and z is minimal in $h_2 \backslash h_1$, or vice-versa.*

The problem with this definition is that the histories h_1 and h_2 may well split at multiple locations, as in Figure 2.1. According to our naive definition, the set $\{x_1, y_2\}$ in the figure would count as a choice pair, which is strange. We need some way of expressing that two points $x \in h_1 \backslash h_2$ and $y \in h_2 \backslash h_1$ are in the same "place". We will follow the terminology of Placek and refer to such elements as *hot pairs*, which are defined as follows.

Definition 2.2. *Let x, y be two distinct elements of W, and h_1, h_2 be histories of W. We say that $\{x, y\}$ is a hot pair for h_1 and h_2 iff there is an upper-bounded, non-empty chain $C \subset h_1 \cap h_2$ such that $sup_{h_1}(C) = x$ and $sup_{h_2}(C) = y$ (or vice versa). We denote the collection of all hot pairs of h_1 and h_2 by $\mathfrak{H}(h_1, h_2)$.*

---

[1] This means that the results of 1.12 and 1.13 still hold in $(W, \leq)$.

[2] Roughly speaking, we could form $W := W' \cup \{y\}$ and define a new causal order $\leq'$ so that $\leq'$ is equal to $\leq$ once restricted to W, and for all $w \in W$, defining $w \leq' y$ iff $w \leq x$, and $y \leq' w$ iff $x \leq w$.

[3] Instead they would be isomorphic to a Minkowski spacetime with two origins.

[4] This means that instead of a choice point being the last element in the intersection $h_1 \cap h_2$, a choice pair would consist of the *first* elements that lie in the set-theoretic differences $h_1 \backslash h_2$ and $h_2 \backslash h_1$.

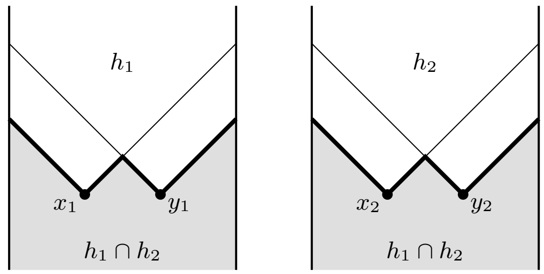

Fig. 2.1: The intuition for a BST92* model. In contrast to Figure 1.5, here $x_1, y_1 \in h_1 \backslash h_2$ and $x_2, y_2 \in h_2 \backslash h_1$. The only choice pairs of this model ought to be $\{x_1, x_2\}$ and $\{y_1, y_2\}$.

Observe that hot pairs are well-defined, since such suprema exist by **BST3**. The requirement that x and y be distinct has a number of consequences. First, this removes the possibility of any singletons existing in $\mathfrak{H}(h_1, h_2)$. Second, since suprema in individual histories are unique, it means that hot pairs have to exist within the difference of h_1 and h_2, that is, $\mathfrak{H}(h_1, h_2) \subset (h_1 \backslash h_2) \cup (h_2 \backslash h_1)$. The intuition behind hot pairs is captured in Figure 2.1, where the hot pairs are displayed by the thick lines above the x_i and y_i.

It seems as though requiring choice pairs to also be hot pairs solves the problem outlined previously: in Figure 2.1, the pair $\{x_1, y_2\}$ is not a hot pair, so is excluded as a possible choice pair of the model. We make this explicit by defining choice pairs as follows.

Definition 2.3. *A hot pair $\{x, y\} \in \mathfrak{H}(h_1, h_2)$ is called a choice pair for h_1 and h_2 iff x is minimal in $h_1 \backslash h_2$ and y is minimal in $h_2 \backslash h_1$, or vice-versa. We denote the set of all choice pairs of h_1 and h_2 by $\mathfrak{C}(h_1, h_2)$.*

Now that we have properly defined the concept of a choice pair, we can introduce the final axiom of BST92* , which is a modification of the Prior Choice Principle.

PCP*: Let C be a chain of W such that $C \subset h \backslash h'$. Then there is some choice pair $\{x, y\} \in \mathfrak{C}(h, h')$ such that $x \in h$ and $x \leq C$.

Observe the similarity between **PCP** and **PCP***. The only difference is that now, we are requiring that the prior choice is witnessed by an element of a choice pair, instead of a single choice point.

We should remark that Placek defines choice pairs differently, and we ought to show that our definition coincides. Observe first that the set of hot pairs $\mathfrak{H}(h_1, h_2)$ has a natural ordering given by:

$$\{x, y\} \preceq \{x', y'\} \text{ iff } x \leq x' \text{ and } y \leq y'.$$

It is not hard to be convinced that $\preceq$ is a partial order on $\mathfrak{H}(h_1, h_2)$, since all of the properties are inherited from the ordering $\leq$ on W.[5] We then have the following.

Lemma 2.4. *If $\{x, y\} \in \mathfrak{C}(h_1, h_2)$ then $\{x, y\}$ is an $\preceq$-minimal element of $\mathfrak{H}(h_1, h_2)$.*

Proof. By definition it is the case that $\{x, y\} \in \mathfrak{H}(h_1, h_2)$. Suppose that there is some $\{w, z\} \in \mathfrak{H}(h_1, h_2)$ such that $w \in h_1$ and $z \in h_2$, and $\{w, z\} \preceq \{x, y\}$. Then $w \leq x$ and $z \leq y$. If $\{w, z\} \neq \{x, y\}$ then we would get $w < x$ or $z < y$, or both. In any case, this would contradict x and y as minimal elements of $h_1 \backslash h_2$ and $h_2 \backslash h_1$ respectively. ☐

In the next section we will prove the converse to the above result.

2.1.2 Basic Facts About Hot Pairs and Choice Pairs

As previously mentioned, there is only one resource for the theory BST92* . As such, there is little in the way of a basic development of the theory. In this section, we will take the opportunity to provide some basic facts about hot pairs and choice pairs. Throughout this section, we will assume that the tuple $(W, \leq)$ is a model of BST92* . We start by proving some basic facts about hot pairs.

Lemma 2.5 (Basic facts about hot pairs). *Let h_1, h_2 be histories, and $x, y \in W$ such that $\{x, y\} \in \mathfrak{H}(h_1, h_2)$. Denote by H_x and H_y the collections of histories of W that contain x and y, respectively.*

1. *If $x \in h_3$, then $\{x, y\} \in \mathfrak{H}(h_3, h_2)$*
2. *$\{x, y\} \in \mathfrak{H}(h_3, h_4)$ for all $h_3 \in H_x$ and $h_4 \in H_y$*
3. *$x \not\leq y$ and $y \not\leq x$*
4. *H_x and H_y are disjoint.*

Proof. Without loss of generality we may assume that $x \in h_1$ and $y \in h_2$. By assumption there is a chain $C \subseteq h_1 \cap h_2$ such that $x = sup_{h_1}(C)$ and $y = sup_{h_2}(C)$.

1. Since h_3 is downwards closed and $C \leq x$, it follows that $C \subseteq h_3$, and thus $C \subseteq h_3 \cap h_2$. Suppose towards a contradiction that $x \neq sup_{h_3}(C)$. By assumption $x = sup_{h_1}(C)$, so $C \leq x$. It follows from the definition of suprema that $sup_{h_3}(C) < x$. By Prop REF, h_1 is downwards-closed, so $sup_{h_3}(C)$ lies in h_1, which contradicts $x = sup_{h_1}(C)$. Thus $x = sup_{h_3}(C)$, from which we may conclude that $\{x, y\} \in \mathfrak{H}(h_2, h_3)$.

[5] Since $x \leq x$ and $y \leq y$, the order $\preceq$ is reflexive. Suppose that $\{x, y\} \preceq \{w, z\} \preceq \{x, y\}$. Without loss of generality let $x \leq w$ and $y \leq z$. If $w \leq x$ then antisymmetry of $\leq$ is violated, and if $w \leq y$ then $x \leq y$ which contradicts Prop 1.14.6. Now suppose that there is some $\{u, v\}$ such that $\{x, y\} \preceq \{w, z\} \preceq \{u, v\}$. There are four cases to consider (namely $x \leq w \leq u$, $x \leq w \leq v$, $x \leq z \leq u$ and $x \leq z \leq v$), and transitivity follows from the transitivity of $\leq$.

2. Follows immediately from two applications of the previous item.
3. Suppose towards a contradiction that $x \leq y$. We can proceed as in item 1, this time using the downward-closure of h_2 to conclude that $C < x < sup_{h_2}(C) = y$, which is a contradiction since by definition, hot pairs consist of distinct elements.
4. Suppose towards a contradiction that there is some history h_3 containing both x and y. Since h_3 is downwards closed, $C \leq x$ implies that $C \subseteq h_3$. By axiom **BST3**, the chain C has a unique supremum $z := sup_{h_3}(C)$. Note that $z \neq x$, since $z = x$ and $C \leq y$ implies that $x \leq y$, contradicting item 3. However, since h_1 is downwards closed, it follows that $z \in h_1$ and thus $C \leq z < x$, which contradicts $x = sup_{h_1}(C)$. □

It is quite often the case that the histories of a branching spacetime (be it in BST92 or BST92*) are pairwise isomorphic. In such a situation, the models are very well-behaved. The following lemma illustrates this.

Lemma 2.6. *Let $h_1, h_2 \in H(W)$ and $\{x, y\} \in \mathfrak{H}(h_1, h_2)$, and let $f : h_1 \to h_2$ be an order-isomorphism.*

1. *f maps x to y.*
2. *If $D \subset h_1 \cap h_2$ is a chain, then $sup_{h_1}(D) = x$ iff $sup_{h_2}(D) = y$.*
3. *$\{x, z\} \notin \mathfrak{H}(h_1, h_2)$ for all z in h_2 distinct from y.*
4. *For every $z \in h_1 \cap h_2$, $z \leq x$ iff $z \leq y$.*

Proof. 1. Since f acts as the identity on the intersection $h_1 \cap h_2$, it follows that any chain C witnessing $\{x, y\}$ as a hot pair will be mapped to itself under f. Since order-isomorphisms preserve suprema, it must be the case that $f(x) = f(sup_{h_1}(C)) = sup_{h_2}(f(C)) = sup_{h_2}(C) = y$.
2. Suppose that $sup_{h_1}(D) = x$. By the previous item, the order-isomorphism f maps x to y. Hence $y = f(x) = f(sup_{h_1}(D)) = sup_{h_2}(D)$. A similar argument holds for the converse, except this time we use the inverse map f^{-1} instead of f.
3. Suppose that there is some $z \in h_2$ and some chain $D \subset h_1 \cap h_2$ such that $sup_{h_1}(D) = x$ and $sup_{h_2}(D) = z$. Since $\{x, y\}$ is a hot pair for h_1 and h_2, it follows from the previous item that $sup_{h_2}(D) = y$, and thus $z = y$.
4. Suppose first that $z \leq x$. We can use the order-isomorphism f to conclude that $f(z) \leq f(x)$. Then item 1 and the fact that f acts as the identity on $h_1 \cap h_2$ imply that $z \leq y$. A symmetric argument holds for when $z \leq y$. □

Now we move on to choice pairs. By definition, all choice pairs are hot pairs, so the results of the previous two lemmas also apply to choice pairs. We also have the following facts specific to choice pairs.

Lemma 2.7 (Basic facts about Choice pairs). *Let $x, y \in W$ and $h_1, h_2 \in H(W)$ such that $\{x, y\} \in \mathfrak{C}(h_1, h_2)$, where $x \in h_1$ and $y \in h_2$. Then:*

1. *If $x \in h_3$, then $\{x, y\} \in \mathfrak{C}(h_3, h_2)$.*
2. *If $x \in h_3$ and $y \in h_4$, then $\{x, y\} \in \mathfrak{C}(h_3, h_4)$.*

3. *If* $\{w, z\} \in \mathfrak{C}(h_3, h_4)$ *where* $x < z$ *and* $x < w$, *then* $\{x, y\} \in \mathfrak{C}(h_3, h_2)$ *and* $\{x, y\} \in \mathfrak{C}(h_4, h_2)$.
4. *If* $\{w, z\} \in \mathfrak{C}(h_1, h_3)$ *and* $w < x$, *then* $\{w, z\} \in \mathfrak{C}(h_2, h_3)$.

Proof. 1. Suppose towards a contradiction that $\{x, y\} \notin \mathfrak{C}(h_3, h_2)$. By Lemma 2.5, we have that $\{x, y\} \in \mathfrak{H}(h_3, h_2)$, since all choice pairs are hot pairs. So, the only way that $\{x, y\}$ does not form a choice pair is if x is not minimal in $h_3 \backslash h_2$, i.e. there is some $z \in h_3 \backslash h_2$ such that $z < x$. However, since $x \in h_1$ and histories are downward closed, it follows that $z \in h_1 \backslash h_2$ which contradicts x as a minimal element of $h_1 \backslash h_2$.
2. Follows from two applications of the previous item.
3. Since histories are downwards closed, x is in both h_3 and h_4. Hence item 1 implies that $\{x, y\}$ is in both $\mathfrak{C}(h_3, h_2)$ and $\mathfrak{C}(h_4, h_2)$.
4. Suppose that $\{w, z\} \in \mathfrak{C}(h_1, h_3)$ where $w \neq x$ and $w < x$. Since h_1 is downwards closed, $w \in h_1$. Also, $w \in h_2$, since otherwise $w \in h_1 \backslash h_2$ and $w < x$ would contradict x as a minimal element of $h_1 \backslash h_2$. Since $w \in h_2$, we can use item 1 to conclude that $\{w, z\} \in \mathfrak{C}(h_2, h_3)$. □

We will now prove the converse of Lemma 2.4.

Lemma 2.8. *If* $\{x, y\}$ *is an* $\preceq$*-minimal element of* $\mathfrak{H}(h_1, h_2)$, *then* $\{x, y\}$ *forms a choice pair for* h_1 *and* h_2.

Proof. Suppose towards a contradiction that x is not minimal in $h_1 \backslash h_2$. Then there is some $z \in h_1 \backslash h_2$ such that $z < x$. Since $z \notin h_2$, by **PCP*** there is some choice pair $\{a, b\}$ for h_1 and h_2 such that $a \leq z$. Since both $\{x, y\}$ and $\{a, b\}$ are hot pairs for h_1 and h_2, it follows from Lemma 2.6.1 that $f(x) = y$ and $f(a) = b$ under the order-isomorphism $f : h_1 \to h_2$. Then $a \leq x$ implies that $b = f(a) \leq f(x) = y$, and thus $\{a, b\} \preceq \{x, y\}$, contradicting $\{x, y\}$ as an $\preceq$-minimal element of $\mathfrak{H}(h_1, h_2)$. Hence no such z exists, and thus x is minimal in $h_1 \backslash h_2$. The case for y is similar. □

The above result, together with Lemma 2.4, shows that our notion of a choice pair is definitionally-equivalent to the one discussed by Placek. We have assumed that histories are pairwise order-isomorphic, though from a practical perspective this is not too important, since we will mostly be working with Minkowskian BSTs (whose histories are always isomorphic to some M^n). We finish our introduction to BST92* with a result that will be useful in Chapter 4.

Lemma 2.9. *Let* W *be a BST92* model. Then every pair of histories of* W *have non-empty intersection.*

Proof. Let h_1 and h_2 be histories of W. Pick some $x \in h_1 \backslash h_2$ (which exists by Prop 1.12.4). By **PCP*** there is some choice pair $\{w, x\}$ for h_1 and h_2 such that $w \leq x$. Since choice pairs are also hot pairs, by definition there is some chain $C \subset h_1 \cap h_2$ whose history-relative suprema are equal to w and z. In particular, C is non-empty, and lies in $h_1 \cap h_2$. □

2.2 Minkowskian Branching Spacetimes

In this section we will construct our Minkowskian BSTs. These are nothing new – in the BST92 setting, MBSTs have been discussed extensively, see e.g. [13], [14], [18] and [24]. As a framework for our construction, we will closely follow the work of Müller found in [13], and simply adapt it to the BST92* setting. We will start by introducing the basic models known as *simple MBSTs*.

2.2.1 A Motivating Example: Simple MBSTs

Simple MBSTs are models in which all the histories split from each other at the origin. Following the convention of Müller [24], we will denote by M^n_m the n-dimensional simple MBST with m-many histories.

To construct M^n_m, we first consider m-many disjoint copies of M^n, i.e.

$$\bigsqcup_{i=1}^{m} M^n_i := \{(x,i) \mid x \in M^n,\ i = 1, ..., m\}.$$

Observe that this set naturally inherits a causal ordering from $\leq^M$. We can then define a relation $\sim$ on this disjoint union as follows:

$$(x,i) \sim (y,j) \text{ iff } \begin{cases} x = y \text{ and } i = j & \text{whenever } 0 \leq^M x \\ x = y & \text{whenever } 0 \not\leq^M x \end{cases}$$

that is, $\sim$ is the reflexive relation that identifies every element (x,i) of M^n_i with its counterparts $(x,j) \in M^n_j$, unless (x,i) is in the future lightcone of $(0,i)$, in which case we only identify (x,i) to itself. It should be clear that $\sim$ is an equivalence relation, with equivalence classes equal to:

$$[(x,i)] := \begin{cases} \{(x,i)\} & \text{if } x \leq 0 \text{ in } \mathbb{R}^n \\ \{(x,j) \mid j = 1, ..., m\} & \text{if } x \not\leq 0 \text{ in } \mathbb{R}^n \end{cases}$$

Throughout this thesis we will abuse notation and denote the equivalence class of a point (x,i) simply by $[x,i]$. We can now define the tuple $(M^n_m, \leq)$, where:

$$M^n_m = \left(\bigsqcup_{i=1}^{m} M^n_i\right)\Big/\!\sim \text{ and } [x,i] \leq [y,j] \ \text{ iff } \ (x,j) \in [x,i] \text{ and } x \leq^M y$$

We will refer to $(M^n_m, \leq)$ as a *simple MBST*, and will typically denote it by M^n_m. Figure 2.2 depicts the construction of the simplest non-trivial[6] model, namely M^2_2.

[6] We have that M^n_1 is just $(\mathbb{R}^n, \leq)$, and in the case of M^1_m, we have a variety of dense branching time models with one branching point.

Strictly speaking, we ought to show that the M^n_m are actually models of BST92* , however we will prove this for the general case in the next section. For now, we remark that the histories of M^n_m will be the images of M^n under the natural embeddings $\iota_i : M^n \to M^n_m$, defined by $\iota_i(x) = [x, i]$.

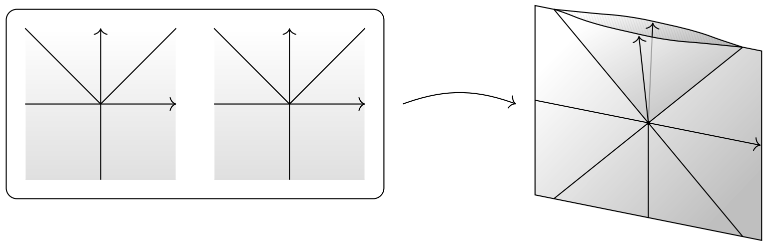

Fig. 2.2: Constructing a binary Minkowski BST in 2 dimensions. Observe that the two future lightcones of the origin are in some sense "on top" of one another.

Remark 2.10. Observe that when defining the equivalence relation $\sim$, it is possible to instead use the strict ordering $<^M$. This would mean that every copy $(0, i)$ is identified. The results of this section would still apply, and the resultant structures will be BST92 models.

2.2.2 The Construction of Minkowskian BSTs

In this section we construct our Minkowskian BSTs. The idea is the same as in the simple case above, except for a few deviations. First, we will not assume that the copies of M^n only split at the origin. This means that for every pair of copies M^n_i and M^n_j, there are a collection of points $C_{ij} \subseteq M^n$ at which these two copies split from each other. In the case of simple MBSTs, each C_{ij} is equal to the singleton $\{0\}$. Second, we will relax the assumption that there are finitely-many copies of M^n. As such, we will assume that there is a potentially-infinite set I that indexes the various copies of M^n. Of course in the simple case, this indexing set is equal to m.

Aside from this, the construction is essentially the same as in the simple MBST case. Specifically, we take I-many copies of some M^n, i.e. $\bigsqcup_{i \in I} M^n_i$, and identify elements (x, i) and (x, j) in the case that x lies outside of the causal future $J^+(C_{ij})$ in M^n. We will show that the resulting space is a BST92* model, and the histories of this model are precisely the images of M^n under the natural embeddings ι_i, as in the simple case. So, we will now construct these models formally.

We start our construction by discussing the nature of the splitting-sets C_{ij}. The collection of all such C_{ij}, which we will denote by $\mathcal{C}$, needs to meet at least some criteria in order to exclude certain pathologies. The following definition makes this precise.

Definition 2.11. *A set $\mathcal{C} := \{C_{ij} \mid i, j \in I\} \subset \mathcal{P}(M^n)$ is called splitting data for M^n iff the set $\bigcup \mathcal{C}$ is finite, and every element C_{ij} of $\mathcal{C}$ satisfies the following conditions:*

C1) *For all $a, b \in C_{ij}$ it is the case that $a \not\leq^M b$ and $b \not\leq^M a$*
C2) $C_{ij} = C_{ji}$
C3) *For each $k \neq i, j$, and for every $a \in C_{ij}$, there exists some $b \in C_{ik} \cup C_{jk}$ such that $b \leq^M a$.*
C4) $C_{ii} = \emptyset$

The splitting data $\mathcal{C}$ should encode the ways in which histories of an MBST split from each other. The various requirements are stipulated in order to capture the intuitions of a branching.

- The requirement that $\bigcup \mathcal{C} = \bigcup_{i,j \in I} C_{ij}$ be finite is imposed in order to cohere with Remark 1.14. Observe that this condition implies that each C_{ij} is finite.
- The condition **C1** is needed for the same reason that **PCP*** is imposed: requiring C_{ij} to be spacelike means that the copies M^n_i and M^n_j cannot split, and then split again at some point afterwards.
- Condition **C2** translates to the intuition that splitting is a symmetric relation: if M^n_i splits from M^n_j at a certain point, then M^n_j splits from M^n_i at the same place.
- Condition **C3** takes inspiration from Lemma 2.5.4 in that there ought to be no history that contains both elements of a (soon-to-be) choice pair.
- Condition **C4** is needed to ensure that no copy of M^n splits from itself.

We can now begin constructing our Minkowskian BSTs. We start by fixing some M^n and some splitting data $\mathcal{C}$. Consider I-many disjoint copies of M^n, i.e. $\bigsqcup_{i \in I} M^n_i$. From this, we can define a relation $\approx$ by:

$$(x, i) \approx (y, j) \text{ iff } x = y \text{ and for all } a \in C_{ij}, \text{ it is not the case that } a \leq^M x.$$

We would like to proceed as in the construction of simple MBSTs and quotient the collection of M^n's under $\approx$. However, it is not immediately clear that the relation $\approx$ is an equivalence relation. The following result confirms that this is indeed the case.

Lemma 2.12. *The relation $\approx$ is an equivalence relation on $\bigsqcup_i M^n_i$.*

Proof. The fact that $\approx$ is a reflexive, symmetric relation is fairly trivial: reflexivity follows from condition **C4**, and symmetry follows immediately from condition **C2**. So, it suffices to show that $\approx$ is transitive.

Suppose that $(x,i) \approx (y,j) \approx (z,k)$. This means that $x = y = z$, so without loss of generality we can relabel these elements as $(x,i),(x,j)$ and (x,k). In the cases that $i = j$, or $j = k$, or $i = k$, the result is trivial[7], so suppose that $i \neq j \neq k \neq i$. Suppose towards a contradiction that $(x,i) \not\approx (x,k)$. Then there must be some $a \in C_{ik}$ such that $a \leq^M x$. Applying condition **C3** to this element a, it follows that there exists some $b \in C_{ij} \cup C_{jk}$ such that $b \leq^M a$. Observe that the transitivity of $\leq^M$ implies that $b \leq^M x$. However, if $b \in C_{ij}$, then $b \leq^M x$ implies that $(x,i) \not\approx (x,j)$, and if $b \in C_{jk}$ then $b \leq^M x$ implies that $(x,j) \not\approx (x,k)$. In either case we contradict our assumption that $(x,i) \approx (x,j) \approx (x,k)$. We can thus conclude that $(x,i) \approx (x,k)$, from which it follows that $\approx$ is transitive. ☐

The above result allows us to define the *MBST subordinate to $\mathcal{C}$*, which we will denote by $(M^n_{\mathcal{C}}, \leq)$. This is defined as follows:

$$M^n_{\mathcal{C}} := \left(\bigsqcup_{i \in I} M^n_i \right) \Big/ \approx, \text{ and } [x,i] \leq [y,j] \text{ iff } (x,j) \in [x,i] \text{ and } x \leq^M y.$$

We also make the following observation, which follows near-immediately from the definition of $\leq$.

Proposition 2.13. *The natural embeddings $\iota_i : M^n \to M^n_{\mathcal{C}}$ given by $x \mapsto [x,i]$ are order-embeddings.*

The remainder of this chapter can be seen as one big proof that the tuple $(M^n_{\mathcal{C}}, \leq)$ is a model of BST92* . We begin with our first result, which is, amongst other things, a confirmation that each $M^n_{\mathcal{C}}$ satisfies the axiom **BST1**.

Lemma 2.14. *The tuple $(M^n_{\mathcal{C}}, \leq)$ is a dense partial order with no maxima.*

Proof. This is fairly routine to verify, since all of these properties are inherited from $\leq^M$. For the interested reader, the proofs are below.

- *Reflexivity:* Since $(x,i) \in [x,i]$ and $x \leq^M x$, it is always the case that $[x,i] \leq [x,i]$.
- *Antisymmetry:* Suppose that $[x,i] \leq [y,j]$. Then $(x,j) \in [x,i]$ and $x \leq^M y$. If $[y,j] \leq [x,i]$, then $(y,i) \in [y,j]$ and $y \leq^M x$. Since $\leq^M$ is antisymmetric, it follows that $x = y$. Since $(x,i) \sim (x,j)$, it follows that $[x,i] = [x,j] = [y,j]$ as required.
- *Transitivity:* Suppose that $[x,i] \leq [y,j] \leq [z,k]$. Then $(x,j) \in [x,i]$ and $(y,k) \in [y,j]$ and $x \leq^M y \leq^M z$. Since $\leq^M$ is transitive, it follows that $x \leq^M z$. We now show that $(x,k) \in [x,i]$. Suppose not. Then $(x,i) \not\approx (x,k)$. It follows from the transitivity of $\approx$ that $(x,j) \not\approx (x,k)$. Hence there exists some $a \in C_{jk}$ such that $a \leq^M x$. However, $x \leq^M y$ implies that $a \leq^M y$, and thus $(y,j) \not\approx (y,k)$, which contradicts $(y,k) \in [y,j]$.

[7] If $i = j$ then $(x,i) \approx (y,i)$ implies that $x = y$ (since $C_{ii} = \emptyset$). Hence $(x,i) = (y,j) \approx (z,k)$ implies that $(x,i) \approx (z,k)$. If $j = k$ then a similar argument holds, and if $i = k$, then $(y,j) \approx (z,i)$ implies that $y = z$, and $(x,i) \approx (y,j)$ implies that $x = y$, hence $(x,i) = (z,k)$ and the result follows from reflexivity.

- *Density*: Suppose that $[x,i] \leq [y,j]$. Then $(x,j) \in [x,i]$ and $x \leq^M y$. Since $\leq^M$ is dense, there exists some $z \in M^n$ such that $x \leq^M z \leq^M y$. The element $[z,j]$ will then lie between $[x,i]$ and $[y,j]$. Indeed: $[x,i] \leq [z,j]$, since $(x,j) \in [x,i]$ and $x \leq^M z$, and $[z,j] \leq [y,j]$ since $(z,j) \in [z,j]$ and $z \leq^M y$.
- *No maxima*: Let $[x,i]$ be any element of $M^n_{\mathcal{C}}$. Since $\leq^M$ has no maxima, there exists some $y \in M^n$ such that $x \leq^M y$. It follows from this that $[x,i] \leq [y,i]$. □

The following result shows that events in the future of splitting data are causally-disconnected in $M^n_{\mathcal{C}}$.

Lemma 2.15. *Let $[x,i]$ and $[y,j]$ be elements of $M^n_{\mathcal{C}}$ such that there exists some $a \in C_{ij}$ where both $a \leq^M x$ and $a \leq^M y$. Then no element of $M^n_{\mathcal{C}}$ upper-bounds both $[x,i]$ and $[y,j]$.*

Proof. If $[x,i] \leq [z,k]$ and $[y,j] \leq [z,k]$, then it would be the case that both $(x,k) \approx (x,i)$ and $(y,k) \approx (y,j)$. However, an application of **C3** gives us the existence of some $b \in C_{jk} \cup C_{ik}$ such that $b \leq^M a$. If $b \in C_{ik}$, then $b \leq^M a \leq^M x$ implies that $b \leq^M x$ and thus $(x,i) \not\approx (x,k)$, and if $b \in C_{jk}$, then $b \leq^M y$ implies that $(y,j) \not\approx (y,k)$. In either case we arrive at a contradiction, thus we may conclude that no such $[z,k]$ exists. □

Corollary 2.16. *If $[x,i] \neq [x,j]$ then $[x,i]$ and $[x,j]$ have no common upper bound.*

Proof. If $[x,i] \neq [x,j]$, then $(x,i) \not\approx (x,j)$, and thus there is some $a \in C_{ij}$ such that $a \leq^M x$. The result then follows from the previous lemma. □

2.2.3 Characterisation of Basic Features

We will now describe the nature of the histories, hot pairs, and choice pairs of Minkowskian BSTs.

Characterisation of Histories

We will now show that, as with the case of simple MBSTs, the histories of $M^n_{\mathcal{C}}$ turn out to coincide with the images of M^n under the order-embeddings ι_i defined as in Proposition 2.13. We will need a way to refer to these images before proving that they are histories, so as an intermediate definition we will follow Müller and refer to the sets $L_i := \iota_i(M^n)$ as *layers*. We ought to also remark that a corollary of Lemma 2.14 is that the MBST $M^n_{\mathcal{C}}$ satisfies the axiom **BST1**, and as such, the results of Proposition 1.12 apply. Our first result shows that each layer is a history.

Lemma 2.17. *Let $M^n_{\mathcal{C}}$ be an MBST. Then each L_i is a maximal directed subset of $M^n_{\mathcal{C}}$.*

Proof. Directedness follows immediately from Prop. 2.13 and the fact that $\leq^M$ is directed. For maximality, suppose that there is some directed subset h of $M^n_{\mathcal{C}}$ such that $L_i \subsetneq h$. Then there is some element $[x, j]$ of $M^n_{\mathcal{C}}$ such that $[x, j] \in h \backslash L_i$. Since h is directed and $[x, i]$ lies in L_i, there is some element of h that is above both $[x, i]$ and $[x, j]$. However, this contradicts Corollary 2.16. Thus we may conclude that that no such subset h exists, and consequently L_i is maximal. □

The converse of Lemma 2.17 is more difficult to prove. We will follow in the footsteps of Müller [13], and appeal to our assumption that our splitting data is finite.[8] We have the following result.

Lemma 2.18. *Suppose $M^n_{\mathcal{C}}$ is an MBST with a finite indexing set I. Then every history h is of the form $\iota_i(M^n)$ for some i in I.*

Proof. Let h be some maximal, directed subset of $M^n_{\mathcal{C}}$, and suppose towards a contradiction that h is distinct from each L_i. Then it cannot be the case that $h \subseteq L_i$ for some i in I, since the maximality of h and Lemma 2.17 would imply that $h = L_i$, contradicting our supposition. Thus for each i there is some element $[x_i, j_i] \in h \backslash L_i$. For each i, pick one such element and let $X = \{[x_i, j_i] \mid i \in I\}$. Observe that the set X is finite, since we are only picking one witness per i, and we have assumed that I is finite. Since the set X is a *finite* subset of h, we use Prop 1.12.5 to conclude that there is some element $a \in h$ such that $X \leq a$. This element a, also being an element of $M^n_{\mathcal{C}}$, must be of the form $[y, k]$ for some k in I. Since $[y, k]$ lies above all elements of X, in particular $[x_k, j_k] \leq [y, k]$. However, since L_k is a history, in particular L_k is downwards-closed, and thus $[x_k, j_k]$ lies in L_k. This contradicts the fact that $[x_k, j_k] \in h \backslash L_k$. We may thus conclude that no such h exists, whence every history of $M^n_{\mathcal{C}}$ is of the form L_i for some i. □

Of course, we can only use the argument above in the case that I is finite, no analogue of Prop. 1.12.5 holds for infinite subsets of a history.[9] The case for when I is countably-infinite is slightly more complicated, but still holds true.

Lemma 2.19. *Suppose $M^n_{\mathcal{C}}$ is an MBST with a countably-infinite indexing set I. Then every history h is of the form $\iota_i(M^n)$ for some i in I.*

Proof. Let h be some maximal directed subset of $M^n_{\mathcal{C}}$, and again suppose towards a contradiction that $h \not\subset L_i$ for all i in I. We will now recursively construct an infinite subset X of $\bigcup \mathcal{C}$. This will suffice for a contradiction,

[8] It should be noted that this is not the only way to prove that the converse to Lemma 2.17 – Placek and Wronski have showed that with a topological assumption, the converse of Lemma 2.16 will hold when the splitting data is infinite. This is done in the context of BST92, so a similar argument might be possible in the BST92* case.

[9] As a counterexample: think of a maximal chain.

since by definition 2.11, the set $\bigcup \mathcal{C}$ is finite. We start the construction by fixing some i in I.

Base Step: By our supposition $h \nsubseteq L_i$, so there is some $[x, j]$ in h such that $[x, j] \neq [x, i]$. It follows that $(x, j) \not\approx (x, i)$ and thus there is some $a_0 \in C_{ij}$ such that $a_0 \leq^M x$. It is also the case that $h \nsubseteq L_j$, and thus there is some $[y, k] \in h$ such that $[y, k] \neq [y, j] \in L_j$. As such, there is some $a_1 \in C_{jk}$ such that $a_1 \leq^M y$. Since h is directed, there is some $[z, l] \in h$ that lies above both $[x, j]$ and $[y, k]$. It follows from the definition of $\leq$ that $(x, l) \approx (x, j)$ and $(y, l) \approx (y, k)$, where both $x \leq^M z$ and $y \leq^M z$.

Suppose towards a contradiction that $a_0 = a_1$. Then $a_0 \in C_{jk}$, and thus by **C3** there is some $a' \in C_{jl} \cup C_{kl}$ such that $a' \leq^M a_0$. If $a' \in C_{jl}$, then $a' \leq^M a_0 \leq^M x$ implies that $(x, l) \not\approx (x, j)$, and if $a' \in C_{kl}$ then $a' \leq^M a_0 = a_1 \leq^M y$ implies that $(y, l) \not\approx (y, k)$. Either of these leads to a contradiction, thus we may conclude that a_0 and a_1 are distinct. We then set $X_0 := \{a_0, a_1\}$. Observe that since $a_0 \in C_{ij}$ and $a_1 \in C_{jk}$, it follows that $X_0 \subseteq \bigcup \mathcal{C}$.

Recursive Step: Suppose now that we have the subset $X_n = \{a_0, ..., a_n\}$ consisting of distinct members of $\bigcup \mathcal{C}$, such that for each $a_\alpha \in X$ there exists some $j_\alpha \in I$ such that $[a_\alpha, j_\alpha] \in h$. Since h is maximally-directed, and $\{[a_\alpha, j_\alpha] \mid \alpha = 0, ..., n\}$ is a finite subset of h, by Prop. 1.12.5 there is some element $[y, k]$ that upper bounds all of the $[a_\alpha, j_\alpha]$. It follows from the definition of $\leq$ that

$$(a_\alpha, k) \approx (a_\alpha, j_\alpha) \text{ and } a_\alpha \leq^M y \text{ for every } \alpha = 0, .., n.$$

By assumption $h \nsubseteq L_k$, so there is some $[z, l] \in h \backslash L_k$. Hence there is some $a_{n+1} \in C_{kl}$ such that $a_{n+1} \leq^M z$. Moreover, since h is directed, there is some $[w, m]$ that lies above both $[z, k]$ and $[y, l]$. Since $\leq$ is transitive and $[y, l]$ lies above all of the $[a_\alpha, j_\alpha]$, it is also the case that $[w, m]$ lies above all of the $[a_\alpha, j_\alpha]$. As such, it follows that $(z, m) \approx (z, l)$ and $(y, m) \approx (y, k)$, and $(a_\alpha, j_\alpha) \approx (a_\alpha, m)$ for every α.

We now show that this a_{n+1} is distinct from all of the a_α. Suppose towards a contradiction that $a_{n+1} = a_\beta$ for some fixed $\beta \in \{0, ..., n\}$. Then $a_\beta \in C_{kl}$, thus by **C3** there is some $a' \in C_{km} \cup C_{lm}$ such that $a' \leq^M a_\beta$. If $a' \in C_{km}$, then $a' \leq^M a_\beta \leq^M y$ implies that $(y, k) \not\approx (y, m)$, and if $a' \in C_{lm}$ then $a' \leq^M a_\beta = a_{n+1} \leq^M z$ implies that $(z, l) \not\approx (z, m)$. Either of these leads to a contradiction, hence it follows that $a_{n+1} \neq a_\beta$. Since a_β was fixed but arbitrary, we may conclude that a_{n+1} is distinct from all the a_α. We then set $X_{n+1} = X_n \cup \{a_{n+1}\}$.

The construction of X is then complete by setting $X := \bigcup_{n \in \mathbb{N}} X_n$. We then see that this set X consists of infinitely-many distinct elements of $\bigcup \mathcal{C}$, which contradicts Definition 2.11. We may thus conclude that $h \subseteq L_i$ for some i, and since h is maximal, Lemma 2.17 implies that $h = L_i$. $\square$

The above two results combine to give a converse to Lemma 2.17. We summarise these results in the following theorem.

Theorem 2.20. *Let $M^n_{\mathcal{C}}$ be an Minkowskian Branching Spacetime. A subset h of $M^n_{\mathcal{C}}$ is a history iff h is of the form $L_i := \iota_i(M^n)$ for some i in I.*

Characterising Hot pairs and Choice Pairs

Now that we know each history of $M^n_{\mathcal{C}}$ is of the form $L_i := \iota_i(M^n)$, we can provide a characterisation of hot pairs and choice pairs. We start with the hot pairs, for which we have the following result.

Lemma 2.21. *The hot pairs for histories L_i and L_j are precisely the elements of the form $\{\iota_i(x), \iota_j(x)\}$, where $x \in E^+(C_{ij})$.*

Proof. Let $x \in J^+(C_{ij}) \backslash I^+(C_{ij})$, and consider some $y \in I^-(x)$. By Prop FIX1.10 there is a a $\leq^M$-chain C in M^n that is contained in $I^-(x)$ and whose supremum is equal to x. Consider now $\iota_i(C)$. Since ι_i is an order-embedding, $\iota_i(C)$ is a $\leq$-chain in $(M^n_{\mathcal{C}}, \leq)$. Moreover, it is also the case that $\iota_i(C) = \iota_j(C)$. Indeed: if it were the case that $\iota_i(c) \neq \iota_j(c)$ for some c in C, then $(c, i) \not\approx (c, j)$, and thus there is some $b \in C_{ij}$ such that $b \leq^M c$, i.e. $c \in J^+(b)$. Since $c \in I^-(x)$, it follows from Prop 1.4.1 that $b \in I^+(x)$, contradicting our assumption that $x \notin I^+(C_{ij})$. Thus $\iota_i(c) = \iota_j(c)$ for every c in C. Since order-embeddings preserve suprema, it follows that $sup_{L_i}(\iota_i(C)) = \iota_i(sup_{M^n}(C)) = \iota_i(x)$ and $sup_{L_j}(\iota_j(C)) = \iota_j(sup_{M^n}(C)) = \iota_j(x)$. Thus we have a chain $\iota_i(C)$ that lies in the intersection $L_i \cap L_j$, and whose suprema in L_i and L_j are equal to $[x, i]$ and $[x, j]$ respectively. It follows that $\{[x, i], [x, j]\}$ is a hot pair for L_i and L_j.

We now show that all the hot pairs are of this form. Suppose that $\{[x, i], [y, j]\}$ is a hot pair for L_i and L_j, and let $C \subset L_i \cap L_j$ be a $\leq$-chain in $M^n_{\mathcal{C}}$ such that $sup_{L_i}(C) = [x, i]$ and $sup_{L_j}(C) = [y, j]$. Since ι_i is an order-isomorphism once restricted to its image L_i, it follows that $sup_{M^n}(\iota_i^{-1}(C)) = \iota_i^{-1}([x, i]) = x$, and similarly, $sup_{M^n}(\iota_j^{-1}(C)) = y$. Since suprema in M^n are unique, it follows that $x = y$. Observe that by assumption, C lies in the intersection $L_i \cap L_j$, so the preimages $\iota_i^{-1}(C)$ and $\iota_j^{-1}(C)$ coincide.

Since $[x, i]$ and $[x, j]$ form a hot pair, and by definition hot pairs consist of distinct elements of $M^n_{\mathcal{C}}$, it must be the case that $(x, i) \not\approx (x, j)$. Thus there is some $a \in C_{ij}$ such that $a \leq^M x$, and hence $x \in J^+(C_{ij})$. We now show that $x \notin I^+(C_{ij})$. Suppose towards a contradiction that there is some $a' \in C_{ij} \cap I^-(x)$. Then by Prop 1.10.1 there is some $c \in C$ such that $a' \leq^M c$. However, it then follows that $[c, i] \neq [c, j]$, which contradicts $C \subseteq L_i \cap L_j$. Thus we may conclude that $I^-(x) \cap C_{ij} = \emptyset$, which completes the proof. □

The above result captures the intuition that hot pairs ought to be the horismotic rims of light cones of the sets C_{ij}. We also have the following result, which confirms that choice pairs of $M^n_{\mathcal{C}}$ correspond to the images of the splitting data $\mathcal{C}$ under the embeddings of 2.13.

Lemma 2.22. *The choice pairs for L_i and L_j are precisely the elements of the form $\{\iota_i(a), \iota_j(a)\}$ where $a \in C_{ij}$.*

Proof. Let $a \in C_{ij}$ and consider the pair $\{[a,i],[a,j]\}$. Since $a \in J^+(a)$ it follows that $a \in J^+(C_{ij})$. It should be fairly clear that $a \notin I^+(C_{ij})$. Indeed – if there were some $b \in C_{ij}$ such that $b \in I^+(a)$, then $a \leq^M b$, which contradicts condition **C1**. By the previous lemma it follows that $[a,i]$ and $[a,j]$ form a hot pair for histories L_i and L_j.

Suppose towards a contradiction that $[a,i]$ is not minimal in $L_i \backslash L_j$. Then there is some $[x,i] \in L_i \backslash L_j$ such that $x <^M a$. If $[x,i] \notin L_j$, then in particular $(x,i) \not\approx (x,j)$, and thus there is some $b \in C_{ij}$ such that $b \leq^M x$. However, then it would be the case that $b <^M a$, and this contradicts **C1**. Thus $[a,i]$ is a minimal element of $L_i \backslash L_j$. By a similar argument it can be shown that $[a,j]$ is minimal in $L_j \backslash L_i$, from which we can conclude that $\{[a,i],[a,j]\}$ is a choice pair for L_i and L_j.

We now show that every choice pair for L_i and L_j is of the form $\{[a,i],[a,j]\}$, where $a \in C_{ij}$. Suppose that $[x,i]$ and $[x,j]$ form a choice pair for L_i and L_j. Since $[x,i]$ and $[x,j]$ also form a hot pair, in particular they are distinct. Then $(x,i) \not\approx (x,j)$, and thus there is some $a \in C_{ij}$ such that $a \leq^M x$. Since $a \in C_{ij}$, it is the case that $(a,i) \not\approx (a,j)$, thus $[a,i] \in L_i \backslash L_j$ and $[a,j] \in L_j \backslash L_i$. Since $a \leq^M x$, it follows that both $[a,i] \leq [x,i]$ and $[a,j] \leq [x,j]$. By assumption $[x,i]$ and $[x,j]$ are minimal elements of the differences $L_i \backslash L_j$ and $L_j \backslash L_i$. Thus $[a,i] = [x,i]$ and $[a,j] = [x,j]$ as required. □

2.2.4 Minkowskian BSTs as Models of BST92*

We now complete this chapter by confirming that each $M^n_{\mathcal{C}}$ satisfies the axioms of BST92* .

Theorem 2.23. *The MBST $(M^n_{\mathcal{C}}, \leq)$ is a model of BST92* .*

Proof. That **BST1** is satisfied follows as a corollary of Lemma 2.14. For axiom **BST2**: let C be an upper-bounded chain in a history h of $M^n_{\mathcal{C}}$. By Thm 2.20 this history h is of the form $L_i = \iota_i(M^n)$ for some i. Then $\iota_i^{-1}(C)$ is a chain in M^n that is upper-bounded. Thus Prop 1.3 implies that there the supremum $sup_{M^n}(\iota_i^{-1}(C))$ exists. Map this element back to L_i under ι_i, and the resulting element is equal to $sup_{L_i}(C)$ (since order-embeddings preserve suprema of chains). This confirms that $M^n_{\mathcal{C}}$ satisfies axiom **BST2**. The argument for **BST3** is similar.

We now verify **PCP***. Suppose towards a contradiction that $C \in L_i \backslash L_j$ is an $\leq$-chain that is lower-bounded by some $[x,i]$, but there is no $a \in C_{ij}$ such that $[a,i] \leq C$. Let $[c_1,i] \in C$. Since $[c_1,i] \notin L_j$, there is some $a_1 \in C_{ij}$ such that $[a_1,i] \leq [c_1,i]$. By assumption $[a_1,i] \not\leq C$, so there is some $[c_2,i]$ in C such that $[a_1,i] \not\leq [c_2,i]$. Note that $[c_2,i] \leq [c_1,i]$. Similarly, there is some $a_2 \in C_{ij}$ such that $[a_2,i] \leq [c_2,i] \leq [c_1,i]$. Thus $a_2 \neq a_1$, since $[a_1,i] \not\leq [c_2,i]$. Repeating this argument, we can inductively define a subset $\{a_n \mid n \in \mathbb{N}\}$ of C_{ij}, which contradicts our assumption that each C_{ij} is finite. We may conclude that there is some $a \in C_{ij}$ such that $[a,i] \leq C$, and thus **PCP*** is satisfied. □

3

Adjunction Spaces

In the previous chapter, we constructed the Minkowskian BSTs of the modified theory BST92* by appropriately gluing together copies of Minkowski spacetime, viewed as a causal structure $(M^n, \leq^M)$. This gluing was performed by taking a quotient of a disjoint collection of the same space. In the topological setting, this type of construction is known as an *adjunction space*. The idea is essentially the same – we can glue two topological spaces, say X and Y, together by quotienting their disjoint union $X \sqcup Y$. The information on where to glue X to Y is encoded within a subspace A of X, and a continuous map $f : A \to Y$, as pictured in Figure 3.1. The adjunction space based on this information is denoted by $X \cup_f Y$.

In Chapter 4, we will show that the MBSTs constructed in Section 2.2.2 can be naturally endowed with a topological structure. We will do this by representing MBSTs as a topological adjunction space. Before we do anything of the sort, we will first need to introduce a basic theory of adjunction spaces. Throughout this chapter, we will do just that.

There are two main ways in which we will deviate from a standard introduction to adjunction spaces found in say, [25]. First, we will explore adjunction spaces formed from a subset A of X that is topologically open.[1] The reason for this is that when constructing MBSTs, we glued the i^{th} and j^{th} copies of M^n together along the subsets $M^n \backslash J^+(C_{ij})$. Since the set $J^+(C_{ij})$ is topologically closed in M^n (see Prop. 1.4.4), this means that we are gluing along *open* subsets of M^n. Second, we will introduce a generalised theory in which it is possible to glue arbitrarily-many topological spaces to each other.

[1] This is in contrast to the standard theory, which typically assumes that A is a closed subset of X. This is done to ensure that the adjunction of two Hausdorff spaces is still Hausdorff.

3.1 The Binary Case

The idea behind a binary adjunction space is fairly simple – we take two topological spaces, say X and Y, and specify the parts which are to be glued to each other. Formally, (binary) adjunction spaces are defined as follows.

Definition 3.1. *Let X, Y be topological spaces and $A \subseteq X$. Given a continuous function $f : A \to Y$, we define the adjunction space of X and Y under f as:*

$$X \cup_f Y := {}^{X \sqcup Y}\!\big/\!_{\sim}$$

where $\sim$ is the reflexive transitive closure of the relation that identifies $(a, 1)$ and $(f(a), 2)$ for each $a \in A$. The topology on $X \cup_f Y$ induced from the quotient of the disjoint union topology on $X \sqcup Y$ is called the adjunct topology, and we will denote this topology by τ_A.

The construction of $X \cup_f Y$ can be described with the following diagram,

$$\begin{tikzcd}
A \arrow[r, "id_X"] \arrow[d, "f"'] & X \arrow[d, "\varphi_1"] \arrow[ddr, dashed, bend left, "\phi_X"] & \\
Y \arrow[r, "\varphi_2"'] \arrow[drr, dashed, bend right, "\phi_Y"'] & X \sqcup Y \arrow[dr, "q"] & \\
 & & X \cup_f Y
\end{tikzcd}$$

where the maps ϕ_X and ϕ_Y are simply the compositions of the canonical inclusions φ_i and the quotient map q associated to the equivalence relation $\sim$, i.e. $\phi_X = q \circ \varphi_1$ and $\phi_Y = q \circ \varphi_2$. Throughout this thesis we will refer to maps such as ϕ_X and ϕ_Y as *canonical maps*. Observe that since the maps φ_i and q are continuous, so are the canonical maps ϕ_X and ϕ_Y. We will follow standard practice and abbreviate the above diagram to the following:

$$\begin{tikzcd}
A \arrow[r, "id_X"] \arrow[d, "f"'] & X \arrow[d, "\phi_X"] \\
Y \arrow[r, "\phi_Y"'] & X \cup_f Y
\end{tikzcd}$$

The diagram above is commutative, which means that on the subset A, the maps ϕ_X and $\phi_Y \circ f$ are equal. Observe also that the points of $X \cup_f Y$ are *equivalence classes*, and be described explicitly as:

- $[x, 1] = \{(x, 1)\}$ if $x \notin A$
- $[y, 2] = \{(y, 2)\}$ if $y \notin f(A)$

- $[a, 1] = \{(f(a), 2)\} \cup \{(a', 1) \mid a' \in f^{-1}(f(a))\}$ for each $a \in A$
- $[y, 2] = \{(y, 2)\} \cup \{(a, 1) \mid a \in f^{-1}(y)\}$ for each $y \in f(A)$

We also have the following observation, which follows immediately from the definition of the topology τ_A.

Proposition 3.2. *Let U be a subset of $X \cup_f Y$. Then U is open in $X \cup_f Y$ iff $\phi_X^{-1}(U)$ and $\phi_Y^{-1}(U)$ are open in their respective spaces.*

The following lemma shows that the adjunction space $X \cup_f Y$ possesses a certain universal property.

Lemma 3.3. *Let Z be a topological space where $\psi_X : X \to Z$ and $\psi_Y : Y \to Z$ are continuous maps such that $\psi_X(a) = \psi_Y \circ f(a)$ for each $a \in A$. Then there is a unique continuous map $g : X \cup_f Y \to Z$ that makes the following diagram commute.*

$$\begin{tikzcd}
A \arrow[r, "id_X"] \arrow[d, "f"'] & X \arrow[d, "\phi_X"'] \arrow[ddr, bend left, "\psi_X"] & \\
Y \arrow[r, "\phi_Y"] \arrow[drr, bend right, "\psi_Y"'] & X \cup_f Y \arrow[dr, dashed, "g"] & \\
& & Z
\end{tikzcd}$$

Proof. We define g by:

$$g([x, i]) = \begin{cases} \psi_X(x) & \text{if } i = 1 \\ \psi_Y(x) & \text{if } i = 2 \end{cases}$$

Observe first that this is well-defined: if $[x, 1] = [y, 2]$, then $f(x) = y$ and thus:

$$g([x, 1]) = \psi_X(x) = \psi_Y \circ f(x) = \psi_Y(y) = g([y, 2]),$$

We now show that this map is continuous. Let $U \subseteq Z$ be open and consider the preimage $g^{-1}(U)$. According to Prop. 3.2, it suffices to show that both $\phi_X^{-1}(g^{-1}(U))$ and $\phi_Y^{-1}(g^{-1}(U))$ are open in their respective spaces. It is not hard to see that $\phi_X^{-1}(g^{-1}(U)) = \psi_X^{-1}(U)$, which is open in X since we assumed ψ_X was continuous. A similar situation holds for Y, from which we can conclude that $g^{-1}(U)$ is open in $X \cup_f Y$, and thus g is continuous.

To see that g is unique, suppose we have some other continuous map $g' : X \cup_f Y \to Z$ such that $\psi_X = g' \circ \phi_X$ and $\psi_Y = g' \circ \phi_Y$. Consider an element $[x, 1]$ in $X \cup_f Y$. Then $g([x, 1]) = g \circ \phi_X(x) = \psi_X(x) = g' \circ \phi_X(x) = g'([x, 1])$ and thus $g([x, 1]) = g'([x, 1])$ for all $x \in X$. The argument for Y is similar, and from this we can conclude that $g = g'$. $\square$

It will be useful to know under what circumstances X and Y topologically embed into $X \cup_f Y$ under the canonical maps ϕ_X and ϕ_Y.[2] The following lemma summarises a number of results.

Lemma 3.4 (Basic facts about the canonical maps). *Let $X \cup_f Y$ be an adjunction space.*

1. *ϕ_Y is always an injection.*
2. *If f is an injection, then so is ϕ_X.*
3. *If A is open, then ϕ_Y is an open map.*
4. *If f is injective and open map, then ϕ_X is open.*
5. *ϕ_Y is always a topological embedding.*
6. *If f is an injective, open map, then ϕ_X is a topological embedding.*

Proof. See Appendix A1. □

In the above result we have omitted certain well-known results for the case when A is assumed to be closed, since these are not useful for our purposes. The following is an immediate corollary of the above lemma, and is a useful identification of some sufficient conditions required to turn the canonical maps ϕ_X and ϕ_Y into open embeddings.

Corollary 3.5. *If A is an open subset of X and f is an injective, open map, then ϕ_X and ϕ_Y are both open topological embeddings.*

It will also be useful to know which properties the adjunction space $X \cup_f Y$ naturally inherits from X and Y. The following lemma collects some results regarding the preservation of various properties, the proofs of which can be found in the Appendices.

Lemma 3.6. *Let $X \cup_f Y$ be an adjunction space.*

1. *If $\mathcal{B}^X$ and $\mathcal{B}^Y$ are bases for X and Y respectively, and ϕ_X and ϕ_Y are open maps, then the collection*

$$\mathcal{B} = \{\phi_X(U) \mid U \in \mathcal{B}^X\} \cup \{\phi_Y(V) \mid V \in \mathcal{B}^Y\}$$

   forms a basis for the topology τ_A.
2. *If X and Y are connected and A is non-empty, then $X \cup_f Y$ is connected.*
3. *If X and Y are compact, then so is $X \cup_f Y$.*

3.2 The General Case

We will now set about generalising the binary adjunction spaces to the setting in which multiple spaces can be adjoined.

[2] Recall that a topological embedding is an injective, continuous map between topological spaces that acts as a homeomorphism once restricted to its image.

As a motivating example, suppose that we have three topological spaces X_1, X_2 and X_3 that we would like to adjoin. In order to join all three spaces together, we need information telling us how to glue the spaces pairwise, namely for each $i = 1, 2, 3$, a collection of subsets A_{ij} of X_i and attaching maps $f_{ij} : A_{ij} \to X_j$. We would like there to be some topological space, which we will tentatively denote by $\bigcup_{\mathcal{F}} X_i$, that is the colimit of the system of X_i, A_{ij} and f_{ij}. The result may be depicted as in the diagram below.

$$
\begin{tikzcd}
& X_1 \arrow[dd, dashed, "\phi_1"] & \\
A_{12} \arrow[ur, "id_{X_1}"] \arrow[dd, "f_{12}"'] & & A_{31} \arrow[ul, "f_{31}"'] \arrow[dd, "id_{X_3}"] \\
& \bigcup_{\mathcal{F}} X_i & \\
X_2 \arrow[ur, dashed, "\phi_2"] & & X_3 \arrow[ul, dashed, "\phi_3"'] \\
& A_{23} \arrow[ul, "id_{X_2}"] \arrow[ur, "f_{23}"'] &
\end{tikzcd}
$$

This can be generalised by fixing some index set I, and defining a tuple of sets $\mathcal{F} := (\mathbf{X}, \mathbf{A}, \mathbf{F})$, where the set $\mathbf{X}$ is a collection of topological spaces X_i, the set $\mathbf{A}$ is a collection of sets A_{ij} such that $A_{ij} \subseteq X_i$ for all $j \in I$, and the set $\mathbf{F}$ is a collection of continuous maps $f_{ij} : A_{ij} \to X_j$. We will now define what it means for such an $\mathcal{F}$ to be an *adjunction system*.

Definition 3.7. *The tuple $\mathcal{F} = (\boldsymbol{X}, \boldsymbol{A}, \boldsymbol{F})$, is called an adjunction system if it satisfies the following conditions for all $i, j \in I$.*

A1) $A_{ii} = X_i$ *and* $f_{ii} = id_{X_i}$
A2) $A_{ji} = f_{ij}(A_{ij})$, *and* $f_{ij}^{-1} = f_{ji}$
A3) $f_{ik}(a) = f_{jk} \circ f_{ij}(a)$ *for each* $a \in A_{ij} \cap A_{ik}$.

The conditions listed above in some sense resemble familiar constructions in differential geometry. Condition **A2** ensures that each f_{ij} is a homeomorphism from A_{ij} to A_{ji}, and thus condition **A3** resembles the *cocycle condition* used when constructing vector bundles from local data.[3]

From an adjunction system, we can then define the adjunction space $\bigcup_{\mathcal{F}} X_i$ as the space obtained from quotienting the disjoint union $\bigsqcup_i X_i$ under the relation $\cong$, where $(x, i) \cong (y, j)$ iff $f_{ij}(x) = y$. The following result confirms that this space is well-defined.

Proposition 3.8. *The relation $\cong$ described above is an equivalence relation.*

[3] Cf. the "Bundle Chart Lemma", listed as Lemma A.14 in Chapter 5.

Proof. Reflexivity and symmetry follow immediately from conditions **A1** and **A2** respectively. For transitivity, suppose that there are three elements $(x,i),(y,j)$ and (z,k) such that $(x,i) \cong (y,j) \cong (z,k)$. Then by definition $f_{ij}(x) = y$ and $f_{jk}(y) = z$. By condition **A2**, it follows that $y \in A_{ji}$ and $f_{ji}(y) = x$. Then $y \in A_{ji} \cap A_{jk}$, hence by **A3** it follows that

$$z = f_{jk}(y) = f_{ik} \circ f_{ji}(y) = f_{ik}(x),$$

thus $(x,i) \cong (z,k)$. □

The above result shows that as a set, the adjunction space $\bigcup_{\mathcal{F}} X_i$ is a collection of equivalence classes of the form

$$[x,i] = \{[y,j] \mid f_{ij}(x) = y\}.$$

As in the previous section, we define the topology on $\bigcup_{\mathcal{F}} X_i$ to be the quotient topology induced from the disjoint union topology $\bigsqcup_i X_i$ by the quotient map q associated to $\cong$. We then have a collection of canonical maps $\phi_i : X_i \to \bigcup_{\mathcal{F}} X_i$ defined as the composition $q \circ \varphi_i$, where φ_i is the canonical embedding into the disjoint union. We make the following observation, which follows immediately by definition.

Proposition 3.9. *A subset U of $\bigcup_{\mathcal{F}} X_i$ is open in the adjunction topology iff $\phi_i^{-1}(U)$ is open in X_i for all i in I.*

We also have the following result, which confirms that the adjunction space is indeed the colimit of the diagram formed from the adjunction system $\mathcal{F}$.

Lemma 3.10. *Let $\psi_i : X_i \to Y$ be a collection of continuous maps from each X_i to some topological space Y, such that for every $i,j \in I$ it is the case that $\psi_i(x) = \psi_j(f_{ij}(x))$. Then there is a unique continuous map g such that $g : \bigcup_{\mathcal{F}} X_i \to Y$ and $\psi_i = g \circ \phi_i$ for all i in I.*

Proof. We define the map g by $g([x,i]) = \psi_i(x)$, that is, $g = \psi_i \circ \phi_i^{-1}$. To see that this defines a function, we need to confirm that g preserves equivalence classes. Suppose that $[x,i] = [y,j]$, i.e. $x = f_{ij}(y)$. Then:

$$g([x,i]) = \psi_i(x) = \psi_j(f_{ij}(x)) = \psi_j(y) = g([y,j])$$

as required. We now show that g is continuous. Let U be open in Y, and consider the set $g^{-1}(U) = \phi_i \circ \psi_i^{-1}(U)$. Recall the set $g^{-1}(U)$ is open in $\bigcup_{\mathcal{F}} X_i$ iff for each $i \in I$, the set $\phi_i^{-1}(g^{-1}(U))$ is open in X_i. Observe that:

$$\phi_i^{-1}(g^{-1}(U)) = \phi_i^{-1} \circ \phi_i \circ \psi_i^{-1}(U) = \psi_i^{-1}(U)$$

which is open since ψ_i is continuous. It follows that $g^{-1}(U)$ is open in $\bigcup_{\mathcal{F}} X_i$, and thus g is continuous. To see that g is unique, we can use the same argument as in 3.3. □

It will be useful to identify under which conditions the maps ϕ_i are open, and are topological embeddings. Note that each ϕ_i is always continuous and injective – continuity follows immediately from the definition of canonical maps, and injectivity follows as a consequence of condition **A1**. We also have the following result, which is an analogue of Corollary 3.5.

Lemma 3.11. *If each f_{ij} is a topological embedding and each A_{ij} is an open subset of X_i, then ϕ_i is an open topological embedding.*

Proof. Fix some ϕ_i. We have already seen that each ϕ_i is injective and continuous. Since every open, continuous, injective map is a topological embedding, it suffices to show that ϕ_i is an open map. Let U be an open subset of X_i, and consider $\phi_i(U)$. By Prop. 3.9 this set is open in $\bigcup_{\mathcal{F}} X_i$ iff for every $j \in I$, the preimage $\phi_j^{-1} \circ \phi_i(U)$ is open in X_j. Observe that $\phi_j^{-1} \circ \phi_i(U) = f_{ij}(U \cap A_{ij})$.[4] Since U is open in X_i, the set $U \cap A_{ij}$ is open in A_{ij} (equipped with the subspace topology). By assumption, $f_{ij} : A_{ij} \to X_j$ is a topological embedding. In particular, it is a homeomophism onto its image, that is, $f_{ij} : A_{ij} \to A_{ji}$ is a homeomorphism. Thus the set $f_{ij}(U \cap A_{ij})$ is open in A_{ji} (again equipped with the subspace topology). By assumption A_{ji} is an *open* subspace, thus $f_{ij}(U \cap A_{ij})$ is also open in X_j.[5] Since U was arbitrary, we may conclude that ϕ_i is an open map, from which the result follows. □

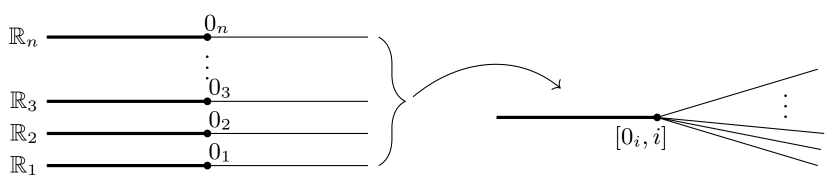

Fig. 3.1: The construction of the n-branched real line.

Example 3.12 (The n-branched real line). If we take I be to be of size n, set each X_i equal to $\mathbb{R}$, each A_{ij} equal to the set $(-\infty, 0)$, and each f_{ij} to be the identity function on $\mathbb{R}$, then the associated tuple $\mathcal{F}$ forms an adjunction system. The adjunction space subordinate to $\mathcal{F}$ will be the real line with n-branches as pictured in Figure 3.1. Of course, it is also possible to relax the requirement that I be finite. In this case, we could obtain a real line with arbitrarily-many branches at the origin.

[4] If $\phi_j(y) \in \phi_i(U)$, then $\phi_j(y) = \phi_i(u)$ for some $u \in U$. Thus $[y, j] = [u, i]$, hence $(y, j) \cong (u, i)$, and $f_{ij}(u) = y$. Thus $y \in f_{ij}(U \cap A_{ij})$. Conversely, if $y \in f_{ij}(U \cap A_{ij})$ then there is some $u \in U \cap A_{ij}$ such that $f_{ij}(u) = y$. Hence $(y, j) \cong (u, i)$, so $[y, j] = [u, i]$ and thus $\phi_j(y) = \phi_i(u)$, i.e. $y \in \phi_j^{-1} \circ \phi_i(u)$.

[5] In an open subspace A of X, $V \subset A$ is open in A iff $V = A \cap U$ where U and A are open in X, so V is the finite intersection of open sets, thus is open in X.

3.2.1 Preservation of Properties

We now prove analogues to the binary case. For our first result, we confirm that bases can be transferred into the adjunction space, provided that each ϕ_i is open.

Lemma 3.13. *Suppose that for each i in I, the collection $\mathcal{B}_i$ forms a basis for X_i. If each ϕ_i is an open map, then the collection $\mathcal{B} = \{\phi_i(B) \mid B \in \mathcal{B}_i\}$ forms a basis for $\bigcup_{\mathcal{F}} X_i$.*

Proof. By assumption each $\phi_i(B)$ is open in $\bigcup_{\mathcal{F}} X_i$. So, it suffices to show that every open set of $\bigcup_{\mathcal{F}} X_i$ can be represented as union of elements of $\mathcal{B}$. Let U be some open subset of $\bigcup_{\mathcal{F}} X_i$, and consider the subset $\phi_i^{-1}(U)$ for some fixed i (which is open in X_i since ϕ_i is continuous). Since $\mathcal{B}_i$ forms a basis for X_i, the set $\phi_i^{-1}(U)$ can be represented as:

$$\phi_i^{-1}(U) = \bigcup_{\alpha \in \mathcal{A}_i} B_\alpha^i$$

where each $B_\alpha^i \in \mathcal{B}_i$. It then follows that:

$$\phi_i(X_i) \cap U = \phi_i(\phi_i^{-1}(U)) = \phi_i\left(\bigcup_{\alpha \in \mathcal{A}_i} B_\alpha^i\right) = \bigcup_{\alpha \in \mathcal{A}_i} \phi_i(B_\alpha^i)$$

and, since the $\phi_i(X_i)$ cover $\bigcup_{\mathcal{F}} X_i$, it follows that:

$$U = U \cap \bigcup_{i \in I} \phi_i(X_i) = \bigcup_{i \in I} (U \cap \phi_i(X_i)) = \bigcup_{i \in I} \bigcup_{\alpha \in \mathcal{A}_i} B_\alpha^i$$

Thus U can be represented as a union of elements of $\mathcal{B}$, as required. □

Recall that a topological space is *second-countable* iff it has a countable basis. The following result is a natural continuation of the previous lemma.

Corollary 3.14. *Suppose each X_i is second-countable. If I is countable and each ϕ_i is an open map, then $\bigcup_{\mathcal{F}} X_i$ is also second-countable.*

Proof. We use the collection $\mathcal{B}$ as in the previous lemma, and since I is countable, the basis $\mathcal{B}$ will be a countable union of countable sets, i.e. a countable basis for $\bigcup_{\mathcal{F}} X_i$. □

We also have the following result regarding the connectedness of the adjunction space, which is an analogue to Lemma 3.6.2.

Lemma 3.15. *If each X_i is connected and each A_{ij} is non-empty, then $\bigcup_{\mathcal{F}} X_i$ is connected.*

Proof. Let $\psi : \bigcup_{\mathcal{F}} X_i \to \{0,1\}$ be a continuous map. Suppose towards a contradiction that ψ is not constant, i.e. suppose there are elements $[x,i]$ and $[y,j]$ of the adjunction space $\bigcup_{\mathcal{F}} X_i$ such that $\psi([x,i]) = 0$ and $\psi([y,j]) = 1$. Since ϕ_i and ϕ_j are continuous, it follows that the maps $\psi_i := \psi \circ \phi_i : X_i \to \{0,1\}$ and $\psi_j := \psi \circ \phi_j$ are continuous maps. By assumption X_i and X_j are both connected, so it follows that ψ_i and ψ_j are both constant. Observe that $\psi_i(x) = \psi([x,i]) = 0$, hence it follows that $\psi_i(X_i) = 0$, and by a similar argument, $\psi_j(X_j) = 1$. However, by assumption the set A_{ij} is non-empty. Then for any element a in A_{ij}, we have that:

$$\psi_j(f_{ij}(a)) = \psi \circ \phi_j \circ f_{ij}(a) = \psi \circ \phi_i(a) = 0$$

which contradicts $\psi_j(X_j) = 1$. We can conclude that the map ψ is constant, and thus $\bigcup_{\mathcal{F}} X_i$ is indeed connected. □

It is not true in general that the adjunction of infinitely-many compact spaces is compact. As such, we have a restricted analogue to Lemma 3.6.3.

Lemma 3.16. *If I is finite and each X_i is compact, then so is $\bigcup_{\mathcal{F}} X_i$.*

Proof. We can use that compactness is preserved in finite disjoint unions and under quotients, as in Lemma 3.6.3, together with [26, Thm 3.2.3]. □

Lemma 3.17. *If $K \subseteq X_i$ is compact, then $\phi_i(K)$ is a compact subset of $\bigcup_{\mathcal{F}} X_i$.*

Proof. This follows immediately from the fact that the continuous image of a compact set is compact.[6] □

Remark 3.18. In general, the maps ϕ_i are not proper, that is, the preimages of compact subsets of $\bigcup_{\mathcal{F}} X_i$ are not necessarily compact in the X_i. To see this, consider again the n-branched real line as in Example 3.12. By the Heine-Borel theorem[7], the subset $[-1,1]$ is a compact subset of $\mathbb{R}_i$. By the previous lemma, the image $\phi_i([-1,1])$ is compact in the space $\bigcup_{\mathcal{F}} \mathbb{R}_i$. Consider now any ϕ_j where j is distinct from i. Then the preimage of this compact subset is:

$$\phi_j^{-1} \circ \phi_i([-1,1]) = f_{ij}([-1,0)) = [-1,0)$$

which is not compact in $\mathbb{R}_j$ since it is not closed.

Recall that a topological space X is *locally-Euclidean* if every point in X has a neighbourhood that is homeomorphic to some Euclidean space $\mathbb{R}^n$. The following result shows that this property can be transferred to an adjunction space.

[6] Let $A \subseteq X$ be compact, and suppose that $\{U_i\}$ forms an open cover of $f(A)$ in Y. Since f is continuous, the sets $\{f^{-1}(U_i)\}$ form an open cover for A. Since A is compact, there is a finite subcover $\{f^{-1}(U_1), ..., f^{-1}(U_n)\}$. Thus $\{U_1, ..., U_n\}$ is a finite subcover of $\{U_i\}$.

[7] The Heine-Borel theorem states that a subset of Euclidean space is compact iff it is closed and bounded. See [27, Thm 27.3].

Lemma 3.19. *If each X_i is locally-Euclidean, and each ϕ_i is an open topological embedding, then $\bigcup_{\mathcal{F}} X_i$ is also locally-Euclidean.*

Proof. Let $[x,i] \in \bigcup_{\mathcal{F}} X_i$. Then $x \in X_i$. Since X_i is locally-Euclidean, there is a chart (U, φ) of X_i at x. We show that $(\phi_i(U), \varphi \circ \phi_i^{-1})$ is a chart for $\bigcup_{\mathcal{F}} X_i$ at $[x,i]$. By assumption ϕ_i is open, thus $\phi_i(U)$ is an open subset of the adjunction space. Since we assumed that ϕ_i is a topological embedding, it is homeomorphic onto its image. Then the restriction of ϕ_i to an open subset is also a homeomorphism, as is its inverse. Thus the map $\varphi \circ \phi_i^{-1} : \phi_i(U) \to \varphi(U)$ is a composition of homeomorphisms, thus is also a homeomorphism. It follows that $(\varphi(U), \varphi \circ \phi_i^{-1})$ is a chart for $\bigcup_{\mathcal{F}} X_i$ at $[x,i]$. $\square$

4

The Bartha Topology on BST92* Models

In this chapter we will evaluate the topological properties of the models of BST92*. We will start with a general treatment of the issue, and provide a BST92*-analogue to the results found in Lemma 1.20. WE will see that the results are very similar, however in the BST92* context the histories of a model are open in the Bartha topology. This is a crucial observation that underpins various arguments throughout the remainder of this thesis. One important consequence, which we will explore in Section 4.2, is that the Bartha topology on a given model can be canonically reconstructed as an adjunction of its history-relative Bartha topologies. In Section 4.3, we will complete our discussion by computing the Bartha topology on the Minkowskian BSTs. The main result here is that the Bartha topology coincides with a topology proposed by Müller in [16]. This encouraging result unifies the order-theoretic approach of Placek et. al. with the topological approach of Müller.

4.1 Comparison with the BST92 Bartha Topology

In this section we evaluate the Bartha topology on BST92* models. The framework for our exploration is the results found in Lemma 1.20 – we would like to see how the Bartha topology interacts with its history-relative counterparts, whether the spaces are connected, and so on. Throughout this section, we fix $(W, \leq)$ as a BST92* model, and let h_1 and h_2 be arbitrary histories of W. The Bartha topology on W is defined as in Definition 1.16.[1] Again we will denote by τ_B^W the Bartha topology on W, and by τ_B^h the Bartha topology relative to the history h of W.

[1] Namely, a subset U of W is open in the Bartha topology iff for each point in U and each maximal chain C passing through x, there are elements c_1 and c_2 in C such that $c_1 < x < c_2$ and the causal diamond $d_{c_1 c_2}$ is contained within U (see Fig. 1.6 for the intuition).

Our first result is an analogue to Lemma 1.18. We omit the proof, since it is identical to the BST92 case.

Proposition 4.1. *Let* $(W, \leq)$ *and* $(W', \leq')$ *be BST92* models. If* $f : W \to W'$ *is an order-isomorphism, then* f *is also a homeomorphism between the topological spaces* (W, τ_B^W) *and* $(W', \tau_B^{W'})$.

Before proving any analogues to the results of Lemma 1.20, we first need a few facts regarding maximal $\leq$-chains.

Lemma 4.2. *Suppose that* C *is a maximal chain of* W*, and* h_1 *and* h_2 *are histories of* W *such that* $C \cap h_1 \neq \emptyset$ *and* $C \not\subseteq h_1$ *and* $C \subseteq h_2$*. Then*

$$sup_{h_2}(C \cap h_1) = inf(C \backslash h_1).$$

Proof. Denote $a := sup_{h_2}(C \cap h_1)$ and $b := inf(C \backslash h_1)$. We show that both a and b are members of C.

Consider first a. Let $c \in C$. If $c \in C \cap h_1$, then by definition $c \leq a$. If $c \notin h_1$, then $c \in C \backslash h_1$, and thus $C \cap h_1 \leq c$ (otherwise, $C \cap h_1 \not\leq c$ implies that there is some $d \in C \cap h_1$ such that $d \leq c$, hence $c \leq d$ and thus $c \in h_1$). Hence $a \leq c$. It follows that $C \cup \{a\}$ is a chain, and since C is maximal, we may conclude that $a \in C$. Consider now b, and let $c \in C$ be arbitrary. If $c \in C \backslash h_1$ then $b \leq c$ by definition. If $c \in h_1$ that $c \in C \cap h_1$, and thus $c \leq C \backslash h_1$ (if not, $d \leq c$ for some $d \in C \backslash h_1$, hence $d \in h_1$ follows from the downward closure of h_1). It follows that $C \cup \{b\}$ is a chain, and hence $b \in C$.

Suppose towards a contradiction that $a \neq b$. Observe that $a \leq C \backslash h_1$, since otherwise any element $d \in C \backslash h_1$ such that $d < a$ would contradict $a = sup_{h_2}(C \cap h_1)$. Hence $a < b$, and since W is dense (see **BST1**), there is some c in W such that $a < c < b$. We now show that this element c lies in the chain C. Consider some $d \in C$. If $d \in h_1$, then $d \leq a < c$ and if $d \notin h_1$, then $d \in C \backslash h_1$ and $c < b \leq d$. It follows that $C \cup \{c\}$ is a chain, hence $c \in C$. However this leads to a contradiction – if $c \in h_1$, then $c \in C \cap h_1$ thus $c \leq a$, and if $c \notin h_1$, then $c \in C \backslash h_1$ thus $b \leq c$. In either case, we contradict $a < c < b$. We may then conclude that $a = b$, as required. ☐

From the above result, we can identify hot pairs using maximal chains. This is done as follows.

Lemma 4.3. *If* C *is a maximal chain of* W*, and* h_1 *and* h_2 *are two histories such that* $C \cap h_1$ *and* $C \not\subset h_1$ *and* $C \subseteq h_2$*, then the pair*

$$\{sup_{h_1}(C \cap h_1), sup_{h_2}(C \cap h_1)\}$$

forms a hot pair for h_1 *and* h_2.

Proof. If $C \not\subseteq h_1$ then there is some $c \in C$ such that $c \notin h_1$. Then this element upper-bounds the subchain $C \cap h_1$. By axiom **BST3** it follows that there exists (possibly two) history-relative suprema $a := sup_{h_1}(C \cap h_1)$ and $b := sup_{h_2}(C \cap h_1)$. We show that $\{a, b\} \in \mathfrak{H}(h_1, h_2)$.

Observe first that $C \cap h_1 \subset h_1 \cap h_2$. So, to show that a and b form a hot pair for h_1 and h_2, it suffices to show that $a \neq b$. Suppose towards a contradiction that $a = b$, and consider the subchain $C \backslash h_1$. This is a chain in the difference $h_2 \backslash h_1$, so by **PCP*** there exists some choice pair $\{x, y\} \in \mathfrak{C}(h_1, h_2)$ such that $x \leq C \backslash h_1$ (where here $y \in h_1$ and $x \in h_2$). By the previous lemma, $b = inf(C \backslash h_1)$, hence $x \leq b = a$. Since we have supposed that $a = b$, and a is by definition a member of h_1, we may use the downwards-closure of histories to conclude that $x \in h_1$. However, this means that both elements of the choice pair $\{x, y\}$ lie in the history h_1, which contradicts Lemma 2.5.4. Thus $a \neq b$, and we may conclude that $\{a, b\} \in \mathfrak{H}(h_1, h_2)$. □

Using the above lemma, we can prove the following result, which says that every open set in τ_B^h is also an open set in τ_B^W.

Lemma 4.4. *Let h be a history of W and τ_B^h its associated Bartha topology. Then $\tau_B^h \subseteq \tau_B^W$.*

Proof. Let U be some element of τ_B^h. Consider some element x in U and some chain C that is a maximal chain of W. If $C \subset h$, then C is also a maximal chain of h, and the result follows immediately. So, suppose that $C \not\subseteq h$. Then there is some $c \in C$ such that $c \notin h$. By Lemma 4.3, it follows that $\{sup_h(C \cap h), sup_{h'}(C \cap h)\}$ forms a hot pair for the histories h and h', where h' is some history such that $C \subset h'$. Since $sup_h(C \cap h)$ is compatible with every element of $C \cap h$, we can extend $(C \cap h) \cup \{sup_h(C \cap h)\}$ to a chain D that is maximal in h. By assumption U is open in the Bartha topology on h, and since $x \in C \cap h \subset D$, and D is a maximal chain in h, it follows that there are $c_1, c_2 \in D$ such that $c_1 < x < c_2$ and $d_{c_1 c_2} \subset U$. Without loss of generality we can assume that $c_2 < sup_h(C \cap h)$ (i.e. $c_2 \in C \cap h$), since otherwise we can take some $c_3 \in C \cap h$ such that $x < c_3 < sup_h(C \cap h) \leq c_2$ and use that $d_{c_1 c_3} \subset d_{c_1 c_2} \subset U$. Since c_1 and c_2 are members of C, it follows that $U \in \tau_B^W$ as required. □

An immediate corollary of the above is the following.

Corollary 4.5. *Every history h of W is open in the topology τ_B^W.*

This is in stark contrast to the BST92 case (cf. Lemma 1.20.2), and this is one of the key reasons that we have adopted BST92* over BST92 in this thesis. We also have the following useful characterisation of τ_B^h, which is suspected to be false in the BST92 setting.

Lemma 4.6. *The Bartha topology τ_B^h on a history coincides with the subspace topology $\tau_S^h := \{U \cap h \mid U \in \tau_B^W\}$.*

Proof. Let $V \in \tau_B^h$. Then by Lemma 4.4 $V \in \tau_B^W$, hence $V = V \cap h \in \tau_S^h$, and consequently $\tau_B^h \subseteq \tau_S^h$. For the converse, suppose that $V \in \tau_S^h$, i.e. $V = U \cap h$ where $U \in \tau_B^W$. Let $x \in V$ and C a maximal chain in h. Since every maximal chain in h is also a maximal chain in W, it follows from our assumption that there are c_1 and c_2 in C such that $x \in d_{c_1 c_2} \subset U$. Since $U \in \tau_B^W$ and $x \in U$, there are c_1 and c_2 in C such that $x \in d_{c_1 c_2} \subset U$. Since $c_2 \in C \subset h$, the downward closure of H implies that $d_{c_1 c_2} \subset h$, and thus $d_{c_1 c_2} \subset U \cap h = V$. It follows that V is an open set of τ_B^h, from which the result follows. □

Corollary 4.7. *A subset U of W is open in the Bartha topology on W iff for each history h of W, the set $U \cap h$ is open in the history-relative Bartha topology τ_B^h.*

Proof. The direction from left-to-right follows from the previous lemma. So, suppose that $U \cap h$ is open in τ_B^h for all h in $H(W)$. By Lemma 4.4, the sets $U \cap h$ are also open in τ_B^W. Then $U = \bigcup_{h \in H(W)} (U \cap h)$ is a union of open sets, and is thus open in τ_B^W as well. □

From the above result we can conclude that the analogue of Lemma 1.20.3 holds in the BST92* case.

We will now move on to discussing the connectedness of the Bartha topology. In the next two results, we will prove the BST92*-analogue to Lemma 1.20.4. We start with a proof that all histories of a BST92* model are connected. Before introducing the argument, it should be noted that the following proof is an adaptation of that found in [19, Fact 53].[2]

Lemma 4.8. *The topology (h, τ_B^h) is connected.*

Proof. Suppose towards a contradiction that there are two open subsets $U, V \in \tau_B^h$ such that $U \cap V = \emptyset$ and $U \cup V = h$. Since both U and V are empty, there are two elements $x, y \in h$ such that $x \in U$ and $y \in V$. Since h is directed, there is some $z \in h$ that upper-bounds both x and y. By assumption U and V cover h, so z must lie in one of these sets. Without loss of generality, suppose that $z \in U$. Since $y \leq z$, we can extend $\{y, z\}$ to a chain C that is maximal in h. We can restrict the chain C to the segment $D := C \cap \downarrow z \cap \uparrow y$, which can be seen as a *maximal chain with endpoints y and z*.[3]

Consider now the subchain $D \cap U$. Since this is a chain that is lower-bounded by y, by **BST3** there exists an infimum $inf(D \cap U)$. We now show

[2] Unfortunately, a modification of [19, Fact 53] is quite necessary. This is because Placek applies the axiom **BST3** to a chain without first confirming that the chain in question is lower-bounded. As it turns out, the chain in question *is* lower-bounded, but this complicates things sufficiently as to warrant a reshaping of the argument.

[3] Here we do not use "maximal" in our normal sense, but instead we mean that D is a maximal element of the poset of bounded chains in h whose endpoints are y and z (ordered by inclusion.)

that the element $inf(D \cap U)$ is an element of the chain D. Let a be some element in D. We proceed via case distinction:

Case 1: $a \in U$. Then $a \in D \cap U$, so by definition we have that $inf(D \cap U) \leq a$.
Case 2: $a \in V$ and a lower-bounds the sub-chain $D \cap U$. Since $inf(D \cap U)$ is the *greatest* lower-bound of $D \cap U$, we have that $a \leq inf(D \cap U)$.
Case 3: $a \in V$ and a does not lower-bound $D \cap U$. Then there is some $b \in D \cap U$ such that $b < a$. Since $inf(D \cap U)$ lower-bounds $D \cap U$, we have that $inf(D \cap U) \leq b < a$, whence $inf(D \cap U) < a$.

We can conclude from the above cases that $inf(D \cap U)$ is comparable with every element of D. Moreover, the first case above ensures that $inf(D \cap U) < z$, and the second case above ensures that $y \leq inf(D \cap U)$. We then have that $D \cup \{inf(D \cap U)\}$ is a chain whose endpoints are y and z, from which it follows from the maximality of D that $inf(D \cap U) \in D$.

Now that we have proved that $inf(D \cap U)$ lies in D, consider again our original chain C, which is maximal h. Since $inf(D \cap U) \in D \subseteq C$, we can now apply the Bartha condition in the relevant places to obtain our contradiction. Since U and V cover h, it must be the case that $inf(D \cap U)$ lies in one of these two covering sets.

If $inf(D \cap U) \in U$, then we can apply the Bartha condition to C and U to conclude that there are $c_1, c_2 \in C$ such that $c_1 < inf(D \cap U) < c_2$ and $d_{c_1 c_2} \subset U$. Observe that $y < c_1$ (since otherwise, $c_1 < y < inf(D \cap U) < c_2$ implies that $y \in U$). Then c_1 is an element of the subchain $D \cap U$ that is strictly below $inf(D \cap U)$, which is a contradiction.

Similarly, if $inf(D \cap U) \in V$, then we can again apply the Bartha condition to C and V to conclude that there are $e_1, e_2 \in C$ such that $e_1 < inf(D \cap U) < e_2$ and $d_{e_1 e_2} \subset V$. Observe that e_2 lower-bounds the sub-chain $D \cap U$ (since otherwise, there would be some element $a \in D \cap U$ such that $inf(D \cap U) \leq a < e_2$, so $a \in V$). Then $inf(D \cap U) < e_2 < D \cap U$, which is again a contradiction.

In either case, we arrive at a contradiction. Thus we may conclude that no such U and V exist, and therefore h is indeed connected. □

Lemma 4.9. *The topology* (W, τ_B^W) *is connected.*

Proof. Suppose towards a contradiction that there are two non-empty, open subsets $U, V \subseteq W$ that cover W. Since both U and V are non-empty, there are elements $x \in U$ and $y \in V$. By Lemma 1.12.3, there are histories h_1 and h_2 of W such that $x \in h_1$ and $y \in h_2$.

Consider first the history h_1. As witnessed by x, the set $h_1 \cap U$ is non-empty. Moreover, since U is open in W, by Lemma 4.7 the set $h_1 \cap U$ is open in the history-relative Bartha topology on h_1. It then has to be the case that $h_1 \cap V = \emptyset$, else we contradict the connectedness of h_1. Since h_1 and V are disjoint, and U, V cover W, we may conclude that $h_1 \subseteq U$.

We can perform a similar argument for h_2 to conclude that $h_2 \subseteq V$. By assumption U and V are disjoint, which means that h_1 and h_2 are also disjoint. However, this contradicts Lemma 2.9. □

We will now finish this section with a look at the Hausdorff property on W. The following result tells us that hot pairs are Hausdorff-violating.

Lemma 4.10. *Let h_1 and h_2 be histories of W. Then every hot pair $\{x, y\} \in \mathfrak{H}(h_1, h_2)$ violates the Hausdorff property.*

Proof. Without loss of generality suppose that $x \in h_1$ and $y \in h_2$. Let U, V be open sets of W such that $x \in U$ and $y \in V$. We show that U and V are not disjoint. Since x and y form a hot pair, there is some chain $C \subset h_1 \cap h_2$ such that $sup_{h_1}(C) = x$ and $sup_{h_2}(C) = y$. Let C_1 and C_2 be two maximal extensions[4] of C containing x and y respectively. Observe[5] that $C = \{z \in C_1 \mid z < x\}$ and similarly $C = \{z \in C_2 \mid z < y\}$.

Since U is open, $x \in U$ and $x \in C_1$, by definition there are two points $x_1, x_2 \in C_1$ such that $x_1 < x < x_2$ and $d_{x_1 x_2} \subset U$. Since $x_1 < x$, it must be the case that $x_1 \in C$, hence $x_1 \in C \cap U$. Similarly, since $x < x_2$ it must be the case that $x_2 \in C_1 \backslash C$. We can repeat the same argument for y, C_2 and V to conclude that there is some $y_1, y_2 \in C_2$ such that $y_1 < y < y_2$, with $y_1 \in C \cap V$ and $d_{y_1 y_2} \subset V$.

We would like to show the existence of at least some point in the intersection $U \cap V$. In fact, any point $z \in C$ that lies above both x_1 and y_1 will suffice[6]. Indeed, we have $x_1 \leq z < x < x_2$ and $y_1 \leq z < y < y_2$, thus both $z \in d_{x_1 x_2} \subset V$ and $d_{y_1 y_2} \subset V$, and thus $z \in U \cap V$ as required. Since we chose U and V arbitrarily, it follows that x and y violate the Hausdorff property. $\square$

The following theorem summarises the results of this section, and serves as a direct comparison to Lemma 1.20.

Theorem 4.11. *Let $(W, \leq)$ be a model of BST92*, and h a history of W. Then*

1. $\tau_B^h \subseteq \tau_B^W$
2. *All histories h of W are open.*
3. *U is open in W iff $U \cap h$ is open in (h, τ_B^h) for each h in $H(W)$.*
4. *Both W and h are connected.*
5. *If W is a multi-history model, then W is not Hausdorff.*
6. *The history-relative Bartha topology τ_B^h coincides with the subspace topology induced from τ_B^W.*

[4] These always exist by Zorn's lemma.

[5] If $z \in C$ then $z < x$ ($z \neq x$ since $x \notin h_2$ and $C \subset h_1 \cap h_2$). If $z \notin C$ then $C < z$ (otherwise $z \leq c$ and then $z \in C$) hence $x \leq z$.

[6] An element z such as this always exists: since x_1 and y_1 are both in C, they are comparable, so we can take $z = y_1$ in the case that $x_1 \leq y_1$, or $z = x_1$ in the case that $y_1 \leq x_1$.

4.2 Naturality of the Bartha Topology

Recall that in Section 1.2.1 we identified three *naturality criteria* for structural extensions of BST92(*). When the structure in question is the Bartha topology, these criteria manifest as:

N1) Under certain assumptions, the history-relative Bartha topologies are homeomorphic to a Minkowski spacetime.
N2) The Bartha topology possesses a certain universal property, in that it can be canonically reconstructed from its history-relative substructures.
N3) The Bartha topology is compatible with any pre-existing BST92* concepts.

In this section we will argue that conditions **N2** and **N3** are satisfied. We will start with **N3**, since it is the most straightforward, and then we will tackle **N2**.

4.2.1 Compatibility with Pre-Existing Structure

There is not much interaction between the topological and the order-theoretic properties of BST92* models. As Placek et. al observe, the only potential compatibility issue is between the Bartha-topological convergence of sequences, and the order-theoretic concepts of suprema and infima. In this section, we will verify that (where relevant) these two notions of convergence are in fact one and the same. In order to do this, we first need a small lemma.

Lemma 4.12. *Let h be a history of W. Then the sets:*

$$\uparrow x := \{y \in h \mid x \leq y\} \text{ and } \downarrow x := \{y \in h \mid y \leq x\}$$

are closed in the Bartha topology on h.

Proof. Consider first $\uparrow x$. We show that the complement is $h \backslash \uparrow x$ is open. Let $z \in h \backslash \uparrow x$, and let C be some maximal chain in h that passes through z. In order to show that the Bartha condition is satified by z and C, it suffices to show the existence of some c in C such that $z < c$ and $c \in h \backslash \uparrow x$. Consider the subchain

$$C' := C \cap \uparrow x = \{c \in C \mid x \leq c\}.$$

then $C \subseteq h \backslash \uparrow x$, and we can pick any element above z. So, suppose that C' is non-empty. Since x lower-bounds C', thus by axiom **BST3** the infimum $inf(C')$ exists. This infimum must be distinct from z; otherwise $x \leq z$, which contradicts $z \in h \backslash \uparrow x$. Since C is a maximal chain, in particular it is dense, so we can pick an element $c \in C$ such that $z < c < inf(C')$. Whatever this element is, it is strictly below $inf(C')$, so lies in the complement $h \backslash \uparrow x$. Since C is a maximal chain, there exists some element b below z in C. Thus we have the existence of two elements c and b of C such that $b < z < c$ and the diamond d_{bc} is fully contained within $h \backslash \uparrow x$ (see Figure 4.1 for the idea). Since we picked both z and C arbitrarily, it follows that $h \backslash \uparrow x$ is open in the Bartha topology on h, and and thus its complement $\uparrow x$ is closed. The argument for $\downarrow x$ is similar. □

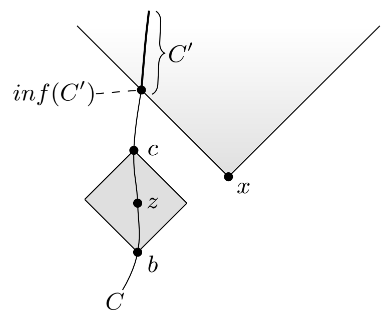

Fig. 4.1: The proof of Lemma 4.12

Using the above result, we can now show that on the level of histories, the topological and the order-theoretic notions of convergence coincide. This result comes in the form of the next two lemmas, where will show the compatibility for suprema and infima separately.

Lemma 4.13. *Let h be a history, and $C \subseteq h$ a countable chain indexed by the natural numbers, so that $c_n \leq c_m$ in h iff $n \leq m$ in $\mathbb{N}$. Then $sup_h(C) = x$ iff C, when viewed as a sequence, converges to x in the history-relative Bartha topology τ_B^h.*

Proof. Suppose first that $sup_h(C) = x$, and let U an open neighbourhood of x. Let D be a maximal extension of the chain C. Since U is open, there are points b_1 and b_2 on D such that $b_1 < x < b_2$ and $d_{b_1 b_2} \subset U$. Since $b_1 < x$ and both elements lie in D, it must be the case that $C \not\leq b_1$, since otherwise this would contradict x as the *least* upper bound of C. Hence there is some $c_N \in C$ such that $b_1 \leq c_N$. Since C is order-isomorphic to $\mathbb{N}$, it follows that $b_1 \leq c_m < x$ for all $m > N$. Since $d_{b_1 b_2} \subset U$, it follows that all such c_m are contained within U, as required. Since U was arbitrary, it follows that C converges to x in τ_B^h.

Conversely, suppose that C converges to x. We first show that $C \leq x$. Suppose towards a contradiction that $C \not\leq x$, that is, there exists some $c_n \in C$ such that $c_n \not\leq x$. Then $x \notin \uparrow c_n$, and thus $x \in h \backslash c_n$. It follows from Lemma 4.12 that $h \backslash \uparrow c_n$ is an open set containing x. Since C converges to x, there exists some $N \in \mathbb{N}$ such that $c_m \in h \backslash \uparrow c_n$ for all $m > N$. Fix such a c_m. Since $c_m \in h \backslash \uparrow c_n$, it follows that $c_n \not\leq c_m$, and thus $c_m \leq c_n$ (since C is a chain). However since C is order-isomorphic to $\mathbb{N}$, it follows that $N < m < n$

and thus $c_n \in h\backslash\uparrow c_n$, a contradiction. We may thus conclude that no such c_n exists, and thus $C \leq x$.

We now show that x is the least upper bound of C. Suppose towards a contradiction that there is some $y \in h$ such that $C \leq y < x$. Since x and y are distinct, there is some $z_1 \in h$ such that $y < z_1 < x$. Since h has no maxima, there is some z_2 in h such that $x < z_2$. Consider now the set $(h\backslash\uparrow z_2) \cap (h\backslash\downarrow z_1)$, which by Lemma 4.12 is an open set. Since $C \leq y < z_1$, that is, $C \subset \downarrow z_1$, it follows that the set $(h\backslash\uparrow z_2) \cap (h\backslash\downarrow z_1)$ is an open neighbourhood of x disjoint from C, which contradicts our assumption that C converges to x. Thus we may conclude that there is no such y, from which it follows that $sup_h(C) = x$. □

We can now show the dual result to the above. Conceptually speaking, the proof is very similar, but here will replace up-arrows with down-arrows.

Lemma 4.14. *Let $C = \{c_n \mid n \in \mathbb{N}\}$ be a countable, lower-bounded chain contained in some history h that is inversely order-isomorphic to $\mathbb{N}$. Then $inf(C) = x$ iff C, when viewed as a sequence, converges to x in the history-relative Bartha topology τ_B^h.*

Proof. Suppose first that $inf(C) = x$, and let U be an open subset of h containing x. Again we can extend C to a maximal chain that passes through x and proceed as in the previous lemma to conclude that C converges to x.

Conversely, suppose that C converges to x. If $x \not\leq C$, then there is some $c_n \in C$ such that $x \not\leq C$. Then $x \in h\backslash\downarrow c_n$, thus there is some $N \in \mathbb{N}$ such that $c_m \in h\backslash\downarrow c_n$ for all $m > N$. Then any such c_m will have to be above c_n, and since C is inversely order-isomorphic to $\mathbb{N}$, it follows that $m < n$ and thus $c_n \in h\backslash\downarrow c_n$, a contradiction. To see that x is the greatest lower bound of C, we can use an argument symmetric to that of the previous lemma. Indeed – if there is some $y \in h$ such that $x < y \leq C$ then again we can pick two elements z_1, z_2 in h such that $z_1 < x < z_2 < y \leq C$, and use that $(h\backslash\uparrow z_2) \cap (h\backslash\downarrow z_1)$ is an open subset containing x that is disjoint from C. □

Now that we have verified that convergence coincides with suprema and infima on the history-relative Bartha topologies, we will do the same for the Bartha topology τ_B^W on W.

Lemma 4.15. *Let $C = \{c_n \mid n \in \mathbb{N}\}$ be a countable, upper-bounded chain contained in some history h that is order-isomorphic to $\mathbb{N}$. Then $sup_h(C) = x$ iff C, when viewed as a sequence, converges to x in the Bartha topology τ_B^W.*

Proof. Suppose first that $sup_h(C)$. By Lemma 4.13 we know that C converges to x in τ_B^h. Since this is a subspace of τ_B^W, C also converges to x in τ_B^W. For the converse, suppose that C converges to x in W, and let h be some history that contains x. By Lemma 4.6, τ_B^h can also be viewed as an open subspace of W. Since $x \in h$, it follows that C also converges to x in τ_S^h. By Lemma 4.6

the subspace topology τ_S^h is equal to τ_B^h, so it follows[7] that C converges to x in h equipped with the history-relative Bartha topology. Thus by Lemma 4.13 it follows that $sup_h(C) = x$, as required. □

It should be clear that a result analogous to Lemma 4.14 also holds. From all of this, we can conclude that on both τ_B^h and τ_B^W, the notion of topological convergence and infima/suprema coincide. The results of this section are then grounds to accept that condition **N3** is satisfied by the Bartha topology on BST92* models.

4.2.2 Construction from History-Relative Topologies

We will now show that the Bartha topology τ_B^W is equivalent to an appropriate adjunction of its history-relative Bartha topologies. In this section we will fix an enumeration $H(W) = \{h_i \mid i \in I\}$, where (in accordance with Remark 1.14), I is assumed to be at most countably-infinite. The choice of such an "appropriate adjunction" should be fairly obvious – consider the tuple $\mathcal{F} = (\mathbf{X}, \mathbf{A}, \mathbf{F})$, where:

- X consists of the topological spaces $(h_i, \tau_B^{h_i})$,
- Each A_{ij} is the intersection $h_i \cap h_j$
- Each $f_{ij} : A_{ij} \to h_j$ is the identity map.

It is not hard to see that the tuple $\mathcal{F}$ defines an adjunction system – the various conditions of Definition 3.7 follow trivially from the properties of the identity map and intersections. Thus we can form the adjunction space $\bigcup_{\mathcal{F}} h_i$, whose elements are equivalence classes of the form

$$[x, i] = \{(y, j) \mid y = x \text{ and } x \in h_i \cap h_j\}.$$

In this context, the canonical maps $\phi_i : h_i \to \bigcup_{\mathcal{F}} X_i$ send each x in h_i to the equivalence class $[x, i]$. We remark that the space $\bigcup_{\mathcal{F}} h_i$ is essentially a copy of W in which elements of W are indexed by the histories that they lie in. It should then come as no surprise that the following holds.

Theorem 4.16. *The space $\bigcup_{\mathcal{F}} h_i$ equipped with the adjunction topology τ_A is homeomorphic to (W, τ_B^W).*

Proof. Consider the inclusion maps $\psi_i : h_i \to W$. These are clearly continuous maps that meet the sufficient conditions of Lemma 3.10. As such, there is a unique continuous map $g : \bigcup_{\mathcal{F}} h_i \to W$. In this context, the map g sends

[7] Let a_n be a sequence that converges to x in X, and let A be an open subspace of X that contains x. Let $U \subset A$ be open in A. Then $U = A \cap V$, where V is open in X. Since A is also open, U is open in X. Thus there is some $N \in \mathbb{N}$ such that $a_n \in U$ for all $n > N$. Since U was arbitrary, it follows that a_n converges to x in A (also, we can effectively view a_n as a sequence lying in A since A is open an contains x so falls into the scope of the convergence property).

elements $[x, i]$ to x in W. We now show that g is a bijective open map. That g is bijective is fairly obvious: if $[x, i] = [y, j]$ then $x = y$, hence $g([x, i]) = x = y = g([y, j])$, and if $x \in W$ then x lies in some history h_i, thus $[x, i] \in \bigcup_{\mathcal{F}} h_i$ will be mapped to x under g.

Suppose now that U is some open subset of $\bigcup_{\mathcal{F}} h_i$. Then for each i, the set $\phi_i^{-1}(U)$ is open in h_i (equipped with the topology τ_B^h). By Corollary 4.5 and Lemma 4.6, the histories h_i are open subspaces of W, thus the inclusion maps ψ_i are open[8]. As such, each $\psi_i(\phi_i^{-1}(U))$ is open in W, and thus so is the set $\bigcup_{i \in I} \phi_i^{-1}(U)$. But this is precisely the set $g(U)$. Indeed: $x \in \phi_i^{-1}(U)$ for some i implies that $[x, i] \in U$, thus $g([x, i]) = x$. Conversely, if $x \in g(U)$ then $[x, i] \in U$ for some i, so $x \in \phi_i^{-1}(U)$. Hence $g(U) = \bigcup_{i \in I} \phi_i^{-1}(U)$ is open in W, and since we chose U arbitrarily, it follows that g is an open map. The proof is complete by observing that all bijective maps that are both continuous and open are homeomorphisms[9]. □

The above theorem suggests that the Bartha topology on any BST92* model can be canonically expressed as an adjunction of its history-relative Bartha topologies. As such, the naturality condition **N2** is satisfied.

4.3 Topological Properties of Minkowskian BSTs

In this section we will discuss the Bartha-topological properties of the Minkowskian BSTs constructed in Section 2.2.2. In particular, we will show that on each $M^n_{\mathcal{C}}$, the Bartha topology is equivalent to an adjunction of the Euclidean topology on M^n (which is equivalent to a topology proposed by Müller in [16]).

Before doing anything of the sort, we first make the following observation. In the previous section we saw that the naturality criteria **N2** and **N3** were satisfied by the Bartha topology. The following result confirms that condition **N1** is satisfied.

Lemma 4.17. *The space $(L_i, \tau_B^{L_i})$ is homeomorphic to M^n equipped with the Bartha topology.*

Proof. We have seen in the form of Prop. 2.13 that the canonical embeddings $\iota_i : M^n \to L_i$ are order-isomorphisms. Thus we can apply Prop. 4.1 to conclude that the ι_i act as homeomorphisms. □

8 This is a general result: if $U \subset A$ is open in an open subspace A of a topological space X, then $U = A \cap V$ where V is open in X. Since A is open, U is the intersection of two sets open in X, hence it is also open in X. Then the inclusion map $\iota : A \to X$ is open.

9 Again, this is a general result. Let $f : X \to Y$ be a bijective, continuous and open map between two topological spaces X and Y. Since f is bijective, the inverse f^{-1} is well-defined. Let $U \subseteq X$ be open. Then: $(f^{-1})^{-1}(U) = f(U)$ is open since f is an open map. Thus f^{-1} is continuous, whence f is a homeomorphism from X to Y.

From this result, we may conclude that the Bartha topology is a natural extension of the order-theoretic structure of BST92* models.

4.3.1 Equivalence of the Bartha and Müller's Topology

In his paper [16], Müller defines a topology on Minkowskian BSTs that is generated by the open balls in each layer L_i. We will now show that such a topology can be equivalently defined as an appropriate topology on an adjunction space.

Müller's Topology Described as an Adjunction Space

Although Müller only defines his topology on simple MBSTs, his definition naturally generalises to the case of arbitrary MBSTs. Consider the system $\mathcal{G} := (\mathbf{Y}, \mathbf{B}, \mathbf{G})$, where:

- each Y_i is equal to M^n equipped with the Euclidean topology,
- each B_{ij} is equal to the set $M^n \backslash J^+(C_{ij})$, and
- each $g_{ij} : B_{ij} \to M^n$ is the inclusion map.

It should be clear that this defines an adjunction system.[10] Observe the (intended) similarity between the above system $\mathcal{G}$ and the construction of $M^n_{\mathcal{C}}$ as in Section 2.2.2. We are essentially performing the same construction, though this time at the topological level. Observe that on the set-theoretic level, the set $\bigcup_{\mathcal{G}} M^n_i$ is equal to the set $M^n_{\mathcal{C}}$, and as functions, the canonical maps $\phi_i : M^n \to \bigcup_{\mathcal{G}} M^n_i$ are equal to the canonical embeddings ι_i defined as in Prop. 2.13.[11]

We will momentarily break our notation and refer to the adjunction topology associated to $\bigcup_{\mathcal{G}} M^n_i$ as the *adjoined Euclidean topology* on $M^n_{\mathcal{C}}$, and we will denote this by τ^E_A. We also remark that by construction, each B_{ij} is an open subset of M^n and each g_{ij} is an open map, so by Lemma 3.14 the canonical maps ϕ_i are open, topological embeddings. As such, we can use Lemma 3.16 to conclude that

$$\mathcal{B} := \{\phi_i(U) \mid U \in \mathcal{B}_{M^n}, i \in I\}$$

is a basis for τ^E_A, where $\mathcal{B}_M$ denotes the basis for M^n consisting of open balls of rational radii centered at rational coordinates.

---

[10] This is the same idea as with the tuple $\mathcal{F}$ defined in Section 4.2.2 – the conditions **A1-3** are satisfied due to basic properties of intersections and inclusion maps.

[11] However, in order to emphasise that the canonical maps are continuous maps between topological spaces, we will keep with the notation of adjunction spaces, and denote them by ϕ_i.

Proving the Equivalence

Before showing that the adjoined Euclidean topology is equivalent to the Bartha topology, we need to make a key observation.

Theorem 4.18. *On M^n, the Bartha topology coincides with the Euclidean topology.*

Proof. We show that the Bartha topology on M^n has a basis comprised of sets of the form $I^+(x) \cap I^-(y)$. To do this, we need to show that every such set is open in the Bartha topology, and that every set U open in the Bartha topology can be represented as a union of sets of this form.

We first show that all subsets of the form $I^+(x) \cap I^-(y)$ are open in the Bartha topology on M^n. Suppose that $z \in I^+(x) \cap I^-(y)$, and let C be a maximal $\leq^M$ chain that passes through z. It follows from Prop 1.10 that there are elements c_1 and c_2 in C such that $x \ll^M c_1 <^M z <^M c_2 \ll^M y$. We can then use the causal diamond $d_{c_1 c_2}$ to witness the Bartha condition at z. Since we chose x, y and z arbitrarily, it follows that all sets of the form $I^+(x) \cap I^-(y)$ are open in the Bartha topology on M^n.

We now show that every subset of M^n that is open in the Bartha topology can be represented as a union of elements of the form $I^+(x) \cap I^-(y)$. So, let U be some subset of M^n open in the Bartha topology. By definition, for each x and each maximal $\leq^M$-chain C, there are $a, b \in C$ such that $d_{ab} = J^+(a) \cap J^-(b) \subset U$, where $a <^M x <^M b$. Let A be the collection of ordered pairs (a, b) such that a and b are witnesses to the Bartha condition for some $x \in U$ and some maximal $\leq^M$-chain C containing x. We will now show that

$$U = \bigcup_{(a,b) \in A} I^+(a) \cap I^-(b).$$

The inclusion from right-to-left is immediate – if $x \in I^+(a) \cap I^-(b)$ for some $(a, b) \in A$, then since $I^+(a) \cap I^-(b) \subset J^+(a) \cap J^-(b) = d_{ab} \subset U$, it follows that $x \in U$. For the converse, suppose that $x \in U$. Pick any timelike curve γ passing through x. By Prop 1.5, γ induces a maximal $\leq^M$-chain C passing through x in which the elements of C are pairwise timelike related (that is, $\ll^M$-related). Since U is open, there are elements a and b of C such that $d_{ab} \subset U$ and $a <^M x <^M b$. Since we chose C to be timelike, it follows that $a \ll^M x \ll^M x$, and thus $x \in I^+(a) \cap I^-(b)$. Hence $x \in \bigcup_{(a,b) \in A} I^+(a) \cap I^-(b)$, from which we may conclude that $U = \bigcup_{(a,b) \in A} I^+(a) \cap I^-(b)$. From all of this we may conclude that the collection

$$\mathcal{B} := \{I^+(x) \cap I^-(y) \mid x, y \in M^n\}$$

forms a basis for the Bartha topology on M^n. However, this is precisely the definition of the Alexandrov topology on M^n. A necessary and sufficient condition for a spacetime to be strongly-causal is that the Alexandrov topology and the manifold topology coincide (see e.g. [23, Sec. 3.6.1.]). It is also well-known that Minkowski spacetime is strongly-causal[12]. From this we can conclude that on M^n, the Bartha topology coincides with the Euclidean topology. $\square$

The above result is the crucial observation that allows us to prove the equivalence of the Bartha and the adjoined Euclidean topologies. The argument is given as follows.

Theorem 4.19. *The Bartha topology on $M^n_{\mathcal{C}}$ is equal to adjoined Euclidean topology.*

Proof. Let U be some subset of $M^n_{\mathcal{C}}$. Suppose first that U is open in the Bartha topology on $M^n_{\mathcal{C}}$. Then by Lemma 4.6, this implies that $U \cap L_i$ is open in the history-relative Bartha topology $\tau_B^{L_i}$. In order to show that U is open in the adjoined Euclidean topology τ_A^E, by Prop. 3.9 it suffices to show that the preimages $\phi_i^{-1}(U)$ are open in the Euclidean topology on M^n. Subject to the remarks made in Section 4.2.2, the canonical maps $\phi_i : M^n \to M^n_{\mathcal{C}}$ are precisely equal to the canonical embeddings ι_i defined as in Prop. 2.13. By Lemma 4.17, the maps ι_i act as homeomorphisms from M^n to L_i, where both are equipped with their respective Bartha topologies. Since $U \cap L_i$ is open in the history-relative Bartha topology $\tau_B^{L_i}$, it follows that the set $\iota_i^{-1}(U)$ is open in $(M^n, \tau_B^{M^n})$. By Theorem 4.18, it follows that $\iota_i^{-1}(U)$ is also open in the Euclidean topology on M^n. Since we chose i to be arbitrary, we can apply Prop. 3.9 to conclude that U is open in the adjoined Euclidean topology τ_A^E, and thus $\tau_B^{M^n_{\mathcal{C}}} \subseteq \tau_A^E$.

The converse is a near-symmetric argument – if U is open in the adjoined Euclidean topology τ_A^E, then by Prop. 3.9 the preimages $\iota_i^{-1}(U)$ are open in the Euclidean topology on M^n. Thus by Theorem 4.18 the preimages $\iota_i^{-1}(U)$ is open in $\tau_B^{M^n}$, hence $\iota_i(\iota_i^{-1}(U)) = U \cap L_i$ is open in the history-relative Bartha topology $\tau_B^{L_i}$. Then $U \in \tau_B^{M^n_{\mathcal{C}}}$ by Corollary 4.7 from which it follows that $\tau_B^{M^n_{\mathcal{C}}} \subseteq \tau_A^E$, whence equality. $\square$

The above result unifies the bottom-up, order-theoretic Bartha topology of Placek et. al with the top-down, pseudo-Euclidean topology of Müller. As a consequence, we may deduce the following.

Theorem 4.20. *Let $M^n_{\mathcal{C}}$ be a Minkowskian BST of the theory* BST92**. Then the Bartha topology on $M^n_{\mathcal{C}}$ is second-countable, connected, and locally-Euclidean. Moreover, $M^n_{\mathcal{C}}$ is Hausdorff iff $M^n_{\mathcal{C}}$ is a single-historied model.*

Proof. When we view $M^n_{\mathcal{C}}$ as the adjoined Euclidean topology, the canonical maps ϕ_i are open, topological embeddings. Thus we can use Corollary 3.14 to conclude that $M^n_{\mathcal{C}}$ is second-countable, Lemma 3.15 to conclude that it is connected, and Lemma 3.19 to conclude that it is locally-Euclidean. If $M^n_{\mathcal{C}}$ has a single history, then by Prop 4.17, $M^n_{\mathcal{C}}$ is homeomorphic to M^n, which is Hausdorff. Conversely, if $M^n_{\mathcal{C}}$ has multiple histories, then by Prop 2.21 these histories possess hot pairs, which by Lemma 4.10 violate the Hausdorff property.

[12] In fact, Minkowski spacetime is globally-hyperbolic, a far-stronger condition.

4.3.2 Topological Characterisations of Hot Pairs

We finish this chapter with an observation regarding hot pairs. We saw in the form of Lemma 4.10 that for any BST92* model W, every hot pair violates the Hausdorff property on its Bartha topology. We will now show that on Minkowskian BSTs the converse also holds.

Lemma 4.21. *Every Hausdorff-violating pair $\{w, z\}$ of $M^n_{\mathcal{C}}$ is a hot pair for $M^n_{\mathcal{C}}$.*

Proof. Suppose $w, z \in M^n_{\mathcal{C}}$ violate the Hausdorff property. Since each history L_i is homeomorphic to M^n (see 4.17,4.18), it cannot be the case that w and z lie in the same layer. So, let $w := [x, i] \in L_i \backslash L_j$ and $z := [y, j] \in L_j \backslash L_i$. Since $[x, i] \in L_i \backslash L_j$, in particular $(x, i) \not\approx (x, j)$, thus there is some $a \in C_{ij}$ such that $a \leq^M x$. Similarly, there is some $b \in C_{ij}$ such that $b \leq^M y$. It follows that both x and y are elements of $J^+(C_{ij})$.

Suppose towards a contradiction that $x \neq y$. Since M^n is Hausdorff, there are two disjoint open subsets U and V of M^n such that $x \in U$ and $y \in V$. We know from Lemma 4.17 that the maps $\iota_i : M^n \to L_i$ are homeomorphisms. In particular, the ι_i are open maps. The sets $\iota_i(U)$ and $\iota_j(U)$ are then open in L_i, hence are also open in $M^n_{\mathcal{C}}$ by Lemma 4.6. Observe that the sets $\iota_i(U)$ and $\iota_j(V)$ are disjoint in $M^n_{\mathcal{C}}$, since by assumption U and V are disjoint in M^n. Thus we have two disjoint open sets separating the elements $[x, i]$ and $[y, j]$, which contradicts our assumption that $[x, i]$ and $[y, j]$ violate the Hausdorff property. We may then conclude that $x = y$.

We now show that $x \notin I^+(C_{ij})$. Suppose towards a contradiction that $x \in I^+(C_{ij})$, and denote by $c_1, ..., c_n$ the elements of C_{ij} such that $x \in \bigcap_{\alpha=1}^n I^+(c_\alpha)$ (note that there are at most finitely-many by our convention that $\mathcal{C}$ is finite). By Prop 1.4, it follows that $x \in \bigcap_{\alpha=1}^n I^+(c_\alpha)$ is open in the Euclidean topology on M^n, and is non-empty since it contains x. Thus there is some neighbourhood U of x such that $U \subseteq \bigcap_{\alpha=1}^n I^+(c_\alpha)$. We can always pick U small enough so that the inclusion is proper. Since ι_i and ι_j are open maps, it follows that $\iota_i(U)$ is open in L_i and $\iota_j(U)$ is open in L_j, hence by Lemma 4.4 these sets are also open in $M^n_{\mathcal{C}}$. However, since U is contained entirely within $J^+(C_{ij})$, it follows that $\iota_i(U) \subset L_i \backslash L_j$ and $\iota_j(U) \subset L_j \backslash L_i$. Thus $\iota_i(U)$ and $\iota_j(U)$ are disjoint open sets of $M^n_{\mathcal{C}}$ that separate $[x, i]$ and $[x, j]$, which contradicts our assumption that this pair violates the Hausdorff property. We may conclude that there is no $c \in C_{ij}$ such that $x \in I^+(c)$, and thus $x \notin I^+(C_{ij})$. It follows that $x \in E^+(C_{ij})$, and the result follows as an application of Lemma 2.21. □

The above result, combined with Lemma 4.10, provides a topological characterisation of hot pairs in $M^n_{\mathcal{C}}$, namely the hot pairs are precisely the pairs of points that violate the Hausdorff property. In combination with Theorem 4.20, the above result then suggests that the violation of the Hausdorff property is deeply connected to the branching (i.e. indeterminism) of Minkowskian BSTs.

Part II

Lorentzian Branching Spacetimes

5

Adjoined Manifolds and Their Vector Bundles

In Chapter 3 we identified some conditions under which a collection of topological spaces could be glued along open subspaces. In this chapter, we will extend this idea to the setting of smooth manifolds. In the first section, we will identify conditions under which a collection of smooth manifolds M_i of the same dimension can be glued. The resulting glued space will then be a (typically non-Hausdorff) smooth manifold in which the canonical maps ϕ_i act as smooth embeddings. We will call such manifolds *adjoined manifolds*, and denote them by $\bigcup_{\mathcal{F}} M_i$ (as opposed to the X_i-notation of Chapter 3).

In Section 6.2, we will discuss the nature of vector bundles over adjoined manifolds. We will start by providing a natural method for adjoining vector bundles, and then we will move on to show that the tangent bundle of an adjoined manifold $\bigcup_{\mathcal{F}} M_i$ is naturally isomorphic to an adjunction of the bundles TM_i. After this, we will show that *any* vector bundle F over an adjoined manifold $\bigcup_{\mathcal{F}} M_i$ is naturally isomorphic to an adjunction of its pullback bundles $\phi_i^* F$.

5.1 Adjoining Smooth Manifolds

We will now discuss the adjunction spaces of smooth manifolds. It is fairly common within differential geometry to glue smooth manifolds together to obtain a new one. A particularly useful construction is that of the *connected sum* of manifolds M_1 and M_2, which is typically denoted by $M_1 \# M_2$. This is constructed by taking taking open balls U_i of M_i, removing them, gluing the $M_i \backslash U_i$ along their diffeomorphic boundary, and then defining an appropriate smooth structure on the resulting topological space.

This is essentially how we will approach the problem, with a few small changes. To begin with, we would like to glue our manifolds along regions that are topologically *open*, not closed as in the standard approach. The reason we do this should be clear by now – we are eventually going to call these smooth manifolds "branches" and form a branching spacetime from them. In accordance with the standard BST92* theory outlined in Chapter 2, and the topological properties of the MBSTs found in Chapter 4, we would like the boundaries of these to-be-glued open regions to be left well alone, so that the resulting space will be non-Hausdorff, and these Hausdorff-violating pairs will be hot pairs.

We will now start to define our "adjoined manifolds". Suppose that we have an adjunction system $\mathcal{F}$, where in addition to the conditions outlined in Definition 3.7, we have that the indexing set I is countable, and

- $\mathbf{X}$ consists of smooth n-manifolds M_i,
- each A_{ij} is an open submanifold of M_i, and
- each $f_{ij} : A_{ij} \to M_j$ is a smooth embedding.

Observe that the last item above ensures that the map f_{ij} acts as a diffeomorphism from A_{ij} to A_{ji} (see Lemma 3.11). Put differently, we will be gluing smooth manifolds along diffeomorphic open submanifolds. By definition smooth embeddings are also topological embeddings, so we can apply Lemma 3.11 to conclude that the canonical maps ϕ_i are open topological embeddings into the adjunction space $\bigcup_{\mathcal{F}} M_i$. As such, we can use Lemmas 3.19 and 3.14 to conclude that $\bigcup_{\mathcal{F}} M_i$ is locally-Euclidean and second-countable, that is, the adjunction space $\bigcup_{\mathcal{F}} M_i$ is at least a topological manifold. We will now set about defining a smooth structure.

5.1.1 Defining a Smooth Structure

In this section we will show that the smooth structure of each M_i can be transferred to the adjunction space $\bigcup_{\mathcal{F}} M_i$ in a natural way, so that the canonical maps ϕ_i become smooth. Our first result describes the smooth structure of $\bigcup_{\mathcal{F}} M_i$.

Lemma 5.1. *Let $\mathcal{A}_i$ be the smooth structure of M_i. Then the collection*

$$\mathcal{A} := \{(\phi_i(U), \varphi \circ \phi_i^{-1}) \mid (U, \varphi) \in \mathcal{A}_i\}$$

induces a smooth structure on $\bigcup_{\mathcal{F}} M_i$.

Proof. We begin by showing that $\mathcal{A}$ consists of pairwise-compatible elements. So, let $(\phi_i(U), \varphi \circ \phi_i^{-1})$ and $(\phi_j(V), \psi \circ \phi_j^{-1})$ be elements of $\mathcal{A}$, whose overlap $\phi_i(U) \cap \phi_j(V)$ is non-empty. In order to show that these charts are compatible, we need to compute their transition functions and show that these are smooth in the Euclidean sense. Before doing so, we observe that by assumption, any element in the overlap $\phi_i(U) \cap \phi_j(V)$ will be of the form $[p, i] = [f_{ij}(p), j]$, where f_{ij} is the diffeomorphism mapping A_{ij} to A_{ji}, which are both endowed with the open-submanifold smooth structure induced from their respective parent spaces. By assumption, the tuples (U, φ) and (V, ψ) are members of the atlases $\mathcal{A}_i$ and $\mathcal{A}_j$ respectively. Since these are open subsets of their respective spaces, it follows that $(U \cap A_{ij}, \varphi|_{A_{ij}})$ and $(V \cap A_{ji}, \psi|_{A_{ji}})$ are charts for A_{ij} and A_{ji} respectively. The transition functions for the charts $(\phi_i(U), \varphi \circ \phi_i^{-1})$ and $(\phi_j(V), \psi \circ \phi_j^{-1})$ can be computed as follows

$$(\varphi \circ \phi_i^{-1}) \circ (\psi \circ \phi_j^{-1})^{-1} = \varphi \circ \phi_i^{-1} \circ \phi_j \circ \psi^{-1} = \varphi \circ f_{ji} \circ \psi^{-1}$$

$$(\psi \circ \phi_j^{-1}) \circ (\varphi \circ \phi_i^{-1})^{-1} = \psi \circ \phi_j^{-1} \circ \phi_i \circ \varphi^{-1} = \psi \circ f_{ij} \circ \varphi^{-1}$$

where the last equalities follow from the definition of an adjunction space. The transition maps above are essentially a local representation of the maps f_{ij} and f_{ji}, restricted to an open subset. Since we have assumed that these are diffeomorphisms (i.e. smooth), it follows that the transition maps are smooth and thus the charts $(\phi_i(U), \varphi \circ \phi_i^{-1})$ and $(\phi_j(V), \psi \circ \phi_j^{-1})$ are compatible. We can then apply Prop. A.3 and extend the atlas $\mathcal{A}$ to a smooth structure. $\square$

From now on, we will assume that the adjunction space $\bigcup_{\mathcal{F}} M_i$ is equipped with the smooth structure $\mathcal{A}$ of Lemma 5.1. Our next result says that this smooth structure $\mathcal{A}$ turns the canonical maps ϕ_i into smooth embeddings.

Lemma 5.2. *Each ϕ_i is a smooth embedding into $\bigcup_{\mathcal{F}} M_i$.*

Proof. As previously remarked, our choice of $\mathcal{F}$ and Lemma 3.11 imply that each ϕ_i is already a topological embedding. Thus it suffices to show that both $\phi_i : M_i \to \phi_i(M_i)$ and its left-inverse $\phi_i^{-1} : \phi_i(M_i) \to M_i$ are smooth, since the result will then follow from an application of Lemma A.12.

We start with ϕ_i. Let $p \in M_i$ and consider $\phi_i(p) = [p, i]$ in $\bigcup_{\mathcal{F}} M_i$. Let (U, φ) be any chart for M_i at p. In order to show that ϕ_i is smooth at p, it suffices to show that ϕ_i is smooth in some local representation. We will use the chart $(\phi_i(U), \varphi \circ \phi_i^{-1})$. Since ϕ_i is an open map, the set $\phi_i(M_i)$ forms an open submanifold of $\bigcup_{\mathcal{F}} M_i$. As such, the chart $(\phi_i(U), \varphi \circ \phi_i^{-1})$ is also a chart for $\phi_i(M_i)$. It should be clear[1] that the local representation of ϕ_i is equal to the identity map restricted to $\varphi(U)$, which is clearly smooth. Since we picked the point p arbitrarily, we may use Prop. A.7 to conclude that ϕ_i is indeed smooth.

The case for the inverse map $\phi_i^{-1} : \phi_i(M_i) \to M_i$ is identical in spirit. Suppose we have a point $[p, i]$ in ϕ_i^{-1}. Again we can pick the chart (U, φ) of M at p and $(\phi_i(U), \varphi \circ \phi_i^{-1})$ of $\phi_i(M_i)$ at $[p, i]$. Again, the local representation of ϕ_i^{-1} will be the identity map on $\mathbb{R}^n$, from which it follows that ϕ_i^{-1} is smooth. $\square$

A nice observation from the above argument is that at every point p in M_i, the local expression of ϕ_i can always be made to coincide with the identity map of $\mathbb{R}^n$. This shouldn't be that surprising – when we perform an adjunction, the only possible place that we might deform any of the M_i is on the

[1] Indeed: $(\varphi \circ \phi_i^{-1}) \circ \phi_i \circ \varphi^{-1} = \varphi \circ \varphi^{-1}$.

to-be-glued regions A_{ij}. However, we are assuming that these are *diffeomorphically* mapped into the M_j, and as such our adjunctions of manifolds are constructed by mere rearrangement. There will be an analogous result when we discuss vector bundles in the next section.

The following theorem summarises the previous two results, and will serve as a useful point of reference.

Theorem 5.3. *Let $\mathcal{F} = (\boldsymbol{X}, \boldsymbol{A}, \boldsymbol{f})$ be an adjunction space with a countable indexing set I, in which:*

1. *$\boldsymbol{X}$ consists of smooth (Hausdorff) manifolds M_i,*
2. *each A_{ij} is open submanifold of M_i, and*
3. *each $f_{ij} : A_{ij} \to X_j$ is a smooth embedding.*

Then the adjunction space $\bigcup_{\mathcal{F}} M_i$ is a smooth manifold in which the canonical maps $\phi_i : M_i \to \bigcup_{\mathcal{F}} M_i$ are open, smooth embeddings.

We will refer to any adjunction space $\bigcup_{\mathcal{F}} M_i$ satisfying the conditions of Theorem 5.3 as an *adjoined manifold*.

5.1.2 Smooth Maps on Adjoined Manifolds

We will now discuss the smooth maps on such spaces. Our first result says that we can push certain smooth maps out to the adjoined manifold $\bigcup_{\mathcal{F}} M_i$.[2]

Lemma 5.4. *Suppose that $\psi_i : M_i \to N$ are smooth maps such that $\psi_i = \psi_j \circ f_{ij}$. Then the map $\tilde{f} : \bigcup_{\mathcal{F}} M_i \to N$ defined by $\tilde{f}([p, i]) = \psi_i(p)$ is a well-defined smooth map.*

Proof. By Prop. A.7.1 the smooth maps ψ_i are continuous, so it follows from our assumption and the universal property of topological adjunction spaces (as in Lemma 3.10) that the map $\tilde{f}$ is a well-defined, continuous map from $\bigcup_{\mathcal{F}} M_i$ to N. To see that $\tilde{f}$ is smooth, let $[p, i]$ be an element in $\bigcup_{\mathcal{F}} M_i$ and consider the restriction of $\tilde{f}$ to the open set $\phi_i(M_i)$. Then $\tilde{f}|_{\phi_i(M_i)} : \phi_i(M_i) \to N$ is equal to $\psi_i \circ \phi_i^{-1}$. By Theorem 5.3 ϕ_i is a smooth embedding, thus Lemma A.12 implies that ϕ_i is a diffeomorphism onto its image. In particular ϕ_i^{-1} is smooth, hence the restriction $\tilde{f}|_{\phi_i(M_i)} = \psi_i \circ \phi_i^{-1}$ is a composition of smooth maps, thus is smooth. Since we chose $[p, i]$ arbitrarily, we can apply Prop. A.7.2 to conclude that $\tilde{f}$ is smooth. $\square$

An immediate application of the above lemma is that certain real-valued functions can be transferred to $\bigcup_{\mathcal{F}} M_i$.

[2] In the statement of Lemma 5.4, we have use a tilde to denote the pushed-out map $\tilde{f}$. From now on we will keep this convention and denote with a tilde any objects that are pushed out to adjunction spaces.

Corollary 5.5. *Let $\{f_i \mid f_i \in C^\infty(M_i)\}$ be a collection of real-valued functions such that $f_i(p) = f_j(f_{ij}(p))$ for all $i, j \in I$. Then the function $\tilde{f} : \bigcup_{\mathcal{F}} M_i \to \mathbb{R}$ defined by $[p, i] \mapsto f_i(p)$ is a member of $C^\infty(\bigcup_{\mathcal{F}} M_i)$.*

5.2 Vector Bundles Over Adjoined Manifolds

In this section we will discuss the nature of vector bundles over an adjoined manifold $\bigcup_{\mathcal{F}} M_i$. We will start by showing that given vector bundles E_i over the M_i, we can glue these together to form an adjoined bundle $\bigcup_{\mathcal{G}} E_i$ over $\bigcup_{\mathcal{F}} M_i$. Our approach is in some sense the obvious one – we will require that the E_i are pairwise isomorphic on their restricted bundles $E_i|_{A_{ij}}$, and we will glue the E_i along these subbundles, making sure to keep the vector-space structure of the fibres intact. Figure 5.1 depicts this intuition behind an adjunction of vector bundles. Once we have verified that our construction works, we will discuss the nature of the tangent bundle $T(\bigcup_{\mathcal{F}} M_i)$. We will show that this is bundle-isomorphic to an adjunction of the bundles TM_i. Finally, we show that by starting with a vector bundle F over $\bigcup_{\mathcal{F}} M_i$, it is possible to reconstruct F from an adjunction of its pullback bundles $\phi_i^* F$. Throughout this section we will assume that $\bigcup_{\mathcal{F}} M_i$ is an adjoined manifold as in Theorem 5.3.

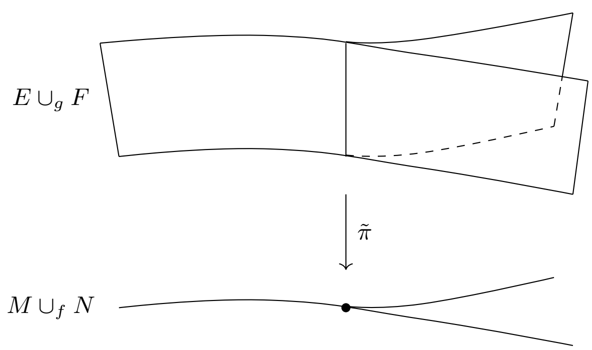

Fig. 5.1: The intuition behind an adjunction of vector bundles $E \to M$ and $F \to N$.

5.2.1 Adjoining Vector Bundles

Suppose that we have a collection of vector bundles $\{(E_i, \pi_i, M_i)\}_{i \in I}$. We saw in Section 6.1 that by requiring the regions A_{ij} to be open submanifolds, and the gluing maps f_{ij} to preserve the smooth structure of A_{ij}, the adjunction space $\bigcup_{\mathcal{F}} M_i$ became a smooth manifold in which the canonical maps ϕ_i act as open, smooth embeddings. We can repeat this "gluing maps preserve structure" mantra once again, and consider an adjunction system $\mathcal{G} = (\mathbf{E}, \mathbf{B}, \mathbf{g})$, where:

- $\mathbf{E}$ consists of the rank-k bundles E_i,
- each $B_{ij} \subseteq E_i$ is equal to the restriction bundle $E_i|_{A_{ij}}$, i.e. $B_{ij} = \pi_i^{-1}(A_{ij})$, and
- each $g_{ij} : B_{ij} \to E_j$ is an injective bundle morphism covering f_{ij}, with every g_{ii} equal to the identity map.

We will denote the adjunction space subordinate to $\mathcal{G}$ by $\bigcup_{\mathcal{G}} E_i$, and we will denote the canonical embeddings by $\chi_i : E_i \to \bigcup_{\mathcal{G}} E_i$. We now confirm that the objects above ensure that $\mathcal{G}$ is well-defined.

Lemma 5.6. *The adjunction space $\bigcup_{\mathcal{G}} E_i$ is an adjoined manifold.*

Proof. We have assumed that $\mathcal{G}$ is an adjunction system to begin with,[3] so it suffices to prove that conditions of Theorem 5.3 are satisfied. Clearly the first condition is met, since all vector bundles are defined to be smooth manifolds. The second condition follows from Lemma A.23 (and our assumption that the A_{ij} are open submanifolds of M_i). For the third condition, we can observe that g_{ij} restricted to its image, namely $g_{ij}| : B_{ij} \to B_{ji}$ is a *bijective* bundle morphism covering the diffeomorphism $f_{ij} : A_{ij} \to A_{ji}$. Thus by Lemma A.25, $g_{ij} : B_{ij} \to B_{ji}$ is a bundle-isomorphism. In particular, it is a diffeomorphism. Since B_{ji} is an open submanifold of E_j, and all open submanifold are embedded submanifolds (see e.g. [20, Prop. 5.2]), it follows from Lemma A.12 that $g_{ij} : B_{ij} \to E_j$ is a smooth embedding of manifolds, as required. $\square$

Now that we have verified that the adjunction space $\bigcup_{\mathcal{G}} E_i$ is a smooth manifold, we can begin to describe its bundle structure. As a rough idea, we would like to do this in such a way that the canonical maps $\chi_i : E_i \to \bigcup_{\mathcal{G}} E_i$ become smooth bundle morphisms covering the ϕ_i. Our first step to achieving this is to define an appropriate projection map $\tilde{\pi} : \bigcup_{\mathcal{G}} E_i \to \bigcup_{\mathcal{F}} M_i$. The natural way to do this is to define $\tilde{\pi}$ in such a way that it behaves like the π_i on their respective bundles, that is, we would like to define $\tilde{\pi}$ in such a way that the diagram

$$\begin{tikzcd} E_i \arrow[r, "\chi_i"] \arrow[d, "\pi_i"'] & \bigcup_{\mathcal{G}} E_i \arrow[d, dashed, "\tilde{\pi}"] \\ M_i \arrow[r, "\phi_i"'] & \bigcup_{\mathcal{F}} M_i \end{tikzcd}$$

commutes for each i. Since our manifolds M_i and our bundles E_i are glued along pairwise isomorphic substructures, we also need $\tilde{\pi}$ to be consistent on the overlaps, so that the following diagram commutes for each i and j.

[3] For the reader unsatisfied by our choice to surmise that such adjunction systems exist – don't worry, we will see plenty of examples in the next two sections.

$$\begin{tikzcd}
 & E_i|_{A_{ij}} \arrow[rr] \arrow[dl, "g_{ij}"'] \arrow[dd] & & E_i \arrow[dl, "\chi_i"'] \arrow[dd, "\pi_i"] \\
E_j \arrow[rr, "\chi_j"] \arrow[dd, "\pi_j"'] & & \bigcup_{\mathcal{G}} E_i \arrow[dd, dashed, "\tilde{\pi}"] & \\
 & A_{ij} \arrow[rr] \arrow[dl, "f_{ij}"] & & M_i \arrow[dl, "\phi_i"] \\
M_j \arrow[rr, "\phi_j"'] & & \bigcup_{\mathcal{F}} M_i &
\end{tikzcd}$$

Now that we have fixed our intuitions, we can set about defining the projection map $\tilde{\pi}$. Recall that since $\bigcup_{\mathcal{G}} E_i$ is an adjunction space, its elements are equivalence classes of the form

$$[u, i] = \{(v, j) \mid g_{ij}(u) = v\}.$$

As such, a natural definition for the projection map $\tilde{\pi}$ would be to map each $[u, i]$ to the element $[\pi_i(u), i]$ in $\bigcup_{\mathcal{F}} M_i$. The following lemma confirms that this map is a well-defined projection map.

Lemma 5.7. *The map $\tilde{\pi}$ defined by $\tilde{\pi}([u, i]) = [\pi_i(u), i]$ is a smooth surjection from $\bigcup_{\mathcal{G}} E_i$ to $\bigcup_{\mathcal{F}} M_i$.*

Proof. To see that the map $\tilde{\pi}$ is well-defined as a function, suppose that there are two elements $u \in E_i$ and $v \in E_j$ such that $[u, i] = [v, j]$, i.e. such that $g_{ij}(u) = v$. Since we required that g_{ij} covers f_{ij}, the following diagram commutes

$$\begin{tikzcd}
B_{ij} \arrow[r, "g_{ij}"] \arrow[d, "\pi_i"'] & B_{ji} \arrow[d, "\pi_j"] \\
A_{ij} \arrow[r, "f_{ij}"'] & A_{ji}
\end{tikzcd}$$

that is, $f_{ij}(\pi_i(u)) = \pi_j(v)$. Hence $[\pi_i(u), i] = [\pi_j(v), j]$, from which it follows that $\tilde{\pi}([u, i]) = [\pi_i(u), i] = [\pi_j(v), j] = \tilde{\pi}([v, j])$, and thus $\tilde{\pi}$ is a function. To see that $\tilde{\pi}$ is a surjection, let $[p, i]$ be an element of $\bigcup_{\mathcal{F}} M_i$. We can pick any element u in the fibre E_p, since then $\tilde{\pi}([u, i]) = [\pi_i(u), i] = [p, i]$ as required.

We now show that $\tilde{\pi}$ is smooth. Observe first that the restriction of $\tilde{\pi}$ to the subspace $\chi_i(E_i)$ is equal to $\tilde{\pi}|_{\chi_i} = \phi_i \circ \pi_i \circ \chi_i^{-1}$. Since this is a composition of smooth maps, it follows from Prop. A.7.2 that $\tilde{\pi}$ is smooth. □

Now that we have confirmed that $\tilde{\pi}$ is a projection map, we can discuss the fibres of $\bigcup_{\mathcal{G}} E_i$, i.e. the preimages $\tilde{\pi}^{-1}([p, i])$. The following result confirms that these fibres have a vector space structure.

Lemma 5.8. *The fibres $\tilde{\pi}^{-1}([p,i])$ of $\bigcup_{\mathcal{G}} E_i$ are real-valued vector spaces of dimension k.*

Proof. The idea is fairly simple – for every p in M_i and every $i \in I$, we are just going to transfer the structure from the fibres $(E_i)_p$ into $\tilde{\pi}^{-1}([p,i])$. In order to do this, we need to show that the canonical map χ_i is a fibrewise bijection, since then we can use a standard argument to induce a vector-space structure on $\tilde{\pi}^{-1}([p,i])$.

We begin by showing that $\tilde{\pi}^{-1}([p,i])$ is equal to the image of $(E_i)_p$ under the canonical map χ_i, that is, we show

$$\tilde{\pi}^{-1}([p,i]) = \chi_i((E_i)_p) := \{[u,i] \mid u \in (E_i)_p\}.$$

Suppose that there is some element u in the fibre $(E_i)_p$. Then $\tilde{\pi}([u,i]) = [\pi_i(u),i] = [p,i]$, so $\chi_i((E_i)_p)$ is a subset of $\tilde{\pi}^{-1}([p,i])$. For the other direction, suppose that we have some $[v,j] \in \tilde{\pi}^{-1}([p,i])$. Then $\tilde{\pi}([v,j]) = [\pi_j(v),j] = [p,i]$, so $f_{ij}(p) = \pi_j(v)$. Since $p \in A_{ij}$, it follows that $v \in B_{ji}$ and thus there is some $g_{ji}(v) \in B_{ij}$, hence $[g_{ji}(v),i] = [v,j]$. By assumption g_{ji} is a bundle map, so $g_{ji}(v) \in (E_i)_{f_{ji}(\pi_j(v))} = (E_i)_{f_{ji}\circ(f_{ij}(p))} = (E_i)_p$. Thus $[v,j] = [g_{ji}(v),i] \in \chi_i \circ (E_i)_p$ as required.

Now that we have confirmed that $\chi_i(\pi_i^{-1}(p)) = \tilde{\pi} \circ \phi_i(p)$, we can use the fact that χ_i is a fibrewise bijection to transfer the vector space structure of $(E_i)_p$ into the fibre $\tilde{\pi}^{-1}([p,i])$. Suppose that we have two elements $[u,j]$ and $[v,k]$ of $\tilde{\pi}^{-1}([p,i])$. By the above argument, it must be the case that $[u,j] = [g_{ji}(u),i]$ and $[v,k] = [g_{ki}(v),i]$. We can then define the operations of addition and scalar multiplication on $\tilde{\pi}^{-1}([p,i])$ by:

$$[u,j] + [v,k] = [g_{ji}(u),i] + [g_{ki}(v),i] := [g_{ji}(u) + g_{ki}(v),i], \text{ and}$$

$$r \cdot [u,j] = r \cdot [g_{ji}(u),i] := [r \cdot g_{ji}(u),i].$$

where we have used the $g_{..}$ maps to shift the representative of the equivalence classes $[u,j]$ and $[v,k]$ to the element coming from $(E_i)_p$.[4] It should also be noted that the maps $g_{..}$ ensure that the operations are well-defined, and stay within the fibre $\tilde{\pi}^{-1}([p,i])$. Indeed, if for example $[u,j] = [w,l]$, i.e. $g_{lj}(w) = u$, then

[4] This is not *that* necessary to do, since there are other representations, for instance

$$\begin{aligned}
[u,j] + [v,k] &= [g_{ji}(u) + g_{ki}(v),i] \\
&= [g_{ij}(g_{ji}(u) + g_{ki}(v)),j] \\
&= [g_{ij} \circ g_{ji}(u) + g_{ij} \circ g_{ki}(v),j] \\
&= [u + g_{kj}(v),j]
\end{aligned}$$

and similarly $[u,j] + [v,k] = [g_{kj}(u) + v,k]$ and $r \cdot [u,j] = [r \cdot u,j]$.

$$\begin{aligned}[w,l]+[v,k] &= [w+g_{kl}(v),l] \\ &= [g_{lj}(w+g_{kl}(v)),j] \\ &= [g_{lj}(w)+g_{lj}\circ g_{kl}(v),j] \\ &= [u,g_{kj}(v),j]\end{aligned}$$

and similarly:

$$r\cdot[w,l]=[r\cdot g_{li}(w),i]=[r\cdot g_{ji}\circ g_{lj}(w),i]=[r\cdot g_{ji}(u),i]=r\cdot[u,j].$$

This argument can be reproduced for whenever $[v,k]$ contains more than two members. Thus these operations of addition and scalar multiplication are well-defined as functions. It should also be clear that the vector space axioms are satisfied, though this is cumbersome to prove, so we will omit this part of the proof. □

The previous two results confirm that the adjoined manifold $\bigcup_{\mathcal{G}} E_i$ satisfies the first few conditions of Def.A.13. The following result confirms that the space $\bigcup_{\mathcal{F}} E_i$ is locally-trivialisable, from which we may .

Lemma 5.9. *The space $\bigcup_{\mathcal{G}} E_i$ is locally-trivialisable.*

Proof. The proof of this lemma is similar in spirit to Lemma 5.1 in that we will show that local trivialisations of the E_i can be tranferred into $\bigcup_{\mathcal{G}} E_i$. Let $[p,i]$ be an arbitrary member of $\bigcup_{\mathcal{F}} M_i$. Since $p\in M_i$ and E_i is a vector bundle over M_i, we can pick a local trivialisation (U,Φ) of E_i at p.

We now show that $\phi_i(U)$ induces a local trivialisation of $\bigcup_{\mathcal{G}} E_i$. Observe first that since ϕ_i is a diffeomorphism from M_i to $\phi_i(M)$, it follows from A.7.3 that the restriction of ϕ_i to U is a diffeomorphism onto its image, i.e. U and $\phi_i(U)$ are diffeomorphic. As such, we can define a map $\zeta: U\times\mathbb{R}^k\to\phi_i(U)\times\mathbb{R}^k$ by $\zeta(p,r)=(\phi_i(u),r)$, that is, componentwise, ζ is equal to (ϕ_i,id). Since the components of this map are smooth, we can use Prop. A.8.4 to conclude that ζ is smooth. Moreover, ζ has an inverse map given by $\zeta^{-1}=(\phi_i^{-1},id)$, and again this is manifestly smooth. Thus $\phi_i(U)\times\mathbb{R}^k$ is diffeomorphic to $U\times\mathbb{R}^k$. We will now define $\Psi:\tilde{\pi}^{-1}(\phi_i(U))\to\phi_i(U)\times\mathbb{R}^k$ in such a way that the following diagram commutes.

$$\begin{tikzcd}
U\times\mathbb{R}^k \arrow[d, "p_1"'] \arrow[rrr, bend left, "\zeta"] & \pi_i^{-1}(U) \arrow[l, "\Phi"'] \arrow[r, "\chi_i"] \arrow[d, "\pi_i"'] & \tilde{\pi}^{-1}(\phi_i(U)) \arrow[r, dashed, "\Psi"] \arrow[d, "\tilde{\pi}"] & \phi_i(U)\times\mathbb{R}^k \arrow[d, "p_1"] \\
U & U \arrow[l, "id"] \arrow[r, "\phi_i"'] & \phi_i(U) \arrow[r, "id"'] & \phi_i(U)
\end{tikzcd}$$

We define $\Psi = \zeta \circ \Phi \circ \chi_i^{-1}$. This is clearly well-defined, and is a composition of diffeomorphisms, so is also a diffeomorphism. It follows that the tuple $(\phi_i(U), \Psi)$ acts as a local trivialisation for $\bigcup_{\mathcal{G}} E_i$ at the point $[p, i]$. Since we chose $[p, i]$ arbitrarily, it follows that $\bigcup_{\mathcal{G}} E_i$ is locally-trivialisable. □

The following theorem is a summary of the results covered thus far.

Theorem 5.10. *Let (E_i, π_i, M_i) be a collection of vector bundles, $\mathcal{F}$ an adjunction system as in Theorem5.3, and $\mathcal{G} = (\boldsymbol{E}, \boldsymbol{B}, \boldsymbol{g})$ an adjunction system where additionally:*

1. *$\boldsymbol{E}$ consists of the bundles E_i,*
2. *each B_{ij} is equal to the restriction bundle $E_i|_{A_{ij}}$, and*
3. *each $g_{ij} : B_{ij} \to B_{ji}$ is a bundle isomorphism covering f_{ij}.*

Then the space $\bigcup_{\mathcal{G}} E_i$ forms a smooth generalised vector bundle over the space $\bigcup_{\mathcal{F}} M_i$ in which the canonical maps $\chi_i : E_i \to \bigcup_{\mathcal{G}} E_i$ are bundle morphisms and open, smooth embeddings.

Proof. We know from Lemmas 5.6, 5.7, 5.8 and 5.9 that $\bigcup_{\mathcal{G}} E_i$ is a vector bundle over $\bigcup_{\mathcal{F}} M_i$. It also follows from our construction that χ_i is a bundle map covering ϕ_i – the remarks above Lemma 5.7 confirms that χ_i covers ϕ_i, Lemma 5.6 confirms that χ_i is smooth, and the construction of the fibres as in the proof of Lemma 5.8 confirms that χ_i is a fibrewise linear map. □

We will now show that bundles such as $\bigcup_{\mathcal{G}} E_i$ possess a certain universal property.

Lemma 5.11. *Let E_i be vector bundles over M_i and F be a vector bundle over $\bigcup_{\mathcal{F}} M_i$. If $\psi_i : E_i \to F$ is a collection of vector-bundle morphisms covering ϕ_i, such that $\psi_i(u) = \psi_j(g_{ij}(u))$ for every $i, j \in I$, then there is a unique bundle morphism $\xi : \bigcup_{\mathcal{G}} E_i \to F$.*

Proof. We define the map $\xi : \bigcup_{\mathcal{G}} E_i \to F$ by $[u, i] \mapsto \psi_i(u)$. This is well-defined, since $[u, i] = [v, j]$ implies that $g_{ij}(u) = v$ and thus $\xi([u, i]) = \psi_i(u) = \psi_j(g_{ij}(u)) = \psi_j(v) = \xi([v, j])$. We now show that ξ is smooth. Consider some element $[u, i]$ of $\bigcup_{\mathcal{G}} E_i$, and its image $\psi_i(u)$. Since $\bigcup_{\mathcal{G}} E_i$ has its smooth structure induced from the smooth structures of the E_i, without loss of generality we can pick a chart $(\chi_i(U), \varphi \circ \chi_i^{-1})$ of $\bigcup_{\mathcal{G}} E_i$ at $[u, i]$, where (U, φ) forms a chart for E_i at u. We then let (V, ψ) be any chart of F at $\psi_i(u)$. The local representation of ξ is then:

$$\psi \circ \xi \circ (\varphi \circ \chi_i^{-1})^{-1} = \psi \circ \xi \circ \chi_i \circ \varphi^{-1} = \psi \circ (\psi_i \circ \chi_i^{-1}) \circ \chi_i \circ \varphi^{-1} = \psi \circ \psi_i \circ \varphi^{-1}$$

which is smooth, since we assumed ψ_i is smooth. We now show that ξ is fibrewise linear. Consider the fibrewise restriction $\xi| : \left(\bigcup_{\mathcal{G}} E_i\right)_{[u,i]} \to F_{\psi_i(u)}$. Recall that $\xi = \psi_i \circ \chi_i^{-1}$, so it follows from our assumption and Theorem 5.10 that the restriction $\xi|$ is the composition of linear maps, so is also linear.

Our proof is complete by showing that ξ covers the identity map on $\bigcup_{\mathcal{F}} M_i$. Fortunately, this follows from our assumption that ψ_i covers ϕ_i. Indeed:

$$\begin{aligned}
\pi_F \circ \xi([u,i]) &= \pi_F(\psi_i(u)) && \text{since } \xi = \psi_i \circ \chi_i^{-1} \\
&= \phi_i \circ \pi_i(u) && \text{since } \psi_i \text{ covers } \phi_i \\
&= \tilde{\pi} \circ \chi_i(u) && \text{since } \chi_i \text{ covers } \phi_i \\
&= \tilde{\pi}([u,i]).
\end{aligned}$$

We have shown that ξ is a smooth, fibrewise-linear map from $\bigcup_{\mathcal{G}} E_i$ to $\bigcup_{\mathcal{F}} M_i$ that commutes with the identity map on $\bigcup_{\mathcal{F}} M_i$, i.e. ξ is a bundle morphism as required. We note that uniqueness follows from an argument that is near-identical to that found in 6.7. $\square$

A corollary of Theorem 5.10 is that the restriction of the bundle $\bigcup_{\mathcal{G}} E_i$ to the subspace $\phi_i(M_i)$ is naturally bundle-isomorphic to E_i. Similarly, we also make the following observation.

Lemma 5.12. *E_i is bundle-isomorphic to the pullback bundle $\phi_i^*\left(\bigcup_{\mathcal{G}} E_i\right)$ for all $i \in I$.*

Proof. We know from Section 5.2.4 that the bundle morphism $\chi_i : E_i \to \bigcup_{\mathcal{G}} E_i$ has to factor through the pullback bundle $\phi_i^*\left(\bigcup_{\mathcal{G}} E_i\right)$, that is, there is some bundle morphism $\xi : E_i \to \phi_i^*\left(\bigcup_{\mathcal{G}} E_i\right)$ that makes the following diagram commute.

$$\begin{tikzcd}
E_i \arrow[r, dashed, "\xi"] \arrow[rr, bend left, "\chi_i"] \arrow[d, "\pi_i"'] & \phi_i^*\left(\bigcup_{\mathcal{G}} E_i\right) \arrow[r, "p_2"] \arrow[d, "p_1"'] & \bigcup_{\mathcal{G}} E_i \arrow[d, "\tilde{\pi}"] \\
M_i \arrow[r, equal] & M_i \arrow[r, "\phi_i"'] & \bigcup_{\mathcal{F}} M_i
\end{tikzcd}$$

Recall that ξ sends each u in E_i to $(\pi_i(u), \chi_i(u))$. The map ξ is a smooth bundle morphism with inverse given by $\xi^{-1} = \chi_i^{-1} \circ p_2$. Thus we can apply Lemma A.17 to conclude that ξ is a bundle isomorphism. $\square$

We will finish this section with a useful lemma regarding the sections of an adjoined bundle.

Lemma 5.13. *Let $\{s_i \mid s_i \in \Gamma(E_i)\}$ be a collection of sections such that the diagram*

$$\begin{tikzcd}
B_{ij} \arrow[r, "g_{ij}"] & B_{ji} \\
A_{ij} \arrow[u, "s_i"] \arrow[r, "f_{ij}"'] & A_{ji} \arrow[u, "s_j"']
\end{tikzcd}$$

commutes for each i and j. Then the map $\tilde{s} : \bigcup_{\mathcal{F}} M_i \to \bigcup_{\mathcal{G}} E_i$ defined by $\tilde{s}([p,i]) = [s_i(p), i]$ is a section of $\bigcup_{\mathcal{G}} E_i$.

Proof. We first show that $\tilde{s}$ is well-defined as a function. Suppose that $[p,i]$ is an element of $\bigcup_{\mathcal{G}} E_i$, and consider $[f_{ij}(p), j]$. Then

$$\tilde{s}([f_{ij}(p), j]) = [s_j \circ f_{ij}(p), j] = [g_{ij} \circ s_i(p), j] = [s_i(p), i] = \tilde{s}([p,i])$$

as required. We now show that $\tilde{s}$ is smooth. To do this, we can use Prop.A.7.2, and show that every point has an open neighbourhood on which $\tilde{s}$ is smooth. This is fairly simple: for $[p,i]$ in $\bigcup_{\mathcal{F}} M_i$, we can pick the neighbourhood $\phi_i(M_i)$, and observe that $\tilde{s}|_{\phi_i(M_i))} = \chi_i \circ s_i \circ \phi_i^{-1}$, which is a composition of smooth maps. We complete the proof by showing that $\tilde{s}$ is actually a right-inverse of $\tilde{\pi}$. Let $[p,i]$ be a point in $\bigcup_{\mathcal{F}} M_i$, and consider $\tilde{\pi}(\tilde{s}([p,i])) = \tilde{\pi}([s_i(p), i])$. It follows from the definition of $\tilde{\pi}$ that $\tilde{\pi}([s_i(p), i]) = [\pi_i(s_i(p)), i]$, and since we assumed that s_i is a section of E_i, we may conclude that $[\pi_i \circ s_i(p), i] = [p, i]$. □

5.2.2 The Tangent Bundle of an Adjoined Manifold

We will now show that given an adjoined manifold $\bigcup_{\mathcal{F}} M_i$ as in Theorem 5.3, the tangent bundles TM_i can be adjoined, and moreover the resulting space is always isomorphic to the tangent bundle of $T\left(\bigcup_{\mathcal{F}} M_i\right)$. To do this, we will form an adjoined bundle as in the previous section, and then exploit the universal property of such a bundle to obtain an isomorphism. To begin with, consider the tuple $\mathcal{G} = (\mathbf{E}, \mathbf{B}, \mathbf{g})$, where:

- each E_i is equal to the tangent bundle TM_i,
- each B_{ij} is equal to the restricted bundle $TM_i|_{A_{ij}}$, and
- each g_{ij} is equal to the differential df_{ij}.

The following result confirms that first and foremost, the tuple $\mathcal{G}$ induces an adjoined bundle.

Lemma 5.14. *The tuple $\mathcal{G}$ described above is an adjunction system that satisfies the criteria of Theorem 5.10.*

Proof. It is fairly straightforward to see that $\mathcal{G}$ is an adjunction system: condition **A1** is trivial, and conditions **A2** and **A3** follow from the basic facts about differentials (as in Prop. A.10). Indeed, for **A2** we have:

$$g_{ij}^{-1} = (df_{ij})^{-1} = d(f_{ij}^{-1}) = df_{ji} = g_{ji},$$

and for **A3** we have:

$$g_{ik} = df_{ik} = d(f_{jk} \circ f_{ij}) = df_{jk} \circ df_{ij} = g_{jk} \circ g_{ij}.$$

As for the conditions of Theorem 5.10, observe that the first two are trivially met by construction. So, it suffices to show that each g_{ij} is a bundle isomorphism covering f_{ij}. Since we have assumed that $\mathcal{F}$ satisfies the conditions of Theorem 5.3, each f_{ij} is a smooth embedding from A_{ij} into X_j. Thus by Lemma A.12, f_{ij} is a diffeomorphism onto its image, i.e. A_{ij} is diffeomorphic to $f_{ij}(A_{ij}) = A_{ji}$. It follows from Lemma A.18 that the differential $df_{ij} : TA_{ij} \to TA_{ji}$ is a bundle isomorphism covering f_{ij}. Moreover, since each A_{ij} is an open submanifold, we can use Lemma A.24 and its subsequent remarks to make the identification $TA_{ij} = TM_i|_{A_{ij}}$, and thus df_{ij} is also a bundle morphism from B_{ij} to B_{ji}. □

Now that we have confirmed that $\bigcup_{\mathcal{G}} TM_i$ is a well-defined adjoined bundle over $\bigcup_{\mathcal{F}} M_i$, we can prove our desired result.

Theorem 5.15. *The bundles $T\left(\bigcup_{\mathcal{F}} M_i\right)$ and $\bigcup_{\mathcal{G}} TM_i$ are isomorphic.*

Proof. We would like to use the universal property of $\bigcup_{\mathcal{G}} TM_i$. In order to do so, we need a collection of bundle maps from the bundles TM_i to $T\left(\bigcup_{\mathcal{F}} M_i\right)$ that commute on their overlaps. We claim that the differential maps $d\phi_i$ are suitable candidates.

Since $\phi_i : M_i \to \bigcup_{\mathcal{F}} M_i$ is smooth, it follows from Lemma A.18 that the differential $d\phi_i$ is a bundle morphism covering ϕ_i. It also follows from Prop. A.18 that

$$d\phi_i = d(\phi_j \circ f_{ij}) = d\phi_j \circ df_{ij} = d\phi_j \circ g_{ij}$$

and thus the maps $d\phi_i$ commute on overlaps. We may then apply Lemma 5.11 to conclude that there is some bundle morphism $\xi : \bigcup_{\mathcal{G}} TM_i \to T\left(\bigcup_{\mathcal{F}} M_i\right)$. We will now show that the map ξ is a bijection. Before we do this, observe that by definition, the map ξ sends $[(p,u), i]$ to $([p,i], (d\phi_i)_p(u))$.

To see that ξ is injective, let $[(p,u), i]$ and $[(q,v), j]$ be two elements of $\bigcup_{\mathcal{G}} TM_i$ such that

$$\xi([(p,u),i]) = ([p,i], (d\phi_i)_p(u)) = ([q,j], (d\phi_j)_q(v)) = \xi([(q,v),j]).$$

It follows immediately that $[p,i] = [q,j]$ (i.e. $f_{ij}(p) = q$) and $(d\phi_i)_p(u) = (d\phi_j)_q(v)$. Using the various properties of the pointwise differential (as in Prop. A.10), we have that

$$\begin{aligned}
(df_{ij})_p(u) &= d(\phi_j^{-1} \circ \phi_i)_p(u) \\
&= (d\phi_j^{-1})_{\phi_i(p)} \circ (d\phi_i)_p(u) \\
&= (d\phi_j^{-1})_{\phi_j(q)} \circ (d\phi_j)_q(v) \\
&= d(\phi_j^{-1} \circ \phi_j)_q(v) \\
&= d(id_{M_j})_q(v) \\
&= v.
\end{aligned}$$

It then follows that $g_{ij}(p,u) = (f_{ij}(p), (df_{ij})_p(u)) = (q,v)$, and thus $[(p,u),i] = [(q,v),j]$ as required.

To show surjectivity, suppose that we have some element of the tangent bundle $T\left(\bigcup_{\mathcal{F}} M_i\right)$. Whatever this element is, it must be of the form $([p,i],u)$ for some $p \in M_i$. We claim the element $[(p,(d\phi_i^{-1})_{\phi_i(p)}(u)),i]$ maps to $([p,i],u)$ under ξ. Observe first that this element is a well-defined member of $\bigcup_{\mathcal{G}} TM_i$ – since ϕ_i is a diffeomorphism, the differential map $(d\phi_i)_p$ is a bijective linear map (see Lemma A.10), so it has a well-defined inverse

$$(d\phi_i^{-1})_{\phi_i(p)} : T_{\phi_i(p)}\left(\bigcup_{\mathcal{F}} M_i\right) \to T_pM_i.$$

As such, the element u maps to $(d\phi_i^{-1})_{\phi_i(p)}(u)$ in T_pM_i, and consequently the element $[(p,(d\phi_i^{-1})_{\phi_i(p)}(u)),i]$ exists as a member of $\bigcup_{\mathcal{G}} TM_i$. It then follows from Lemma A.10 that

$$\begin{aligned}
\xi\left(\left[(p,(d\phi_i^{-1})_{\phi_i(p)}(u)),i\right]\right) &= \left([p,i],(d\phi_i)_p \circ (d\phi_i^{-1})_{\phi_i(p)}(u)\right)\\
&= \left([p,i],d(\phi_i \circ \phi_i^{-1})_{\phi_i(p)}(u)\right)\\
&= \left([p,i],d(id)_{\phi_i(p)}(u)\right)\\
&= ([p,i],u)
\end{aligned}$$

and thus we may conclude that ξ is surjective. Now that we have show that ξ is a bijective bundle morphism, the result then follows from an application of Lemma A.17. $\square$

5.2.3 Canonicity of Adjoined Bundles

We saw in the form of Lemma 5.12 that the pullback of an adjoined bundle $\bigcup_{\mathcal{G}} E_i$ along the canonical map ϕ_i is isomorphic to E_i. We will now work in the opposite direction, and show that any vector bundle F over $\bigcup_{\mathcal{F}} M_i$ can be expressed as an appropriate adjunction of its pullback bundles ϕ_i^*F. In order to do this, we will need the adjunction system $\mathcal{H} = (\mathbf{F},\mathbf{D},\mathbf{h})$, where:

- $\mathbf{F}$ consists of the pullback bundles ϕ_i^*F,
- each D_{ij} is the restriction of the pullback to A_{ij}, i.e. $D_{ij} := (\phi_i^*F)|_{A_{ij}} = p_1^{-1}(A_{ij})$, and
- each $h_{ij} : D_{ij} \to D_{ji}$ is the map defined by $(p,u) \mapsto (f_{ij}(p),u)$.

The reasoning behind the definition of $\mathcal{H}$ is that we would like the diagram

$$\begin{tikzcd}
& (\phi_i^* F)|_{A_{ij}} \arrow[rr] \arrow[dl, "h_{ij}"'] \arrow[dd] & & \phi_i^* F \arrow[dl, dashed] \arrow[dd, "p_2^i"] \\
\phi_j^* F \arrow[rr, dashed] \arrow[dd, "p_2^j"'] & & \bigcup_{\mathcal{H}} \phi_i^* F \arrow[dd, dashed] & \\
& A_{ij} \arrow[rr] \arrow[dl, "f_{ij}"'] & & M_i \arrow[dl, "\phi_i"] \\
M_j \arrow[rr, "\phi_j"'] & & \bigcup_{\mathcal{F}} M_i &
\end{tikzcd}$$

to commute for every i and j. Our first task is to show that $\mathcal{H}$ is well-defined.

Lemma 5.16. *The tuple $\mathcal{H}$ defines an adjunction system.*

Proof. It follows from the definition of the pullback bundle, and the fact that $\pi_F(u) = [p, i] = [f_{ij}(p), j]$ that the h_{ij} are injective functions. It should also be clear condition **A1** is met, since $D_{ii} = \pi_i^{-1}(A_{ii}) = \pi_i^{-1}(M_i) = \phi_i^* F$, and each h_{ii} acts as the identity map, since $f_{ii} = id_{M_i}$. We also have that

$$h_{ij}(D_{ij}) = \{(f_{ij}(p), u) \mid p \in D_{ij}\} = \{(q, u) \mid q \in f_{ij}(A_{ij}) = A_{ji}\} = D_{ji}$$

and that $h_{ji} \circ h_{ij}(p, u) = h_{ji}(f_{ij}(p), u) = (f_{ji} \circ f_{ij}(p), u) = (p, u)$, and consequently $h_{ij}^{-1} = h_{ji}$. Thus condition **A2** is also satisfied. To see that condition **A3** is met, let $(p, u) \in D_{ij} \cap D_{ik}$. Then:

$$h_{ik}(p, u) = (f_{ik}(p), u) = (f_{jk} \circ f_{ij}(p), u) = h_{jk}(f_{ij}(p), u) = h_{jk} \circ h_{ij}(p, u)$$

as required. □

A corollary of the above result is that the adjunction space $\bigcup_{\mathcal{H}} \phi_i^* F$ is a well-defined topological space. The next lemma confirms that the space $\bigcup_{\mathcal{H}} \phi_i^* F$ is an adjoined bundle.

Lemma 5.17. *The space $\bigcup_{\mathcal{H}} \phi_i^* F$ is an adjoined bundle over $\bigcup_{\mathcal{F}} M_i$.*

Proof. We would like to show that $\mathcal{H}$ satisfies the conditions of Theorem 5.10. Of course, the first two conditions of are met by construction, so it suffices to show that each h_{ij} is a bundle isomorphism covering f_{ij}.

We first show that h_{ij} is smooth. It is always possible to find a local representation of h_{ij} that has component functions f_{ij} and $id_{\mathbb{R}^k}$.[5] Of course

[5] Indeed, for any p in A_{ij} we can let (U, Φ) be a local trivialisation of F at $[p, i] = [f_{ij}(p), j]$. We can then use the local trivialisation $\Psi_i : \pi_i^{-1}(\phi_i^{-1}(U)) \to \phi_i^{-1}(U) \times$

these are both smooth and local trivialisations are equipped with the product-manifold structure, so we can use A.8 to conclude that h_{ij} is smooth. Since $h_{ij}^{-1} = h_{ji}$, it follows that each h_{ij} is a diffeomorphism from D_{ij} to D_{ji}. We finish the proof by showing that h_{ij} is a fibrewise linear map. Fix some point p in M_i, and let (p, u) and (p, v) be elements of the fibre $(\phi_i^* F)|_p$. It follows from the definition of the fibres of the pullback bundle (as in Def.A.21) that:

$$h_{ij}(p, u + v) = (f_{ij}(p), u + v) = (f_{ij}(p), u) + (f_{ij}(p), v) = h_{ij}(p, u) + h_{ij}(p, v)$$

and

$$h_{ij}(p, ru) = (f_{ij}(p), ru) = r(f_{ij}(p), u) = rh_{ij}(p, u)$$

as required. It should be clear that h_{ij} covers f_{ij}, since $f_{ij} \circ p_1(p, u) = f_{ij}(p) = p_1(f_{ij}(p), u) = p_1 \circ h_{ij}(p, u)$.

We have shown that each h_{ij} is a diffeomorphism that is fibrewise linear, and covers f_{ij}, i.e. h_{ij} is a bundle isomorphism from D_{ij} to D_{ji}. We can thus apply Theorem 5.10 to conclude that the adjunction space $\bigcup_{\mathcal{G}} \phi_i^* F$ is a well-defined vector bundle over $\bigcup_{\mathcal{F}} M_i$ in which the canonical maps ψ_i are open, smooth embeddings, and are also bundle morphisms covering the f_{ij}. □

We finish this chapter with the result that *all* vector bundles over an adjoined manifold are isomorphic to an adjoined bundle.

Theorem 5.18. *Let F be a vector bundle over $\bigcup_{\mathcal{F}} M_i$. Then F is isomorphic to the adjunction bundle $\bigcup_{\mathcal{H}} \phi_i^*(F)$.*

Proof. The proof of this result revolves around the fact that the adjunction bundle possesses a certain universal property (see Lemma 5.11). We know from Def.A.21 that for each pullback bundle $\phi_i^* F$, the projection map $p_2^i : \phi_i^* F \to F$ acts as a bundle morphism covering ϕ_i. Moreover, for any $p \in A_{ij}$ it is also the case that $p_2^i(p, u) = u = p_2^j(f_{ij}(p), u) = p_2^j \circ h_{ij}(p, u)$. Hence we can apply Lemma 5.11 to conclude that there is some (unique) bundle morphism $\xi : \bigcup_{\mathcal{H}} \phi_i^* F \to F$.

In the proof of Lemma 5.11 we saw that this map is described by $[(p, u), i] \mapsto p_2^i(p, u)$, that is, $\xi = p_2^i \circ \psi_i^{-1}$. We already know that the map ξ is a bundle morphism, i.e. it is smooth, fibrewise linear, and covers the identity map on $\bigcup_{\mathcal{F}} M_i$, so the proof is complete once we show that ξ is bijective.

To see that ξ is injective, suppose that we have two distinct elements $[(p, u), i]$ and $[(q, v), j]$ of $\bigcup_{\mathcal{G}} \phi_i^* F$, where $\xi([(p, u), i]) = u$ and $\xi([(q, v), j]) = v$. Suppose towards a contradiction that $u = v$. Since both elements are distinct,

---

$\mathbb{R}^k$ of D_{ij} at p given by $(q, u) \mapsto (q, p_2 \circ \Phi(u))$, and we can similarly define Ψ_j as a local trivialisation for D_{ji} at $f_{ij}(p)$. The representation of h_{ij} in these local triviliasations is given by

$$\Psi_j \circ h_{ij} \circ \Psi_i^{-1}(q, r) = \Psi_j \circ h_{ij}(q, u) = \Psi_j(f_{ij}(q), u) = (f_{ij}(q), r).$$

it must be the case that $h_{ij}(p,u) = (f_{ij}(p), u) \neq (q,v) = (q,u)$. By definition, this means that $f_{ij}(p) \neq q$, and thus $[p,i] \neq [q,j]$. Since we assumed $u = v$, we have that $\pi_F(u) = [p,i]$ and $\pi_F(u) = [q,j]$, which contradicts π_F as a function. Thus $u \neq v$, from which it follows that ξ is injective. To see that ξ is surjective, let u be some element of F. Then u lies in some fibre $F_{[p,i]}$ where $p \in M_i$. Then the element $[(p,u),i]$ is mapped to u under ξ. Thus ξ is bijective. We can then apply Lemma A.17 to conclude that ξ is a bundle isomorphism from $\bigcup_{\mathcal{H}} \phi_i^* F$ to F. $\square$

6

The Lorentzian Structure of Branching Spacetimes

In the previous chapter we identified conditions under which a collection of smooth Hausdorff manifolds M_i may be glued together to form an adjoined manifold $\bigcup_{\mathcal{F}} M_i$. We also saw that a collection of objects (e.g. functions, sections of bundles) defined on the M_i could be transferred to $\bigcup_{\mathcal{F}} M_i$, provided that the f_{ij} preserved the structure of the objects. In this chapter, we will apply these ideas in the situation that each M_i is a spacetime. A notable result, listed as Theorem 6.6, identifies conditions under which a collection of spacetimes (M_i, g_i) can glued together to form an *adjoined spacetime*.

In Section 7.2, we apply all of the theory outlined in Part 2 spaces (in particular, Theorem 6.6) to the Minkowskian BSTs of Part 1. The approach is similar to that outlined in Section 4.3, however now we will express the construction of MBSTs at the level of the *metric* structure. We will see that the MBSTs can be expressed as an adjoined spacetime, and moreover this extra mathematical structure meets all the criteria for a natural extension of the order-theoretic BST92* structure.

Finally, we will generalise this approach to the case of arbitrary spacetimes. We will define a new class of adjoined spacetimes called *Lorentzian Branching Spacetimes*, which are constructed similarly to MBST. Drawing inspiration from the fact that spacetimes are locally-Minkowskian, we will show that Lorentzian BSTs are locally isometric to MBSTs.

6.1 Constructing Adjoined Spacetimes

We will now show that under some conditions, spacetimes in the sense of Definition A.32 can be naturally adjoined. Throughout this section, we will fix a collection of spacetimes (M_i, g_i) and an adjunction system $\mathcal{F}$ such that the space $\bigcup_{\mathcal{F}} M_i$ is an adjoined manifold as in Theorem 5.3. We will also assume that each A_{ij} is non-empty, since we can then apply Lemma 3.15 conclude that $\bigcup_{\mathcal{F}} M_i$ is connected.

Given these assumptions, there are two things to do: we need to define a Lorentzian metric on $\bigcup_{\mathcal{F}} M_i$, and we need to define a time-orientation. We will treat these tasks separately.

6.1.1 Defining a Lorentzian Metric

Recall that a Lorentzian metric on a manifold M is a section g of the tensor bundle $T^2(T^*M)$ that is everywhere symmetric and non-degenerate, with signature $(-,+,...,+)$. In order to construct a Lorentzian metric on $\bigcup_{\mathcal{F}} M_i$ we first need to describe the bundle $T^2(T^* \bigcup_{\mathcal{F}} M_i)$. In the same spirit as Theorem 5.15, we will show that the bundle $T^2(T^* \bigcup_{\mathcal{F}} M_i)$ is isomorphic to an adjunction of the bundles $T^2(T^*M_i)$.

To begin with, consider the tuple $\mathcal{G} = (\mathbf{E}, \mathbf{B}, \mathbf{g})$, where:

- each E_i is equal to the bundle $T^2(T^*M_i)$,
- each B_{ij} is equal to the restricted bundle $(T^2(T^*M_i))|_{A_{ij}}$, and
- each g_{ij} is the map sending (p, α) to $(f_{ij}(p), (f_{ij})_*\alpha)$, where $(f_{ij})_*\alpha$ is the pushforward of α along f_{ij} defined as in Section 5.2.3.

Observe the similarity between this system and the system defined in Section 6.2.2. In fact, lots of the following results follow the same conceptual framework as our discussion of adjoined tangent bundles. Our next result is an analogue of Lemma 5.14.

Lemma 6.1. *The tuple $\mathcal{G}$ defined above is an adjunction system that satisfies the conditions of Theorem 5.10.*

Proof. That $\mathcal{G}$ is an adjunction system follows routinely from the basic facts of pointwise pushforwards and pullbacks – condition **A1** follows from Prop.A.19.1, condition **A2** follows from Prop.A.19.2, and condition **A3** follows from Prop.A.19.4.

By construction, $\mathcal{G}$ satisfies the first two conditions of Thm. 6.REF. To see that the third condition is met, observe that the definition of the g_{ij} is essentially the same as the map used in Lemma A.20. We can use the same arguments found there to conclude that the g_{ij} are bundle isomorphisms covering the maps f_{ij}. □

The above result, together with Theorem 5.10, confirms that $\bigcup_{\mathcal{G}} T^2(T^*M_i)$ is an adjoined bundle over $\bigcup_{\mathcal{F}} M_i$. We will now show that this adjoined bundle is isomorphic to the bundle of bilinear forms on the adjoined manifold $\bigcup_{\mathcal{F}} M_i$.

Theorem 6.2. *The bundles $\bigcup_{\mathcal{G}} T^2(T^*M_i)$ and $T^2(T^* \bigcup_{\mathcal{F}} M_i)$ are isomorphic.*

Proof. Consider the maps $\psi_i : T^2(T^*M_i) \to T^2(T^* \bigcup_{\mathcal{F}} M_i)$ defined by $\psi_i(p,\alpha) = ([p,i],(\phi_i)_*\alpha)$. It should be clear that these maps are well-defined – by Thm. 5.3 the canonical maps ϕ_i are smooth embeddings, thus they are diffeomorphic onto their images. Moreover, we can use the same argument as in Lemma A.20 to conclude that the ψ_i are bundle morphisms covering the ϕ_i. We now show that $\psi_i = \psi_j \circ g_{ij}$ where defined. Suppose that (p,α) is some element of $T^2(T^*A_{ij})$. Then:

$$\begin{aligned}\psi_j \circ g_{ij}(p,\alpha) &= \psi_j(f_{ij}(p),(f_{ij})_*\alpha)\\ &= ([f_{ij}(p),j],(\phi_j)_* \circ (f_{ij})_*\alpha)\\ &= ([f_{ij}(p),j],(\phi_j \circ f_{ij})_*\alpha)\\ &= ([p,i],(\phi_i)_*\alpha)\end{aligned}$$

where the third equality follows from Prop. A.19.4. We can then use Lemma 5.11 to conclude that there is some smooth bundle morphism

$$\xi : \bigcup_{\mathcal{G}} T^2(T^*M_i) \to T^2(T^* \bigcup_{\mathcal{F}} M_i), \text{where } \xi([(p,\alpha),i]) = ([p,i],(\phi_i)_*\alpha).$$

We will now show that ξ is bijective. For injectivity, suppose that we have two distinct elements $[(p,\alpha),i]$ and $[(q,\beta),j]$ of the adjoined bundle $\bigcup_{\mathcal{G}} T^2(T^*M_i)$. Then $g_{ij}(p,\alpha) = (f_{ij}(p),(f_{ij})_*\alpha) \neq (q,\beta)$. In the case that $f_{ij}(p) \neq q$, i.e. $[p,i] \neq [q,j]$, we immediately have that

$$\xi([(p,\alpha),i]) = ([p,i],(\phi_i)_*\alpha) \neq ([q,j],(\phi_j)_*\beta) = \xi([(q,\alpha),j]).$$

So, suppose that $f_{ij}(p) = q$. Thus it must be the case that $(f_{ij})_*\alpha \neq \beta$. Then $g_{ij}(p,alpha) = (f_{ij}(p),(f_{ij})_*(\alpha)) = (q,(f_{ij})_*(\alpha)$. So, in order for $[(p,\alpha),i]$ and $[(q,\beta),j]$ to be distinct, it must be the case that $(f_{ij})_*(\alpha) \neq \beta$. This means that there are elements v,w in T_qM_j such that $(f_{ij})_*\alpha(v,w) \neq \beta(v,w)$. Since f_{ij} is a diffeomorphism, by Prop. A.10.4 the differential $(df_{ij})_p$ is a bijective map from T_pM_i to T_qM_j. Thus the elements v and w are of the form $v = (df_{ij})_p(v')$ and $w = (df_{ij})_p(w')$, where v' and w' are elements of T_pM_i. Consider now the elements $(d\phi_i)_p(v')$ and $(d\phi_i)_p(w')$ in the tangent space $T_{[p,i]} \bigcup_{\mathcal{F}} M_i$. Then:

$$\begin{aligned}(\phi_i)_*\alpha\left((d\phi_i)_p(v'),(d\phi_i)_p(w')\right) &= \alpha(v',w')\\ &= (f_{ij})_*\alpha((df_{ij})_p(v'),(df_{ij})_p(w'))\\ &\neq \beta\left((df_{ij})_p(v'),(df_{ij})_p(w')\right)\\ &= (\phi_j)_*\beta\left((d\phi_j)_q \circ (df_{ij})_p(v'),(d\phi_j)_q \circ (df_{ij})_p(w')\right)\\ &= (\phi_j)_*\beta\left(d(\phi_j \circ f_{ij})_p(v),d(\phi_j \circ f_{ij})_p(w)\right)\\ &= (\phi_j)_*\beta\left((d\phi_i)_p(v),(d\phi_i)_p(w)\right)\end{aligned}$$

that is, $(\phi_i)_*\alpha \neq (\phi_j)_*\beta$. Again we have that $\xi([(p,\alpha),i]) \neq \xi([(q,\alpha),j])$, from which we may conclude that ξ is indeed injective.

To see that ξ is surjective, suppose that we have some element $([p,i],\alpha)$ in the bundle $T^2(T^* \bigcup_{\mathcal{F}} M_i)$. Consider the element $[(p,\phi_i^*\alpha),i]$ in the adjoined bundle $\bigcup_{\mathcal{G}} T^2(T^*M_i)$. It is not hard to see that ξ will map this element to $([p,i],\alpha)$. Indeed, we have:

$$\xi([(p,\phi_i^*\alpha),i]) = ([p,i],(\phi_i)_*(\phi_i^*\alpha)) = ([p,i],\alpha)$$

where the final equality follows from Prop. A.19.2. We may then conclude that ξ is surjective. The result then follows from an application of Lemma A.17. $\square$

Now that we have a nice description of the bundle of bilinear forms on $\bigcup_{\mathcal{F}} M_i$, we can transfer the Lorentzian metrics of each M_i. Of course, this can not be done arbitrarily – we need to make the assumption that the g_i are compatible on their overlaps. The following theorem says that if this assumption is made, then the Lorentzian metrics of the spacetimes M_i can be pushed into $\bigcup_{\mathcal{F}} M_i$.

Theorem 6.3. *Suppose that each M_i is equipped with a Lorentzian metric g_i, and each f_{ij} is an isometric embedding. Then there is a well-defined Lorentzian metric $\tilde{g}$ on $\bigcup_{\mathcal{F}} M_i$ that turns the canonical maps ϕ_i into isometric embeddings.*

Proof. Since each f_{ij} is an isometry, we have that $g_j \circ f_{ij} = g_{ij} \circ g_i$, that is, the diagram

$$\begin{array}{ccc} T^2(T^*A_{ij}) & \xrightarrow{g_{ij}} & T^2(T^*A_{ji}) \\ \uparrow{\scriptstyle g_i} & & \uparrow{\scriptstyle g_j} \\ A_{ij} & \xrightarrow[f_{ij}]{} & A_{ji} \end{array}$$

commutes. Thus we can apply Lemma 5.13 to define a section $\hat{g}$ of the adjoined bundle $\bigcup_{\mathcal{G}} T^2(T^*M_i)$, where in this case $\hat{g}$ maps each $[p,i]$ to $[g_i(p),i]$. By the previous result, the adjoined bundle $\bigcup_{\mathcal{G}} T^2(T^*M_i)$ is isomorphic to the bundle $T^2(T^* \bigcup_{\mathcal{F}} M_i)$ under the mapping ξ. We can then use xi to push the section $\hat{g}$ forward to $T^2(T^* \bigcup_{\mathcal{F}} M_i)$, by defining the map $\tilde{g} := \xi \circ \hat{g}$. We now show that this new section $\tilde{g}$ is a well-defined Lorentzian metric on $\bigcup_{\mathcal{F}} M_i$. Clearly $\tilde{g}$ is a smooth section – it is smooth since both $\hat{g}$ and ξ are smooth, and it is the case that $\pi \circ \xi \circ \hat{g} = \tilde{\pi} \circ \hat{g} = id$ since ξ covers the identity map on $\bigcup_{\mathcal{F}} M_i$ and $\hat{g}$ is section of $\bigcup_{\mathcal{G}} T^2(T^*M_i)$.

To see that $\tilde{g}$ is everywhere symmetric and non-degenerate, let $[p,i]$ be arbitrary and consider two tangent vectors v,w in $T_{[p,i]} \bigcup_{\mathcal{F}} M_i$. By the results of Prop. A.10, we know that the differential $(d\phi_i)_p$ is a bijective linear map from T_pM_i to the fibre $T_{[p,i]} \bigcup_{\mathcal{F}} M_i$. Thus there are two unique elements v', w'

in T_pM_i such that $v = (d\phi_i)_p(v')$ and $w = (d\phi_i)_p(w)$. Moreover, by definition we have that:

$$\tilde{g}([p,i]) = \xi \circ \hat{g}([p,i]) = \xi([(p, g_i(p)), i]) = ([p,i], (\phi_i)_*(g_i(p))) \qquad (*)$$

where $(\phi_i)_*(g_i(p))(v,w) = g_i(p)(v',w')$. Then:

$$(\phi_i)_*(g_i(p))(v,w) = g_i(p)(v',w') = g_i(p)(w',v') = (\phi_i)_*(g_i(p))(w,v)$$

and thus the bilinear form $(\phi_i)_*(g_i(p))$ is symmetric. To see that $\tilde{g}$ is also non-degenerate, suppose now that v is an element of $T_{[p,i]} \bigcup_{\mathcal{F}} M_i$ such that $(\phi_i)_*(g_i(p))(v,w) = 0$ for all w in the same fibre. Then $v' \in T_pM_i$ will have the same property. Indeed, for any $w' \in T_pM_i$, we would have that:

$$g_i(p)(v',w') = (\phi_i)_*(g_i(p))(v,w) = 0.$$

Since g_i is non-degenerate, this can only be the case iff v' is equal to the zero element of T_pM_i. Since the pointwise differential $(d\phi_i)_p$ is a *linear* map, it follows that $v = (d\phi_i)_p(v') = (d\phi_i)_p(0) = 0$ as required. We now show that $\tilde{g}$ turns the ϕ_i into isometric embeddings. Of course, by Theorem 5.3 we know that the ϕ_i are smooth embeddings, so it suffices to show that

$$g_i(v,w) = \tilde{g}((d\phi_i)_p(v), (d\phi_i)_p(w))$$

for all v, w in T_pM_i. However, this is precisely how we *defined* $\tilde{g} := \xi \circ \hat{g}$, as seen in the expression $(*)$. Thus the ϕ_i are isometric embeddings. Since isometries preserve the signature of a metric (see Prop. A.28), it follows that the signature of $\tilde{g}$ is equal to the signature of each g_i, and thus $\tilde{g}$ is a Lorentzian metric. □

6.1.2 Defining a Time-Orientation

Now that we have shown that our adjoined manifold $\bigcup_{\mathcal{F}} M_i$ has well-defined Lorentzian metric $\tilde{g}$, we can discuss time-orientations. Again, it is not the case in general that every collection τ_i of time-orientations can induce a time-orientation on $\bigcup_{\mathcal{F}} M_i$. However, if we assert that the τ_i are pairwise compatible, then we obtain the following result.

Lemma 6.4. *If the f_{ij} preserve time-orientations, then there is a time-orientation $\tilde{\tau}$ on $\bigcup_{\mathcal{F}} M_i$ such that the ϕ_i preserve time-orientations.*

Proof. Recall that f_{ij} preserves time-orientations whenever $df_{ij}(\tau_i(p)) = \tau_j(f_{ij}(p))$ for every p in A_{ij}. We define the time orientation $\tilde{\tau}$ by $\tilde{\tau}([p,i]) = d\phi_i(\tau_i(p))$. Observe first that this is well-defined, since whenever $[f_{ij}(p), j] = [p,i]$, we have that

$$d\phi_j(\tau_j(f_{ij}(p))) = d\phi_j \circ df_{ij}(\tau_i(p)) = d(\phi_j \circ f_{ij})(\tau_i(p)) = d\phi_i(\tau_i(p)).$$

We now check that $\tilde{\tau}$ is smooth, i.e. we check that for every point in $\bigcup_{\mathcal{F}} M_i$ there is an open neighbourhood and a locally-defined vector field that agrees with $\tilde{\tau}$. So, let $[p,i]$ be a point in $\bigcup_{\mathcal{F}} M_i$, and consider p in M_i. By assumption the time-orientation τ_i on M_i is smooth, so there is some open neighbourhood U of p and a vector field X_U on U such that for each p' in U, it is the case that $X_U(p') \in \tau_i(p)$. Consider now the subset $\phi_i(U)$, which is an open neighbourhood of $[p,i]$ (since ϕ_i is an open map). Since ϕ_i is a smooth embedding, in particular it is a diffeomorphism onto $\phi_i(M_i)$ and thus by Lemma A.9 it is a diffeomorphism once restricted to U. Thus we can pushforward the vector field X_U along ϕ_i to obtain the vector field $(\phi_i)_* X_U$. Recall that the pushforward of a vector field along a diffeomorphism is defined to be:

$$(\phi_i)_* X_U([p',i]) := d\phi_i \circ X_U(\phi_i^{-1}([p',i])) = d\phi_i \circ X_U(p').$$

We now show that this vector field will suffice to witness the smoothness of $\tilde{\tau}$ at $[p,i]$. Consider some element $[q,j]$ of $\phi_i(U)$. Since $\phi_i(U) \subseteq \phi_i(M_i)$, it must be the case that $[q,j] = [f_{ij}(p'),j] = [p',i]$ for some $p' \in A_{ij} \cap U$. Then $\tilde{\tau}([p',i]) = d\phi_i(\tau_i(p'))$. It follows from the definition of $(\phi_i)_* X_U$ that $(\phi_i)_* X_U([p',i]) = d\phi_i \circ X_U(p')$. By assumption $X_U(p') \in \tau_i(p')$, thus $d\phi_i \circ X_U(p') \in d\phi_i(\tau_i(p')) = \tilde{\tau}([p',i])$, and we may conclude that $\tilde{\tau}$ is indeed smooth. Observe that the canonical maps ϕ_i preserve time-orientations follows immediately from the construction of $\tilde{\tau}$. □

We suggested in Section 5.3.3 that when a spacetime is non-Hausdorff, it may not be the case that the two definitions of a time-orientation coincide. As such, we cannot conclude from the above result that the existence of a time-oreintation $\tilde{\tau}$ entails the existence of a globally-defined timelike vector field. As a small aside, the following is a sufficient condition to guarantee the existence of globally-defined timelike vector field on the adjoined manifold $\bigcup_{\mathcal{F}} M_i$.

Lemma 6.5. *Let $\{X_i\}$ be a collection of globally-defined, timelike vector fields for the M_i such that the diagram*

$$\begin{tikzcd} TA_{ij} \arrow[r, "df_{ij}"] & TA_{ji} \\ A_{ij} \arrow[u, "X_i"] \arrow[r, "f_{ij}"'] & A_{ji} \arrow[u, "X_j"'] \end{tikzcd}$$

commutes for each i and j. Then $(\bigcup_{\mathcal{F}} M_i, g)$ has a globally-defined timelike vector field.

Proof. By assumption, we can use Lemma 5.13 to define a section $\hat{X}$ of the bundle $\bigcup_{\mathcal{G}} TM_i$. By Thm.5.15, there is a bundle isomorphism $\Psi : \bigcup_{\mathcal{G}} TM_i \to$

$T(\bigcup_{\mathcal{F}} M_i)$. We define the vector field $\tilde{X} : \bigcup_{\mathcal{F}} M_i \to T(\bigcup_{\mathcal{F}} M_i)$ to be $\tilde{X} := \Psi \circ \hat{X}$. This means that

$$\tilde{X}([p,i]) = \Psi([(p, X_i), i]) = ([p,i], (d\phi_i)_p(X_i(p))).$$

It is not hard to see that this is indeed a globally-defined section of $T(\bigcup_{\mathcal{F}} M_i)$ – the argument is similar to that of Thm.6.3. To see that $\tilde{X}$ is timelike, suppose we have some $[p,i] \in \bigcup_{\mathcal{F}} M_i$, and consider $\tilde{X}([p,i]) \in T_{[p,i]} \bigcup_{\mathcal{F}} M_i$. Then $\tilde{X}([p,i]) = ([p,i], (d\phi_i)_p(X_i(p)))$. By definition of the Lorentzian metric $\tilde{g}$, we have that

$$\tilde{g}_{[p,i]}((d\phi_i)_p(X_i(p)), (d\phi_i)_p(X_i(p))) = (g_i)_p(X_i(p), X_i(p)).$$

which is timelike since we assumed that X_i is timelike on M_i. □

The next theorem is a summary of the results thus far.

Theorem 6.6. *Let $\{(M_i, g_i)\}$ be a countable collection of Hausdorff spacetimes, and $\mathcal{F} = (\boldsymbol{X}, \boldsymbol{A}, \boldsymbol{f})$ an adjunction system in which:*

1. *$\boldsymbol{X}$ consists of the spaces M_i,*
2. *each A_{ij} is an open Lorentzian submanifold of M_i, and*
3. *each $f_{ij} : A_{ij} \to M_j$ is an isometric embedding that preserves time-orientation.*

Then the adjunction space $\bigcup_{\mathcal{F}} M_i$ possesses a Lorentzian metric that makes the canonical maps ϕ_i act as open, isometric embeddings that preserve time-orientation.

Proof. We can use Theorem 6.3 to conclude that $\bigcup_{\mathcal{F}} M_i$ has a Lorentzian metric, and we can use Lemma 6.4 to define a time-orientation $\tilde{\tau}$ on $\bigcup_{\mathcal{F}} M_i$ that makes every ϕ_i preserve time-orientations. □

We will refer to adjunction spaces satisfying the conditions of Theorem 6.6 as *adjoined spacetimes*. We complete this section by showing that the adjoined spacetime $(\bigcup_{\mathcal{F}} M_i, \tilde{g})$ possesses a certain universal property.

Lemma 6.7. *Let $\mathcal{F}$ be as in Theorem 6.6. Suppose that there are smooth isometric embeddings $\psi_i : M_i \to N$ that commute on overlaps. Then there is a unique isometric embedding from $\bigcup_{\mathcal{F}} M_i$ to N.*

Proof. We know from Lemma 5.11 that there is at least a unique smooth embedding $\xi : \bigcup_{\mathcal{F}} M_i \to N$ where $[p,i] \mapsto \psi_i(p)$. To see that ξ is an isometry, we need to show that $\tilde{g}_{[p,i]}(v,w) = g^N_{\xi([p,i])}((d\xi)_{[p,i]}(v), (d\xi)_{[p,i]}(w))$, however this follows from our assumption that each ψ_i is an isometric embedding. Indeed:

$$\begin{aligned}\tilde{g}_{[p,i]}(v,w) &= (g_i)_p((d\phi_i^{-1})_{[p,i]}(v),(d\phi_i^{-1})_{[p,i]}(w))\\ &= (g^N)_{\psi(p)}((d\psi_i)_p\circ(d\phi_i^{-1})_{[p,i]}(v),(d\psi_i)_p\circ(d\phi_i^{-1})_{[p,i]}(w))\\ &= (g^N)_{\psi(p)}(d(\psi_i\circ\phi_i^{-1})_{[p,i]}(v),d(\psi_i\circ\phi_i^{-1})_{[p,i]}(w))\\ &= (g^N)_{\xi([p,i])}((d\xi)_{[p,i]}(v),(d\xi)_{[p,i]}(w))\end{aligned}$$

Thus ξ is an isometric embedding from $\bigcup_{\mathcal{F}} M_i$ into N. □

6.1.3 Causal Properties of Adjoined Spacetimes

Given an adjoined spacetime $(\bigcup_{\mathcal{F}} M_i, \tilde{g})$, we can associate a causal ordering $\leq^g$, defined by: $[p,i] \leq^{\tilde{g}} [q,j]$ iff there is a future directed, causal curve $\gamma : [0,1] \to \bigcup_{\mathcal{F}} M_i$ such that $\gamma(0) = [p,i]$ and $\gamma(1) = [q,j]$. Since the maps ϕ_i are so well-behaved, it will be interesting to see whether the causal properties of $\bigcup_{\mathcal{F}} M_i$ are inherited from the M_i. In what follows we will denote by $\leq^i$ the causal ordering of the i^{th} spacetime (M_i, g_i). Our first result shows that the canonical maps ϕ_i are monotone with respect to the orderings $\leq^i$ and $\leq^{\tilde{g}}$.

Lemma 6.8. *Let $p \leq^i q$ in M_i. Then $[p,i] \leq^{\tilde{g}} [q,i]$ in $\bigcup_{\mathcal{F}} M_i$.*

Proof. By Thm. 6.6, the canonical map ϕ_i is an isometric embedding that is time-orientation preserving. The result then follows from an application of Lemma A.34 and the observation that a causal curve in $\phi_i(M_i)$ is also a causal curve in $\bigcup_{\mathcal{F}} M_i$. □

At the very least, we have that the ϕ_i are order-preserving. We will now set about proving a restricted dual to the above lemma. Specifically, we will show that whenever the A_{ij} are past-sets, then two elements $[p,i]$ and $[q,j]$ that are $\leq^{\tilde{g}}$-related must also be $\leq^j$-related. Recall that a subset A of a spacetime (M,g) is called a *past-set* if it is downwards-closed under $\leq^g$ (or equivalently, if $J^-(A) = A$). This will mean that the ordering $\leq^{\tilde{g}}$ will have a downwards-closure property similar to that of Prop. 1.12.1. Before we prove our restricted converse to Lemma 6.8, we need a small result.

Lemma 6.9. *Let $\gamma : [0,1] \to \bigcup_{\mathcal{F}} M_i$ be a future-directed causal curve such that $\gamma(1) = [p,i]$. If each A_{ij} is a past-set, then $\gamma \subset \phi_i(M_i)$.*

Proof. Suppose towards a contradiction that $\gamma \not\subset \phi_i(M_i)$. Since γ is a curve, in particular it is continuous. Hence the preimages $\gamma^{-1}(\phi_j(M_j))$ are all open subsets of $[0,1]$. By our supposition the set $\gamma^{-1}(\phi_i(M_i))$ is an open, non-empty, *proper* subset of $[0,1]$.

We would like to pick a (maximal) open set $U \subset [0,1]$ such that $U \subseteq \gamma^{-1}(\phi_i(M_i))$ and $1 \in U$. This is done as follows:

$$U := \bigcup\{U_\alpha \mid U_\alpha := (a_\alpha, 1] \subseteq \gamma^{-1}(\phi_i(M_i))\}$$

This set is clearly open, and it is connected since it is the union of connected subsets that are pairwise non-disjoint[1]. Since the set U is a connected, open subset of $[0,1]$, it must[2] be of the form $U := (a,1]$.

Consider now the element $a \in [0,1]$. It should be clear that $a \notin \gamma^{-1}(\phi_i(M_i))$. Indeed, suppose towards a contradiction that $a \in \gamma^{-1}(\phi_i(M_i))$. Since the set $\gamma^{-1}(\phi_i(M_i))$ is open and contains a, we can pick some open interval $(a-\epsilon, a+\epsilon)$ contained within $\gamma^{-1}(\phi_i(M_i))$. Then it would be the case that $(a-\epsilon, a+\epsilon) \cup (a,1] = (a-\epsilon, 1]$ is a connected, open set contained within $\gamma^{-1}(\phi_i(M_i))$, and thus $a \in U = (a,1]$, a contradiction. Thus $a \notin \gamma^{-1}(\phi_i(M_i))$. Denote by $[y,j]$ the image of a under γ, i.e. $\gamma(a) = [y,j]$.

Since $a \notin \gamma^{-1}(\phi_i(M_i))$, it must be the case that $[y,j] \notin \phi_i(M_i) \cap \phi_j(M_j) = \phi_j(A_{ji})$. Consider now the set $\gamma^{-1}(\phi_j(M_j))$, which is an open subset of $[0,1]$. Since $\gamma^{-1}(\phi_j(M_j))$ is open and a lies in this set, it follows that there is some open subset V of $[0,1]$ such that $a \in V \subseteq \gamma^{-1}(\phi_j(M_j))$. Without loss of generality we can assume V is of the form (b_1, b_2).[3] Pick any element $c \in (b_1,b_2) \cap (a,1]$. Observe that $c \in (b_1,b_2) \subseteq \gamma^{-1}(\phi_j(M_j))$ and $c \in (a,1] \subseteq \gamma^{-1}(\phi_i(M_i))$, so c lies in their intersection, and thus $\gamma(c) \in \phi_i(M_i) \cap \phi_j(M_j) = \phi_i(A_{ij})$. Denote this element by $\gamma(c) = [z,j]$.

Observe now that the restriction of γ to the closed (connected) interval $[a,c] \subset [0,1]$ is a future directed curve that lies entirely within $\phi_j(M_j)$, so can effectively be seen as a curve on the subspace $(\phi_i(M_i), \tilde{g}|_{\phi_i(M_i)})$. Since ϕ_j^{-1} is an isometry, the curve $\gamma' := \phi_j^{-1} \circ \gamma$ is a future-directed causal curve connecting y to z. Thus $y \leq^j z$. Recall that $[z,j]$ lies in $\phi_j(A_{ji})$. Thus $z \in A_{ji}$. We have assumed that A_{ji} is a past-set, thus $y \leq^j z$ and $z \in A_{ji}$ implies that $y \in A_{ji}$. Thus $[y,j] \in \phi_j(A_{ji})$. However, this contradicts the fact that $[y,j] \notin \phi_i(M_i)$. □

We can now use the above lemma to prove the following result, which is a restricted dual to Lemma 6.8.

Lemma 6.10. *Suppose that every A_{ij} is a past set. Then $[p,i] \leq^{\tilde{g}} [q,j]$ iff $[p,i] \in \phi_i(A_{ij})$ and $f_{ij}(p) \leq^j q$.*

Proof. Suppose first that $[p,i] \leq^{\tilde{g}} [q,j]$. Then there is a future-directed, causal curve γ in $\bigcup_{\mathcal{F}} M_i$ such that $\gamma(0) = [p,i]$ and $\gamma(1) = [q,j]$. It follows from

[1] This is a fairly well-known fact. Let A_i be connected subsets of a topological space X such that $A_i \cap A_j \neq \emptyset$ for every i,j. If $\bigcup_i A_i$ were disconnected, then there are two disjoint open subsets U, V of X such that $U \cup V = \bigcup_i A_i$. Let $x \in U \cup V$. Then $x \in A_i$. Since A_i connected, it must be the case that $A_i \subseteq U$, since otherwise $A_i = (A_i \cap U) \cup (A_i \cap V)$, a contradiction. Since $A_i \cap A_j \neq \emptyset$, every A_k is a subset of U, and thus V is empty, a contradiction.

[2] The connected subsets of $[0,1]$ are of the form $(a,b), [a,b), (a,b]$ or $[a,b]$. Thus the open connected subsets are of the form $(a,b), [0,a)$, or $(a,1]$.

[3] If V is open then it is a union of basis elements, which are of the form $(r,r') \cap [0,1]$. So, $a \in V = [0,1] \cap \bigcup (r,r')$ implies that $a \in (r,r') \cap [0,1]$, so just pick a local basis for a, say $(a-\epsilon, a+\epsilon)$ small enough that its contained within the interval (r,r').

the previous lemma that $\gamma \subseteq \phi_j(M_j)$, and thus $[p,i] = [f_{ij}(p),j]$. Hence $[f_{ij}(p),j] \leq^{\tilde{g}} [q,j]$, and since $\phi_j^{-1} : \phi_j(M_j) \to M_j$ is a time-orientation-preserving isometry, it follows from Lemma A.34 that $f_{ij}(p) \leq^j q$, as required. Conversely, suppose that $p \in A_{ij}$ and $f_{ij}(p) \leq^j q$. By Lemma 6.8, it follows that $[f_{ij}(p),j] \leq^{\tilde{g}} [q,j]$, and thus $[p,i] \leq^{\tilde{g}} [q,j]$. □

The above result is fairly strong, and it allows us to transfer potential causal properties of the M_i to the adjoined spacetime $\bigcup_{\mathcal{F}} M_i$. Our first result says that whenever the M_i possess compatible global time functions, then these can be pushed into $\bigcup_{\mathcal{F}} M_i$.

Lemma 6.11. *Suppose that each A_{ij} is a past-set. If each M_i has a global time function t_i such that $t_j \circ f_{ij}(p) = t_i(p)$ for each j, then $\bigcup_{\mathcal{F}} M_i$ has a well-defined global time function.*

Proof. Define $\tilde{t} : \bigcup_{\mathcal{F}} M_i \to \mathbb{R}$ by $[p,i] \mapsto t_i(p)$. By Cor. 5.5 this is a well-defined, real-valued function on $\bigcup_{\mathcal{F}} M_i$. To see that $\tilde{t}$ is monotone, suppose that we have some $[p,i]$ and $[q,j]$ in $\bigcup_{\mathcal{F}} M_i$ such that $[p,i] \leq^{\tilde{g}} [q,j]$. By Lemma 6.10, this is the case iff $[p,i] \in \phi_j(A_{ji})$ and $f_{ij}(p) \leq^j q$. Since t_j is a global time function, it follows that:

$$\tilde{t}([p,i]) = \tilde{t}([f_{ij}(p),j]) = t_j(f_{ij}(p)) \leq^j t_j(q) = \tilde{t}([q,j]),$$

and consequently $\tilde{t}$ is a global time function on $\bigcup_{\mathcal{F}} M_i$. □

We finish this section with a corresponding result for global hyperbolicity.

Lemma 6.12. *Let I be a finite indexing set. If each M_i is globally-hyperbolic and each A_{ij} is a past-set, then every causal diamond in $\bigcup_{\mathcal{F}} M_i$ is compact.*

Proof. Let $[p,i] \leq^{\tilde{g}} [q,j]$, and consider the set

$$J^+([p,i]) \cap J^-([q,j]) := \{[r,k] \mid [p,i] \leq^{\tilde{g}} [r,k] \leq^{\tilde{g}} [q,j]\}.$$

We can apply Lemma 6.10 to conclude that the above set is contained as a subset of $\phi_j(M_j)$. It should be clear[4] that the following equivalence holds:

$$J^+([p,i]) \cap J^-([q,j]) = \phi_j(J^+(f_{ij}(p)) \cap J^-(q)).$$

By assumption, the set $J^+(f_{ij}(p)) \cap J^-(q)$ is compact in M_j. By assumption I is finite, so we can apply Lemma 3.16 to conclude that the set $\phi_j(J^+(f_{ij}(p)) \cap J^-(q))$ is compact in $\bigcup_{\mathcal{F}} M_i$. □

[4] Specifically, $[p,i] \leq^{\tilde{g}} [w,k] \leq^{\tilde{g}} [q,j]$ iff $f_{ij}(p) \leq^j f_{ik}(w) \leq^j q$ iff $f_{ij}(w) \in J^+(f_{ij}(p)) \cap J^-(q)$ iff $[w,k] \in \phi_j(J^+(f_{ij}(p)) \cap J^-(q))$

6.2 Branching Spacetimes

It is now time to combine our theory of adjunction spaces with the branching spacetimes of Part 1. We will start by expressing our Minkowskian BSTs as an adjoined manifold, and justifying this extra structure as natural. After this, we will generalise this construction and effectively replace the Minkowski spacetime with an arbitrary one. This generalisation, which is the main contribution of this thesis, enables the definition of a new class of BSTs, which we will call *Lorentzian Branching Spacetimes*.

6.2.1 The Lorentzian Structure of Minkowskian BSTs

We have already discussed the construction of Minkowskian BSTs at two levels of structure – in Section 2.2 we constructed MBSTs as models of BST92*, and in Section 4.3 we constructed MBSTs from a topological perspective. In this section, we will do this one last time, and construct MBSTs as adjoined spacetimes.

The obvious way to do this is to invoke Theorem 6.6. In order to do so, we first need to define an appropriate adjunction system $\mathcal{F} = (\mathbf{X}, \mathbf{A}, \mathbf{f})$. We can take inspiration from Section 4.3.1, where now we require that:

- each member of $\mathbf{X}$ consists of the n-dimensional Minkowski spacetime (M^n, η) viewed as a Lorentzian manifold,
- $A_{ij} = M^n \backslash J^+(C_{ij})$, viewed as an open Lorentzian submanifold of M^n, and
- each $f_{ij} : A_{ij} \to M^n$ is the inclusion map.

The underlying set of the adjunction space $\bigcup_{\mathcal{F}} M_i$ is of course equal to $M^n_{\mathcal{C}}$ (constructed as in Section. 2.2.2). It should also be clear that the tuple $\mathcal{F}$ meets the criteria of Theorem 6.6 – all of the conditions either follow by construction, or from the basic properties of the inclusion map. So, we can conclude that there is a well-defined Lorentzian metric $\tilde{\eta}$ on the adjunction space $M^n_{\mathcal{C}}$. In order to distinguish this object from the standard MBSTs defined in Section 2.2.2, we will denote the adjoined spacetime subordinate to the system $\mathcal{F}$ by $(M^n_{\mathcal{C}}, \tilde{\eta})$.

Naturality of the Structure

In Section 1.2.1 we identified three general criteria for the naturality of structures defined on BST92(*) models. These were:

N1) The structure, once restricted to a single-historied model, should agree with the Minkowski version.
N2) The structure should possess a certain universal property, in that it can be canonically reconstructed from its history-relative substructures.

N3) The structure should be compatible with any pre-existing BST92* concepts.

We will now argue that the structure $(M^n_{\mathcal{C}}, \tilde{\eta})$ meets all of these criteria. Of course, we have already discussed the nature of the topological structure of $M^n_{\mathcal{C}}$ (cf. Section 4.3.1), so we will only need to discuss the differentiable and metric structures.

We start with condition **N1**. We know from Theorem 2.17 that the histories of $M^n_{\mathcal{C}}$ are precisely equal to the layers $L_i := \phi_i(M^n)$. We also know from Theorem 6.6 that the canonical maps ϕ_i are isometric embeddings that preserve time-orientation. This means that each ϕ_i acts as an isometry from M^n to the history $\phi_i(M^n)$. As such, the spaces $\phi_i(M^n)$ are isomorphic to the Minkowski spacetime (M^n, η) (at least from the perspective of Lorentzian geometry).

To see that condition **N2** is met, can use the results of 5.15, 5.18 and 6.2 to confirm that all of the structures discussed can be reconstructed as an adjunction of their history-relative substructures, and the results of 5.4, 5.11, and 6.7 to conclude that they possess the relevant universal properties.

To complete our argument, we need to confirm that $(M^n_{\mathcal{C}}, \tilde{\eta})$ is compatible with all pre-existing BST92* structure. Fortunately, in BST92* there are no notions pertaining to differentiable structures nor vector bundles. Nevertheless, the elephant in the room is clearly the causal structure of $(M^n, \tilde{\eta})$, i.e. we need to show that the causal relation $\leq^{\tilde{\eta}}$ coincides with the relation $\leq$ defined as in Section 2.2.2.[5] The following result confirms this.

Theorem 6.13. *Let $(M^n_{\mathcal{C}}, \tilde{\eta})$ be an MBST, viewed as an adjoined spacetime. Then $\leq^{\tilde{\eta}}$ coincides with $\leq$.*

Proof. Each $A_{ij} = M^n \backslash J^+(C_{ij})$ is closed under $\leq^M$, thus is a past-set. We know from Lemma 6.10 that $[x,i] \leq^{\tilde{\eta}} [y,j]$ iff $[x,i] \in \phi_i(A_{ij})$ and $f_{ij}(x) \leq^M y$. Since f_{ij} is equal to the inclusion map on M^n, it follows that $[x,i] \leq^{\tilde{\eta}} [y,j]$ iff $[x,i] = [x,j]$ and $x \leq^M y$, which is precisely the definition of $\leq$. □

We also have the following result, which can be seen as an analogue to Proposition 1.5.

Lemma 6.14. *Let $[x,i], [y,j]$ be elements of $M^n_{\mathcal{C}}$ such that $[x,i] \leq [y,j]$. Then every causal curve from $[x,i]$ to $[y,j]$ induces a maximal $\leq$-chain C.*

Proof. Since $[x,i] \leq^{\tilde{\eta}} [y,j]$. Then $x \leq^M y$, so by Prop.1.5 there is a maximal $\leq^M$-chain C connecting x to y. Since ϕ_j acts as an order-isomorphism between M^n and $\phi_j(M^n)$, it follows that $\phi_j(C)$ is a maximal $\leq$-chain in $\phi_j(M_j)$. Thus $\phi_j(C)$ is also a maximal $\leq$-chain in $M^n_{\mathcal{C}}$. □

[5] Recall that the causal order $\leq^{\tilde{\eta}}$ induced from $\tilde{\eta}$ is defined by: $[x,i] \leq^{\tilde{\eta}} [y,j]$ iff there is a future-directed causal curve $\gamma : [0,1] \to M^n_{\mathcal{C}}$ from $[x,i]$ to $[y,j]$.

The above results suggest that the causal structure of $\tilde{\eta}$ is essentially the same as the causal structure $\leq$. Since all of our conditions for naturality are met, we have a strong argument to suggest that the Lorentzian structure defined on $M^n_{\mathcal{C}}$ is indeed a natural extension of the order-theoretic structure of MBSTs, viewed as models of BST92*.

6.2.2 Lorentzian Branching Spacetimes

We will now generalise the approach of the previous section, by defining a new class of models which we will call *Lorentzian Branching Spacetimes*. We start by fixing a Hausdorff spacetime (M, g), and defining what it means for M to have splitting data.

Definition 6.15. *A set* $\mathcal{C} := \{C_{ij} \mid i, j \in I\} \subset \mathcal{P}(M)$ *is called splitting data for* M *iff the set* $\bigcup \mathcal{C}$ *is finite, and every element* C_{ij} *of* $\mathcal{C}$ *satisfies the following conditions:*

- ***C1)*** *For all* $a, b \in C_{ij}$ *it is the case that* $a \not\leq^g b$ *and* $b \not\leq^g a$
- ***C2)*** $C_{ij} = C_{ji}$
- ***C3)*** *For each* $k \neq i, j$, *and for every* $a \in C_{ij}$ *there exists some* $b \in C_{ik} \cup C_{jk}$ *such that* $b \leq^g a$.
- ***C4)*** $C_{ii} = \emptyset$

The above definition is essentially a rehash of the splitting data introduced in Definition 2.11, and as such enjoys the same justification.

Suppose now that we have some fixed splitting data $\mathcal{C}$ for the spacetime (M, g). We can then proceed as in the previous section, and define an adjunction system $\mathcal{F} = (\mathbf{X}, \mathbf{A}, \mathbf{f})$, where:

- every member of $\mathbf{X}$ is a copy of (M, g), viewed as a spacetime,
- each A_{ij} is equal to the set $M \backslash Cl(J^+(C_{ij}))$, viewed as an open Lorentzian submanifold,[6] and
- each $f_{ij} : A_{ij} \to M$ is the inclusion map.

Observe that $\mathcal{F}$ is well-defined. Indeed, that $\mathcal{F}$ is an adjunction system follows from a routine argument – all of the conditions of Definition 3.7 follow from the properties of the splitting data $\mathcal{C}$ and the fact that each f_{ij} is equal to the identity map on M. Moreover, it follows that by construction, $\mathcal{F}$ meets the first two criteria of Theorem 6.6, and the third criterion is met since the inclusion map of an open Lorentzian submanifold always preserves distances and time-orientations.

Since $\mathcal{F}$ is a well-defined adjunction system meeting the criteria of Theorem 6.6, we may conclude that the adjunction space $\bigcup_{\mathcal{F}} M_i$ is an adjoined

[6] The reason we use the closure of $J^+(C_{ij})$ is to guarantee that A_{ij} is open, since in the general setting, it is not guaranteed that causal futures $J^+(x)$ are topologically closed (as an example, see [22, Fig. 11]).

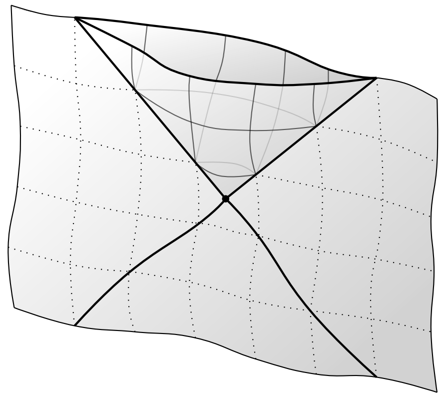

Fig. 6.1: A simple Lorentzian BST constructed in a similar manner to the MBST M_2^n.

spacetime. We will denote the adjoined spacetime $(\bigcup_{\mathcal{F}} M_i, \tilde{g})$ by $M_{\mathcal{C}}$, and we will call the class of all such $M_{\mathcal{C}}$ *Lorentzian Branching Spacetimes* (Lorentzian BSTs, or LBSTs for short). Figure 6.1 depicts the type of branching that LBSTs exhibit.

The following result shows that whenever M is a locally-flat spacetime (that is, whenever it is locally Minkowskian), the Lorentzian BSTs built from M are locally isomorphic to Minkowskian BSTs.

Theorem 6.16. *Let $M_{\mathcal{C}}$ be a Lorentzian BST, where M is a locally-flat spacetime. Then at every point in $M_{\mathcal{C}}$, there is an open neighbourhood that is locally isometric to some Minkowskian BST.*

Proof. Let $[x, i]$ be some element of $M_{\mathcal{C}}$, and consider the element x in (M, g). Since M is locally-Minkowskian, there is some open neighbourhood U of x and an isometric embedding $\varphi : U \to M^n$ that preserves time-orientations. The idea is to look at the fragment of the splitting data $\mathcal{C}$ of M that is contained within U, and then use this to define some new splitting data $\mathcal{D}$ for M^n.

We first need to define the indexing set and splitting data for $M^n_{\mathcal{D}}$. Consider the set

$$J := \{i \in I \mid \exists j \in I(U \cap C_{ij} \neq \emptyset)\}$$

that is, J is the subset of I that has splitting data contained within U. We can then define the splitting data $\mathcal{D}$ on M^n by using the sets

$$D_{ij} := \varphi(C_{ij}) = \{\varphi(a) \mid a \in C_{ij}\}.$$

To see that the collection $\mathcal{D} := \{D_{ij} \mid i, j \in J\}$ defines splitting data for M^n, we need to confirm that the conditions of Definition 2.11 (or 6.15) are satisfied. However, these follow routinely from the fact that $\mathcal{C}$ forms splitting data for M, and the observation that φ, once restricted to its image, is an isometry that preserves time-orientation (which by Lemma A.34 implies that φ can be seen as a "causal isomorphism" between U and $\varphi(U)$).

1. For **C1**, suppose towards a contradiction that there are $a, b \in D_{ij}$ such that $a \leq^M b$. Since $D_{ij} = \varphi(C_{ij})$, and φ acts as a causal isomorphism, it follows that $\varphi^{-1}(a) \leq^g \varphi^{-1}(b)$, which contradicts C_{ij} as spacelike. We may conclude that D_{ij} is also spacelike in M^n.
2. For condition **C2**, we have: $D_{ij} = \varphi(C_{ij}) = \varphi(C_{ji}) = D_{ji}$.
3. For condition **C3**, suppose that $a \in D_{ij}$. Since $D_{ij} = \varphi(C_{ij})$, there is some element $a' \in C_{ij}$ such that $a = \varphi(a')$. Since $\mathcal{C}$ forms splitting data for (M, g), there is some $b' \in C_{ik} \cup C_{jk}$ such that $b' \leq^g a'$. Since the map φ preserves causal orders, it follows that $\varphi(b') \in D_{ik} \cup D_{jk}$ and $\varphi(b') \leq^M a$ as required.
4. For condition **C4**, we have that $D_{ii} = \varphi(C_{ii}) = \varphi(\emptyset) = \emptyset$.

We may then conclude that the set $\mathcal{D}$ forms splitting data for M^n. As such, we can construct the Minkowskian BST subordinate to $\mathcal{D}$, which we will denote by $(M^n_{\mathcal{D}}, \tilde{\eta})$. In order to distinguish the adjunction space $M^n_{\mathcal{D}}$ from the space $(M_{\mathcal{C}}, \tilde{g})$, we will denote the to-be-glued regions of M^n by B_{ij} (recall that $B_{ij} := M^n \backslash J^+(D_{ij})$), and the canonical maps by $\chi_i : M^n \to M^n_{\mathcal{D}}$. We will use the standard notation for (M, g), that is, we denote the canonical maps and to-be-glued regions of (M, g) by ϕ_i and A_{ij}, respectively.

We will now show that the set $\tilde{U} := \bigcup_{j \in J} \phi_j(U)$ is isometric to $M^n_{\mathcal{D}}$. In order to do this, we will define an isometry $\psi : \tilde{U} \to M^n_{\mathcal{D}}$ so that the following diagram commutes.

$$\begin{tikzcd}
(\tilde{U}, \tilde{g}|_U) \arrow[r, dashed, "\psi"] & (M^n_{\mathcal{D}}, \tilde{\eta}) \\
(U, g|_U) \arrow[u, "\cdots \phi_j \cdots"] \arrow[r, "\varphi"'] & (M^n, \eta) \arrow[u, "\cdots \chi_j \cdots"]
\end{tikzcd}$$

We define the map ψ by

$$\psi([y, j]) = \chi_j \circ \varphi \circ \phi_j^{-1}([y, j]) = [\varphi(x), j].$$

To see that ψ is well-defined, suppose we have some element $[z, k]$ such that $[y, j] = [z, k]$. By construction of $M_{\mathcal{C}}$, it must be the case that $y = z$ and $x \in A_{jk}$. We now show that $\varphi(y) \in B_{jk} := M^n \backslash J^+(D_{jk})$. Suppose towards a contradiction that $\varphi(x) \notin B_{jk}$. There there is some element $a \in D_{jk}$ such that $a \leq^M \varphi(y)$. Since $D_{jk} = \varphi(C_{jk})$, the element a must be of the form $a = \varphi(a')$, where $a' \in C_{jk}$. Since φ preserves causal orders, it follows that $a' \leq^g x$, thus

$x \in J^+(C_{jk})$, and consequently, $x \notin A_{jk}$, a contradiction. We may then conclude that $\varphi(x) \in B_{jk}$. It follows that $\psi([z,k]) = [\varphi(z), k] = [\varphi(y), j]$, and thus ψ is well-defined as a function.

To see that ψ is a diffeomorphism, we can pick the open neighbourhoods $\phi_j(U)$ and use Prop. A.7.2 in the obvious way. We now show that ψ is an isometry, that is, for all v, w in $T_{[y,j]}(\tilde{U})$, we show that

$$\tilde{g}|_{\tilde{U}}(v,w) = \tilde{\eta}\left(d\psi)_[y,j](v), \tilde{\eta}(d\psi)_[y,j](w)\right).$$

This follows from the fact that φ and the canonical maps ϕ_j and χ_j act as isometries on their respective domains of definition. Indeed, we have that

$$\begin{aligned}
\tilde{g}|_{\tilde{U}}(v,w) &= (g|_U)\left((d\phi_j^{-1})_{[y,j]}(v), (d\phi_j^{-1})_{[y,j]}(w)\right) \\
&= \eta\left((d\varphi)_y \circ (d\phi_j^{-1})_{[y,j]}(v), (d\varphi)_y \circ (d\phi_j^{-1})_{[y,j]}(w)\right) \\
&= \tilde{\eta}\big((d\chi_j)_{\varphi(y)} \circ (d\varphi)_y \circ (d\phi_j^{-1})_{[y,j]}(v), \\
&\qquad\qquad (d\chi_j)_{\varphi(y)} \circ (d\varphi)_y \circ (d\phi_j^{-1})_{[y,j]}(w)\big) \\
&= \tilde{\eta}\left(d(\chi_j \circ \varphi \circ \phi_j^{-1})_{[y,j]}(v), d(\chi_j \circ \varphi \circ \phi_j^{-1})_{[y,j]}(w)\right) \\
&= \tilde{\eta}\left((d\psi)_{[y,j]}(v), (d\psi)_{[y,j]}(w)\right)
\end{aligned}$$

as required. It follows by a similar argument that the map ψ preserves time-orientations. We may then conclude that $\psi : \tilde{U} \to M^n_{\mathcal{D}}$ is a well-defined, smooth isometry that preserves time-orientations. □

Before moving on, we should remark that in the above theorem, it is not the case that the Lorentzian BST $M_{\mathcal{C}}$ is locally-isometric to some *fixed* Minkowskian BST. Moreover, we can always take the open region U small enough so that the MBSTs are always *simple* MBSTs.

Causal Properties of Lorentzian BSTs

We finish this section by commenting on the potential causal properties of models such as $M_{\mathcal{C}}$. For a first result, we make the interesting observation that in LBSTs, the two definitions of time-orientability once again coincide.

Theorem 6.17. *Let $M_{\mathcal{C}}$ be a Lorentzian BST built from the Hausdorff spacetime (M, g). Then $M_{\mathcal{C}}$ is time-orientable iff there exists a global timelike vector field $\tilde{X}$ on $M_{\mathcal{C}}$.*

Proof. The direction from left-to-right is trivial – we can define the time-orientation τ by picking the pointwise values of $\tilde{X}$. Conversely, suppose that we have a time-orientation τ on $\bigcup_{\mathcal{F}} M_i$. By Theorem 6.6 the canonical maps ϕ_i preserve time-orientations, so by definition this means that the pullback of τ along ϕ_i is equal to τ_i, the time-orientation of (the i^{th} copy of) M. Since M is Hausdorff, we can apply Lemma A.29 to conclude that there is a globally-defined timelike vector field X_i on M. If we do this for every i, we will have a collection of X_i's that will agree on the overlaps A_{ij} (since the time orientations τ_i and τ_j agree on $A_{ij} = A_{ji}$). Thus we can apply 6.5 and push out the vector fields X_i to a vector field $\tilde{X}$ on $M_{\mathcal{C}}$. By Theorem 6.6 the maps ϕ_i preserve time-orientation, so we apply the same idea as in Figure A.4 to conclude that the vector field $\tilde{X}$ will be future-directed. □

We also have the following result, which suggests that a Lorentzian BST $M_{\mathcal{C}}$ may be causally well-behaved, provided M is.

Theorem 6.18. *Let (M, g) be a Hausdorff spacetime, and $M_{\mathcal{C}}$ a Lorentzian BST built from M.*

1. *If M is causal, then so is $M_{\mathcal{C}}$.*
2. *If M has a global time function, then so does $M_{\mathcal{C}}$.*
3. *If the indexing set I is finite and M has compact causal diamonds, then so does $M_{\mathcal{C}}$.*

Proof. If there were some CTC γ in $M_{\mathcal{C}}$, then we could use Lemma 6.10 to conclude that γ lies in some layer $\phi_i(M)$. Then we could pull γ back to M along ϕ_i, and this would be a CTC in M. Items 2 and 3 follow from Lemmas 6.11 and 6.12, respectively. □

We finish this chapter with an observation that is mostly beyond the scope of this thesis. Using item 3 in the above theorem, and the fact that Minkowksi spacetimes are globally-hyperbolic, we can obtain a non-Hausdorff counterexample to Geroch's Splitting Theorem.[7] For a concrete example, any MBST will do, take for instance M_2^2 (pictured as in Figure 2.2) endowed with the appropriate Lorentzian structure. This spacetime is globally-hyperbolic (i.e. it has compact causal diamonds), but it is clearly not topologically-tubular (that is, M_2^2 is not homeomorphic to $\mathcal{S} \times \mathbb{R}$ for any Cauchy surface $\mathcal{S}$). Moreover, such a model would contain *non-homeomorphic* (and thus, non-isometric) Cauchy surfaces.

[7] See A.37 for the statement of this result.

7 Conclusion

In the introduction to this thesis we identified two limitations of Belnap's approach to the unification of relativity and indeterminism. These were:

1. the intended models of BST92, namely the Minkowskian BSTs, are an inadequate representation of the structure of a relativistic model, and
2. BST92 can only deal with special-relativistic branching, and cannot capture the branching of arbitrary spacetimes.

In this thesis we have removed both these limitations by building up to the construction of the Lorentzian BSTs outlined in Section 7.2.2. These branching spacetimes display a rich-enough mathematical structure to be considered by general relativists. Indeed – we know from results like Theorem 6.6 that our LBSTs are smooth manifolds, possessing a Lorentzian metric and a time-orientation; results from Sections 6.2 and 6.3 suggest that the vector bundles over $M_{\mathcal{C}}$ can be naturally constructed by adjoining bundles over M, and Theorem 6.18 suggests that at least some causality conditions are preserved during the construction of $M_{\mathcal{C}}$. The only apparent anomaly is the violation of the Hausdorff property. However, the discussion of Section 4.3.2 suggests that the Hausdorff-violating pairs are intricately connected to the points of indeterminism in $M_{\mathcal{C}}$. The intention of branching is to capture a models indeterminism, so this violation of the Hausdorff property is somewhat justifiable.

We were able to draw such a conclusion by a serendipity – the use of BST92*, together with a type-2 construction of its MBSTs catalysed our development of adjoined spacetimes. It should be clear why we opted for the modified theory BST92* over BST92. In Chapter 4, we saw that the histories of a BST92* model are open in the Bartha topology, and thus we could express the construction of MBSTs as a topological adjunction space in which the histories of a given $M^n_{\mathcal{C}}$ were glued to each other along homeomorphic open subspaces. This approach of gluing along open regions could then be transferred to the Lorentzian manifold setting. Moreover, the results of Chapters 6 and 7 were strongly motivated by the fact that histories of a BST92* model were topologically open. This heavily simplified some potentially problematic arguments, and in particular allowed us to transfer objects defined on histories to the branching spacetime.

In fairness, it may well be possible recreate the results of this thesis in the BST92 setting, however it is not immediately clear how this could be done. For instance, it is known that the histories of a BST92 model are specifically *not* open, so we cannot use the results of Section 3.2.1, and thus would have to develop an even more general theory of adjunction spaces in order to argue that the BST92 Bartha topology were natural. Moreover, if we cannot glue the MBSTs of BST92 along smooth, open submanifolds (i.e. along the intersection of histories), then we could not prove the results of Chapters 6 and 7.

Future Work

The results of thesis suggest a number of interesting developments. We will finish the thesis by highlighting some potential avenues of exploration.

Extending the Branching of LBSTs

Although we have alleviated Belnap's approach of some limitations, there are still some lurking throughout this thesis. In particular, the remarks of 1.14 mean that MBSTs and LBSTs can only have splitting data of *finite* size, and can only have at most *countably-many* branches in total. Recall that we used the finitude of splitting data in order to prove that histories of MBSTs are precisely equal to the layers L_i (see 2.20), and we used that the countability of the indexing set I to ensure that the adjoined spacetimes were second-countable. A next step would be to relax these assumptions, and to see whether the resulting structures are still well-behaved.

Generalising BST92*

We saw in Section 7.2.2 that from a Lorentzian perspective, we can build Lorentzian BSTs from splitting data, in the same way that we did for Minkowskian BSTs. However, in the case of MBSTs, we have an associated logical theory, i.e. BST92*. A pertinent question to ask is: what is the associated logical theory of Lorentzian BSTs? It is probably the case our LBSTs are precisely the class of intended models of Placek's recent theory found in [19], though this needs to be checked. Perhaps a good measure of the correct theory would be to consider the intended class of models that the theory permits, and then to employ some form of metric recovery, and compare the resulting BSTs with our Lorentzian BSTs.[1]

[1] This is a slightly limited approach – typically you can only recover a metric up to a conformal factor, however recent results by Martin/Panangaden [28] suggest that

Establishing the Relationship between BST92 and BST92*

Another interesting avenue of research would be to explore the relationship between models of BST92 and BST92*.[2] We mentioned in Chapter 2 that a useful criterion for motivating the definition of a choice pair would be that, once identified and endowed with an appropriate quotient ordering, the resulting space should be a BST92 model. What will be interesting to see is if there is a general procedure for passing between models of BST92 and BST92*.

Finding the Mathematical Characterisation of MBST Histories

By definition, a history of an MBST $M^n_{\mathcal{C}}$ is a maximal directed subset. We saw in the form of Theorem 2.20 that the histories of $M^n_{\mathcal{C}}$ are precisely the layers $\iota_i(M^n)$. These layers are also equal to the images $\phi_i(M^n)$, which we know from the results of Chapter 4 are open, Hausdorff subspaces of $M^n_{\mathcal{C}}$. A pertinent question is the following – is there a topological property that only the $\phi_i(M^n)$ possess? A result is known in the case of simple MBSTs – Müller has shown that histories of a simple MBST are maximal *H-submanifolds*, see [24, Sec. 3.2.2] for the details. However, it is not immediately clear how such arguments generalise to the setting of arbitrary MBSTs (and I have tried). Taking this idea one step further: what is the Lorentzian characterisation of a history in an LBST? We know that histories *should* equal the images $\phi_i(M)$, however it is not immediately clear what the characterising properties of these subspaces are. For instance – we know as a corollary of Theorem 6.6 that the images $\phi_i(M)$ are open, Hausdorff Lorentzian submanifolds of $M_{\mathcal{C}}$, and they are probably maximal with respect to these properties, but is there more?

Tensor Bundles of Adjoined Manifolds

We saw in the form of Theorem 6.2 that the bundle of bilinear forms of the adjoined manifold $\bigcup_{\mathcal{F}} M_i$ is naturally isomorphic to an adjunction of the bundles $T^2(T^*M_i)$. The crux of the proof was really Lemma A.20. This result is actually an instance of a more general result – if two (Hausdorff) manifolds are diffeomorphic, then their tensor bundles of all orders agree. As such, the construction of Section 7.1.1 easily generalises to the case of arbitrary tensor bundles on $\bigcup_{\mathcal{F}} M_i$.

Defining Global Objects Using Partitions of Unity

We saw in the form of Theorem 6.17 that in Lorentzian BSTs the two notions of time-orientations coincide. This was proved by pulling the time-orientation

---

the full metric can be recovered in the case of a globally-hyperbolic (Hausdorff) spacetime. However, all of this is done with Hausdorff spacetimes, and it is not necessarily true that such a recovery applies in the non-Hausdorff setting.

[2] This was mentioned to me by Thomas Müller in private correspondence.

back to the Hausdorff spacetimes, and then appealing to an argument involving partitions of unity. The approach of 6.17 is actually a special case of something deeper, which could potentially enable the construction of globally-defined objects on adjoined manifolds. As a sketch: we could potentially define an object on adjoined manifold by using partitions of unity on the Hausdorff manifolds M_i. With the extra condition that these M_i-relative objects agree on overlaps, it could then be possible to use something akin to the topologist's Gluing lemma to obtain an object that is globally-defined. This idea is promising in that it would suggest that adjunctions of Hausdorff smooth manifolds are in some sense well-behaved, compared to class of all non-Hausdorff manifolds.

The Consistency of LBSTs and General Relativity

From a coarse perspective, Lorentzian BSTs adhere to the basic mathematical tenets of General Relativity (with the exception of the Hausdorff property, though its violation is linked to the points of indeterminism). However, before any real conclusions can be made about the utility of Lorentzian BSTs, the question of their physical reasonableness must be answered. In particular, there may well be deeper reasons as to why LBSTs are not suitable objects of study for general relativists. A next step will be to evaluate some physically-reasonable conditions in the context of Lorentzian BSTs. We have already made some progression – in Section 7.2.2 we showed that the causal properties of LBSTs *can* be reasonable. However, it follows from a result of Hajicek [29, Thm. 4] that stably-causal LBSTs must contain bifurcating geodesics. This may well activate the criticisms of Earman [5, Sec. 3.4] against non-Hausdorff spacetimes, though this needs to be verified. This is perhaps the most crucial area of future work.

Applications of LBSTs to Quantum Mechanics

A popular criticism (e.g. Earman [5]) of branching spacetimes is that there is currently no dynamical theory for the branching process. Authors such as McCall [3] and Douglas [4] (and probably others) suggest that the branching of spacetime occurs at a measurement of a quantum system, though there is no consensus on these ideas. Of course this is wildly beyond the scope of this thesis, but there are some potentially non-trivial applications of LBSTs to quantum mechanics. The best guess at an application of branching spacetimes would be the *consistent histories approach* introduced by Isham.[3] Müller has made the first steps to combining Isham's approach within a branching framework – in [30] he provides a branching-time interpretation of Isham's approach. A next step would be to use our Lorentzian BSTs as a generalisation of Müller's approach to the spacetime setting.

[3] This is an unfortunate name – here "histories" are not the histories in the BST92(*) sense of the word, but a use that is particular to Isham's formalism.

A

Appendix: Manifolds, Vector Bundles and Spacetimes

In the first part of this thesis, we saw that the models of BST92* can be enriched with a topological structure, namely the Bartha topology. We argued that this topology was natural in the sense of the criteria outlined in Section 4.2. We also saw that the Minkowskian BSTs of BST92* are well-behaved with regards to their Bartha topology. In particular, we expressed their Bartha topologies as adjunctions of the Euclidean topology on Minkowski spacetime. In this part of the thesis, we will extend this idea further, by adjoining Minkowski spacetimes at the level of its smooth structure, and eventually, at the level of the metric structure. The key result of this part will be that Minkowskian BSTs possess a natural pseudometric that is obtained by gluing η appropriately.

Before we do anything of the sort, we will first remind the reader of the precise definition of a spacetime, in the physicist's sense of the word. In short, a *spacetime* is a connected, smooth manifold possessing a Lorentzian metric and a time-orientation. In this chapter, we will slowly unpack the various terms in this definition. Our treatment of spacetimes is all fairly standard, with one exception: we will not assume that our smooth manifolds are Hausdorff. This is because we would eventually like to describe Minkowskian BSTs as spacetimes, and we know from Theorem 4.20 that MBSTs are typically non-Hausdorff. Unfortunately the Hausdorff property is used throughout the standard literature on Lorentzian geometry,[1] so we will have to be very careful when proving results that Hausdorffness is not assumed. As a convention, we will label any results that assume Hausdorffness with a dagger (†).

[1] For instance, the Hausdorff property is used to define *partitions of unity*, which are a tool for constructing global objects from local data.

A.1 Smooth Manifolds

We start by introducing some of the basic theory of smooth manifolds. We will use Lee's *Introduction to Smooth Manifolds* [20] as our resource, and make minor alterations depending on the role of the Hausdorff property. As a convention we will refer to the proofs in [20] where possible, and any facts stated without proof are proved in the Appendix A.2.

A.1.1 Basic Notions

We will call a topological space M a *topological manifold* whenever it is second-countable and locally-Euclidean. It is often useful to equip a topological manifold with a *smooth structure*. This allows the various techniques of calculus (i.e. differentiation, integration etc.) to be generalised to the locally-Euclidean setting. We will start by defining this concept formally.

A topological manifold M is a locally-Euclidean space, so M possesses a collection of charts (U, φ), where φ is a homeomorphism from the open subset U to some open subset of $\mathbb{R}^n$.[2] In the situation that two charts (U, φ) and (V, ψ) overlap, i.e. when $U \cap V \neq \emptyset$, we can use the homeomorphisms φ and ψ to obtain two maps $\varphi \circ \psi^{-1}$ and $\psi \circ \varphi^{-1}$, which are homeomorphisms between open subsets of $\mathbb{R}^n$. Maps of this type are known as *transition maps*, since they allow us to switch between different local representations of M.[3] This motivates the following definition.

Definition A.1. *We say that two charts (U, φ) and (V, ψ) of M are smoothly compatible iff the maps $\varphi \circ \psi^{-1} : \psi(V) \to \varphi(U)$ and $\psi \circ \varphi^{-1} : \varphi(U) \to \psi(V)$ are smooth in the Euclidean sense.*[4]

Smoothly compatible charts then allow us to jump between local representations of objects defined on M without destroying any sense of differentiability. We will refer to a collection $\mathcal{A}$ of charts of M as a *smooth atlas* whenever the charts of $\mathcal{A}$ cover M, and they are pairwise smoothly compatible. Note that

---

[2] Whatever the value of n, it is constant throughout M, and we call this the *dimension of X*.

[3] As an example – suppose we have a curve γ on M, and a point p in M that lies in the intersection of the two charts $U \cap V$. These charts give us two ways of representing γ within $\mathbb{R}^n$, namely the curves $\varphi \circ \gamma$ and $\psi \circ \gamma$. In the overlap $U \cap V$, we can switch between these two local representations by using the transition maps, e.g. $(\varphi \circ \psi^{-1}) \circ (\psi \circ \gamma) = \varphi \circ \gamma$.

[4] Recall that a map $f : \mathbb{R}^n \to \mathbb{R}^n$ is smooth in the Euclidean sense iff the component maps $f^i : \mathbb{R}^n \to \mathbb{R}$ have continuous partial derivatives of all orders. For more on this see [20, App. 3]

the existence and uniqueness of smooth atlases is not guaranteed, and sometimes two smooth atlases can induce the same structure.[5] To avoid this latter issue, smooth structures are defined as follows.

Definition A.2. *A smooth structure on M is a smooth atlas $\mathcal{A}$ that is maximal with respect to inclusion.*

We will call a topological manifold M a *smooth manifold* if it is equipped with a fixed smooth atlas $\mathcal{A}$. We also have the following fact regarding the construction of smooth structures.

Proposition A.3. *Every smooth atlas can be extended to a unique smooth structure.*

We remark that the proof of the above proposition does not require the Hausdorff property (see [20, Prop. 1.17] for the full argument). We finish this section with two useful examples of smooth manifolds.[6]

Example A.4 (Open Submanifolds). Given a smooth manifold M with atlas $\mathcal{A}$, and an open subset U of M, there is a natural way in which U inherits a smooth structure from M. We can define an atlas $\mathcal{A}_U$ by:

$$\mathcal{A}_U := \{(U \cap V, \varphi|_{U\cap V}) \mid (V, \varphi) \in \mathcal{A}\}$$

that is, $\mathcal{A}_U$ is simply the restriction of charts in M to U. One can verify (see e.g. [20, Example 1.26]) that U equipped with this atlas is a smooth manifold of the same dimension as M. The set U, viewed as a smooth manifold, is then called an *open submanifold* of M.

Example A.5 (Product Manifolds). Given m-many smooth manifolds $M_1, ..., M_m$ of dimensions $n_1, ..., n_m$ respectively, there is a natural way in which the Cartesian product $M_1 \times ... \times M_n$ can be endowed with a smooth structure. We can define an atlas by

$$\mathcal{A} = \{(U_1 \times ... \times U_n, \varphi_1 \times ... \times \varphi_n) \mid (U_i, \varphi) \in \mathcal{A}_i\}$$

where $\mathcal{A}_i$ is the smooth structure of M_i, and

$$\varphi_1 \times ... \times \varphi_n : U_1 \times ... \times U_n \to \mathbb{R}^{n_1+...+n_m}$$

maps $(p_1, ..., p_n)$ to the concatenation of the tuple $(\varphi_1(p_1), ..., \varphi_n(p_n))$.

[5] Here by “same structure” we mean that two smooth atlases may determine the same class of smooth, real-valued functions on M.

[6] For more examples of Hausdorff smooth manifolds, see [20, Pg. 17].

A.1.2 Smooth Maps

Suppose we have a map $f : M \to N$ between smooth manifolds M and N. We can express f in terms of local coordinates by using two charts (U, φ) for M at p, and (V, ψ) of N at $f(p)$, and then by looking at the map $\psi \circ f \circ \varphi^{-1}$. Intuitively speaking, we will say that the map f is *smooth* if it is smooth in all of its local expressions. The following is a more-precise formulation.

Definition A.6. *A map $f : M \to N$ between smooth manifolds is smooth at p iff there are charts (U, φ) at p and (V, ψ) at $f(p)$ such that the map $\psi \circ f \circ \varphi^{-1} : \varphi(U) \to \psi(V)$ is smooth in the Euclidean sense. We say that f is smooth iff it is smooth at every p in M.*

It can be shown that this definition of smoothness is independent of the choice of local expression. We will then say that the map f is *smooth* if it is smooth at every point p in M. We have the following useful facts about smooth maps.

Proposition A.7 (Basic facts about smooth maps). *Let $f : M \to N$ be a map between smooth manifolds.*

1. *If f is smooth, then f is continuous.*
2. *If every point p has an open neighbourhood U for which the restricted map $f|_U : U \to N$ is smooth, then f is smooth.*
3. *If f is smooth, then the restriction of f to an open submanifold is smooth.*
4. *If there is a collection of open maps U_α such that each $f_\alpha : U_\alpha \to N$ is smooth and $f_\alpha|U_\alpha \cap U_\beta = f_\beta|_{U_\alpha \cap U_\beta}$, then $f = \bigcup_\alpha f_\alpha$ is a well-defined smooth map from M to N.*

Proof. See [20, Pg.35] and thereafter for the various proofs. □

The following is a collection of maps that are known to be smooth.

Proposition A.8. *Let M, N and P be smooth generalised manifolds.*

1. *The identity map id_M is smooth.*
2. *If U is an open submanifold of M, then the inclusion map $\iota : U \to M$ is smooth.*
3. *If $f : M \to N$ and $g : N \to P$ are smooth, then so is $g \circ f : M \to P$.*
4. *A map $h : M \to N_1 \times ... \times N_k$ is smooth iff the components $h_i := p_i \circ h : M \to N_i$ are smooth.*

Proof. See [20, Prop. 2.10, 2.12]. □

A homeomorphism between topological spaces is a bijective, continuous map whose inverse is also continuous. The analogous notion for smooth manifolds is that of a *diffeomorphism.* A map $f : M \to N$ is called a diffeomorphism if it is a bijective, smooth map whose inverse is smooth. In such a situation, we will call M and N *diffeomorphic.* We finish this section by recalling some useful properties of diffeomorphisms.

Proposition A.9. *Let $f : M \to N$ be a diffeomorphism.*

1. *f is a homeomorphism and an open map.*
2. *If U is an open submanifold of M, then $f|_U : U \to f(U)$ is a diffeomorphism.*
3. *If $g : N \to P$ is a diffeomorphism, then so is the composition $g \circ f : M \to P$.*

Proof. See [20, Prop. 2.15]. □

A.1.3 Tangent Spaces

We will eventually discuss the notion of embeddings of smooth manifolds. But before doing this, we need to introduce the notion of a the tangent space of a manifold. For our purposes, we will never need to appeal to the precise definition of a tangent space, so we will save some time and simply give an intuitive sketch of the idea.

Recall that in the Euclidean setting, the tangent to a curve γ at a point p is obtained by evaluating its derivative at p. Analogously, if the curve γ and the point p live on a smooth manifold M, we can compute the derivative of γ in local coordinates (by using some chart (U, φ) at p) and evaluate it at p to obtain a derivative. The collection of all the derivatives[7] of all the curves passing through p forms the *tangent space* of M at p, which we denote by T_pM. The space T_pM can then be seen as the collection of all velocity vectors of curves passing through p.[8] It turns out that the tangent space T_pM forms a real-valued vector space of dimension equal to $dim(M)$.

We will now introduce the notion of a differential of a smooth map. The idea is that every smooth map $f : M \to N$, induces a map $df_p : T_pM \to T_{f(p)}N$ at every point p in M, as in Figure A.1.[9] We will refer to this map as the *pointwise differential of f at p*. Using this idea, it can be shown (see [20, Pg. 60]) that from a chart (U, φ) at p, the collection of directional derivatives:

$$\left.\frac{\partial}{\partial \varphi^i}\right|_p := (d\varphi_p)^{-1}\left(\left.\frac{\partial}{\partial \varphi^i}\right|_{\varphi(p)}\right) = (d\varphi^{-1})_{\varphi(p)}\left(\left.\frac{\partial}{\partial \varphi^i}\right|_{\varphi(p)}\right)$$

[7] Actually, two different curves might have the same derivative at p, so in practice we would identify them and work with equivalence classes of derivatives of curves.

[8] There are a number of other ways to define the space T_pM. For instance, Lee defines tangent vectors $v \in T_pM$ as *derivations*, which are linear maps $v : C^\infty(M) \to \mathbb{R}$ that satisfy a type of product rule (see [20, Chpt. 3] for more information). It can be shown that both these definitions of T_pM are equivalent, so for the purposes of this thesis it doesn't really matter.

[9] In the velocity-vector construction of T_pM, the differential df_p maps the derivative of the curve γ in M to the derivative of the composite curve $f \circ \gamma$ in N. In the derivation construction of T_pM, the map df_p maps a derivation v to the derivation $df_p(v)$ such that $v(g) = df_p(v)(g \circ f)$.

forms a basis for the vector space T_pM (where here φ^i is the i^{th} component function of a local expression φ of p). We also have the following facts about df_p.

Proposition A.10 (Facts about df_p). *Let $f : M \to N$ and $g : N \to L$ be smooth maps. Then:*

1. $df_p : T_pM \to T_{f(p)}N$ *is a linear map*
2. $d(id_M)_p = id_{(T_pM)}$
3. $d(g \circ f)_p = dg_{f(p)} \circ df_p$
4. *If f is a diffeomorphism then df_p is an isomorphism for each p, and* $(df_p)^{-1} = d(f^{-1})_{f(p)}$

Proof. See [20, Prop. 3.6]. □

We finish this section by remarking that all of this discussion is done at the local level, so tangent spaces and pointwise differentials can also be defined in the case where M is not Hausdorff.

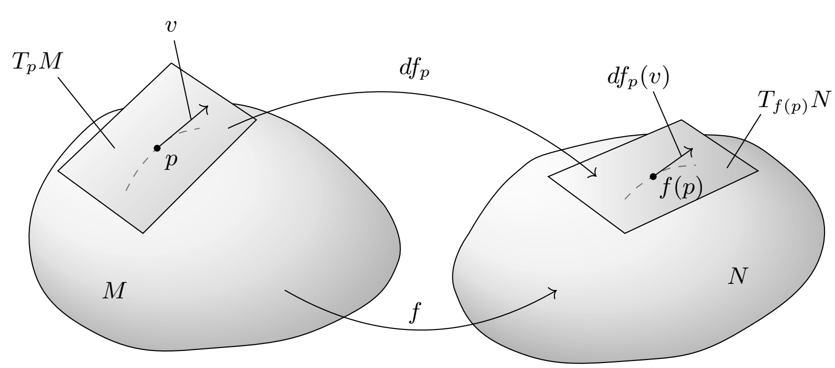

Fig. A.1: The differential of f at p.

A.1.4 Embedded Submanifolds

Now that we have introduced the pointwise differential df_p of a smooth map, we can make precise the notion of a *smooth embedding*.[10] These are defined as follows.

Definition A.11. *A smooth map $f : M \to N$ is called a smooth embedding if it is a topological embedding and the maps df_p are injective for every p in M.*

[10] In most literature, smooth embeddings are usually referred to as simply "embeddings". However, since we are working with embeddings on different structures, we can not use this convention.

Note that the requirement that df_p be injective is needed to ensure that the image $f(M)$ is a smooth submanifold of N. We will call a subset A of a smooth manifold M an *embedded submanifold* whenever A is equipped with the subspace topology and a smooth structure such that the inclusion map $\iota : A \to M$ is a smooth embedding.[11] We finish this section with a characterisation of smooth embeddings.[12]

Lemma A.12. *Let $f : M \to N$ be a smooth map. Then the following are equivalent.*

1. *f is a smooth embedding.*
2. *$f(M)$ is an embedded submanifold of N, and f acts as a diffeomorphism from M to $f(M)$.*

A.2 Vector Bundles

A vector bundle E over a smooth manifold M is a collection of vector spaces E_p, one for each point p in M, that is endowed with a smooth structure, effectively turning E into a smooth manifold in its own right. Each vector space E_p is referred to as a *fibre* of E, and there is a natural projection map $\pi : E \to M$ sending every vector in E_p to p. Vector bundles are comprised of vector spaces, and as such are not only locally-Euclidean, but also locally equal to a product manifold. Figure A.2 depicts the intuition behind a vector bundle.

We will start by introducing the formal definition of a vector bundle, as well as its morphisms and sections. After this, we will discuss some particular bundles used in this thesis – namely the tangent bundle, the bundle of bilinear forms, pullback bundles, and restricted bundles.

A.2.1 Basic Notions

The formal definition of a vector bundle may seem obscure at first, but it essentially all that is required to formally capture the intuition described in Fig. A.2. The definition is as follows.

Definition A.13. *A vector bundle of rank k is a tuple (E, π, M), where E and M are smooth manifolds, and $\pi : E \to M$ is a smooth, surjective map satisfying the following properties.*

[11] It does not have to be the case that A is of the same dimension as M. The difference in dimensions of A and M is known as the *codimension* of A, i.e. $codim(A) = dim(M) - dim(A)$. It should come as no shock that the open submanifolds are the embedded submanifolds of codimension 0 (see [20, Prop 5.1] for the proof of this).

[12] Lemma A.12 is analogous to the topological result that a topological embedding $f : X \to Y$ acts as a homeomorphism from X to $f(X)$.

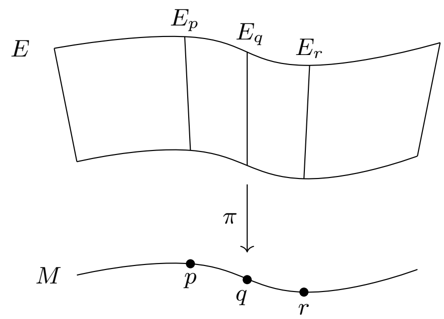

Fig. A.2: A vector bundle E over a smooth manifold M

1. *For each p in M, the pre-image $\pi^{-1}(p) \subset E$ is a k-dimensional real valued vector space, and*
2. *For every p in M there is a neighbourhood U of p and a diffeomorphism $\Phi : \pi^{-1}(U) \to U \times \mathbb{R}$ such that the following diagram commutes (where p_1 is the map that projects onto the first factor),*

$$\begin{tikzcd}
\pi^{-1}(U) \arrow[r, "\Phi"] \arrow[d, "\pi"'] & U \times \mathbb{R}^k \arrow[d, "p_1"] \\
U \arrow[r, equal] & U
\end{tikzcd}$$

and for each $q \in U$, the restriction $\Phi|_q : E_q \to \{q\} \times \mathbb{R}^k$ is an isomorphism of vector spaces.

The map Φ is known as a *local trivialisation*, and allows us to compute a local expression of objects of E in terms of vector spaces. These are analogous to the open charts of a smooth manifold, and there is also an associated notion of compatibility between local trivialisations. Suppose that we have two local trivialisations (U_α, Φ_α) and (U_β, Φ_β) such that the intersection $U_\alpha \cap U_\beta$ is non-empty. We then have the following diagram,

$$\begin{tikzcd}
U_{\alpha\beta} \times \mathbb{R}^k \arrow[d, "p_1"'] \arrow[rr, bend left, dashed, "\Phi_\beta \circ \Phi_\alpha^{-1}"] & \pi^{-1}(U_{\alpha\beta}) \arrow[l, "\Phi_\alpha"'] \arrow[r, "\Phi_\beta"] \arrow[d, "\pi"] & U_{\alpha\beta} \times \mathbb{R}^k \arrow[d, "p_1"] \\
U_{\alpha\beta} \arrow[r, equal] & U_{\alpha\beta} \arrow[r, equal] & U_{\alpha\beta}
\end{tikzcd}$$

where we have used $U_{\alpha\beta}$ as a shorthand for the intersection $U_\alpha \cap U_\beta$. Since both Φ's are diffeomorphisms, the map $\Phi_\beta \circ \Phi_\alpha^{-1}$ is well-defined, and sends $(u, r) \mapsto (u, g_{\alpha\beta}(u)(r))$, where the map $g_{\alpha\beta}$ maps each element u of $U_{\alpha\beta}$ to some linear transformation[13] on $\mathbb{R}^k$. The maps $g_{\alpha\beta}$ are known as *transition functions*, and are similar to the transition maps of compatible charts of a smooth manifold (cf. Def. A.1). We also have the following useful result, which allows to construct vector bundles from local data.

Lemma A.14 (Bundle Chart Lemma). *Let M be a smooth manifold, and suppose that for each p in M we are given a real-valued vector space E_p of some fixed dimension k. Let $E := \bigsqcup_p E_p$, and let $\pi : E \to M$ be the map that takes each element of E_p to the point p. Suppose furthermore that we are given the following data:*

1. *an open cover $\{U_\alpha\}_{\alpha \in A}$ of M*
2. *for each $\alpha \in A$, a bijection $\Phi_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^k$ whose restriction to each E_p is a vector space isomorphism from E_p to $\{p\} \times \mathbb{R}^k$*
3. *for each $\alpha, \beta \in A$ with $U_\alpha\beta \neq \emptyset$, a smooth map $g_{\alpha\beta} : U_{\alpha\beta} \to GL_k(\mathbb{R})$ such that the map $\Phi_\alpha \circ \Phi_\beta^{-1}$ from $U_{\alpha\beta}$ to itself has the form $\Phi_\alpha \circ \Phi_\beta^{-1}(p, v) = (p, g_{\alpha\beta}(p)(v))$.*

Then E has a unique topology and smooth structure making it into a vector bundle over M, with π as its projection map, and (U_α, Φ_α) as local trivialisations.

Proof. See [20, Lem. 10.6] for the full proof. Note that nowhere is the Hausdorff property used (except when it is shown that E is Hausdorff, provided M is). □

We will now define the notion of a section of a bundle. This is fairly straightforward – a section of a vector bundle E is formed by smoothly choosing a single element from each fibre E_p, as in Figure A.3. Formally, sections are defined as follows.

Definition A.15. *Let E be a vector bundle over M. A section is a smooth map $s : M \to E$ such that $\pi \circ s = id_M$. We denote the space of all sections of E by $\Gamma(E)$.*

A section is essentially a smooth right-inverse of the projection map π. This means that it assigns to every p in M some vector v in E_p, in a smooth manner. Sections are pretty useful objects, we will see that they can be used to define smooth *fields* of objects on a manifold.

[13] Actually, it is a map $g_{\alpha\beta} : U_{\alpha\beta} \to GL_k(\mathbb{R})$, where $GL_k(\mathbb{R})$ is the *general linear group of degree k* over $\mathbb{R}$. This is the set of all invertible $k \times k$ matrices with entries in $\mathbb{R}$, and is a Lie group. The Lie group associated to the maps $g_{\alpha\beta}$ is known as the *structure group* of the bundle E.

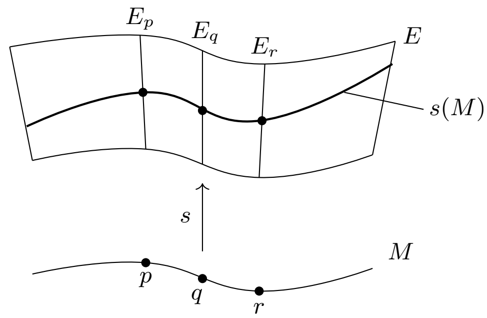

Fig. A.3: A smooth section s of E

We complete this section by discussing what it means for a map between bundles to be structure-preserving. Such maps are called *bundle morphisms*, and are defined as follows.

Definition A.16. *Let $f : M \to N$ be a smooth map, and let (E, π_E, M) and (F, π_F, N) be vector bundles. A smooth map $g : E \to F$ is said to be a bundle morphism covering f iff the diagram*

$$\begin{tikzcd} E \arrow[r, "g"] \arrow[d, "\pi_E"'] & F \arrow[d, "\pi_F"] \\ M \arrow[r, "f"'] & N \end{tikzcd}$$

commutes, and g is a fibrewise linear map.

Observe that a bundle morphism is a smooth map (i.e. a morphism of manifolds) and a linear map (i.e. a morphism of vector spaces) once restricted to fibres. The requirement that the map covers f amounts to requiring that g maps each fibre E_p to $F_{f(p)}$. We say that g is a *bundle isomorphism* iff it is bijective, and its inverse map g^{-1} is also a bundle morphism. We have the following useful condition for identifying bundle isomorphisms.

Lemma A.17. *Let E and F be vector bundles over the same base manifold M, and $f : E \to F$ be a bundle morphism. If f is bijective, then f is a bundle morphism.*

Proof. Lee leaves this as an exercise (see [20, Prop. 10.26]), and the full proof can be found in [31, Lem. 5.3.12]. □

A.2.2 The Tangent Bundle

We saw earlier that the tangent space to M is well-defined at every point. It is possible to define a vector-bundle structure on the set

$$TM := \bigsqcup_{p \in M} T_p M.$$

Lee proves this in the Hausdorff case (see [20, Prop. 3.18]), though at no point is the Hausdorff property used in the proof (except when it is shown that TM is Hausdorff whenever M is). As such, tangent bundles can still be constructed without assuming the Hausdorff property.

We should comment on the local trivialisations of the bundle TM. We saw in Section 5.1.3 that the tangent spaces T_pM have as bases the directional derivatives of the components of a coordinate representation. This allows us to express the local trivialisations as follows. Suppose we have some point p in M and a chart (U, φ) for M at p, and consider the coordinate functions $\varphi^i : U \to \mathbb{R}$. The map $\Phi : \pi^{-1}(U) \to U \times \mathbb{R}^n$ is defined by

$$\Phi(v) = \Phi\left(v_i \frac{\partial}{\partial \varphi^i}\right) = (p, v_1, ..., v_n).$$

In fact, we can go one step further and use the map Φ to describe a chart for TM using $\pi^{-1}(U)$, by mapping $(p, v_i \frac{\partial}{\partial \varphi^i}) \mapsto (\varphi^1(p), ..., \varphi^n(p), v_1, ..., v_n)$.

It is also possible to define the differential of a smooth map $f : M \to N$ using the pointwise differentials df_p. The differential of f is denoted by df, and is defined as

$$df(p, v) = (f(p), df_p(v)).$$

We have the following useful facts about the differential map.

Proposition A.18. *Let $f : M \to N$ and $g : N \to L$ be smooth maps. Then*

1. *$df : TM \to TN$ is a bundle morphism covering f*
2. *$d(g \circ f) = dg \circ df$*
3. *If f is a diffeomorphism then df is a bundle isomorphism, with $(df)^{-1} = d(f^{-1})$.*
4. *If $f : M \to N$ is a diffeomorphism then $df : TM \to TN$ is a bundle isomorphism.*

Proof. See [20, Prop. 2.12] and thereafter. □

A.2.3 The Bundle of Bilinear Forms

Recall that a bilinear form on a real-valued, k-dimensional vector space V is a binary map $\alpha : V \times V \to \mathbb{R}$, such that:

$$\alpha(u + v, w) = \alpha(u, w) + \alpha(v, w)$$

$$\alpha(u, v + w) = \alpha(u, v) + \alpha(u, w)$$

$$\alpha(r \cdot u, v) = r \cdot \alpha(u, v) = \alpha(u, r \cdot v)$$

for every $u, w, v \in V$ and $r \in \mathbb{R}$. We will denote the space of all bilinear forms on V by $T^2(V^*)$.[14] The space $T^2(V^*)$ can be turned into a vector space using the operations of pointwise addition and scalar multiplication of functions.[15] The dimension of the space $T^2(V^*)$ is equal to k^2, with a natural basis given by the collection $\{\epsilon_i \otimes \epsilon_j \mid i, j \leq n\}$, where $\{\epsilon_i\}_{i=1}^n$ forms a basis for the dual space V^*.[16]

In order to eventually define what it means for a smooth manifold M to be a spacetime, we will need to discuss the bilinear forms built from elements TM. In this case, we can use the fibres $T^2(T_p^*M)$ and apply the Bundle chart lemma to equip the set

$$T^2(T^*M) := \bigsqcup_{p \in M} T^2(T_p^*M)$$

with a vector bundle structure. The exact details of this construction are not too important for our purposes, though we will remark that the local trivialisations of $T^2(T^*M)$ are given by:

$$\Phi(p, \alpha) = \Phi(p, r_{ij} d\varphi^i \otimes d\varphi^j) = (p, r_{11}, ..., r_{1n}, ..., r_{n1}, ..., r_{nn}),$$

where we here we denote by $d\varphi^i$ the dual-basis element associated to the directional derivative $\partial/\partial\varphi^i$. Local charts of the bundle $T^2(T^*M)$ are then given by:

$$\pi^{-1}(U) \mapsto U \times \mathbb{R}^{n^2} \mapsto \varphi(U) \times \mathbb{R}^{n^2}$$

in the obvious way, and the resulting space $T^2(T^*M)$ becomes a vector bundle of rank n^2 over M.[17]

Given a smooth map f between smooth manifolds M and N, there is a natural sense in which (some of) the bilinear forms on N can be transferred to M. Suppose we have some element of the form $(f(p), \alpha)$ in $T^2(T_{f(p)}^*N)$. We can then define the object $f^*\alpha$ in $T^2(T_p^*M)$ by:

$$f^*\alpha(v, w) = \alpha(df_p(v), df_p(w))$$

for all $v, w \in T_pM$. Observe that since f is smooth, by Proposition A.10.4 the images of v and w under the differential df_p are indeed elements of $T_{f(p)}N$, so this definition makes sense. Moreover, the form $f^*\alpha$ inherits its bilinearity from α, so $(p, f^*\alpha)$ lies in the fibre $T^2(T_p^*M)$. We will refer to the object $f^*\alpha$ as the (pointwise) pullback of α under f.[18]

---

[14] This space is also commonly denoted as $V^* \otimes V^*$, or $T^{(0,2)}V$, or $T_2^0(V)$.

[15] That is, we define $(\alpha+\beta)(v, w) := \alpha(v, w) + \beta(v, w)$ and $(r \cdot \alpha)(v, w) = r \cdot \alpha(v, w)$.

[16] The dual space for a vector space V is the set of all linear maps $\alpha : V \to \mathbb{R}$. If $\{e_i\}$ is a basis for V, then the collection $\{\epsilon_i\}$ of maps defined by $\epsilon_i(e_j) = 1$ iff $i = j$ and $\epsilon_i(e_j) = 0$ iff $i \neq j$ forms a basis for V^*.

[17] When viewed as a smooth manifold, $T^2(T^*M)$ is of dimension $n + n^2$.

[18] There is a sense in which fields of bilinear forms (and in general, covariant tensor fields) can be pulled back along a smooth map – see [20] Chpt. 12 for more details.

There is also a sense in which the pointwise pullback has a restricted dual. In the case where M and N have the same dimension and f is a smooth embedding (which by Lemma A.12 means that M and $f(M)$ are diffeomorphic), we can also push bilinear forms of M forward into N. Given some $(p, \beta) \in T^2(T_p^*M)$, we can define the map $f_*\beta$ to act as:

$$f_*\beta(v, w) = \beta(v', w'), \text{ where } v = df_p(v') \text{ and } w = df_p(w').$$

Observe that whenever f is a diffeomorphism, the pointwise differential df_p is a bijection (see Prop. A.10.4), and since $f(M)$ is a submanifold of the same dimension as N, we can use the maps as in the proof of Lemma A.12 to conclude that every element of the tangent space $T_{f(p)}N$ is of the form $df_p(v)$ for some vector v in T_pM, and thus $f_*\beta$ is a well-defined bilinear form that lives in the fibre $T^2(T^*_{f(p)}N)$. We have the following facts about pointwise pullbacks and pushforwards.

Lemma A.19. *Let $f : M \to N$ and $g : N \to P$ be smooth embeddings, with $dim(M) = dim(N) = dim(P)$, and let α, β and δ be elements of the fibres $T^2(T_p^*M), T^2(T^*_{f(p)}N)$ and $T^2(T^*_{g\circ f(p)}P)$ respectively.*

1. $(id_M)_*\alpha = (id_M)^*\alpha = \alpha$
2. $f_*\alpha = (f^{-1})^*\alpha$
3. $f^*\beta = (f^{-1})_*\beta$
4. $(f^* \circ g^*)\delta = (g \circ f)^*\delta$
5. $f^*(f_*\alpha) = \alpha$
6. $g_* \circ f_*(\alpha) = (g \circ f)_*\alpha$

We saw in Proposition A.18 that whenever $f : M \to N$ is a smooth map, the map df (defined fibrewise as the pointwise differential maps df_p) becomes a smooth bundle morphism from TM to TN. Similarly, we can also define a map that acts on fibres as the pointwise pullback (pushforward) does. Just as the differential map associated to a diffeomorphism is a bundle isomorphism between tangent bundles, there is an associated result for the bundle of bilinear forms.

Lemma A.20. *If $f : M \to N$ is a diffeomorphism, then the bundles $T^2(T^*M)$ and $T^2(T^*N)$ are isomorphic.*

We remark that in the proof of the above result, the isomorphism is given by the map $\xi : T^2(T^*M) \to T^2(T^*N)$ defined by $\xi^{-1}(p, \alpha) = (f(p), f_*\alpha)$. This will be important in Section 7.1, when we adjoin bundles of bilinear forms.

A.2.4 Restrictions and Pullbacks

We saw in the previous section that in some cases, it is possible to transfer some objects from bundle-to-bundle along a smooth map between the base manifolds. This idea can be taken a step further, and we can actually transfer (part of a) vector bundle's structure along a smooth map. Given a smooth map $f : M \to N$ and a vector bundle (F, π_F, N), we can define a new bundle over M by restricting the structure of F to the fibres that cover $f(M)$. This is called a *pullback bundle*, and is commonly denoted by f^*F. The fibres are given by $(f^*F)_p = F_{f(p)}$. The formal definition is as follows.

Definition A.21. *Let $f : M \to N$ be a smooth map and (F, π_F, N) a vector bundle. The pullback bundle has as elements:*

$$f^*F := \{(m, u) \in M \times F \mid f(m) = \pi_F(u)\},$$

and the projection map of f^*F is given by the projection onto the first factor.

We also make the following useful observation.

Proposition A.22. *Let $f : M \to N$ be a smooth map, and (F, π_F, N) a vector bundle. The map $p_2 : f^*F \to F$ is a bundle morphism, and p_2 is an isomorphism whenever f is a diffeomorphism.*

It can also be shown (see Prop. A.14) that every bundle morphism $g : E \to F$ that covers the map f has to factor through the pullback bundle f^*F, i.e. for every such g there is a unique morphism $h : E \to f^*F$ such that the following diagram commutes.

$$\begin{tikzcd}
E \arrow[r, dotted, "h"] \arrow[rr, bend left, "g"] \arrow[d, "\pi_E"'] & f^*F \arrow[r, "p_2"] \arrow[d, "p_1"'] & F \arrow[d, "\pi_F"] \\
M \arrow[r, equal] & M \arrow[r, "f"'] & N
\end{tikzcd}$$

Given a vector bundle (E, π, M) and an embedded submanifold A of M, we can restrict the bundle E to A by pulling E back along the inclusion map $\iota : A \to M$ (which is smooth by definition). We will denote this bundle by $E|_A$, instead of ι^*E. The bundle $(E|_A, \pi|_A, A)$ is then called a *restricted bundle*. In the special case that A is an open submanifold of M, we have the following useful observation.

Lemma A.23. *If A is an open submanifold of M and E a vector bundle over M, then the restricted bundle $E|_A$ is an open submanifold of E.*

The next result will be needed in Chapter 6, when we glue tangent bundles together.

Proposition A.24 (†). *If M is Hausdorff and A an open submanifold, then TA is isomorphic to $TM|_A$.*

Proof. The full argument can be found as [20, Prop. 3.9], and is proved using partitions of unity, so we will not assume that the same result holds in the non-Hausdorff case. □

We will also follow the convention of Lee and identify the bundles TA and $TM|_A$ whenever A is an open submanifold of M (and when M and A are Hausdorff, of course). Our final result of this section is a generalisation of Lemma A.17.

Lemma A.25. *If $g : E \to F$ is a bijective bundle morphism covering a diffeomorphism $f : M \to N$, then g is a bundle isomorphism.*

A.3 Lorentzian Manifolds and Spacetimes

We complete this chapter by defining just what it means for a smooth manifold to be a spacetime. We will start by defining *Lorentzian metrics* on a manifold. After this we will define time-orientations and spacetimes, and finally we will discuss some elements of causality theory.

A.3.1 Lorentzian Manifolds

We will first define what it means for a smooth manifold to be *Lorentzian.* In order to do this, we need to define a pseudometric field on M which locally mimics the behaviour of the pseudometric η. We saw back in Section 1.1 that η is a symmetric, non-degenerate bilinear form that acts on elements x and y of M^n as

$$\eta(x, y) = -x_0y_0 + x_1y_1 + ... + x_{n-1}y_{n-1}.$$

Although we didn't mention it at the time, the metric η is actually acting on the tangent vectors of M^n, but since each tangent space $T_x\mathbb{R}^n$ is canonically isomorphic to $\mathbb{R}^n$, the identification $T_x\mathbb{R}^n \cong \mathbb{R}^n$ is implicitly made. Observe that the pseudometric η has a signature of $(-, +, ..., +)$.[19]

A Lorentzian metric is a generalisation of η to the smooth-manifold setting. Intuitively speaking, Lorentzian metrics are defined to be a smooth *field* of bilinear forms, that possess the same local properties as η. Formally, Lorentzian metrics are defined as follows.

Definition A.26. *A Lorentzian metric on a smooth manifold M is a section g of the bundle $T^2(T^*M)$ that is everywhere symmetric and non-degenerate, with signature* $(-, +, ..., +)$.

[19] Recall that the signature of a matrix B is the tuple consisting of the signs of the non-zero entries of the diagonalised version of B. In the case of η, the matrix representation is

$$\eta = \begin{pmatrix} -1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}$$

and thus the signature of η is $(-, +, ..., +)$.

It can be shown that the signature of a Lorentzian metric g is constant throughout the manifold.[20] We will typically denote Lorentzian manifolds by (M, g), though we will abbreviate this to simply M where the context is clear. We will also causally refer to the metric g as the "Lorentzian structure" of M.

We will now discuss the embeddings of the Lorentzian world. For a map $f : M \to N$ between Lorentzian manifolds to be structure preserving, it has to preserve topological and smooth structures of M, as well as the behavior of the Lorentzian metric g. Such maps are called *isometric embeddings*, and they are defined as follows.

Definition A.27. *Let (M, g^M) and (N, g^N) be Lorentzian manifolds. A map $f : M \to N$ is called an isometric embedding iff f is a smooth embedding and additionally*

$$g^M(v, w) = g^N(df_p(v), df_p(w))$$

for every v, w in T_pM and p in M.

In the case that f is a diffeomorphism from M to N, then f is called an *isometry*, and we will refer to M and N as *isometric*. We finish this section with the following observation, which will be useful in Chapter 7.

Lemma A.28. *Isometries preserve the signature of metrics.*

Proof. Suppose that $f : M \to N$ is an isometry, and let (U, φ) be a chart of M at p. As in the proof of Lemma A.20, since f is a diffeomorphism we can always pick a local representation of g^M and g^N that have the same coefficient matrix. Since the signature of these metrics is independent from the choice of local expression, we are done. □

A.3.2 Time-Orientations

We will now define what it means for a Lorentzian manifold to be *time-oriented*. We start by fixing a Lorentzian manifold (M, g^M), and an element p in M. In a manner similar to Definition 1.1, we define a tangent vector v in T_pM to be:

- *timelike* iff $g^M(v, v) < 0$,
- *lightlike* iff $g^M(v, v) = 0$, and
- *spacelike* iff $g^M(v, v) > 0$.

We can then define two *timecones* of T_pM, by considering an arbitrary timelike vector v and defining the sets:

$$C(v) := \{u \in T_pM \mid g^M(u, u) < 0 \text{ and } g^M(v, u) < 0\}$$

[20] This is a consequence of *Sylvester's Theorem*, which says that the signature of a matrix is invariant under change of basis. See [20, Pg. 343] for more.

$$C(-v) := \{u \in T_pM \mid g^M(u,u) < 0 \text{ and } g^M(v,u) > 0\}$$

It can be shown (e.g. [32, Pg. 143-5]) that the set of timelike vectors of T_pM is equal to the disjoint union of these two sets.[21]

Denote by M^T the collection of all timecones of all elements of M, and suppose we have some map $\tau : M \to M^T$ such that for each p in M, the map τ chooses one of the timecones of T_pM. We will call this choice *smooth* if for every $p \in M$, there is some neighbourhood U and a vector field X on U such that $X(q) \in \tau(q)$ for every $q \in U$. A *time-orientation* is then defined to be a smooth choice of timecone at each point p in M. We also have the following equivalent characterisation.

Lemma A.29 (†). *Let M be a Hausdorff manifold. Then M is time-oriented iff there exists a globally-defined timelike vector field X on M.*

Proof. See [32, Lem. 32] for the proof. The direction from right-to-left is straightforward, but the other direction is an argument involving a partition of unity, so the same result may not hold in general for non-Hausdorff manifolds. □

Given the above result, some authors simply *define* a time-orientation to be the globally-defined timelike vector field. However in the non-Hausdorff case, it might not true that such an equivalence holds. This will cause a slight inconvenience in Section 7.1 when we define a time orientation on our adjoined manifolds.

One might be wondering how isometries interact with time-orientations. It is not true in general that isometries preserve time-orientations.[22] This motivates the following definition.

Definition A.30. *Let $f : M \to N$ be an isometry between time-oriented Lorentzian manifolds. We say that f preserves time-orientation if for every p in M, the differential df_p maps $\tau_M(p)$ to $\tau_N(f(p))$.*

Remark A.31 (†). Suppose we have two time-oriented, Hausdorff Lorentzian manifolds M and N, with time-orientations τ_M and τ_N respectively. We know from Lemma A.29 that a time-orientation can be characterised by a globally-defined timelike vector field, so there are vector fields X on M and Y on N characterising τ_M and τ_N resp. If $f : M \to N$ preserves time-orientation, the differential df maps $\tau_M(p)$ to $\tau_N(f(p))$. Since f is a diffeomorphism, we can define the pushforward f_*X by:

[21] We have seen this before – in the case of Minkowski spacetime, the timecone is simply the interior of the lightcones pictured in Fig. 1.1.

[22] As a counterexample, consider the map that takes each element $(x_0, x_1, ..., x_{n-1})$ of M^n to $(-x_0, x_1, ..., x_{n-1})$. This map is clearly an isometry, but it reverses the time-orientation of M^n.

$$f_*X(p) = df_p(X(p)),$$

and this is a globally-defined vector field on N. Since f is an isometry, f_*X is timelike in N. So, if f preserves time-orientation, then at every point in N, the pushforward f_*X lies in the same timecone as Y. Figure A.4 depicts this idea.

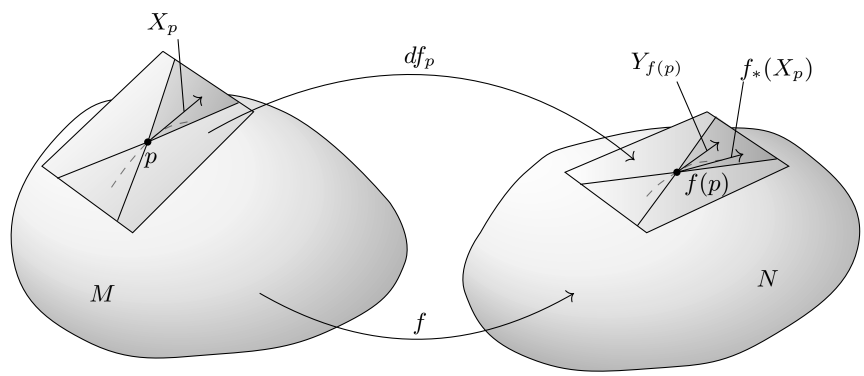

Fig. A.4: An isometry $f : M \to N$ that preserves time-orientation. At every point p, the vector $f_*(X_p)$ lies in the same timecone as $Y_{f(p)}$.

We are now in the position to fix a definition of a *spacetime*.

Definition A.32. *A spacetime is a Lorentzian manifold* (M, g) *that is connected and time-oriented.*

The requirement that M be connected is used as an enjoyable convention, and we see no reason to neglect it. Before moving on to causal properties of spacetimes, we will briefly mention what it means for a spacetime to be *locally flat*. Roughly speaking, in the same way that smooth manifolds are locally Euclidean, locally flat spacetimes are Lorentzian manifolds that are locally *Minkowskian*, as pictured in Figure A.5. This means that at every point p in a locally-flat spacetime (M, g), there is a chart (U, φ) such that $\varphi : U \to M^n$ is an isometric embedding that preserves time-orientations.[23]

Elements of Causality Theory

We will finish this chapter with a brief introduction of *causality theory*, which deals with the causal properties of spacetimes. Of course, there is a large amount that we simply cannot cover – the interested reader is invited to look

[23] This is found as [33, Def. 2.1], and is defined for Hausdorff spacetimes, however the same definition also works for our not-necessarily-Hausdorff spacetimes.

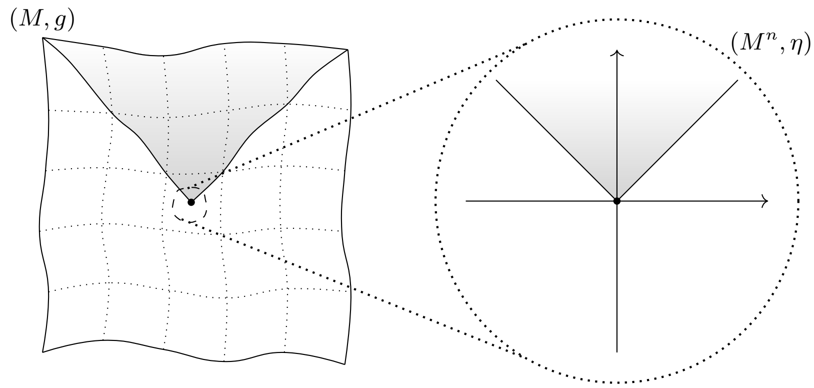

Fig. A.5: The intuition behind a locally-flat spacetime (M, g).

at the excellent survey by Minguzzi and Sanchez [23] for more information on causality theory. Most of our definitions are taken directly from this paper, and some from the first few chapters of [22]. Throughout this section we will assume that (M, g) is a spacetime.

We start by defining the causal orderings of (M, g), as we did for Minkowski spacetimes in Section 1.1. Given two points p and q in M, we say that p *causally (temporally) precedes* q, written $p \leq^g q$ ($p \ll^g q$), if there is a future-directed, causal (timelike) curve γ connecting p to q. Clearly $p \ll^g q$ implies that $p \leq^g q$. We can then define the causal future/past of a point/set as in Definition 1.2.

Definition A.33. *Let p be a point in M, and $S \subseteq M$.*

- *The causal future of p is the set $J^+(p) := \{q \in M \mid p \leq^g q\}$,*
- *The temporal future of p is the set $I^+(p) := \{q \in M \mid p \ll^g q\}$,*
- *The horismotic future of p is the set $E^+(p) := J^+(p) \backslash I^+(p)$,*
- *The causal future of S is the set $J^+(S) := \{q \in M \mid \exists p \in S(p \leq^g q)\}$,*
- *The temporal future of S is the set $I^+(S) := \{q \in M \mid \exists p \in S(p \ll^g q)\}$,*
- *The horismotic future of S is the set $E^+(S) := J^+(S) \backslash I^+(S)$.*

We define the causal, temporal and horismotic past analogously, and denote them by $J^-(\cdot), I^-(\cdot)$ and $E^-(\cdot)$ respectively.

We will now show that any isometry $f : M \to N$ that preserves time orientation will also preserve the causal orders of M and N.

Lemma A.34. *If an isometry preserves time-orientation, then it preserves the causal ordering.*

Proof. let $f : M \to N$ be an isometry that preserves time-orientation, and suppose that $p \leq^M q$. Then there is a future-directed, timelike curve γ :

$[0,1] \to M$. Thus $f \circ \gamma : [0,1] \to N$ is a curve connecting $f(p)$ to $f(q)$. Since f preserves time-orientation, the curve $f \circ \gamma$ is future-directed. Since f is an isometry, $f \circ \gamma$ is also timelike. Indeed, let $f(r) \in f \circ \gamma$. Then:

$$g^N((f \circ \gamma)'(t), (f \circ \gamma)'(t)) = g^N(df(\gamma'(t)), df(\gamma'(t))) = g^M(\gamma'(t), \gamma'(t))$$

A symmetric argument proves the other direction. □

We will finish our discussion by introducing some causally well-behaved spacetimes. There are a number of conditions one might impose on a class of spacetimes in order to manage causal properties. These are known as *causality conditions*, and as it turns out these conditions form a hierarchy. We will focus on three causality conditions, namely causality, stable causality and global hyperbolicity.

We say that a spacetime (M, g) is *causal* if it contains no closed timelike curves.[24] Observe that this is equivalent to asserting that the causal order $\leq^g$ is anti-symmetric. Roughly speaking, a spacetime (M, g) is *stably causal* iff it is causal, and for every Lorentzian metric g' "close" to g, the spacetime (M, g') is also causal.[25] This means that the lightcones of a stably-causal spacetime can be continuously deformed by a small amount, and the resulting cones will still contain no CTCs. The precise formulation of stable causality is not too important for us, so we will focus on the characterising property that stably-causal spacetimes possess a *global time function.* These are defined as follows.

Definition A.35. *Let (M, g) be a spacetime and $f : M \to \mathbb{R}$ a smooth function. Then f is called a global time function if it is monotone with respect to the temporal ordering $\ll^g$, that is, if $p \ll^g q$, then $f(p) \leq f(q)$ in $\mathbb{R}$.*

It can be show (see [23, Thm. 3.56]) that in the Hausdorff setting, stable causality is equivalent to the existence of a global time function.[26] Our final causality condition is global hyperbolicity, which is defined as follows.

Definition A.36. *A spacetime (M, g) is called globally-hyperbolic iff it is causal, and the causal diamonds $J^+(p) \cap J^-(q)$ are all compact.*

Famously, Geroch [34, Thm. 11] showed that global hyperbolicity is equivalent to the existence of a *Cauchy surface*, which is a spacelike hypersurface that

[24] Here by "closed" we mean a curve $\gamma : [0,1] \to M$ whose endpoints are mapped to the same point p in M.

[25] This proximity of Lorentzian metrics is defined by topologizing the space $Lor(M)$ of Lorentzian metrics on M.

[26] Actually, Minguzzi/Sanchez prove this result for "temporal functions", which are defined to be smooth functions whose gradient is everywhere past-directed and timelike (see [23, Def 3.48]). However, as the remarks after Def 3.48 confirm, these temporal functions are also global time functions.

is deterministic in the sense that from the initial conditions on the Cauchy surface, one can determine the past and future of the whole spacetime uniquely. We finish this chapter with a celebrated result regarding global hyperbolicity, namely Geroch's *Splitting Theorem*, which connects the causal structure of M to its global topological structure (see [34, Prop. 7] for the proof).

Theorem A.37 (Geroch's Splitting Theorem). *Let (M, g) be a globally-hyperbolic (Hausdorff) spacetime. Then M is homeomorphic to the product $\mathcal{S} \times \mathbb{R}$, where $\mathcal{S}$ is a Cauchy surface of M.*

B

Appendix: Proofs

B.1 Proofs from Section 3.1

B.1.1 Proof of Lemma 3.4

Proposition B.1. *ϕ_X is injective iff f is injective.*

Proof. Suppose that f is injective, and let $x, x' \in X$ be distinct. Without loss of generality we can assume that x, x' are elements of A, since otherwise the result is trivial[1]. Since f is injective it has a right inverse, i.e. a map $f^{-1} : f(A) \to A$ such that $f^{-1}(f(a)) = a$ for each a in A. Following the description of the equivalence classes given above, for each $a \in A$, the mapping under ϕ_X is given by:

$$\phi_X(a) = [a, 1] = \{(a, 1), (f(a), 2)\}$$

It is not hard to see that any distinct element a' cannot be mapped to this equivalence class.

For the converse direction, suppose that f is *not* injective. Then there is some a, a' such that $f(a) = f(a')$. Hence:

$$\phi_X(a) = [a, 1] = \{(a, 1), (a', 1), (f(a), 2)\} = [a, 2] = \phi_X(a')$$

It follows that ϕ_X is also not injective. By contraposition our result follows.

$\square$

$\square$

Lemma B.2. *If A is open in X then ϕ_Y is an open map.*

Proof. Let U be open in Y. From the previous proposition we know that $\phi_Y(U)$ is open in $X \cup_f Y$ iff both $\phi_X^{-1}(\phi_Y(U))$ and $\phi_Y^{-1}(\phi_Y(U))$ are open in X and Y respectively. The latter is open, since ϕ_Y is always an injection, we

[1] If $x \in X \backslash A$ then $\phi_X(x) = [x, 1] = \{(x, 1)\}$

have that $\phi_Y^{-1}(\phi_Y(U)) = U$. So, it suffices to show that $\phi_X^{-1}(\phi_Y(U))$ is an open subset of X. It is not hard to see[2] that:

$$\phi_X^{-1}(\phi_Y(U)) = \{x \in X \mid \phi_X(x) \in \phi_Y(U)\} = \{a \in A \mid f(a) \in U\}$$

Put differently, $\phi_X^{-1}(\phi_Y(U)) = f^{-1}(U)$. Since f is a continuous map and U open, the set $f^{-1}(U)$ is open in A. Since A is equipped with the subspace topology, this is the case iff $f^{-1}(U) = A \cap V$ where V is some open set in X. By assumption, A is open in X and thus $f^{-1}(U)$ is the intersection of open subsets of X, so is also open in X. It follows that $\phi_Y(U)$ is open in $X \cup_f Y$. Since U was arbitrary, we can conclude that ϕ_Y is an open map. □

Lemma B.3. *If A is open in X and f is an open map, then ϕ_X is an open map.*

Proof. Let U be an open subset of X, and consider $\phi_X(U) \subset X \cup_f Y$. As in the previous lemma, in order to show that $\phi_X(U)$ is open in the adjunction space, it suffices to show that the preimages of $\phi_X(U)$ under ϕ_X and ϕ_Y are open in their respective spaces.

Consider first $\phi_X^{-1}(\phi_X(U))$. Observe[3] that:

$$\begin{aligned}\phi_X^{-1}(\phi_X(U)) &= \{x \in X \mid \phi_X(x) \in \phi_X(U)\} \\ &= U \cup \{a \in A \mid \exists u \in U \cap A(f(u) = f(a)\} \\ &= U \cup f^{-1}(f(A \cap U))\end{aligned}$$

We show that $f^{-1}(f(A \cap U))$ is open in X. Observe that since U is open in X, the set $A \cap U$ is open in A. By assumption, f is an open map, hence $f(A \cap U)$ is open in Y. Since f is continuous, $f^{-1}(f(A \cap U))$ is open in A, i.e. $f^{-1}(f(A \cap U)) = A \cap V$ for some V open in X. By assumption, A is open in X and thus $f^{-1}(f(A \cap U))$ is open in X. We can conclude from all this that $\phi_X^{-1}(\phi_X(U))$ is open in X.

Consider now the set $\phi_Y^{-1}(\phi_X(U))$. Observe[4] that:

---

[2] Something like: $a \in f^{-1}(U)$ implies $f(a) \in U$ so $\phi_X(a) = \phi_Y(f(a))$ implies that $\phi_X(a) \in \phi_Y(U)$, hence $a \in \phi_X^{-1}(\phi_Y(U))$. Conversely, $x \in \phi_X^{-1}(\phi_Y(U))$ implies that $\phi_X(x) \in \phi_Y(U)$, so there is some $u \in U$ such that $\phi_X(x) = \phi_Y(u)$, and this is only the case if $x \in A$ and $f(x) = u$, i.e. $x \in f^{-1}(u) \subseteq f^{-1}(U)$.

[3] If $x \in \phi_X^{-1}(\phi_X(U))$ then $\phi_X(x) = \phi_X(u)$ for some $u \in U$. If $x \notin U$, this can only be the case if x and u are in A and $f(x) = f(u)$, i.e. $x \in f^{-1}(f(u))$. Since $u \in A \cap U$, we get that $x \in f^{-1}(f(A \cap U))$. Conversely, $x \in f^{-1}(f(A \cap U))$ implies that $f(x) = f(u)$ for some $u \in A \cap U$. Hence $[x, 1] = \{(f(x), 2)\} \cup \{(a, 1) \mid f(a) = f(x)\} = [u, 1]$, hence $\phi_X(x) = \phi_X(u)$ and thus $x \in \phi_X^{-1}(\phi_X(u)) \subseteq \phi_X^{-1}(\phi_X(U))$.

[4] If $y \in \phi_Y^{-1}(\phi_X(U))$ then $\phi_Y(y) = \phi_X(u)$ for some $u \in U$, which can only be the case if $u \in A$ and $u \in f^{-1}(y)$. Hence $f(u) = y$ and $y \in f(A \cap U)$. Conversely, $y \in f(A \cap U)$ implies there is some $u \in A \cap U$ such that $f(u) = y$. In particular, $u \in A$, hence $\phi_X(u) = \phi_Y(f(u)) = \phi_Y(y)$. So since ϕ_Y is injective, $y = \phi_Y^{-1}(\phi_X(u)) \subseteq \phi_Y^{-1}(\phi_X(U))$.

$$\begin{aligned}\phi_Y^{-1}(\phi_X(U)) &= \{y \in Y \mid \exists u(u \in U \wedge \phi_Y(y) = \phi_X(u))\} \\ &= \{y \in Y \mid \exists u(u \in U \cap A \wedge f(u) = y)\} \\ &= f(A \cap U)\end{aligned}$$

Since U is open in X, the set $U \cap A$ is open in A. By assumption, f is an open map, hence $f(A \cap U)$ is an open set in Y.

We can conclude that $\phi_X(U)$ is open in $X \cup_f Y$, and thus ϕ_X is an open map. □

Lemma B.4. *If f is injective and open, then ϕ_X is open.*

Proof. Let $U \subseteq X$ and consider $\phi_X(U)$. Since f is injective, it follows from Prop. A.1 that ϕ_X is also injective, and thus $\phi_X^{-1}(\phi_X(U)) = U$, which is clearly open in X. Consider now $\phi_Y^{-1}(\phi_X(U))$. We can as in the previous lemma to conclude that $\phi_Y^{-1}(\phi_X(U))$ is open. It follows that ϕ_X is an open map. □

Lemma B.5. *ϕ_Y is always a topological embedding.*

Proof. Since ϕ_Y is injective, the restriction to its image is a bijection. Similarly, the restriction of a continuous map to its image is also continuous[5] and thus it suffices to show that ϕ_Y is homeomorphic onto it's domain, i.e. the map $\phi_Y : Y \to \phi_Y(Y)$ is open.

Let U be an open subset of Y. Since $\phi_Y(Y)$ is equipped with the subspace topology, in order to show that $\phi_Y(U)$ is open in $\phi_Y(Y)$, we need to show that $\phi_Y(U) = \phi_Y(Y) \cap V$, where V is some open subset of $X \cup_f Y$. Observe that U is open in Y and f is continuous, so the preimage $f^{-1}(U)$ is open in A. Since A is equipped with the subspace topology, this means that $f^{-1}(U) = A \cap W$, where W is some open set in X.

We now argue that the set $\phi_X(W) \cup \phi_Y(U)$ suffices as our choice for V. Note first that:

$$\begin{aligned}\phi_Y(Y) \cap (\phi_X(W) \cup \phi_Y(U)) &= (\phi_Y(Y) \cap \phi_X(W)) \cup (\phi_Y(Y) \cap \phi_Y(U)) \\ &= (\phi_Y(Y) \cap \phi_X(W)) \cup \phi_Y(U)\end{aligned}$$

Let $z \in \phi_X(W) \cap \phi_Y(Y)$. Then $z = \phi_X(w)$ for some $w \in W$, and $z = \phi_Y(y)$ for some $y \in Y$. This can only be the case if $w \in A$ and $f(w) = y$. Since $w \in A$, it follows that $w \in A \cap W = f^{-1}(U)$, hence $f(w) \in U$, and thus $z = \phi_Y(y) = \phi_Y(f(w)) \in U$. We can conclude that $\phi_X(W) \cap \phi_Y(Y) \subseteq U$, and thus $(\phi_Y(Y) \cap \phi_X(W)) \cup \phi_Y(U) = \phi_Y(U)$ as required.

We now show that the set $\phi_X(W) \cup \phi_Y(U)$ is open in $X \cup_f Y$. Observe first that:

[5] Let $f : X \to Y$ be any continuous map between topological spaces, and consider $f|_X : X \to f(X)$. Let $U \subseteq f(X)$ be open in the subspace topology. Then $U = f(X) \cap V$ where V is open in Y. The preimage is then $(f|_X)^{-1}(U) = f^{-1}(f(X) \cap V) = X \cap f^{-1}(V) = f^{-1}(V)$ which is open in X since f is continuous.

$$\begin{aligned}
\phi_X^{-1}(\phi_X(W) \cup \phi_Y(U)) &= \phi_X^{-1}(\phi_X(W)) \cup \phi_X^{-1}(\phi_Y(Y)) \\
&= (W \cup f^{-1}(U)) \cup f^{-1}(U) \\
&= W \\
\phi_Y^{-1}(\phi_X(W) \cup \phi_Y(U)) &= \phi_Y^{-1}(\phi_X(W)) \cup \phi_Y^{-1}(\phi_Y(Y)) \\
&= f(W \cap A) \cup U \\
&= U
\end{aligned}$$

Since W and U are open sets of X and Y respectively, it follows from Prop. 3.2 that $\phi_X(W) \cup \phi_Y(U)$ is open in $X \cup_f Y$. We may thus conclude that $\phi_Y(U) = \phi_Y(Y) \cap (\phi_X(W) \cup \phi_Y(U))$ is open in the subspace topology $\phi_Y(Y)$. Since we picked an arbitrary open set U of Y, it follows that the map $\phi_Y : Y \to \phi_Y(Y)$ is an open map, and thus is a homeomorphism as required. □

Lemma B.6. *If f is an injective open map, then ϕ_X is a topological embedding.*

Proof. Note first that since f is by assumption an injection, it follows from Prop. A.1 that ϕ_X is also injective. Since ϕ_X is continuous, it is continuous once restricted to its image. So, it suffices to show that ϕ_X is homeomorphic onto its image.

Let $U \subseteq X$ be an open set. We want a set V such that $\phi_X(U) = \phi_X(X) \cap V$. The set $V := \phi_X(U) \cup \phi_Y(f(U \cap A))$ will suffice. Observe first that:

$$\phi_X(X) \cap (\phi_X(U) \cup \phi_Y(f(U \cap A))) = ... = \phi_X(U) \cup (\phi_X(X) \cap \phi_Y(f(U \cap A)))$$

Let $z \in \phi_X(X) \cap \phi_Y(f(U \cap A))$. Then there is some $x \in X$ and some $y \in f(U \cap A)$ such that $\phi_X(x) = z = \phi_Y(y)$. This can only be the case if $f(x) = y$ and since $y \in f(U \cap A)$, it must be the case that $x \in U$. It follows that $z \in \phi_X(U)$ and thus $\phi_X(X) \cap \phi_Y(f(U \cap A))$ is a subset of U, whence $\phi_X(U) = \phi_X(X) \cap V$.

We now show that V is open in $X \cup_f Y$, i.e. $\phi_X^{-1}(V)$ and $\phi_Y^{-1}(V)$ are open in X and Y respectively.

Consider first $\phi_X^{-1}(V)$. Observe:

$$\begin{aligned}
\phi_X^{-1}(\phi_X(U) \cup \phi_Y(f(U \cap A))) &= \phi_X^{-1}(\phi_X(U)) \cup \phi_X^{-1}(\phi_Y(f(U \cap A))) \\
&= U \cup f^{-1}(f(A \cap U)) \\
&= U \cup (A \cap U) \\
&= U
\end{aligned}$$

Where in the second and third lines we use that ϕ_X and f are injective. Since U is open, we may conclude that $\phi_X^{-1}(V)$ is open. Now observe that:

$$\begin{aligned}
\phi_Y^{-1}(\phi_X(U) \cup \phi_Y(f(U \cap A))) &= \phi_Y^{-1}(\phi_X(U)) \cup \phi_Y^{-1}(\phi_Y(f(U \cap A))) \\
&= f(A \cap U) \cup f(A \cap U) \\
&= f(A \cap U)
\end{aligned}$$

Since U is open, the set $A \cap U$ is an open set the subspace topology on A. Since we assumed that f is an *open* map, it follows that the set $f(A \cap U) = \phi_Y^{-1}(V)$ is open in Y.

We may conclude that V is open in $X \cup_f Y$, and since $\phi_X(U) = V \cap \phi_X(X)$, it follows that $\phi_X(U)$ is open in $\phi_X(X)$ equipped with the subspace topology. As such, ϕ_X is an open map onto $\phi_X(X)$, and thus $\phi_X : X \to X \cup_f Y$ is a topological embedding. □

B.1.2 Proof of Lemma 3.6

Proposition B.7. *Let $\mathcal{B}^X$ and $\mathcal{B}^Y$ be bases for X and Y respectively. If ϕ_X and ϕ_Y are open maps, then the collection $\mathcal{B} = \{\phi_X(U) \mid U \in \mathcal{B}^X\} \cup \{\phi_Y(V) \mid V \in \mathcal{B}^Y\}$ forms a basis for the adjunction topology τ_A.*

Proof. It suffices to show that $\mathcal{B}$ consists of open sets, and that every open set in $X \cup_f Y$ can be represented as a union of elements of $\mathcal{B}$. To see that $\mathcal{B}$ consists of open sets, we can appeal to our assumption that the canonical maps ϕ_X and ϕ_Y are open. Since every element U of $\mathcal{B}^X$ is open in X, it follows that $\phi_X(U)$ is open in $X \cup_f Y$. The case for Y is similar. Thus $\mathcal{B}$ consists of open sets.

Suppose now that W is an open set in $X \cup_f Y$. By Prop. 3.2, the sets $\phi_X^{-1}(W)$ and $\phi_Y^{-1}(W)$ are open in X and Y, respectively. As such, these preimages can be represented as unions of basis elements, i.e.

$$\phi_X^{-1}(W) = \bigcup_{i \in I} U_i \quad \text{and} \quad \phi_Y^{-1}(W) = \bigcup_{j \in J} V_j$$

where each $U_i \in \mathcal{B}^X$ and each $V_j \in \mathcal{B}^Y$. Since for *any*[6] subset $Z \subseteq X \cup_f Y$ it is the case that $Z = \phi_X(\phi_X^{-1}(Z)) \cup \phi_Y(\phi_Y^{-1}(Z))$, it follows that:

$$\begin{aligned} W &= \phi_X(\phi_X^{-1}(W)) \cup \phi_Y(\phi_Y^{-1}(W)) \\ &= \phi_X\left(\bigcup_{i \in I} U_i\right) \cup \phi_Y\left(\bigcup_{j \in J} V_j\right) \\ &= \left(\bigcup_{i \in I} \phi_X(U_i)\right) \cup \left(\bigcup_{j \in J} \phi_Y(V_j)\right) \end{aligned}$$

and thus W is a union of members of $\mathcal{B}$. We can conclude that $\mathcal{B}$ is indeed a basis. □

[6] Since $X \cup_f Y = \phi_X(X) \cup \phi_Y(Y)$, any subset Z of $X \cup_f Y$ can be represented as $Z = (Z \cap \phi_X(X)) \cup (Z \cap \phi_Y(Y))$. Since ϕ_X is a surjective map once restricted to its image, it has a right-inverse ϕ_X^{-1}. In this case, $\phi_X(\phi^{-1}(W)) = W$ for any subset $W \subseteq \phi_X(X)$. In particular, $Z \cap \phi_X(X)$ is a subset of $\phi_X(X)$, so $Z \cap \phi_X(X) = \phi_X(\phi_X^{-1}(Z \cap \phi_X(X))) = ... = \phi_X(\phi_X^{-1}(Z)) \cap \phi_X(\phi_X^{-1}(\phi_X(X))) = \phi_X(\phi_X^{-1}(Z)) \cap \phi_X(X) = \phi_X(\phi_X^{-1}(Z))$. The case for ϕ_Y is similar, and we then get $Z = (Z \cap \phi_X(X)) \cup (Z \cap \phi_Y(Y)) = \phi_X(\phi_X^{-1}(Z)) \cup \phi_Y(\phi_Y^{-1}(Z))$.

Lemma B.8. *If X, Y are connected and A is non-empty, then $X \cup_f Y$ is connected.*

Proof. Suppose towards a contradiction that $X \cup_f Y$ is disconnected, i.e. there exist disjoint open sets U, V of $X \cup_f Y$ such that U, V are both non-empty and cover $X \cup_f Y$. Consider $\phi_X^{-1}(U)$ and $\phi_X^{-1}(V)$. By assumption, these sets are open in X, and since pre-images preserve disjointedness[7], they are disjoint. It is also the case that they cover[8] X, in order to preserve the connectedness of X, it must be the case that either $\phi_X^{-1}(U)$ or $\phi_X^{-1}(V)$ is empty. Without loss of generality suppose that $\phi_X^{-1}(U) = \emptyset$. It follows from $\phi_X^{-1}(V) = X$ that $\phi_X(X) \subseteq V$ and, since $\phi_X(A) = \phi_Y(f(A))$, that $f(A) \subseteq \phi_Y^{-1}(V)$. Since we have assumed that A is non-empty, this means that $\phi_Y^{-1}(V)$ is also non-empty. However, in order to preserve the connectedness of Y, it cannot be the case that $\phi_Y^{-1}(U)$ is also non-empty. We have thus arrived at our contradiction - if $\phi_Y^{-1}(U) = \emptyset$, then $\phi_Y(Y) \subseteq V$, and thus $X \cup_f Y = \phi_X(X) \cup \phi_Y(Y) \subset V$, contradicting U as a non-empty set[9]. ☐

Lemma B.9. *If both X and Y are compact, then so is $X \cup_f Y$.*

Proof. It is well-known that the disjoint union of two (or, finitely-many) compact spaces is again compact, and that the quotient of a compact space is compact[10]. In our context, this means that $X \sqcup Y$ is compact and thus $X \cup_f Y$ is also compact. ☐

B.2 Proofs from Appendix A

Lemma B.10 (Lemma 5.11). *Let $f : M \to N$ be a smooth map. Then the following are equivalent.*

1. *f is a smooth embedding.*
2. *$f(M)$ is an embedded submanifold of N, and f acts as a diffeomorphism from M to $f(M)$.*

Proof. The proof of $(1 \Rightarrow 2)$ is found as [20, Prop 5.2], and the reader may verify that the Hausdorff property is not used. So, we show $(2 \Rightarrow 1)$. To begin with, we denote by g the map from M to $f(M)$. By assumption $g : M \to f(M)$ is a diffeomorphism, so in particular it is a homeomorphism. Also, f is smooth, thus it is continuous. Hence $f : M \to N$ is a topological embedding. It suffices to show that the maps $df_p : T_pM \to T_{f(p)}N$ are all injective. So, let p be arbitrary, and consider $df_p := d(\iota \circ g)_p$. Since the map $g : M \to f(M)$ is

[7] Simple enough - $\phi_X^{-1}(U) \cap \phi_X^{-1}(V) = \phi_X^{-1}(U \cap V) = \phi_X^{-1}(\emptyset) = \emptyset$
[8] $\phi^{-1}(U) \cup \phi_X(V) = \phi_X^{-1}(U \cup V) = X$
[9] This proof is a bit weak at the end - we want to say that $X \cup_f Y = V$ and thus $U = \emptyset$.
[10] This is a corollary of [26, Thm. 3.2.3.].

a diffeomorphism, by Prop. A.10.4 it follows that $dg_p : T_pM \to T_{f(p)}f(M)$ is an isomorphism of vector spaces. In particular, it is injective. Since $f(M)$ is an embedded submanifold, the inclusion map $\iota : f(M) \to N$ is a smooth embedding. Then $df_p = d(\iota \circ g)_p = d\iota_{f(p)} \circ dg_p$ is a composition of injective maps, so is also injective. Thus f is a smooth embedding. □

Lemma B.11 (Lemma 5.18). *Let $f : M \to N$ and $g : N \to P$ be smooth embeddings, with $dim(M) = dim(N) = dim(P)$, and let α, β and δ be elements of the fibres $T^2(T_p^*M), T^2(T_{f(p)}^*N)$ and $T^2(T_{g\circ f(p)}^*P)$ respectively.*

1. $(id_M)_*\alpha = (id_M)^*\alpha = \alpha$
2. $f_*\alpha = (f^{-1})^*\alpha$
3. $f^*\beta = (f^{-1})_*\beta$
4. $(f^* \circ g^*)\delta = (g \circ f)^*\delta$
5. $f^*(f_*\alpha) = \alpha$
6. $g_* \circ f_*(\alpha) = (g \circ f)_*\alpha$

Proof. It should be noted from the outset that most of these results follow immediately from the properties of the pointwise differential maps (cf. Prop. A.10).

1. $(id_M)^*\alpha(v, w) = \alpha(d(id_M)_p(v), d(id_M)_p(w)) = \alpha(v, w)$
2. This is immediate from the definition of the pushforward. Indeed, $f_*\alpha(v, w) = \alpha(v', w')$ where $v = df_p(v')$ and $w = df_p(w')$. Since f is a diffeomorphism onto its image, the differential df_p is a bijection, with inverse equal to $(df^{-1})_{f(p)}$. This means that $v' = (df^{-1})_{f(p)}(v)$ and $w' = (df^{-1})_{f(p)}(w)$. Hence $(f^{-1})^*\alpha(v, w) = \alpha((df^{-1})_{f(p)}(v), (df^{-1})_{f(p)}(w)) = \alpha(v', w') = f_*\alpha(v', w')$ as required.
3. By the previous item, $(f^{-1})_*\beta = ((f^{-1})^{-1})^*\beta = f^*\beta$.
4. For any $v, w \in T_pM$, we have:
$$\begin{aligned}(f^* \circ g^*)\delta(v, w) &= g^*\delta(df_p(v), df_p(w)) \\ &= \delta(dg_{f(p)} \circ df_p(v), dg_{f(p)} \circ df_p(w)) \\ &= \delta(d(g \circ f)_p(v), d(g \circ f)_p(w)) \\ &= (g \circ f)^*\delta\end{aligned}$$
where we have used A.10.3 in the fourth line.
5. $f^*(f_*\alpha) = f^* \circ (f^{-1})^*\alpha = (f^{-1} \circ f)^*\alpha = id_M^*\alpha = \alpha$.
6. $g_* \circ f_* = (g^{-1})^* \circ (f^{-1})^* = (f^{-1} \circ g^{-1})^* = ((g \circ f)^{-1})^* = (g \circ f)_*$

□

Lemma B.12 (Lemma 5.19). *If $f : M \to N$ is a diffeomorphism, then the bundles $T^2(T^*M)$ and $T^2(T^*N)$ are isomorphic.*

Proof. We define the map $\xi : T^2(T^*N) \to T^2(T^*M)$ fibrewise by $(q, \alpha) \mapsto (p, f^*\alpha)$, where $p = f^{-1}(q)$ and $f^*\alpha$ is the element of $(T^2(T^*M))_p$ such that $f^*\alpha(v, w) = \alpha(df_p(v), df_p(w))$ for all $v, w \in T_pM$. It should be clear that the element $f^*\alpha$ is unique, from which we can conclude that ξ is a well-defined function that covers the inverse map f^{-1}. We will now show that ξ is a smooth, bijective bundle morphism, since the result will then follow from Lemma A.25.

We start by showing that ξ is a linear map once restricted to fibres. This is fairly routine – for any two elements β_1 and β_2 in $T_{f(p)}N$, we have:

$$\begin{aligned}
\xi(\beta_1 + \beta_2)(v, w) &= f^*(\beta_1 + \beta_2)(v, w) \\
&= (\beta_1 + \beta_2)(df_p(v), df_p(w)) \\
&= \beta_1(df_p(v), df_p(w)) + \beta_2(df_p(v), df_p(w)) \\
&= f^*\beta_1(v, w) + f^*\beta_2(v, w) \\
&= \xi(\beta_1)(v, w) + \xi(\beta_2(v, w)
\end{aligned}$$

and similarly, for the scalar multiplication we have:

$$\begin{aligned}
\xi(r \cdot \beta)(v, w) &= f^*(r \cdot \beta)(v, w) = (r \cdot \beta)(df_p(v), df_p(w)) \\
&= r \cdot (\beta(df_p(v), df_p(w))) = r \cdot (f^*\beta(v, w)) = r \cdot \xi(\beta)(v, w)
\end{aligned}$$

for each $r \in \mathbb{R}$ and $\beta \in T_{f(p)}N$.

To see that ξ is injective, suppose that we have two distinct elements α and β in $(T^2(T^*N))_q$. If α and β are distinct, then there is are two elements $v', w' \in T_qN$ such that $\alpha(v', w') \neq \beta(v', w')$. Since f is a diffeomorphism, by Prop. A.10 the differential df_p is a bijection, and thus v and w are of the form $v' = df_p(v)$ and $w' = df_p(w)$ where $v, w \in T_pM$. Then: $f^*\alpha(v, w) = \alpha(v', w') \neq \beta(v', w') = f^*\beta(v, w)$, and thus ξ is injective.

To show surjectivity, suppose that we have some element (p, α) in the fibre $(T^2(T^*M))_p$.We define the map $\beta : T_qN \times T_qN \to \mathbb{R}$ by $\beta(v', w') = \alpha(v, w)$ where $v' = df_p(v)$ and $w' = df_p(w)$. Note that this is well-defined since the differential df_p is bijective, and β inherits its bilinearity from α. It is not hard to see that $f^*\beta = \alpha$. Indeed, for arbitrary $v, w \in T_pM$, we have that: $f^*\beta(v, w) = \beta(df_p(v), df_p(w)) = \alpha(v, w)$. It follows from this that $\xi(q, \beta) = (p, f^*\beta) = (p, \alpha)$ as required.

So far we have shown that ξ is a well-defined, bijective map that acts linearly on fibres. We finish the proof by showing that ξ is a smooth map. So, let $p \in M$ and $q \in N$ be arbitrary such that $f(p) = q$. Let (U, Φ) be a local trivialisation of $T^2(T^*M)$ at p that is small enough[11] so that U also forms a chart for M. Since $f : M \to N$ is a diffeomorphism, the tuple $(f(U), \varphi \circ f^{-1})$ forms a chart for N at q. Without loss of generality we can pick a local trivialisation (V, Ψ) of $T^2(T^*N)$ at q such that $V = f(U)$.[12] The expression of ξ in local trivialisations can be computed by using $\Phi \circ \xi \circ \Psi^{-1}$. This can be computed as:

---

[11] This is always possible – we can pick some chart (U', φ) of M at p, and then take the intersection $U \cap U'$ and restrict Φ and φ accordingly.

[12] Again, pick any local trivialisation (V, Ψ) at q and intersect it with the chart $(f(U), \varphi \circ f^{-1})$.

$$
\begin{aligned}
\Phi \circ \xi \circ \Psi^{-1}(q, r_{11}, ..., r_{nn}) &= \Phi \circ \xi(r_{ij} d\psi^i \otimes d\psi^j)) \\
&= \Phi(r_{ij} d(\psi^i \circ f) \otimes d(\psi^j \circ f)) \\
&= \Phi(r_{ij} d(\varphi^i \circ f^{-1} \circ f) \otimes d(\varphi^j \circ f^{-1} \circ f)) \\
&= \Phi(r_{ij} d\varphi^i \otimes d\varphi^j) \\
&= (p, r_{11}, ..., r_{nn}).
\end{aligned}
$$

Thus locally, we can pick a coordinate representation of ξ in which the component maps are equal to $(f^{-1}, id_{\mathbb{R}}, ..., id_{\mathbb{R}})$.[13] This map is smooth, since f^{-1} is smooth. Since we chose q arbitrarily, it follows that for any point in N, the map ξ has a local representation which is smooth in the Euclidean sense. Thus ξ is smooth.

Since ξ is a bijective bundle morphism covering a diffeomorphism f^{-1}, we can then (albeit prematurely) apply Lemma A.25 to conclude that ξ is a bundle isomorphism. □

Proposition B.13 (Prop. 5.21). *Let $f : M \to N$ be a smooth map, and (F, π_F, N) a vector bundle. The map $p_2 : f^*F \to F$ is a bundle morphism, and p_2 is an isomorphism whenever f is a diffeomorphism.*

Proof. This result and the next are standard facts about the pullback bundle. See Hatcher [35, Sec. 1.2] (particularly Prop. 1.5) for the proofs. □

Proposition B.14. *If $g : E \to F$ is a bundle morphism covering f, then there is a unique morphism $h : E \to f^*F$ such that the following diagram commutes.*

$$
\begin{tikzcd}
E \arrow[r, dotted, "h"] \arrow[rr, bend left, "g"] \arrow[d, "\pi_E"'] & f^*F \arrow[r, "p_2"] \arrow[d, "p_1"'] & F \arrow[d, "\pi_F"] \\
M \arrow[r, equal] & M \arrow[r, "f"'] & N
\end{tikzcd}
$$

Lemma B.15 (Lemma 5.23). *If A is an open submanifold of M and E a vector bundle over M, then the restriction bundle $E|_A$ is an open submanifold of E.*

Proof. We will only provide a sketch, since the argument contains some notions that we have not introduced. It is known that every open submanifold is an embedded submanifold of codimension 0 (cf [20, Lem. 5.1]). It can also be shown that the projection map $\pi : E \to M$ is a *submersion*, and thus π is *transverse* to every embedded submanifold of M (cf. [20, Ex. 10.1]). It can be

---

[13] Actually, we are implicitly assuming that we can diffeomorphically send $U \times \mathbb{R}^{n^2}$ to $\varphi(U) \times \mathbb{R}^{n^2}$ by using the map $(p, r_{11}, ..., r_{1n}, ..., r_{n1}, ..., r_{nn}) \mapsto (\varphi^1(p), ..., \varphi^n(p), r_{11}, ..., r_{1n}, ..., r_{n1}, ..., r_{nn})$.

shown whenever a smooth map is transverse, the preimage of an embedded submanifold of codimension k is an embedded submanifold of codimension k (cf. [20, Cor. 6.1]). Thus $\pi^{-1}(A) := E|_A$ is a submanifold of codimension 0 in E, and hence $E|_A$ is an open submanifold of E. □

Lemma B.16 (Lemma 5.24). *If $g : E \to F$ is a bijective bundle morphism covering a diffeomorphism $f : M \to N$, then g is a bundle isomorphism.*

Proof. We can use the fact that a pullback along a diffeomorphism turns the projection map $p_2 : f^*F \to F$ into a bundle isomorphism. Then the map $g : E \to F$ factors through the pullback f^*F, so there is some bundle morphism $h : E \to f^*F$ such that $g = p_2 \circ h$. Since g is bijective, the map h must also be bijective. Since h is a bijective bundle morphism between two bundles over M, we can apply A.17 to conclude that it is a bundle isomorphism. Then g is a composition of bundle isomorphisms, so is also a bundle isomorphism. □

References

[1] N. Belnap, “Branching space-time, postprint january, 2003”, Jan. 2003. [Online]. Available: `http://philsci-archive.pitt.edu/1003/`.

[2] ——, “Branching space-time”, *Synthese*, vol. 92, no. 3, pp. 385–434, 1992.

[3] S. McCall *et al.*, *A model of the universe: space-time, probability, and decision*. Oxford University Press, 1994.

[4] R. Douglas, “Stochastically branching spacetime topology”, *Time’s arrows today: recent physical and philosophical work on the direction of time*, p. 173, 1997.

[5] J. Earman, “Pruning some branches from “branching spacetimes””, *Philosophy and Foundations of Physics*, vol. 4, pp. 187–205, 2008.

[6] T. Muller, *Probability theory and causation. a branching space-times analysis*, Jan. 2004. [Online]. Available: `http://philsci-archive.pitt.edu/1737/`.

[7] N. Belnap, *Epr-like “funny business” in the theory of branching space-times*, Sep. 2002. [Online]. Available: `http://philsci-archive.pitt.edu/784/`.

[8] ——, *No-common-cause epr-like funny business in branching space-times*, 2002. [Online]. Available: `http://philsci-archive.pitt.edu/670/`.

[9] T. Muller, N. Belnap, and K. Kishida, *Funny business in branching space-times: Infinite modal correlations*, Jun. 2006. [Online]. Available: `http://philsci-archive.pitt.edu/2803/`.

[10] L. Wronski and T. Placek, *On infinite epr-like correlations*, May 2007. [Online]. Available: `http://philsci-archive.pitt.edu/3362/`.

[11] T. Muller, *A branching space-times view on quantum error correction*, May 2007. [Online]. Available: `http://philsci-archive.pitt.edu/3377/`.

[12] L. Wronski, *Branching space-times and parallel processing*, Feb. 2012. [Online]. Available: `http://philsci-archive.pitt.edu/10124/`.

[13] T. Müller, “Branching space-time, modal logic and the counterfactual conditional”, in *Non-locality and modality*, Springer, 2002, pp. 273–291.

[14] L. Wroński and T. Placek, “On minkowskian branching structures”, *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, vol. 40, no. 3, pp. 251–258, 2009.

[15] R. Penrose, “Singularities and time-asymmetry”, in *General relativity*, 1979.

[16] T. Müller, “A generalized manifold topology for branching space-times”, *Philosophy of science*, vol. 80, no. 5, pp. 1089–1100, 2013.

[17] G. McCabe, “The topology of branching universes”, *Foundations of Physics Letters*, vol. 18, no. 7, pp. 665–676, 2005.

[18] T. Placek, N. Belnap, and K. Kishida, “On topological issues of indeterminism”, *Erkenntnis*, vol. 79, no. 3, pp. 403–436, 2014.

[19] T. Placek, “Branching for general relativists”, in *Nuel Belnap on indeterminism and free action*, Springer, Cham, 2014, pp. 191–221.

[20] J. M. Lee, *Introduction to Smooth Manifolds*. Springer, 2013.

[21] M. Visser, *Lorentzian wormholes: from Einstein to Hawking*. 1995.

[22] R. Penrose, *Techniques of differential topology in relativity.* Siam, 1972, vol. 7.

[23] E. Minguzzi and M. Sánchez, “The causal hierarchy of spacetimes”, *Recent developments in pseudo-Riemannian geometry, ESI Lect. Math. Phys*, pp. 299–358, 2008.

[24] T. Muller, “Branching space-times, general relativity, the hausdorff property, and modal consistency”, 2011.

[25] R. Brown, *Topology and groupoids.* 2006.

[26] R. Engelking, *General topology.* Heldermann, 1989.

[27] J. Munkres, *Topology.* 2000.

[28] K. Martin and P. Panangaden, “Spacetime geometry from causal structure and a measurement”, *Mathematical Foundations of Information Flow: Clifford Lectures Information Flow in Physics, Geometry, and Logic and Computation, March 12-15, 2008, Tulane University, New Orleans, Louisiana*, vol. 71, p. 213, 2012.

[29] P. Hajicek, “Causality in non-hausdorff space-times”, *Communications in Mathematical Physics*, vol. 21, no. 1, pp. 75–84, 1971.

[30] T. Müller, “Branch dependence in the consistent histories approach to quantum mechanics”, *Foundations of Physics*, vol. 37, no. 2, pp. 253–276, 2007.

[31] B. I. Dundas, *A Short Course in Differential Topology.* Cambridge University Press, 2018.

[32] B. O’Neill, *Semi-Riemannian geometry with applications to relativity.* Academic press, 1983, vol. 103.

[33] V. Moretti and R. Di Criscienzo, “How can we determine if a spacetime is flat?”, *Frontiers in Physics*, vol. 1, p. 12, 2013, ISSN: 2296-424X. DOI: `10.3389/fphy.2013.00012`.

[34] R. Geroch, “Domain of dependence”, *Journal of Mathematical Physics*, vol. 11, no. 2, pp. 437–449, 1970.

[35] A. Hatcher, *Vector bundles and K-theory.* 2003.