DISSERTATION / DOCTORAL THESIS

Titel der Dissertation /Title of the Doctoral Thesis

„Fundamental constructions for Lorentzian length spaces“

verfasst von / submitted by

Felix Rott, MSc

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Doktor der Naturwissenschaften (Dr. rer. nat.)

Wien, 2024 / Vienna 2024

Studienkennzahl lt. Studienblatt / degree programme code as it appears on the student record sheet:	A 796 605 405
Dissertationsgebiet lt. Studienblatt / field of study as it appears on the student record sheet:	Mathematik
Betreut von / Supervisor:	Univ. Prof. Mag. Mag. Dr. Michael Kunzinger

Fundamental constructions for Lorentzian length spaces

Felix Rott

Abstract

This thesis deals with fundamental constructions and concepts concerning the theory of Lorentzian length spaces, which can be described as an abstraction of Lorentzian geometry in the spirit of metric length spaces and metric geometry. On the one hand, a gluing construction for Lorentzian length spaces is developed. Bulding on this, an analogue to the Reshetnyak gluing theorem as well as the preservation of various other (causal) properties are established. On the other hand, a splitting theorem for spaces with non-negative curvature bounds in the sense of triangle comparison is presented.

Acknowledgements

I want to start chronologically by thanking my parents. My dad for encouraging my mathematical curiosity from the very beginning, and my mum for unbounded support in general.

I want to thank Tobias Beran, who I had both the fortune and privilege to meet on my first day of undergraduate studies, and who has been more than a colleague ever since.

I want to thank Michael Kunzinger, who has been more than a scientific advisor during all of my studies. Even while supervising my Bachelor's thesis, he always had time for my questions, no matter how busy he was, and even at that level treated me like an equal. He always had helpful and encouraging words for me, and I will be forever grateful for everything he enabled for me. He truly justifies the german expression "Doktorvater".

I want to thank Matteo Calisti, Argam Ohanyan, Clemens Sämann and Roland Steinbauer who contribute to the inrcedibly warm, familiar and productive environment that is our research group. I am convinced that spatial separation in the future will not prevent us from keeping in contact.

Introduction

The theory of Lorentzian length spaces is a new and very promising approach to axiomatise and "synthetesise" Lorentzian geometry. The first work on this subject [19] was published only in 2018, and since then the field is developing at an impressive and ever increasing speed. Its creation is inspired by metric geometry and its interplay with Riemannian geometry, where in particular the theory of (metric) length spaces has helped to extract the "metric core" of many results and concepts related to Riemannian manifolds. Metric geometry is both a solidly established and highly active field of research. It is not far fetched to believe that the theory of Lorentzian length spaces will have a similar impact on related subjects, in particular Lorentzian geometry.

Historically, mathematical general relativity has been described in terms of smooth Lorentzian manifolds. It turns out, however, that there are many physically highly relevant situations where one is required to lower the regularity of the metric, such as when modelling different kinds of matter, shock waves, conical singularities or impulsive gravitational waves [33, 38, 14, 36, 35].

Lorentzian length spaces take this a step further and discard the manifold structure entirely. One of the main goals in the initial stages of the theory was to develop an anaolgoue to curvature bounds in the sense of triangle comparison and thus providing Lorentzian versions of Alexandrov spaces or $\mathrm{CAT}(K)$ spaces, the centerpieces of study in metric geometry.

Triangle comparison in a purely Lorentzian context was first worked on in [15], with [2] painting the full picture in the smooth semi-Riemannian setting. The recent interest in Lorentzian length spaces by reseranchers with a background in optimal transport and metric measure geometry gave even more traction to an already rapidly developing field. In particular, the introduction of synthetic lower Ricci curvature bounds into the synthetic Lorentzian setting via a convexity condition on the entropy functional, cf. e.g. [9, 5], is expected to be especially fruitful for problems related to field equations and singularity theorems [29]. In addition, alternative approaches to synthetic Lorentzian geometry in the spirit of Lorentzian length spaces have been developed [30, 22, 28], demonstrating both the value as well as the flexibility of the theory.

In the past, there have been other approaches to abstracting Lorentzian geometry, but they were considerabily less successful. [8] introduced so-called timelike spaces, which, although in the spirit of metric geometry, do not even contain all smooth spacetimes, hence are not a suitable generalisation. The causal spaces of [18], while being very general, are not geometric enough to provide a Lorentzian analogue to metric length spaces.

The theory of Lorentzian length spaces attracts researchers with very diverse backgrounds, such as Lorentzian geometry, general relativity, quantum gravity, Alexandrov geometry and metric measure geometry.

Clearly, this subject is, especially by mathematical standards, quite young, which is why the fundamentals of the theory have not been fully fleshed out yet. One possible point of view on categorising the current work on Lorentzian length spaces is the following: on the one hand, researchers are interested in generalising purely Lorentzian phenomena from the smooth setting of spacetimes. Examples include work on the causal ladder or questions on (in)extendibility [13, 1, 7]. On the other hand, one aims to translate concepts from metric (measure) geometry to a Lorentzian setting [4, 20, 23]. In short, one cares about "synthetesising"

Lorentzian concepts and "Lorentzifying" metric concepts. While there is some obvious overlap, and some works do not fit into this division at all, cf. e.g. [16, 21], it is a good rule of thumb and most works do at least partially fit into one of these categories.

This thesis can confidently be put into the second category, namely developing Lorentzian analogues of useful metric techniques. The main topic of investigation in this thesis was to transport the gluing construction for metric spaces over to Lorentzian length spaces.

Gluing is one of the cornerstone constructions of metric geometry, and provides an extremely general way of producing new spaces out of given ones. The central result about gluing in metric geometry is the Reshetnyak Gluing Theorem, which says that upper curvature bounds (in the sense of triangle comparison) are preserved under gluing. The gluing of metric spaces is a rather natural extension of the gluing of topological spaces, where one equips first the disjoint union and then an appropriate quotient with a suitable metric.

For Lorentzian length spaces, however, things become a bit more complicated. The disjoint union construction essentially works out in complete analogy, but for quotients the lower semi-continuity of the time separation function is all but guaranteed. A possible way of summarising this naively is the following: gluing in the full generality that is present in metric geometry is (to current knowledge) simply not possible while still retaining the lower semi-continuity of the time separation function. So one has to make sacrifices regarding generality, posing limitations on either the spaces themselves or the glued sets. If one values the generality of the spaces, then the identified sets have to, in a way, "behave timelike". For now it does seem like one is unable to avoid the identified sets having a causal character, but if one wants to glue along sets which are not timelike, then these glued sets essentially have to be null or spacelike boundaries of the spaces. Moreover, it is only possible to glue two spaces, with the boundaries being a "future boundary" in the one space and a "past boundary" in the other.

Going back to the gluing theorem itself, this is also substantially more delicate for Lorentzian length spaces than for $\mathrm{CAT}(K)$ spaces. Indeed, due to the missing concept of spacelike distance in the Lorentzian context, it was only possible to formulate the theorem for spacetimes with sectional curvature bounds in the sense of [2]. Concerning the setup of well behaved comparison neighbourhoods, the techniques developed in [25] proved to be crucial. Very recent results [31] introduce a spacelike distance on (certain subsets of) Lorentzian synthetic spaces, which might make a more general version of the gluing theorem feasible. The first two papers included in this thesis are devoted to gluing (general) spaces along timelike sets, while the other possibility is yet to be fully explored in future work.

A purely Lorentzian phenomenon, which came up during the development of these gluing techniques for Lorentzian length spaces, is to investigate which steps of the causal ladder (cf. [27, 26]) survive when gluing. This question is extensively answered in the second work. Furthermore, there are several open questions concerning the gluing of Lorentzian length spaces. As mentioned above, “spacelike gluing” is less explored, and could in particular be crucial for tackling problems related to the inextendibility of spacetimes. Further, while gluing and lower curvature bounds are generally poorly related to each other, it is still possible to glue Alexandrov spaces along their boundary and retain the (lower) curvature bound this way [34]. A similar question can be posed by replacing Alexandrov spaces with RCD spaces [17]. Both of these statements are open questions for Lorentzian length spaces. As a matter of fact, the last question remains unanswered even in the setting of metric measure spaces.

Finally, the third work is concerned with a slightly different topic, although still in the general area of strengthening the fundamental building blocks of the theory, namely establishing a splitting theorem (with respect to triangle comparison) for Lorentzian length spaces. Originally developed in the context of (smooth) Riemannian manifolds, a splitting theorem is a rigidity result and essentially says that under some geometric assumptions, a space splits, i.e., is isometric to a product. These assumptions usually include some form of (lower) curvature bound as well as the existence of a line.

This was first achieved for sectional curvature bounded below by zero and has been further extended to work for nonnegative lower Ricci bounds [37, 10]. Splitting theorems have been Lorentzified to apply to Lorentzian manifolds, both under sectional and Ricci curvature conditions [3, 32], as well as synthetesised to work for Alexandrov spaces using triangle comparison [24, 6], and even for RCD-spaces using synthetic lower Ricci curvature bounds [11, 12]. The splitting theorem displayed in this work is formulated with respect to triangle comparison, with most inspiration drawn from Lorentzian splitting using sectional curvature bounds as well as a proof for splitting of Alexandrov spaces.

List of included papers

1. T. Beran and F. Rott. Gluing constructions for Lorentzian length spaces. *Manuscripta Mathematica.*
   DOI: https://doi.org/10.1007/s00229-023-01469-4

2. F. Rott. Gluing of Lorentzian length spaces and the causal ladder. *Classical Quantum Gravity*, 40(17).
   DOI: https://iopscience.iop.org/article/10.1088/1361-6382/ace585

3. T. Beran, A. Ohanyan, F. Rott, and D. A. Solis. The splitting theorem for globally hyperbolic Lorentzian length spaces with non-negative timelike curvature. *Letters in Mathematical Physics* 113, 48.
   DOI: https://link.springer.com/article/10.1007/s11005-023-01668-w

References

[1] L. Aké Hau, A. J. Cabrera Pacheco, and D. A. Solis. On the causal hierarchy of Lorentzian length spaces. *Classical Quantum Gravity*, 37(21):21, 2020. https://dx.doi.org/10.1088/1361-6382/abb25f.

[2] S. B. Alexander and R. L. Bishop. Lorentz and semi-Riemannian spaces with Alexandrov curvature bounds. *Comm. Anal. Geom.*, 16(2):251–282, 2008. https://doi.org/10.4310/CAG.2008.v16.n2.a1.

[3] J. K. Beem, P. E. Ehrlich, S. Markvorsen, and G. J. Galloway. Decomposition theorems for Lorentzian manifolds with nonpositive curvature. *J. Diff. Geom.*, 22(1), 1985. https://doi.org/10.4310/jdg/1214439719.

[4] T. Beran and C. Sämann. Hyperbolic angles in Lorentzian length spaces and timelike curvature bounds. *Journal of the London Mathematical Society*, 107(5):1823–1880, 2023. https://doi.org/10.1112/jlms.12726.

[5] M. Braun and R. J. McCann. Causal convergence conditions through variable timelike ricci curvature bounds, 2023. https://arxiv.org/abs/2312.17158.

[6] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry, Graduate Studies in Mathematics*, volume 33. American Mathematical Society, 2001.

[7] A. Burtscher and L. García-Heveling. Time functions on Lorentzian length spaces, 2021. https://arxiv.org/abs/2108.02693.

[8] H. Busemann. Timelike spaces. *Dissertationes Math. (Rozprawy Mat.)*, 53:52, 1967.

[9] F. Cavalletti and A. Mondino. Optimal transport in Lorentzian synthetic spaces, synthetic timelike Ricci curvature lower bounds and applications, 2020. https://arxiv.org/abs/2004.08934.

[10] J. Cheeger and D. Gromoll. The splitting theorem for manifolds of nonnegative Ricci curvature. *J. Diff. Geom.*, 6(1):119–128, 1971. https://doi.org/10.4310/jdg/1214430220.

[11] N. Gigli. The splitting theorem in non-smooth context, 2013. https://arxiv.org/abs/1302.5555.

[12] N. Gigli. An overview of the proof of the splitting theorem in spaces with non-negative Ricci curvature. *Anal. Geom. Metr. Spaces*, 2(1):169–213, 2014.

[13] J. D. E. Grant, M. Kunzinger, and C. Sämann. Inextendibility of spacetimes and Lorentzian length spaces. *Ann. Global Anal. Geom.*, 55(1):133–147, 2019.
https://doi.org/10.1007/s10455-018-9637-x.

[14] J. B. Griffiths and J. Podolský. *Exact space-times in Einstein's general relativity.* Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 2009.

[15] S. G. Harris. A triangle comparison theorem for Lorentz manifolds. *Indiana Univ. Math. J.*, 31(3):289–308, 1982.
https://www.jstor.org/stable/24893088.

[16] J. Hedicke. Lorentzian distance functions in contact geometry. *J. Topol. Anal.*, 2022.
https://doi.org/10.1142/S179352532250008X.

[17] V. Kapovitch, C. Ketterer, and K. Sturm. On gluing Alexandrov spaces with lower Ricci curvature bounds, 2020.
https://arxiv.org/abs/2003.06242.

[18] E. H. Kronheimer and R. Penrose. On the structure of causal spaces. *Proc. Cambridge Philos. Soc.*, 67(2):481–501, 1967.
https://doi.org/10.1017/S030500410004144X.

[19] M. Kunzinger and C. Sämann. Lorentzian length spaces. *Ann. Glob. Anal. Geom.*, 54(3):399–447, 2018.
https://doi.org/10.1007/s10455-018-9633-1.

[20] M. Kunzinger and R. Steinbauer. Null distance and convergence of Lorentzian length spaces. *Ann. Henri Poincaré*, 2022.
https://doi.org/10.1007/s00023-022-01198-6.

[21] M. T. Law and J. Lucas. Spacetime representation learning. In *The Eleventh International Conference on Learning Representations*, 2023.
https://openreview.net/forum?id=qV$_{Mr}hYajc$.

[22] R. J. McCann. A synthetic null energy condition, 2023.
https://arxiv.org/abs/2304.14341.

[23] R. J. McCann and C. Sämann. A Lorentzian analogue for Hausdorff dimension and measure. *Pure and Applied Analysis*, 4(2):367–400, 2022.
https://doi.org/10.2140/paa.2022.4.367.

[24] A. D. Milka. Metric structure of a certain class of spaces that contain straight lines. *Ukrain. Geometr. Sb. Vyp.*, 4:43–48, 1967.

[25] E. Minguzzi. Convex neighborhoods for Lipschitz connections and sprays. *Monatsh. Math.*, 177(4):569–625, 2015.

[26] E. Minguzzi. Lorentzian causality theory. *Living Reviews in Relativity*, 22(1):3, 2019.

[27] E. Minguzzi and M. Sánchez. The causal hierarchy of spacetimes. In *Recent developments in pseudo-Riemannian geometry*, ESI Lect. Math. Phys., pages 299–358. Eur. Math. Soc., Zürich, 2008.

[28] E. Minguzzi and S. Suhr. Lorentzian metric spaces and their Gromov–Hausdorff convergence, 2023. https://arxiv.org/abs/2209.14384.

[29] A. Mondino and S. Suhr. An optimal transport formulation of the Einstein equations of general relativity. *J. Eur. Math. Soc. (JEMS)*, 25(3):933–994, 2023.

[30] O. Müller. Lorentzian Gromov-Hausdorff theory and finiteness results. *Gen. Relativity Gravitation*, 54(10):Paper No. 117, 17, 2022.

[31] O. Müller. Gromov-Hausdorff metrics and dimensions of Lorentzian length spaces, 2023. https://arxiv.org/abs/2209.12736.

[32] R. P. Newman. A proof of the splitting conjecture of S.-T. Yau. *Journal of differential geometry*, 31(1):163–184, 1990.

[33] R. Penrose. The geometry of impulsive gravitational waves. In *General relativity (papers in honour of J. L. Synge)*, pages 101–115. Clarendon Press, Oxford, 1972.

[34] A. Petrunin. Applications of quasigeodesics and gradient curves. In *Comparison geometry (Berkeley, CA, 1993–94)*, volume 30 of *Math. Sci. Res. Inst. Publ.*, pages 203–219. Cambridge Univ. Press, Cambridge, 1997.

[35] J. Podolský, C. Sämann, R. Steinbauer, and R. Švarc. The global uniqueness and C^1-regularity of geodesics in expanding impulsive gravitational waves. *Classical Quantum Gravity*, 33(19):23, 2016. https://arxiv.org/abs/1602.05020.

[36] J. Podolský, C. Sämann, R. Steinbauer, and R. Švarc. The global existence, uniqueness and C^1-regularity of geodesics in nonexpanding impulsive gravitational waves. *Classical Quantum Gravity*, 32(2):025003, 23, 2015.

[37] V. Toponogov. The metric structure of Riemannian spaces with nonnegative curvature which contain straight lines. *Amer. Math. Soc. Transl.*, 70:225–239, 1968.

[38] J. A. Vickers and J. P. Wilson. Generalized hyperbolicity in conical spacetimes. *Classical Quantum Gravity*, 17(6):1333–1360, 2000.

manuscripta math. 173, 667–710 (2024)

Tobias Beran · Felix Rott

Gluing constructions for Lorentzian length spaces

Received: 18 February 2022 / Accepted: 28 February 2023 /
Published online: 24 March 2023

Abstract. We introduce an analogue to the amalgamation of metric spaces into the setting of Lorentzian pre-length spaces. This provides a very general process of constructing new spaces out of old ones. The main application in this work is an analogue of the gluing theorem of Reshetnyak for CAT(k) spaces, which roughly states that gluing is compatible with upper curvature bounds. Due to the absence of a notion of spacelike distance in Lorentzian pre-length spaces we can only formulate the theorem in terms of (strongly causal) spacetimes viewed as Lorentzian length spaces.

1. Introduction

The theory of Lorentzian length spaces, introduced in [16], is a new approach to developing a synthetic description of Lorentzian geometry without relying on any differential geometric machinery. It is very much inspired by the relationship between metric geometry and Riemannian geometry, where in particular the theory of length spaces has led to fundamental contributions and essentially has given rise to a purely metric and synthetic point of view of Riemannian manifolds.
Lorentzian length spaces appear to be a very promising approach in this direction and are on the way to becoming an independent field of research, increasingly attracting many established researchers from Lorentzian geometry and general relativity. There have been a variety of interesting results concerning the advancement of Lorentzian length spaces of which we want to mention a few.

- [13] introduces a notion of (in)extendibility for Lorentzian length spaces.
- [3], via generalized cones, introduces an analogue to warped products into the setting of Lorentzian length spaces.
- [11] introduces optimal transport methods in Lorentzian length spaces, defines timelike Ricci curvature bounds via suitable entropy conditions and gives applications to general relativity (synthetic singularity theorems).
- [1] further develops the causal ladder for Lorentzian length spaces.

F. Rott (✉): Faculty of Mathematics, University of Vienna, Vienna, Austria.
e-mail: felix.rott@univie.ac.at
T. Beran
e-mail: tobias.beran@univie.ac.at

Mathematics Subject Classification: 53C23 (primary); 53C50 · 53B30 · 51F99 · 51K10 (secondary)

https://doi.org/10.1007/s00229-023-01469-4

- [15] examines the null distance in Lorentzian length spaces (which was first introduced in [20] for manifolds) and in turn studies Gromov-Hausdorff convergence, establishing first compatibility results with respect to curvature bounds.
- [10] studies (the existence of) time functions on Lorentzian length spaces.
- [17] defines an analogue to Hausdorff measure on Lorentzian length spaces.

1.1. Motivation and summary

Currently, the majority of research in Lorentzian length spaces is concentrated around direct applications to general relativity and only few works result from a purely metric motivation. Indeed, it still seems that many fundamental concepts and constructions from metric geometry have not yet been fully incorporated or are outright missing from the Lorentzian theory.

The main goal of this work is to adapt some of these missing concepts from metric geometry to the Lorentzian setting and in this way contribute to making it an equally applicable and impactful synthetic analogue to the metric theory of length spaces. Leading experts in the field of metric geometry suggest that an idea similar to the amalgamation of metric spaces is essential in this process. Indeed, the amalgamation of metric spaces is *the* fundamental construction for producing new spaces from old ones and thus showcases a significant advantage of metric spaces compared to (Riemannian) manifolds, where gluing is in general only possible along isometric/diffeomorphic boundaries, if at all. Instead, one can usually only consider Cartesian products or submanifolds, both of which offer much less flexibility.

One of the key results concerning gluing in the metric world is the gluing theorem of Reshetnyak: it states that the amalgamation of metric spaces which satisfy an upper curvature bound, so-called CAT(k) spaces, also satisfies the same upper curvature bound. Metric gluing has also found applications in the theory of semi-dispersing billiards, cf. [4,8,9], as well as in geometric group theory, cf. [6].

A gluing process for Lorentzian pre-length spaces turns out to be a more delicate matter than the corresponding process for metric spaces since, roughly speaking, there is much more compatibility one has to respect. In other words, a metric space only consists of a set with a distance function while a Lorentzian pre-length space is both a causal space and a metric space and moreover both have to behave well with respect to the time separation function. Our first task is to translate the metric amalgamation into the Lorentzian setting. In the metric case, it consists of two steps: first forming the disjoint union and then considering the quotient semi-metric with respect to the identifying equivalence relation. The disjoint union can be easily adapted but a "quotient time separation" needs to be treated a bit more carefully.

We continue with the preparations necessary for a Lorentzian analogue of the Reshetnyak gluing theorem, which is the central part of this work. Most important for this goal is to establish a gluing lemma for triangles in the sense of [6, Lemma II.4.10]. This turns out to be a quite technically demanding task. Even worse, without a solid concept of spacelike distance in Lorentzian pre-length spaces, there is no chance to achieve a reasonably general version of the gluing theorem. It does, however, work out when considering manifolds as Lorentzian pre-length spaces, where spacelike distances are well known: this is the content of the last chapter. This is also where the main result of this paper is formulated, namely:

Theorem 5.2.1. (Reshetnyak's gluing theorem, Lorentzian version) Let (X_1, g_1) and (X_2, g_2) be two smooth and strongly causal spacetimes with $\dim(X_1) =: n \geq m := \dim(X_2)$. Let A_1 and A_2 be two closed non-timelike locally isolating subsets of X_1 and X_2, respectively. Let $f : A_1 \to A_2$ be a τ-preserving and $\leq$-preserving locally bi-Lipschitz homeomorphism which locally preserves the signed distance. Suppose A_1 and A_2 are convex in the sense of Remark 5.1.1(iii). Suppose X_1 and X_2 have (sectional) curvature bounded above by $K \in \mathbb{R}$ in the sense of [2]. Then the Lorentzian amalgamation $X := X_1 \sqcup_A X_2$ is a Lorentzian pre-length space with timelike curvature bounded above by K. □

1.2. Outlook

We conclude the introduction by briefly discussing possible applications of gluing constructions in the (synthetic) Lorentzian setting.

- The "causal inheritance": Many steps of the causal ladder for spacetimes have been translated into the synthetic setting, cf. [1]. Given two Lorentzian pre-length spaces that are both situated somewhere on the causal ladder, can the same be said about their amalgamation? If not, are there additional properties that would guarantee the preservation of this property?
- The compatibility of the amalgamation and Gromov-Hausdorff convergence of Lorentzian length spaces with respect to the null distance: convergence of Lorentzian length spaces has been studied in [15]. Given two sequences of Lorentzian length spaces that converge each to a Lorentzian length space, can the same be said about the sequence of the respective amalgamations?
- An analogue to the collision theorem: concerning the theory of semi-dispersing billiards, the collision theorem is a particularly nice application of gluing in the metric world, cf. [4, Theorem 2.6.1]. In the Lorentzian case, this could be useful for investigating particle collisions in general relativity.
- Globalization: the metric version of the gluing lemma is used to globalize upper curvature bounds, see [6, Proposition II.4.9 & Lemma II.4.10]. A Lorentzian version of such a result would certainly be very interesting and might be possible with similar methods.
- General relativity: it is expected that gluing constructions can also be directly applied in various topics from general relativity. Examples include: extending a spacetime (or Lorentzian length space) by gluing, cosmic censorship and gluing at singularities, or matching of spacetimes and impulsive gravitational waves.

2. Preliminaries

By a spacetime (M, g) we mean a smooth manifold M with a Lorentzian metric g and a time orientation. Requiring the spacetime to be C^k means the metric g is C^k. We denote by η the ordinary Minkowski metric on $\mathbb{R}^n$. We write $I(x, z) :=$

$I^+(x) \cap I^-(z) = \{y \mid x \ll y \ll z\}$ for timelike diamonds and $J(x, z)$ for causal diamonds. By a hinge (α, β) we mean a configuration of two (timelike) geodesics α and β and the included (hyperbolic) angle, usually denoted by ω. For basic information regarding Lorentzian pre-length spaces see [16]. For basic information regarding the amalgamation and its compatibility with curvature conditions in the metric case see [6,7]. We will anyways present a very short recap of the most fundamental concepts concerning Lorentzian pre-length spaces and we will briefly describe the amalgamation in the metric picture.

2.1. *A brief introduction to Lorentzian pre-length spaces*

Simply put, a Lorentzian pre-length space encodes certain fundamental properties of a Lorentzian manifold while completely ignoring others. The focus lies on the causality relations and the time separation function, while the Lorentzian metric and the general manifold structure are discarded entirely. Compare this to metric geometry, where length spaces serve as a very useful generalization of Riemannian manifolds.

Definition 2.1.1. *(Lorentzian pre-length space)* A tuple $(X, d, \ll, \leq, \tau)$ is called a Lorentzian pre-length space if it satisfies the following:

(i) $(X, \ll, \leq)$ is a causal space, i.e., $\leq$ is a reflexive and transitive relation on X and $\ll$ is a transitive relation on X contained in $\leq$.
(ii) $\tau : X \times X \to [0, \infty]$ is lower semi-continuous with respect to the metric d.
(iii) τ respects the causal structure in the following way: τ satisfies the reverse triangle inequality for $\leq$-related points and is compatible with $\ll$ in the sense that $\tau(a, b) > 0 \iff a \ll b$.

Note that due to [16, Example 2.11] any smooth spacetime is a Lorentzian pre-length space (where the distance metric is induced by some (complete) Riemannian background metric). Moreover, any continuous causally plain metric on a spacetime yields a Lorentzian pre-length space, cf. [16, Proposition 5.8]

Definition 2.1.2. *(Causal/timelike curves)* Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space.

(i) A locally Lipschitz curve $\gamma : [a, b] \to X$ is called-future directed causal (respectively timelike), if $\gamma(s) \leq \gamma(t)$ (respectively $\gamma(s) \ll \gamma(t)$) for all $s, t \in [a, b]$, $s < t$. Past-directed curves are defined analogously. Unless explicitly stated otherwise, we assume all causal curves to be future-directed.
(ii) The τ-length of a causal curve γ is given as

$$L_\tau(\gamma) := \inf \left\{ \sum_{i=0}^{n} \tau(\gamma(t_i), \gamma(t_{i+1})) \mid a = t_0 < t_1 < \ldots < t_n = b, n \in \mathbb{N} \right\}. \tag{2.1.1}$$

If $\gamma(a) = x$, $\gamma(b) = y$ and $L_\tau(\gamma) = \tau(x, y)$ we say that γ is τ-realizing and we call (the image of) such a curve a geodesic segment.

The main difference between a Lorentzian length space and a Lorentzian pre-length space is in spirit the same as between a length space and a metric space. That is, the time separation function of a Lorentzian length space is intrinsic in the sense that it is given by the (supremum of the) lengths of connecting causal curves. There are also some additional technical assumptions on a Lorentzian length space resembling the existence of small "convex" neighbourhoods. As we will mainly work with Lorentzian pre-length spaces, we only refer to the definition, see [16, Definition 3.22].

The final ingredient we will need is the description of curvature bounds. As in the metric world, triangle comparison replaces the concept of sectional curvature bounds. We denote by M_K the Lorentzian model space of constant sectional curvature K, cf. [16, Definition 4.5]. That is, M_K is either an appropriately scaled version of de Sitter- or anti de Sitter space or the Minkowski plane. We may denote the Lorentzian metric on M_K by $\langle\cdot,\cdot\rangle$. Unless explicitly stated otherwise, we assume all mentioned triangles to satisfy the appropriate size bounds for M_K, cf. [2, Lemma 2.1] or [16, Lemma 4.6].

Definition 2.1.3. *(Timelike curvature bounds)* A Lorentzian pre-length space $(X, d, \ll, \leq, \tau)$ has timelike curvature bounded below (respectively above) by $K \in \mathbb{R}$ if every point in X has a neighbourhood U, called a comparison neighbourhood, which satisfies the following:

(i) $\tau|_{U\times U}$ is finite and continuous.

(ii) For all $x, y \in U$ with $x \ll y$ there exists a τ-realizing curve entirely contained in U.

(iii) Let $\Delta(x, y, z)$ be a timelike triangle in U, i.e., $x \ll y \ll z$ and we have τ-realizing curves connecting these points pairwise. Let $\Delta(\bar{x}, \bar{y}, \bar{z})$ be its comparison triangle in M_K. Then for all $p, q \in \Delta(x, y, z)$ and corresponding comparison points $\bar{p}, \bar{q}$ in the comparison triangle we have $\tau(p, q) \leq \bar{\tau}(\bar{p}, \bar{q})$ (respectively $\tau(p, q) \geq \bar{\tau}(\bar{p}, \bar{q})$).

2.2. A brief introduction to metric amalgamation and the gluing theorem

Here, we collect all metric prerequisites needed for a Lorentzian gluing construction, following [6,7].

Definition 2.2.1. *(Disjoint union metric)* Let $(X_i, d_i)_{i\in I}$ be a family of metric spaces. Let $X := \sqcup_{i\in I} X_i$ be the disjoint union. Then

$$d(x, y) := \begin{cases} d_i(x, y) & x, y \in X_i \\ \infty & \text{else.} \end{cases} \tag{2.2.1}$$

defines a metric on X, called the disjoint union metric.

Definition 2.2.2. *(Quotient semi-metric)* Let (X, d) be a metric space and let $\sim$ be an equivalence relation on X. The quotient semi-metric with respect to $\sim$ is defined as

$$\widetilde{d}([x], [y]) := \inf \left\{ \sum_{i=1}^{n} d(x_i, y_i) \mid x \sim x_1, x_{i+1} \sim y_i, y_n \sim y, n \in N \right\} \tag{2.2.2}$$

Definition 2.2.3. *(Amalgamation)* Let $(X_i, d_i)_{i\in I}$ be a family of metric spaces. Let $(A_i)_{i\in I}$ be a family of closed subspaces each of which is isometric to some metric space A via the isometry $f_i : A \to A_i$. Equip the disjoint union $X := \sqcup_{i\in I} X_i$ with the metric d from above. On X let $\sim$ be the equivalence relation generated by the condition $f_i(a) \sim f_j(a)$ for all $i, j \in I, a \in A$. Then the quotient of X equipped with the quotient semi-metric with respect to $\sim$ is called the amalgamation of the family $(X_i)_{i\in I}$ with respect to A and is denoted by $\sqcup_A X_i$, i.e., $\sqcup_A X_i = (X/\sim, \widetilde{d})$.

Finally, we mention the gluing theorem of Reshetnyak. In the following formulation, it is in fact possible to omit the assumption of each X_i being proper, but the proof then gets significantly more difficult. For a proof, see e.g. [6, Theorem II.11.1 & Theorem II.11.3].

Theorem 2.2.4. (Reshetnyak) *Let $(X_i, d_i)_{i\in I}$ be a family of proper CAT(k) spaces. Let $(A_i)_{i\in I}$ be a family of closed convex and complete subspaces each of which is isometric to some metric space A. Then the amalgamation $\sqcup_A X_i$ is a CAT(k) space.* □

3. Lorentzian structure on a quotient

In this section, we introduce the amalgamation construction for Lorentzian pre-length spaces. To this end we first have to discuss more elementary aspects of gluing.

3.1. Basic gluing preparations

We begin introducing a Lorentzian structure on the quotient of a Lorentzian pre-length space by adapting the definition of the quotient semi-metric with additional causality assumptions.

Remark 3.1.1. (On notation and conventions I) As any distance metric and time separation function takes values only in $[0,\infty]$, we set $\sup\emptyset = 0$ and $\inf\emptyset = \infty$ for the sake of convenience. We will sometimes apply shortcuts commonly used in the theory of metric spaces and just write X for a Lorentzian pre-length space $(X, d, \ll, \leq, \tau)$. By a subspace $A \subseteq X$ we mean a subset viewed as a Lorentzian pre-length space equipped with the restriction of the original metric, causality relations and time separation. Moreover, we will write $[x, y]$ for a geodesic segment between x and y. Either the context or a more detailed description will prevent any ambiguity. We will usually write $\widetilde{X} := X/\sim$ for the (topological) quotient of X with respect to an equivalence relation $\sim$. We will denote the natural projection $x \mapsto [x]$ by $\pi : X \to \widetilde{X}$.

Definition 3.1.2. *(Quotient time separation)* Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space and let $\sim$ be an equivalence relation on X. The quotient time-separation function is defined as $\widetilde{\tau} : \widetilde{X} \times \widetilde{X} \to [0,\infty]$,

$$\widetilde{\tau}([x],[y]) := \sup\left\{\sum_{i=1}^{n} \tau(x_i, y_i) \mid x \sim x_1\right.$$

$$\leq y_1 \sim x_2 \leq y_2 \sim \ldots \sim x_n \leq y_n \sim y, n \in \mathbb{N}\} . \qquad (3.1.1)$$

We call a sequence $(x_1, y_1, \ldots, x_n, y_n)$ as above an n-chain from $[x]$ to $[y]$.

Remark 3.1.3. (Restricting the set of chains) Note that we can always assume without loss of generality that $x_1 = x$ and $y_n = y$. Indeed, suppose $(x_1, y_1, \ldots, x_n, y_n)$ is an n-chain from $[x]$ to $[y]$, then $(x, x, x_1, y_1, \ldots, x_n, y_n, y, y)$ is an $(n+2)$-chain with at least the same length.
Furthermore, we can always assume that $y_i \neq x_{i+1}$. Otherwise by the reverse triangle inequality for τ we could replace

$$\tau(x_i, y_i) + \tau(x_{i+1}, y_{i+1}) = \tau(x_i, y_i) + \tau(y_i, y_{i+1}) \leq \tau(x_i, y_{i+1}) \qquad (3.1.2)$$

to obtain a longer chain. We say that the relation between y_i and x_{i+1} is nontrivial if $y_i \neq x_{i+1}$ and $y_i \sim x_{i+1}$.

We define both causality relations on $\widetilde{X}$ via the quotient time separation.

Definition 3.1.4. *(Quotient causality)* Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space and let $\sim$ be an equivalence relation on X. On $\widetilde{X}$, we define $[x] \mathrel{\widetilde{\ll}} [y] :\iff \widetilde{\tau}([x], [y]) > 0$ and $[x] \mathrel{\widetilde{\leq}} [y] :\iff \{\sum_{i=1}^{n} \tau(x_i, y_i) \mid x \sim x_1 \leq y_1 \sim x_2 \leq y_2 \sim \ldots \sim x_n \leq y_n \sim y, n \in \mathbb{N}\} \neq \emptyset$. Especially the causal relation might be better described in words: we have $[x] \mathrel{\widetilde{\leq}} [y]$ if and only if there exists a chain from $[x]$ to $[y]$. Regarding the timelike relation, we have $[x] \mathrel{\widetilde{\ll}} [y]$ if and only if there exists a chain of positive length from $[x]$ to $[y]$.

Note that without additional assumptions (on e.g. $\sim$) this construction is badly behaved or does not yield a Lorentzian pre-length space at all. This is not very surprising, and in some sense parallels the metric world, where the quotient semi-metric might not be positive definite. There are, however, some properties of a Lorentzian pre-length space that any quotient satisfies.

Proposition 3.1.5. (Quotient causal space) *Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space and let $\sim$ be an equivalence relation on X. Then $(\widetilde{X}, \widetilde{\ll}, \widetilde{\leq})$ is a causal space.*

Proof. The inclusion of $\widetilde{\ll}$ in $\widetilde{\leq}$ is clear from the definition. Suppose $[x] \mathrel{\widetilde{\ll}} [y] \mathrel{\widetilde{\ll}} [z]$. The concatenation of chains with positive length from $[x]$ to $[y]$ and from $[y]$ to $[z]$, respectively, results in a chain with positive length from $[x]$ to $[z]$. Hence $\widetilde{\tau}([x], [z]) > 0$ and so $[x] \mathrel{\widetilde{\ll}} [z]$. By the same argument we have that if there exists a chain from $[x]$ to $[y]$ and a chain from $[y]$ to $[z]$ then there exists a chain from $[x]$ to $[z]$. Thus, $[x] \mathrel{\widetilde{\leq}} [y] \mathrel{\widetilde{\leq}} [z]$ implies $[x] \mathrel{\widetilde{\leq}} [z]$. The reflexivity of $\widetilde{\leq}$ follows from the reflexivity of $\leq$: (x, x) is a valid 1-chain for any $[x] \in \widetilde{X}$ and so $[x] \mathrel{\widetilde{\leq}} [x]$. □

Proposition 3.1.6. (Reverse triangle inequality) *Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space and let $\sim$ be an equivalence relation on X. Then $\widetilde{\tau}$ satisfies the reverse triangle inequality for causally related points, i.e., if $[x] \mathrel{\widetilde{\leq}} [y] \mathrel{\widetilde{\leq}} [z]$, then $\widetilde{\tau}([x], [z]) \geq \widetilde{\tau}([x], [y]) + \widetilde{\tau}([y], [z])$.*

Proof. This follows immediately from the definition: for any chain from $[x]$ to $[y]$ and any chain from $[y]$ to $[z]$, their concatenation results in a chain from $[x]$ to $[z]$. Since there might be chains from $[x]$ to $[z]$ without going through $[y]$, $\widetilde{\tau}([x],[z])$ can only get larger. More precisely, given $\varepsilon > 0$, let $(x_1, y_1, \ldots, x_n, y_n)$ and $(y_1', z_1, \ldots, y_m', z_m)$ be chains from $[x]$ to $[y]$ and from $[y]$ to $[z]$ with lengths at least $\widetilde{\tau}([x],[y]) - \frac{\varepsilon}{2}$ and $\widetilde{\tau}([y],[z]) - \frac{\varepsilon}{2}$, respectively. Then $(x_1, y_1, \ldots, x_n, y_n, y_1', z_1, \ldots, y_m', z_m)$ is a chain from $[x]$ to $[z]$ with length at least $\widetilde{\tau}([x],[y]) + \widetilde{\tau}(y,z) - \varepsilon$ and the claim follows. □

In summary, we obtain the following intuitive properties on any Lorentzian quotient. This can be thought of as the analogue to "gluing can only shrink distances" in the metric case.

Corollary 3.1.7. (Elementary properties of Lorentzian quotients) *Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space and $\sim$ an equivalence relation on X. The quotient Lorentzian structure always satisfies the following properties for all $[x], [y] \in X$.*

(i) $\widetilde{\tau}([x],[y]) \geq \tau(x,y)$.
(ii) $x \ll y \Rightarrow [x] \mathrel{\widetilde{\ll}} [y]$ *and* $x \leq y \Rightarrow [x] \mathrel{\widetilde{\leq}} [y]$.

□

Finally, we observe that the only property that might prevent the quotient of a Lorentzian pre-length space from being a Lorentzian pre-length space itself is that $\widetilde{\tau}$ need not be lower semi-continuous. The following examples show how the lower semi-continuity of $\widetilde{\tau}$ may fail and how this can be prevented.

Example 3.1.8. (Showcasing gluing in the Minkowski plane) Consider the ordinary Minkowski plane $\mathbb{R}^2_1$. Identify two spacelike related points x and y as in Fig. 1. Then the resulting space is not a Lorentzian pre-length space since $\widetilde{\tau}$ is not lower semi-continuous. To see this, let $p \in \partial J^-(x) \setminus J^-(y)$ and $q \in I^+(y) \setminus J^+(x)$. Then $\widetilde{\tau}(p,q) > 0$ (red line). But if we choose a sequence $(p_n)_{n\in\mathbb{N}}$ such that $p_n \to p$ and $p_n \notin J^-(x) \cup I^-(q)$ for all $n \in \mathbb{N}$, then $\widetilde{\tau}(p_n, q) = 0$.
Identifying a closed vertical strip, say $(1,t) \sim (2,t)$ for all $t \in [0,1]$ has exactly the same problem. Doing the identification along an open strip, say $(1,t) \sim (2,t)$ for all $t \in (0,1)$ at first glance seems to eliminate this obstacle. In this case, however, the quotient semi-metric will not be positive definite since for example we would have $\widetilde{d}((1,0),(2,0)) = 0$. Intuitively, we need the identified sets to be (topologically) closed and at the same time always have timelike related points nearby. Identifying $(1,t) \sim (2,t)$ for all $t \in \mathbb{R}$ covers both conditions. This actually turns out to be a Lorentzian pre-length space. The lower semi-continuity immediately follows from the more general proof of below.

3.2. Amalgamation prerequisites

In the spirit of metric amalgamation, we first introduce a very easy construction of formally viewing two Lorentzian pre-length spaces as one.

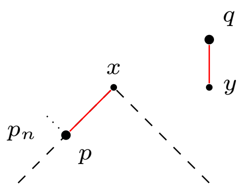

Fig. 1. $\widetilde{\tau}$ is not lower semi-continuous

Definition 3.2.1. *(Lorentzian disjoint union)* Let $(X_1, d_1, \ll_1, \leq_1, \tau_1)$ and $(X_2, d_2, \ll_2, \leq_2, \tau_2)$ be two Lorentzian pre-length spaces and set $X := X_1 \sqcup X_2$. Define $\leq := \leq_1 \sqcup \leq_2$, i.e., "$\leq \subseteq X \times X$" and $x \leq y :\iff \exists i \in \{1, 2\} : x, y \in X_i \wedge x \leq_i y$. Similarly, define $\ll := \ll_1 \sqcup \ll_2$. Let d be the disjoint union metric on X, cf. (2.2.1). Define $\tau : X \times X \to [0, \infty]$ by

$$\tau(x, y) := \begin{cases} \tau_i(x, y) & x, y \in X_i \\ 0 & \text{else.} \end{cases} \tag{3.2.1}$$

We call $(X, d, \ll, \leq, \tau)$ the Lorentzian disjoint union of X_1 and X_2.

Proposition 3.2.2. (Disjoint union Lorentzian pre-length space) *Let* $(X_1, d_1, \leq_1, \ll_1, \tau_1)$ *and* $(X_2, d_2, \ll_2, \leq_2, \tau_2)$ *be two Lorentzian pre-length spaces. Then the Lorentzian disjoint union* $(X, d, \ll, \leq, \tau)$ *is a Lorentzian pre-length space.*

Proof. Clearly, $(X, \ll, \leq)$ is a causal space. The reverse triangle inequality is directly inherited from the respective inequalities in X_1 and X_2, since causal relation can only occur between points coming from the same spaces. Similarly, the lower semi-continuity of τ is inherited in this way, since for $x_n \to x$, say $x \in X_1$, for large enough n_0 we have $x_n \in X_1$ for all $n \geq n_0$. Finally, the compatibility of the causal relations with τ also follows directly from their counterparts in X_1 and X_2. $\square$

Analogous to the metric case, we define the Lorentzian amalgamation as a quotient of the Lorentzian disjoint union where we identify certain subsets. To ensure that this construction actually results in a Lorentzian pre-length space (i.e., that $\widetilde{\tau}$ is lower semi-continuous) we require the following property of the identified subsets.

Definition 3.2.3. *(Local timelike isolation)* A subset A of a Lorentzian pre-length space $(X, d, \ll, \leq, \tau)$ is said to be non-future locally isolating if for all $a \in A$ with $I^+(a) \neq \emptyset$ and for all neighbourhoods $U_a \subseteq A$ of a there exists $b_+ \in U_a$ such that $a \ll b_+$. Similarly, we define a non-past locally isolating set. We say A is non-timelike locally isolating if it satisfies both properties.

Remark 3.2.4. (Examples and comments on local timelike isolation) Clearly, the open image of a timelike curve is non-timelike locally isolating. Furthermore, for any Lorentzian length space X, the set X is non-timelike locally isolating by the sequence lemma, cf. [1, Lemma 2.18]. But note that of course a subset of a Lorentzian length space is not a Lorentzian length space in general (with the restricted structure). We make the additional assumption of $I^{\pm}(a) \neq \emptyset$ to also allow gluing of spaces with "future/past boundary". For example, consider a closed rectangle in the Minkowski plane as a Lorentzian pre-length space and identify two vertical line segments (which are not causally related at all). Then the boundary points of these segments would fail to have the non-isolating property introduced above, but only because there is nothing in the future (respectively in the past) of these points to begin with. In this case also the counterexample of Example 3.1.8 fails: the quotient time separation of a sequence approaching the past of the future boundary point of one segment cannot have a positive value in the limit since on the other segment this is also a future boundary point with empty future.

In the metric case, amalgamation usually occurs along maps which at least locally "preserve structure", i.e., local isometries. This turns out to be not necessary to define our gluing process. So fore the sake of generality, we will formulate the Lorentzian amalgamation with as few assumptions as possible. Not surprisingly, this may be badly behaved, which is why we will almost exclusively work with additional assumptions. Next, we introduce the (for our purposes) correct notion of structure preserving maps in Lorentzian pre-length spaces. The following definition is closely related to corresponding notions in [1,13].

Definition 3.2.5. *(Structure preserving maps)* Let X_1 and X_2 be two Lorentzian pre-length spaces.

(i) A map $f : X_1 \to X_2$ is called τ-preserving if $\tau_1(x, y) = \tau_2(f(x), f(y))$ for all $x, y \in X_1$.
(ii) A map $f : X_1 \to X_2$ is called $\ll$-preserving if $x \ll_1 y \iff f(x) \ll_2 f(y)$ for all $x, y \in X_1$. It is called $\leq$-preserving if $x \leq_1 y \iff f(x) \leq_2 f(y)$ for all $x, y \in X_1$. If f is both $\ll$-preserving and $\leq$-preserving it is called causality preserving.
(iii) A map $f : X_1 \to X_2$ is called locally τ, $\ll$ or $\leq$-preserving if for all $x \in X_1$ there exists a neighbourhood $U \subseteq X_1$ of x such that $f|_U$ is τ, $\ll$ or $\leq$-preserving.

Remark 3.2.6. (Implication and counterexample) One the one hand, it is clear that any τ-preserving map is $\ll$-preserving. On the other hand, a τ-preserving map need not be $\leq$-preserving in general. Indeed, consider in the Minkowski plane a null segment and a spacelike segment, say $\{(s, s) \mid s \in [0, 1]\}$ and $\{(0, s) \mid s \in [0, 1]\}$. Then the map $(s, s) \mapsto (0, s)$ is τ-preserving (it vanishes identically in both cases). But $(s, s) \leq (t, t) \iff s \leq t$ while $(0, s)$ and $(0, t)$ are never causally related.

The following is immediate from the definition.

Corollary 3.2.7. (Inverse is also preserving) *If a bijective map $f : X_1 \to X_2$ between two Lorentzian pre-length spaces is (locally) τ-preserving (respectively $\ll$ or $\leq$-preserving), then so is its inverse. In particular, the neighbourhoods in the local case are compatible in the sense that $f|_U$ is preserving if and only if $f^{-1}|_{f(U)}$ is.* □

As a final prerequisite, we discuss the underlying quotient semi-metric. In metric amalgamation, gluing happens along isometric subsets. In the Lorentzian setting the focus lies on the time separation and the causality relations where as the distance metric plays only a background role as a topological tool. Since we decided to not require any preservation of Lorentzian structure of the identified sets, it is too restrictive to insist on using metric (local) isometries. That is why we decided to use locally bi-Lipschitz homeomorphisms[1] instead. The bi-Lipschitz condition ensures that the quotient semi-metric is positive definite and, when assuming additionally some causality preservation, that a causal curve in one of the identified sets is also a causal curve in the other. Note that in this case, however, there is in general no chance of obtaining a nice representation of the quotient semi-metric á la [6, Lemma I.5.24].

Proposition 3.2.8. (Metric amalgamation with locally bi-Lipschitz maps) *Let X_1 and X_2 be two metric spaces and let $A_1 \subseteq X_2$ and $A_2 \subseteq X_2$ be closed subspaces. Let A be a metric space and let $f_1 : A \to A_1$ and $f_2 : A \to A_2$ be locally bi-Lipschitz homeomorphisms. Let d be the disjoint union metric on $X_1 \sqcup X_2$ and consider the equivalence relation generated by $f_1(a) \sim f_2(a)$ for all $a \in A$. Then the quotient semi-metric $\widetilde{d}$ on $X := (X_1 \sqcup X_2)/\sim$ with respect to $\sim$ is a metric.*

Proof. $\widetilde{d}([x],[x]) = 0$ for all $[x] \in X$ and the symmetry of $\widetilde{d}$ are immediate from the definition. Concerning the triangle inequality, let $(x_1, y_1, \ldots, x_n, y_n)$ be an n-chain from $[x]$ to $[y]$ such that $\sum_{i=1}^n d(x_i, y_i) < \widetilde{d}([x],[y]) + \varepsilon$ and let $(y_1', z_1, \ldots, y_m', z_m)$ be an m-chain from $[y]$ to $[z]$ such that $\sum_{i=1}^m d(y_i', z_i) < \widetilde{d}([y],[z]) + \varepsilon$. Then the concatenation $(x_1, y_1, \ldots, x_n, y_n, y_1', z_1, \ldots, y_m', z_m)$ is a chain from $[x]$ to $[z]$ of length less than $\widetilde{d}([x],[y]) + \widetilde{d}([y],[z]) + 2\varepsilon$. Hence $\widetilde{d}([x],[z]) \leq \widetilde{d}([x],[y]) + \widetilde{d}([y],[z]) + 2\varepsilon$ and the claim follows. It is left to show that $\widetilde{d}$ is positive definite. Let $[x],[y] \in X$. Note that if, say $[x] \in X_1 \setminus A_1$, i.e., $[x] = \{x^1\}$, then any chain that is not (x^1, y^1) has to go through A_1 by similar arguments as in Remark 3.1.3. In both cases we end up with a positive distance: $d(x^1, y^1) > 0$ since d is a metric and $d(x^1, A_1) > 0$ since A is closed. So the only case left to consider is when both points lie in A. To this end, note that since f_1 and f_2 are locally bi-Lipschitz homeomorphisms, it follows that $f := f_2 \circ f_1^{-1} : A_1 \to A_2$ and $f^{-1} : A_2 \to A_1$ are locally bi-Lipschitz homeomorphisms as well. Let $[x] \neq [y], [x],[y] \in A$. Then $x^1, y^1 \in A_1$ and $x^2, y^2 \in A_2$. Let $r > 0$ be such that $y^1 \notin B_r^{A_1}(x^1)$ and $y^2 \notin B_r^{A_2}(x^2)$, where $B_r^{A_i}(x^i)$ is the r-ball of x^i in the metric subspace A_i. By choosing r smaller if necessary, we can assume that $B_r^{A_1}(x^1)$ and $B_r^{A_2}(x^2)$ are bi-Lipschitz neighbourhoods for f and f^{-1} of x^1 and x^2 in A_1 and A_2, respectively. We can furthermore suppose that this is with respect to the same Lipschitz constant $L \geq 1$. That is, we have

$$\frac{1}{L} d_2(f_1(a), f_1(b)) \leq d_1(a,b) \leq L d_2(f(a), f_1(b)) \tag{3.2.2}$$

[1] We call a bijective map $f : X \to Y$ between metric spaces locally bi-Lipschitz if every $x \in X$ has a neighbourhood U such that $f|_U : U \to f(U)$ and its inverse are Lipschitz.

for all $a, b \in B_r^{A_1}(x^1)$ and

$$\frac{1}{L}d_1(f^{-1}(a'), f^{-1}(b')) \le d_2(a', b') \le L d_1(f^{-1}(a'), f^{-1}(b')) \tag{3.2.3}$$

for all $a', b' \in B_r^{A_2}(x^2)$. By the above arguments we can assume that any chain is of the form $(x^1, f(a_1), f^{-1}(a_1), f^{-1}(a_2), f(a_2), f(a_2), \dots, f(a_{n-1}), y^1)$ (it does not matter whether we start in x^1 or x^2 since we could add (x^2, x^2) at the beginning of the chain without increasing its length, and similar for ending in y^1). Given such a chain, by setting $a_0 = x^1$ and $a_n = y^1$, there exists a minimal $j \in \{1, \dots, n\}$ such that $a_i \notin B_r^{A_1}(x^1)$. Similarly, there exists a minimal $k \in \{1, \dots, n\}$ such that $f(a_k) \notin B_r^{A_2}(x^2)$. Without loss of generality assume $j \le k$. Then we compute

$$\begin{aligned}
&d_1(x^1, a_1) + d_2(f(a_1), f(a_2)) + d_1(a_2, a_3) + \ldots + d_1(a_{n-1}, y^1) \ge \\
&d_1(x^1, a_1) + \frac{1}{L}d_1(a_1, a_2) + d_1(a_2, a_3) + \ldots + d_1(a_{n-1}, y^1) \ge \\
&\frac{1}{L}d_1(x^1, a_1) + \frac{1}{L}d_1(a_1, a_2) + \frac{1}{L}d_1(a_2, a_3) + \ldots + d_1(a_{n-1}, y^1) \ge \\
&\frac{1}{L}d_1(x^1, a_3) + \ldots + d_1(a_{n-1}, y^1) \ge \ldots \ge \\
&\frac{1}{L}d_1(x^1, a_{j-1}) + d_1(a_{j-1}, a_j) + \ldots + d_1(a_{n-1}, y^1) \ge \\
&\frac{1}{L}d_1(x^1, a_{j-1}) + \frac{1}{L}d_1(a_{j-1}, a_j) + \ldots + d_1(a_{n-1}, y^1) \ge \\
&\frac{1}{L}d_1(x^1, a_j) + \ldots + d_1(a_{n-1}, y^1) \ge \\
&\frac{1}{L}r + \ldots + d_1(a_{n-1}, y^1) > \frac{1}{L}r > 0.
\end{aligned}$$

Thus, we found a uniform lower bound for any chain from $[x]$ to $[y]$ and so $\widetilde{d}$ is positive definite. □

3.3. Lorentzian amalgamation

We now introduce the central object of this paper, which allows us to create a new Lorentzian pre-length space out of old ones by gluing them together, a process similar to the amalgamation of metric spaces. To avoid pathological counterexamples, we have to make minor additional assumptions on the identified subsets.

Definition 3.3.1. *(Lorentzian amalgamation)* Let $(X_1, d_1, \le_1, \ll_1, \tau_1)$ and $(X_2, d_2, \le_2, \ll_2, \tau_2)$ be two Lorentzian pre-length spaces. Let $(A, d_A, \ll_A, \le_A, \tau_A)$ be a Lorentzian pre-length space and let A_1 and A_2 be closed non-timelike locally isolating subspaces of X_1 and X_2, respectively. Let $f_1 : A \to A_1$ and $f_2 : A \to A_2$ be locally bi-Lipschitz homeomorphisms. Suppose that the causality of A_1 and A_2 are compatible in the following sense: for all $a \in A$ we have $I_1^{\pm}(f_1(a)) \neq \emptyset \iff I_2^{\pm}(f_2(a)) \neq \emptyset$. Let $(X_1 \sqcup X_2, d, \ll, \le, \tau)$ be the Lorentzian disjoint union of X_1 and X_2 and consider the equivalence relation $\sim$ on $X_1 \sqcup X_2$ generated by $f_1(a) \sim f_2(a)$ for all $a \in A$. Then $((X_1 \sqcup X_2)/\sim, \widetilde{d}, \widetilde{\ll}, \widetilde{\leq}, \widetilde{\tau})$ is called the Lorentzian amalgamation (with respect to A) of X_1 and X_2 and is denoted by $X_1 \sqcup_A X_2$.

Remark 3.3.2. (On notation and conventions II) As in the metric case, the space A is usually introduced for easier notation, especially when dealing with an amalgamation of more than two spaces. One can think of A as being any of the identified spaces A_i. For just two spaces, it might be more convenient to only consider a map $f : A_1 \to A_2$ (and its inverse). In fact, we will formulate the gluing theorem for two spaces only. Since most of the results in the current chapter can easily be generalized to several spaces, we will work with an extra space A.
We will use slightly sloppy notation and refer to A as a subset of the amalgamation $X := X_1 \sqcup_A X_2$ and we will view X_1 and X_2 as subsets of X via the identifications with $\pi(X_1)$ and $\pi(X_2)$, respectively. For example, by $[a] \in A$ we mean $[a] = \{f_1(a), f_2(a)\}$. If a point is originally not in one of the identified sets, say $x \in X_1 \setminus A_1$, then it is from its on equivalence class, i.e., $[x] = \{x\}$. If we do not care which space such a singleton is from, we may write $[x] \in X \setminus A$. Occasionally, it will be more convenient to omit the identifying maps f_i. We then denote the origin of a point by a superscript. That is, if $[a] \in A$ we have $[a] = \{a^1, a^2\}$ and if $[x] \in X_1 \setminus A_1 \subseteq X \setminus A$ then $[x] = \{x^1\}$.
Dealing with pasts and futures in the amalgamation can also be a bit tricky notation-wise, even if we require that the identifying maps preserve structure. To this end we denote by a subscript with respect to which causality the set is constructed. For example, we write $J_1^+(x^1) = \{y^1 \in X_1 \mid x^1 \leq_1 y^1\}$ or $I_X([x],[y]) = I_X^+([x]) \cap I_X^-([y]) = \{[p] \in X \mid [x] \widetilde{\ll} [p] \widetilde{\ll} [y]\}$.

Next, we formulate a lemma that is important to show the last remaining property of semi-continuity of $\widetilde{\tau}$.

Lemma 3.3.3. (Restricting to timelike chains) *Let $X := X_1 \sqcup_A X_2$ be the Lorentzian amalgamation of two Lorentzian pre-length spaces X_1 and X_2. If $\widetilde{\tau}([x],[y]) > 0$ then there is a timelike chain, i.e., $x_i \ll y_i$ for all i, from $[x]$ to $[y]$ whose length is arbitrarily close to $\widetilde{\tau}([x],[y])$.*

Proof. Let $[x],[y] \in X$ with $\widetilde{\tau}([x],[y]) > 0$. Given small enough $\varepsilon > 0$ we find a chain $(x_1, y_1, \ldots, x_n, y_n)$ such that $\sum_{i=1}^n \tau(x_i, y_i) > \widetilde{\tau}([x],[y]) - \varepsilon > 0$. By Remark 3.1.3, we can assume that it is of the form $\tau_1(x^1, f_1(a_1)) + \tau_2(f_2(a_1), f_2(a_2)) + \ldots + \tau_2(f_2(a_n), y^2)$. Note that the assumption of $x^1 \in [x]$ does not lose any generality: if this were not the case the chain simply would start out in X_2. Similarly for $y^2 \in [y]$. Since the length of this chain is positive, we have that at least one entry is positive, say $\tau_1(f_1(a_j), f_1(a_{j+1})) > 0$. Then $I_1^+(f_1(a_j)) \neq \emptyset$ and hence also $I_2^+(f_2(a_j)) \neq \emptyset$. Suppose $\tau_2(f_2(a_{j-1}), f_2(a_j)) = 0$. Consider the neighbourhood $S_1 := \{a \in A_1 \mid \tau(a, f_1(a_{j+1})) > \tau(f_1(a_j), f_1(a_{j+1}) - \frac{\varepsilon}{n}\}$ of $f_1(a_1)$ in A_1, which is open by the lower semi-continuity of τ_1. Since f_1 and f_2 are homeomorphisms it follows that $S_2 := f_2(f_1^{-1}(S_1))$ is an open neighbourhood of $f_2(a_j)$ in A_2. Since A_2 is non-timelike locally isolating and $I_2^+(f_2(a_j)) \neq \emptyset$, we

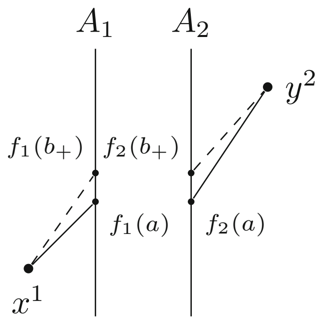

Fig. 2. Moving a to make the null piece connecting x^1 and $f_1(a)$ timelike

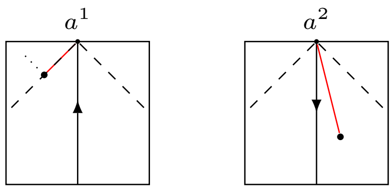

Fig. 3. This amalgamation leads to a $\widetilde{\tau}$ that is not lower semi-continuous. Note that the arrows denote the time orientation and not the gluing identifications

find $f_2(b_+) \in S_2$ such that $f_2(a_j) \ll f_2(b_+)$. Then by the push-up property of τ_1 we have $f_2(a_{j-1}) \ll f_2(b_+)$. Also, $f_1(b_+) \in S_1$. Then

$$\tau_2(f_2(a_{j-1}), f_2(b_+)) + \tau_1(f_1(b_+), f_1(a_{j+1})) > \tau_2(f_2(a_{j-1}), f_2(a_j)) + \tau_1(f_1(a_j), f_1(a_{j+1})) - \frac{\varepsilon}{n}. \tag{3.3.1}$$

We can do this for all other entries of this chain to end up with a timelike chain whose length is at most ε less than the chain we started with. Figure 2 illustrates this process. □

Example 3.3.4. (On minimal assumptions and bad behaviour) In Definition 3.3.1 the compatibility of the causality in A_1 and A_2 via $I_1^{\pm}(f_1(a)) \neq \emptyset \iff I_2^{\pm}(f_2(a)) \neq \emptyset$ for all $a \in A$, which is weaker than f_1 and f_2 being locally $\ll$-preserving, is truly a necessary condition. Let X_1 and X_2 be a closed unit square in the Minkowski plane equipped with the restricted Lorentzian structure. Identify them along a vertical line segment and reverse the time orientation in one of the squares, as is indicated by the arrows in Fig. 3. Then $I_1^+(a^1) = \emptyset$ but $I_2^+(a^2) \neq \emptyset$. Choosing endpoints and a sequence as in Example 3.1.8 leads to a similar failure of the lower semi-continuity of $\widetilde{\tau}$.
However, this does not mean that every Lorentzian pre-length space resulting from the amalgamation process is well behaved. Indeed, similar to the above example let X_1 and X_2 be the whole Minkowski plane and identify them along a vertical line with reversed time orientation in one space. Then although X_1 and X_2 are (much more than) chronological, the amalgamation fails to be chronological and moreover we have $\widetilde{\tau}([x], [y]) = \infty$ for all $[x], [y]$.

Remark 3.3.5. (Convergence in the amalgamation) Suppose $[x_n] \to [x]$. Clearly, if $[x] = \{x^1\}$ then $[x_n] = \{x_n^1\}$ for large enough n. If $[x] \in A$, then at least in one of X_1 or X_2, say X_1, there exists a subsequence $x^1_{n_k} \to x^1$. Indeed, if this were not the case then there would exist neighbourhoods of x^1 and x^2 in X_1 and X_2, respectively, that do not contain any points of the sequence. This is in contradiction to how the quotient topology is defined.

Proposition 3.3.6. (Amalgamation is Lorentzian pre-length space) *Let X be the Lorentzian amalgamation of two Lorentzian pre-length spaces X_1 and X_2. Then X is a Lorentzian pre-length space.*

Proof. We only need to prove that $\widetilde{\tau}$ is lower semi-continuous. Let $[x], [y] \in X$ and let $[x_n] \to [x], [y_n] \to [y]$. Then we need to show $\widetilde{\tau}([x_n], [y_n]) \geq \widetilde{\tau}([x], [y]) - \varepsilon$ for all $\varepsilon > 0$. Note that if $\widetilde{\tau}([x], [y]) = 0$ there is nothing to show, so suppose $\widetilde{\tau}([x], [y]) > 0$. Given a small enough $\varepsilon > 0$, by Remark 3.1.3 and Lemma 3.3.3 we find a timelike chain such that

$$\tau_1(x^1, f_1(a_1)) + \tau_2(f_2(a_1), f_2(a_2)) + \ldots + \tau_2(f_2(a_m), y^2) > \widetilde{\tau}([x], [y]) - \varepsilon > 0. \tag{3.3.2}$$

Assume ε is so small that even $\widetilde{\tau}([x], [y]) - 3\varepsilon > 0$. Then the claim follows from the lower semi-continuity of the original time separation functions: by Remark 3.3.5 and since we assumed $x^1 \in [x]$, $y^2 \in [y]$, we can further assume $x^1_n \to x^1$ and $y^2_n \to y^2$ (otherwise $[x], [y] \in A$ and we find subsequences in the other space). Let

$$\delta := \frac{1}{2} \min\{\tau_1(x^1, f_1(a_1)), \tau_2(f_2(a_m), y^2), \varepsilon\} > 0. \tag{3.3.3}$$

Then since τ_1 and τ_2 are lower semi-continuous there exists $n_0 \in \mathbb{N}$ such that $\tau_1(x^1_n, f_1(a_1)) > \tau_1(x^1, f_1(a_1)) - \delta$ and $\tau_2(f_2(a_m), y^2_n) > \tau_2(f_2(a_m), y^2) - \delta$ for all $n \geq n_0$. Then

$$\begin{aligned}
\widetilde{\tau}([x_n], [y_n]) &\geq \tau_1(x^1_n, f_1(a_1)) + \tau_2(f_2(a_1), f_2(a_2)) + \ldots + \tau_2(f_2(a_m), y^2_n) \\
&> \tau_1(x^1, f_1(a_1)) + \tau_2(f_2(a_1), f_2(a_2)) + \ldots + \tau_2(f_2(a_m), y^2) - 2\delta \\
&> \widetilde{\tau}([x], [y]) - \varepsilon - 2\delta > \widetilde{\tau}([x], [y]) - 3\varepsilon > 0
\end{aligned}$$

and we are done. □

Finally, we note that if we assume global preservation of the Lorentzian structure, the quotient time separation has the following form, which is familiar from the metric case.

Proposition 3.3.7. (Short form of quotient time separation) *Let $X := X_1 \sqcup_A X_2$ be the Lorentzian amalgamation of two Lorentzian pre-length spaces X_1 and X_2 and assume that f_1 and f_2 are τ-preserving and causality preserving. Then the quotient time separation has the following form:*

$$\begin{aligned}
&\widetilde{\tau}([x], [y]) \\
&= \begin{cases} \tau_i(x^i, y^i) & x^i, y^i \in X_i, \\ \sup\limits_{[a] \in J_X([x],[y]) \cap A} \{\tau_i(x^i, f_i(a)) + \tau_j(f_j(a), y^j)\} & x^i \in X_i, y^j \in X_j, i, j \in \{1, 2\}, i \neq j. \end{cases}
\end{aligned} \tag{3.3.4}$$

Proof. Let, say, $x^1 \in [x]$, $y^1 \in [y]$ and let $(x_1, y_1, \ldots, x_n, y_n)$ be an n-chain from $[x]$ to $[y]$. By Remark 3.1.3 we can assume this chain to have nontrivial relations everywhere and $x_1 = x^1$ as well as $y_n = y^1$. Intuitively, the chain starts out at x^1 in X_1, moves to A_1, jumps around between A_1 and A_2 and ends up at y^1 in X_1 again. Since f_1 and f_2 are τ-preserving and causality preserving, we can replace all distances from points in A_2 with corresponding equal distances in A_1. But then by several applications of the reverse triangle inequality for τ_1 we can replace all these distances in A_1 by a single distance to obtain a longer chain. Finally, we can omit the detour through A_1 altogether by replacing it with the direct distance from x^1 to y^1 in X_1.
The second case is symmetric, so suppose $x^1 \in [x]$ and $y^2 \in [y]$. If we have an n-chain $(x_1, y_1, \ldots, x_n, y_n)$ with nontrivial relations, then as above the chain starts out at x^1 in X_1, moves to A_1, jumps around between A_1 and A_2, but this time ends up at y^2 in X_2. Since f_1 and f_2 are τ and causality preserving we can replace, say, all distances in A_1 with equal distances in A_2 and then apply the reverse triangle inequality multiple times to end up with a 2-chain with at least the same length as in the claim. More precisely, we estimate the length of the n-chain

$$\begin{aligned}
&\sum_{i=1}^{n} \tau(x_i, y_i) = \\
&\quad \tau_1(x^1, f_1(a_1)) + \tau_2(f_2(a_1), f_2(a_2)) + \ldots + \tau_1(f_1(a_{n-2})), f_1(a_{n-1}) \\
&\qquad + \tau_2(f_2(a_{n-1}), y^2) = \\
&\quad \tau_1(x^1, f_1(a_1)) + \tau_2(f_2(a_1), f_2(a_2)) + \ldots + \tau_2(f_2(a_{n-2}), f_2(a_{n-1})) \\
&\qquad + \tau_2(f_2(a_{n-1}), y^2) \leq \\
&\quad \tau_1(x^1, f_1(a_1)) + \tau_2(f_2(a_1), y^2).
\end{aligned}$$

□

The following lemma is helpful for the gluing theorem. It says that certain diamonds in the glued space are built from corresponding diamonds in the original spaces.

Lemma 3.3.8. (Amalgamated diamonds) *Let X be the amalgamation of two Lorentzian pre-length spaces X_1 and X_2. Assume that $f_1 : A \to A_1$ and $f_2 : A \to A_2$ are $\leq$-preserving. Then we have the following decomposition of causal diamonds in X:*

(i) If $[x], [y] \in A$, then $J_X([x], [y]) = \pi(J_1(x^1, y^1) \sqcup J_2(x^2, y^2))$.
(ii) Let $i \in \{1, 2\}$. If $x^i, y^i \in X_i \backslash A$ are such that $J_i(x^i, y^i) \cap A_i = \emptyset$, then $J_X([x], [y]) = \pi(J_i(x^i, y^i))$.

If f_1 and f_2 are $\ll$-preserving, we get the same statement for timelike diamonds.

Proof. One inclusion always holds by Corollary 3.1.7: if $q^i \in J_i(x^i, y^i)$, i.e., $x^i \leq_1 q^i \leq_i y^i$, then $[x] \mathrel{\widetilde{\lesssim}} [q] \mathrel{\widetilde{\lesssim}} [y]$, hence $[q] \in J_X([x], [y])$. We show the other inclusion separately.

(i) Let $[q] \in J_X([x],[y])$ and assume without loss of generality $q^1 \in [q]$. Then $[x] \mathrel{\tilde{\leq}} [q]$, i.e., there exists a chain such that $x^1 \leq_1 a_1^1 \sim a_1^2 \leq_2 a_2^2 \sim a_2^1 \leq_1 a_3^1 \sim \ldots \sim a_n^1 \leq_1 q^1$. Since f_1 and f_2 are $\leq$-preserving, we can transfer each relation in X_2 to a corresponding relation in X_1 and ultimately obtain $x^1 \leq_1 q^1$. Analogously, we obtain $q^1 \leq_1 y^1$ and the claim follows.

(ii) Let $[q] \in J_X([x],[y])$. Note that $[x] = \{x^i\}$ and $[y] = \{y^i\}$. Since $J_i(x^i, y^i) \cap A_i = \emptyset$ and f_1 and f_2 are $\leq$-preserving, any chain starting at x^i and ending at y^i must necessarily only contain trivial relations and so has to stay inside X_i. Thus, $[q] = \{q^i\}$, $x^i \leq_i q^i$ and $q^i \leq_i y^i$.

□

4. Gluing preparations

In this section, we revisit the Lorentzian version of Alexandrov's lemma and show a gluing lemma for timelike triangles, which is essential for the proof of the gluing theorem. At first, we recall some terminology from [2], which is of great use to us. For a more detailed analysis of the contents therein, see [14].

Our approach for proving the gluing theorem in the Lorentzian setting is in spirit very close to the metric version, which introduces a so-called gluing lemma for triangles, cf. [6, Lemma II.4.10], whose proof relies on Alexandrov's lemma. While [2, Lemma 2.4] certainly is a very powerful formulation of Alexandrov's Lemma valid in any semi-Riemannian manifold, this is not ideal for our situation since it relies too much on the differential structure present in the model spaces.

4.1. *Signed distance and other techniques*

In this first subsection, we introduce some very general concepts from [2] and collect some useful facts concerning angles.

Definition 4.1.1. *(Signed length and signed distance)* Let (M, g) be a spacetime and let $p \in M$. For $v \in T_pM$, we denote the "norm" of v by $|v| := \sqrt{|g_p(v,v)|}$. We then define the signed length of v as $|v|_\pm := \mathrm{sgn}(v)\sqrt{|g_p(v,v)|} = \mathrm{sgn}(v)|v|$, where

$$\mathrm{sgn}(v) := \begin{cases} 1 & g_p(v,v) \geq 0, \\ -1 & g_p(v,v) < 0. \end{cases} \tag{4.1.1}$$

If $p, q \in M$ are contained in a normal[2] neighbourhood U, then there exists a unique geodesic $\gamma_{pq} : [0,1] \to M$ connecting p and q contained in U. We define the signed distance of p and q as $|pq|_\pm := |\gamma'_{pq}(0)|_\pm$. Note that if $p \leq q$ and M is strongly causal, then

$$\tau(p,q) = -|pq|_\pm = \sqrt{-g_p(\gamma'_{pq}(0), \gamma'_{pq}(0))}. \tag{4.1.2}$$

[2] We follow the notation of [2] where a normal neighbourhood is a (diffeomorphic) exponential image of an open set of the tangent space. More commonly, such neighbourhoods may also be called geodesically convex neighbourhoods or convex normal neighbourhoods.

Definition 4.1.2. *(Hyperbolic and nonnormalized angle)* Let M be a spacetime and let $p \in M$. Let $q, r \in I^{\pm}(p)$ be such that $\gamma'_{pq}(0) =: v$ and $\gamma'_{pr}(0) =: w$ exist. The hyperbolic angle between q and r at p is defined as

$$\measuredangle_p(q, r) := \operatorname{arcosh}\left(\frac{|g_p(v, w)|}{|v||w|}\right). \tag{4.1.3}$$

Let now q, r be any points in a normal neighbourhood of p. The nonnormalized angle between q and r at p is defined as $\angle qpr := g_p(v, w)$.

These two notions of angles are closely related (if the points are timelike related). Keeping the above terminology, one immediately sees that

$$\measuredangle_p(q, r) = \operatorname{arcosh}\left(\frac{|\angle pqr|}{|v||w|}\right). \tag{4.1.4}$$

Note that the nonnormalized angle is much more general in the sense that the points need not be timelike related at all. However, if both angles exist, the nonnormalized angle in some way better captures what "type" of angle we are dealing with, i.e., if the tangent vectors lie in the same timecone or not (recall that for timelike vectors v and w, $g(v, w) < 0$ if both are future or past directed and $g(v, w) > 0$ if they have different time orientation). Observe that the nonnormalized angle is not scale invariant.

Remark 4.1.3. (Implicit inequalities on angles) With this in mind, we want to touch on an implication of inequalities of angles. Keeping the above notation, if say $\measuredangle_p(q, r) \leq \measuredangle_{p'}(q', r')$, $|v| = |v'|$, $|w| = |w'|$ and both tangent vectors point in the same direction, then $\angle qpr \geq \angle q'p'r'$ (and vice versa). If the vectors lie in different timecones, then the inequality of the nonnormalized angles reverses. Note that $g_p(v, w)$ and $g_p(v', w')$ must have the same sign in order to infer some inequality.

One big advantage of signed distance and the nonnormalized angle is the very powerful hinge lemma, of which we will make extensive use. For a proof of the following statement see [2, Lemma 2.2].

Lemma 4.1.4. (Hinge lemma) *Let* $(|pq|_\pm, |qr|_\pm, |pr|_\pm) \in \mathbb{R}^3 \setminus \{0\}$ *be a triple realizable as the sidelengths of a triangle in* M_K. *If we vary the length of the third side* $|pr|_\pm$ *while keeping* $|pq|_\pm$ *and* $|qr|_\pm$ *fixed, then:*

(i) The nonnormalized angle $\angle pqr$ *is a decreasing function of* $|pr|_\pm$.
(ii) The nonnormalized angles $\angle qpr$ *and* $\angle qrp$ *are increasing functions of* $|pr|_\pm$.

□

The following lemma will be essential for one case in the gluing lemma, cf. [14, Lemma 5.1.1].

Lemma 4.1.5. (Bounding function via derivative) *For* $k \in \mathbb{R}$, *let* $f : [0, L] \to \mathbb{R}$ *be a smooth function such that* $f'' + kf \leq 0$, $f(0) = 0$ *and* $f(L) = 0$. *If* $k > 0$, *assume additionally that* $L < \frac{\pi}{\sqrt{k}}$. *Then* $f(t) \geq 0$ *for all* $t \in [0, L]$. □

We note that, on the one hand, by replacing f with $-f$, we can make a similar statement with reversed inequalities. That is, if $f'' + kf \geq 0$, then $f \leq 0$. On the other hand, Lemma 4.1.5 holds as well if we only assume $f(0) \geq 0$ and $f(L) \geq 0$. Indeed, suppose indirectly that this is not the case, then there exists $t_0 \in (0, L)$ such that $f(t_0) < 0$. By the mean value theorem, there must exist $t_1 \in (0, t_0)$, $t_2 \in (t_0, L)$ such that $f(t_1) = f(t_2) = 0$. Then simply apply Lemma 4.1.5 to $[t_1, t_2] \subseteq [0, L]$ to obtain the desired contradiction.

Remark 4.1.6. (Dealing with hyperbolic angles) Before proceeding to the following results, we want to mention some useful properties of the hyperbolic angle. Contrary to the nonnormalized angle, the hyperbolic angle is independent of the length of its sides (when considering it as a hinge). Suppose in M_k (or in fact, any two-dimensional spacetime) we have three geodesics emanating from a point p that all go into the same time direction, say the future. Then we can definitively speak of a "middle segment", say $[p, y]$ lies between[3] $[p, x]$ and $[p, z]$. Then the triangle equality for hyperbolic angles holds:

$$\measuredangle_p(x, y) + \measuredangle_p(y, z) = \measuredangle_p(x, z). \tag{4.1.5}$$

This is immediate when viewing the hyperbolic angle as the area under a hyperbolic segment.

Remark 4.1.7. (Hinge behaviour) Furthermore, we want to highlight a fact valid in any spacetime: consider a timelike triangle $\Delta(q, r, p)$ with $a = \tau(q, p)$, $b = \tau(r, p)$, $c = \tau(q, r)$ and $\omega = \measuredangle_p(q, r)$. Moving the point q further into the past along the geodesic extending $[q, p]$ causes the distance from r to q to increase as well. More precisely, if $q' \ll q$ is such that $[q, p] \subseteq [q', p]$, then $c' := \tau(q', r) \geq \tau(q, r) = c$. This is a simple consequence of the reverse triangle inequality. We want to reformulate this as follows, so that it can be applied similarly to the hinge lemma: let α and β be the (past-oriented) geodesics extending the segments $[q, p]$ and $[r, p]$, respectively. Assume $\alpha(1) = q$. Consider the hinge (α, β) and denote the included hyperbolic angle by ω. Then $\tau(\alpha(t), r)$ is an increasing function[4] in t for all $t \geq 1$. Intuitively, one should think of the behaviour illustrated in Fig. 4: increasing the longest side in a timelike triangle while keeping one of the short sides fixed (and hence also the included hyperbolic angle) causes the other short side to increase as well.

4.2. A visual Lorentzian version of Alexandrov's lemma

As in the metric case, the gluing lemma heavily relies on Alexandrov's lemma. [2] gives an exceptionally general version of Alexandrov's lemma valid in semi-Riemannian comparison theory. The following lemma may in some sense be

[3] The segment $[p, y]$ is said to lie between $[p, x]$ and $[p, z]$ if the following holds: consider the three tangent vectors corresponding to the three segments. These have the same time orientation. Extend them to rays and denote these by R_x, R_y and R_z. Then for all $v \in R_x$, $w \in R_z$ the connecting segment $[v, w]$ intersects the ray R_y.

[4] Note that this of course only makes sense if the geodesic α can actually be extended. Since we are anyways only interested in applying this in model spaces, this is not problematic.

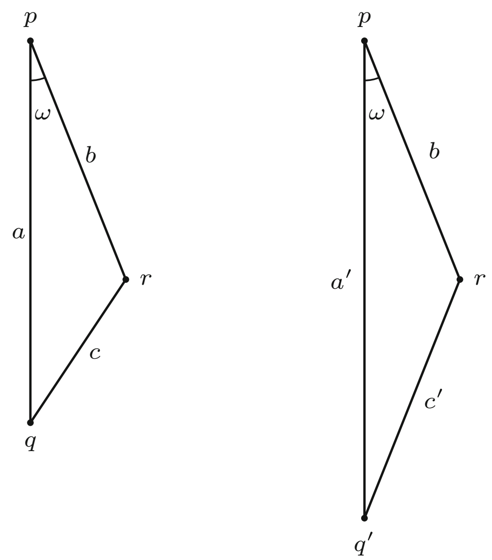

Fig. 4. Increasing a causes c to increase as well

regarded as a special case of [2, Lemma 2.4] but we still feel justified to give a full statement here. On the one hand, our approach is much more visual and hence more in the spirit of the original (Euclidean) version of the lemma. On the other hand, we have to avoid assumptions on the behaviour of nonnormalized angles in the comparison situation since the triangles are originally from Lorentzian pre-length spaces, where currently the concept of (hyperbolic) angles is not fully developed.

Lemma 4.2.1. (Alexandrov's lemma, Lorentzian version) *Let $\Delta(x, p, y)$ and $\Delta(p, y, z)$ be two triangles in M_K arranged in a way such that x and z lie on opposite sides of the geodesic line extending $[p, y]$. Suppose $p, y \in I(x, z)$ and $\tau(x, y) + \tau(y, z) < \tau(x, p) + \tau(p, z)$. Let $\Delta(x', y', z')$ be a (timelike) triangle in M_K such that $\tau(x', y') = \tau(x, y), \tau(y', z') = \tau(y, z)$ and $\tau(x', z') = \tau(x, p) + \tau(p, z)$. In particular, we have to assume that $\Delta(x, p, y)$ and $\Delta(p, y, z)$ are small enough such that the size bounds for $\Delta(x', y', z')$ are satisfied as well. Then $\angle xyz \leq \angle x'y'z'$. Moreover, if $[x, z] \cap [p, y] = \emptyset$ then $\angle pzy \geq \angle p'z'y', \angle pxy \geq \angle p'x'y'$ and $|py|_\pm \leq |p'y'|_\pm$. And if $[x, z] \cap [p, y] \neq \emptyset$ then $\angle pzy \leq \angle p'z'y', \angle pxy \leq \angle p'x'y'$ and $|py|_\pm \geq |p'y'|_\pm$ with equalities everywhere if and only if one of the inequalities is an equality, which happens if and only if $[x, z] \cap [p, y] = \{p\}$.*

Proof. We only show the case $[x, z] \cap [p, y] = \emptyset$, which is also the relevant one for the gluing lemma. The other case is done analogously (see last paragraph). To begin with, we want to explain the assumption

$$\tau(x, y) + \tau(y, z) < \tau(x, p) + \tau(p, z). \tag{4.2.1}$$

On the one hand, this gives the reverse triangle inequality in the "straightened" big triangle, guaranteeing its (nondegenerate) existence. On the other hand, this ensures that we can actually apply Alexandrov's lemma to the gluing lemma below. The condition $[x, z] \cap [p, y] = \emptyset$ serves as an analogue to having an angle greater than π at p in the metric case. So from a Euclidean point of view, the quadrilateral is

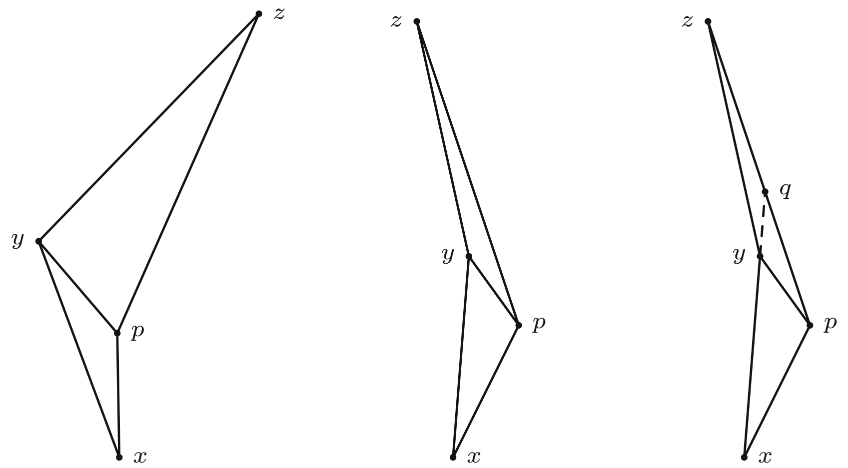

Fig. 5. The left and the middle configuration depict the two possible cases of a concave quadrilateral. The figure on the right illustrates how to rule out the middle case

concave. But if you take a look at Fig. 5, there are two possibilities for this quadrilateral to turn out concave. To put it another way: we know that the quadrilateral is concave, but since we cannot explicitly connect this nonempty intersection with a (Euclidean) large angle at p, it could happen that this large angle appears at y. Suppose we are in this case, then extend the segment $[x, y]$ until it intersects $[p, z]$, say in a point q. Then with the reverse triangle inequality and the fact that $y \in [x, q]$ and $q \in [p, z]$, we compute

$$\begin{aligned}\tau(x, y) + \tau(y, z) &\geq \tau(x, y) + \tau(y, q) + \tau(q, z) = \tau(x, q) + \tau(q, z)\\ &\geq \tau(x, p) + \tau(p, q) + \tau(q, z) = \tau(x, p) + \tau(p, z),\end{aligned}$$

a contradiction to (4.2.1).

Turning now to the actual proof, note that by the reverse triangle inequality we have $\tau(x', z') = \tau(x, p) + \tau(p, z) \leq \tau(x, z)$. For the angle at y, consider the triangles $\Delta(x, y, z)$ and $\Delta(x', y', z')$. They have two sides of equal length, and since $\tau(x', z') \leq \tau(x, z)$, we have $|xz|_\pm \leq |x'z'|_\pm$. Thus, $\angle xyz \geq \angle x'y'z'$ by the hinge lemma.

For the remaining estimates, we follow the same visual approach as the original version of Alexandrov's lemma. Let $\tilde{x}$ be the unique point with $\tau(\tilde{x}, p) = \tau(x, p)$ such that $\tilde{x}$ lies on the extension of the timelike geodesic segment $[p, z]$, see Fig. 6. Since $[x, z] \cap [p, y] = \emptyset$, the segment $[\tilde{x}, p]$ lies between the segments $[x, p]$ and $[p, y]$ (in the sense of Remark 4.1.6). This already implies $\angle xpy \leq \angle \tilde{x}py$. To see this, we observe that $\angle xpy \leq \angle \tilde{x}py$ is equivalent to

$$\langle \gamma'_{p\tilde{x}}(0) - \gamma'_{px}(0), \gamma'_{py}(0)\rangle \geq 0. \tag{4.2.2}$$

Clearly, $\gamma'_{p\tilde{x}}(0) - \gamma'_{px}(0)$ is a spacelike vector as the difference of two past directed timelike vectors of the same length. Recall that the scalar product with a spacelike vector v is non-negative if and only if the other vector lies in the same half-space as v generated by the normal space of v. For better visualization, we apply a Lorentz transformation to view $\gamma'_{p\tilde{x}}(0) - \gamma'_{px}(0)$ as a horizontal vector and p as the origin,

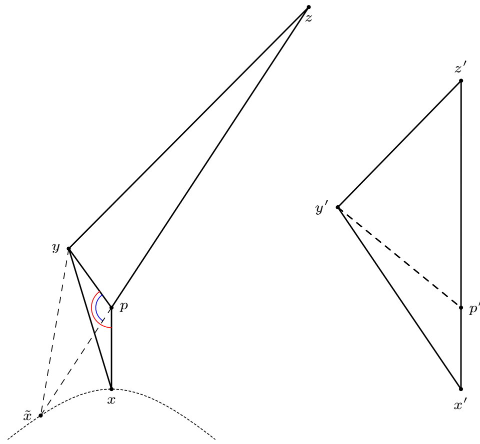

Fig. 6. A visual approach in the spirit of the original version of Alexandrov's lemma

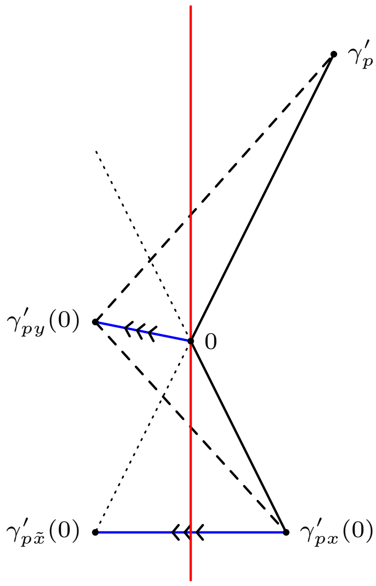

Fig. 7. The two vectors in the scalar product lie in the same half-space

see Fig. 7. Then the normal space of $\gamma'_{p\tilde{x}}(0) - \gamma'_{px}(0)$ is a vertical line through 0. From a Euclidean point of view, $\gamma'_{p\tilde{x}}(0) - \gamma'_{px}(0)$ is at an angle of ninety degrees and symmetric with respect to its normal space. Since $p \in [\tilde{x}, z]$, it follows that $\gamma'_{px}(0)$ and $\gamma'_{p\tilde{x}}(0)$ differ by a Euclidean reflection with respect to the normal space. Together, this implies that $\gamma'_{pz}(0)$ and $\gamma'_{px}(0)$ lie in the same half-space.
Now keep in mind that these directional vectors originally come from the sides of two triangles which together yield a concave situation without self-intersection. In this way, we not only obtain that $\gamma'_{py}(0)$ does not lie between $\gamma'_{px}(0)$ and $\gamma'_{pz}(0)$, but also that $\gamma'_{py}(0)$ has to lie between the mirrored versions of these vectors, that is between $\gamma'_{p\tilde{x}}(0)$ and $-\gamma'_{px}(0)$. Thus, it is clear that both entries in (4.2.2) are in the same half-space with respect to the normal space of the spacelike entry, and so the inequality is true.
Now consider the triangles $\Delta(x, p, y)$ and $\Delta(\tilde{x}, p, y)$. By construction we have $|px|_\pm = |p\tilde{x}|_\pm$, so they have two sides of equal length since they have the segment $[p, y]$ in common. Moreover, since $\measuredangle \tilde{x}py \geq \measuredangle xpy$, the hinge lemma

gives $|\tilde{x}y|_\pm \leq |xy|_\pm$. We continue with the triangles $\Delta(\tilde{x}, y, z)$ and $\Delta(x', y', z')$. Again, we have two pairs of equal length since $|yz|_\pm = |y'z'|_\pm$ and $|x'z'|_\pm = |xp|_\pm + |pz|_\pm = |\tilde{x}p|_\pm + |pz|_\pm = |\tilde{x}z|_\pm$, where the last equality holds since p lies on the segment $[\tilde{x}, z]$. For the third side we know $|\tilde{x}y|_\pm \leq |xy|_\pm = |x'y'|_\pm$ by the above considerations. Thus, the hinge lemma implies $\angle\tilde{x}zy \geq \angle x'z'y'$. And since p and p' correspond to each other on the segments $[\tilde{x}, z]$ and $[x', z']$, respectively, we also have $\angle pzy \geq \angle p'z'y'$.

Analogously, we find a point $\tilde{z}$ that lies on the extension of $[x, p]$ at distance $\tau(p, z)$ from p. A similar argument then implies $\angle pxy \geq \angle y'x'z'$. Finally, the triangles $\Delta(x, p, y)$ and $\Delta(x', p', y')$ have two sides with equal length and we know $\angle pxy \geq \angle p'x'y'$. Consequently, we obtain $|py|_\pm \leq |p'y'|_\pm$ by the hinge lemma.

At last, note that if $[x, z] \cap [p, y] = \{p\}$, then $[x, z]$ in fact is composed of the two sides $[x, p]$ and $[p, z]$ in the triangles and hence $\Delta(x, y, z)$ is isometric to $\Delta(x', y', z')$, so equality in all inequalities in the statement follows immediately. Also if $[x, z] \cap [p, y] = \{q\} \neq \{p\}$, then the procedure from above causes $[x, p]$ to be between $[p, y]$ and $[p, \tilde{x}]$. Then the same calculation yields $\angle\tilde{x}py \leq \angle xpy$ and consequently all the following inequalities are reversed. Any of the inequalities being an equality forces the others to be equalities as well, which then implies $q = p$. $\square$

In the statement above we assumed the configuration of triangles to be such that the subdivision happens along the longest side of $\Delta(x', y', z')$. For a proper application we need to guarantee the desired behaviour in the other cases as well. However, this can be done analogously. We give a rough sketch of the proof.

Lemma 4.2.2. (Alexandrov's Lemma, other constellations) *Let $\Delta(x, p, z)$ and $\Delta(p, y, z)$ be two triangles in M_K such that $[x, p]$, $[x, z]$, $[p, y]$ as well as $[y, z]$ are timelike and the triangles are arranged on opposite sides of the geodesic line extending the segment $[p, z]$, see Fig. 8. Let $\Delta(x', y', z')$ be a timelike triangle such that $\tau(x', y') = \tau(x, p) + \tau(p, y)$, $\tau(y', z') = \tau(y, z)$ and $\tau(x', z') = \tau(x, z)$. Then $\angle pzy \geq \angle p'z'y'$. If $[x, y] \cap [p, z] = \emptyset$, then $\angle pxz \geq \angle p'x'z'$, $\angle pxy \geq \angle p'x'y'$ and $|py|_\pm \leq |p'y'|_\pm$.*

Proof. Note that since $p \ll y \ll z$, also $[p, z]$ is timelike and so we are in fact dealing with two timelike triangles. We have $\tau(x', y') = \tau(x, p) + \tau(p, y) \leq \tau(x, y)$, i.e., $|x'y'|_\pm \geq |xy|_\pm$, hence $\angle pzy \geq \angle p'z'y'$ by the hinge lemma. As for the other inequalities, extend the segment $[p, x]$ to obtain a point $\tilde{y}$ on this extension such that $\tau(p, y) = \tau(p, \tilde{y})$, cf. Fig. 6. Then as in Fig. 7, we obtain $\angle ypz \leq \angle\tilde{y}pz$. Considering the triangles $\Delta(p, \tilde{y}, z)$ and $\Delta(p, y, z)$, we infer $|yz|_\pm \geq |\tilde{y}z|_\pm$ from the hinge lemma. Then consider $\Delta(x, \tilde{y}, z)$ and $\Delta(x', y', z')$ and obtain $\angle pxz \geq \angle p'x'z'$. We can argue similarly to obtain $\angle pyz \geq \angle p'y'z'$. Finally, $|pz|_\pm \leq |p'z'|_\pm$ again by the hinge lemma. $\square$

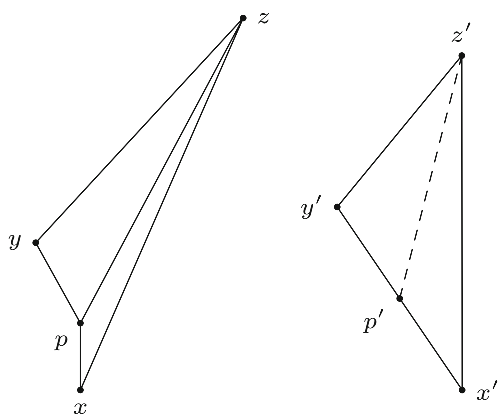

Fig. 8. The other configuration of Alexandrov's Lemma

4.3. The gluing lemma

We now formulate and prove the gluing lemma for timelike triangles. This really is the main tool in the proof of the gluing theorem, both in the metric case and in the Lorentzian setting. The proof is rather long and quite technical.

Lemma 4.3.1. (Gluing lemma for timelike triangles, case I) *Let X be a Lorentzian pre-length space and let $U \subseteq X$ be a subset that satisfies (i) and (ii) in the definition for a comparison neighbourhood in X, cf. Definition 2.1.3. That is, $\tau|_{U\times U}$ is finite and continuous and for all $x, y \in U$ with $x \ll y$ there is a causal curve contained in U with length $\tau(x, y)$. Let $K \in \mathbb{R}$ and let $T_3 := \Delta(x, y, z)$ be a timelike triangle in U satisfying size bounds for M_K. Let $p \in [x, z]$ such that $p \ll y$ (or $y \ll p$). In other words, $T_1 := \Delta(x, p, y)$ and $T_2 := \Delta(p, y, z)$ are again timelike triangles (if $y \ll p$ then the order of the points changes), see Fig. 9. Let $\overline{T_1} := \Delta(\bar{x}, \bar{p}, \bar{y})$ and $\overline{T_2} := \Delta(\bar{p}, \bar{y}, \bar{z})$ be comparison triangles for T_1 and T_2 in M_K, respectively. Suppose T_1 and T_2 satisfy timelike curvature bounds from above for K, i.e., for all $a, b \in T_i$ and corresponding comparison points $\bar{a}, \bar{b} \in \overline{T_i}$, $i = 1, 2$ we have $\tau(a, b) \geq \bar{\tau}(\bar{a}, \bar{b})$. Then T_3 satisfies the same timelike curvature bound from above.*

Proof. Realize the comparison triangles for T_1 and T_2 in such a way that they share the timelike geodesic segment $[\bar{p}, \bar{y}]$ and such that $\bar{x}$ and $\bar{z}$ lie on opposite sides of this segment (as in Alexandrov's lemma). Note that because of the size bounds, either $[\bar{p}, \bar{y}]$ and $[\bar{x}, \bar{z}]$ intersect in a single point or they do not intersect at all. If $[\bar{p}, \bar{y}] \cap [\bar{x}, \bar{z}] = \{p\}$, then as in Lemma 4.2.1, $\overline{T_1}$ and $\overline{T_2}$ together already form a comparison triangle $\overline{T_3}$ for T_3 and we are done. The cases of comparing points which are not immediate from this assumption will be covered later in greater generality. Conversely, if $[\bar{p}, \bar{y}] \cap [\bar{x}, \bar{z}] = \{q\} \neq \{p\}$, then $\bar{\tau}(\bar{x}, \bar{z}) = \bar{\tau}(\bar{x}, \bar{q}) + \bar{\tau}(\bar{q}, \bar{z})$ and, since $\bar{q} \in [\bar{p}, \bar{y}]$, $\bar{q}$ is a comparison point for some $q \in [p, y]$. Moreover, x and q are on the sides of T_1 and q and z are on the sides of T_2. Since T_1 and T_2 satisfy

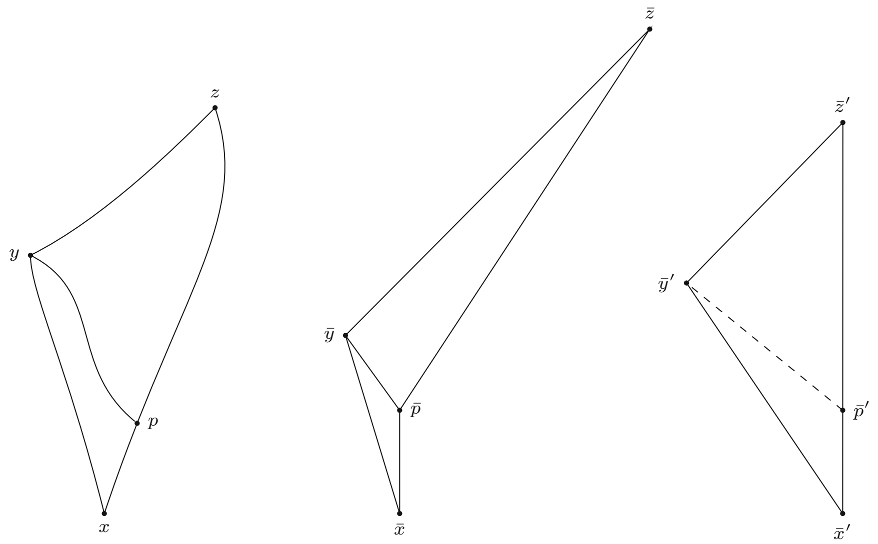

Fig. 9. A timelike triangle in X subdivided into two timelike triangles, the comparison triangles for the smaller triangles and the comparison triangle for the big triangle

timelike curvature bounds, we compute[5]

$$\bar{\tau}(\bar{x}, \bar{z}) = \bar{\tau}(\bar{x}, \bar{q}) + \bar{\tau}(\bar{q}, \bar{z}) \leq \tau(x, q) + \tau(q, z) \leq \tau(x, z). \tag{4.3.1}$$

But this is in contradiction to

$$\tau(x, z) = \tau(x, p) + \tau(p, z) = \bar{\tau}(\bar{x}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}) < \bar{\tau}(\bar{x}, \bar{z}), \tag{4.3.2}$$

so such an intersection cannot occur with an upper curvature bound. Thus, $[\bar{p}, \bar{y}] \cap [\bar{x}, \bar{z}] = \emptyset$ is the only interesting case we have to consider. Moreover, as $\tau(x, p) + \tau(p, z) = \tau(x, z) \geq \tau(x, y) + \tau(y, z)$ by the reverse triangle inequality, we can realize the situation in M_K as in Fig. 9.

What follows now are several applications of Alexandrov's Lemma 4.2.1. First, we "bend" $\overline{T_1}$ and $\overline{T_2}$ in such a way that they form a comparison triangle for T_3, cf. Fig. 6. More precisely, let $\overline{T_3} := \Delta(\bar{x}', \bar{y}', \bar{z}')$, where $\bar{\tau}(\bar{x}', \bar{y}') = \bar{\tau}(\bar{x}, \bar{y}) = \tau(x, y)$, $\bar{\tau}(\bar{y}', \bar{z}') = \bar{\tau}(\bar{y}, \bar{z}) = \tau(y, z)$ and $\bar{\tau}(\bar{x}', \bar{z}') = \bar{\tau}(\bar{x}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}) = \tau(x, p) + \tau(p, z) = \tau(x, z)$. When talking about the comparison triangles for T_1 and T_2 simultaneously, it will be convenient to denote their union, which is a quadrilateral, by $\overline{T_1} \cup \overline{T_2}$. That is, by $\bar{a} \in \overline{T_1} \cup \overline{T_2}$ we mean a point which belongs either to $\overline{T_1}$ or $\overline{T_2}$ (or both if $\bar{a} \in [\bar{p}, \bar{y}]$). We distinguish several cases, depending on which sides the points lie on. Note that for any point on $\overline{T_3}$, we can find a "comparison point" in either $\overline{T_1}$ or $\overline{T_2}$, i.e., a point on the corresponding side of equal time separation to the endpoints (the common edge $[\bar{p}, \bar{y}]$ of $\overline{T_1}$ and $\overline{T_2}$ is the only one not (isometrically) transferred to $\overline{T_3}$). We choose two points $\bar{a}', \bar{b}' \in \overline{T_3}$ and check all possible configurations. The general idea is to at first show that time separation in $\overline{T_3}$ is even smaller than in $\overline{T_1} \cup \overline{T_2}$, i.e., $\bar{\tau}(\bar{a}, \bar{b}) \geq \bar{\tau}(\bar{a}', \bar{b}')$, and then relating

[5] It may a priori not be clear that $x \leq q \leq z$ so one can apply the reverse triangle inequality, but this can be seen as follows: we have $\bar{q} \in [\bar{x}, \bar{z}]$, hence $\bar{x} \ll \bar{q} \ll \bar{z}$ and so $0 < \bar{\tau}(\bar{x}, \bar{q}) \leq \tau(x, q)$ by the curvature condition and thus $x \ll q$. We get a similar estimate for q and z.

this inequality to time separations in X, i.e., $\tau(a,b) \geq \bar{\tau}(\bar{a},\bar{b})$, thus ensuring the desired curvature bound. We may assume without loss of generality that $\bar{a}' \ll \bar{b}'$ always holds, since otherwise there is nothing to show. Since all these (sub-)case distinctions may become a bit confusing, we try to give a "to-do list" summarizing everything. Note that some descriptions might only make sense when reading the proof of the respective cases. The cases (B) and (C) are entirely analogous, hence we omit the descriptions in (C).

(A) $\bar{a}' \in [\bar{x}', \bar{y}']$ and $\bar{b}' \in [\bar{y}', \bar{z}']$ (both points on short sides): this case is easy and also the only one where we do not need any subcases.
(B) $\bar{a}' \in [\bar{x}', \bar{y}']$ and $\bar{b}' \in [\bar{x}', \bar{z}']$ (one point on short and long side each): here, we distinguish whether $\bar{a}'$ and $\bar{b}'$ are in the same triangle or not.
  (1) $\bar{b}' \in [\bar{x}', \bar{p}']$ (same triangle): this case is easy and very similar to (A).
  (2) $\bar{b}' \in [\bar{p}', \bar{z}']$ (different triangle): this needs yet another distinction, namely whether the connecting segment $[\bar{a},\bar{b}]$ stays inside $\overline{T_1} \cup \overline{T_2}$ or not.
    (i) $[\bar{a},\bar{b}]$ stays inside comparison situation: we construct several subtriangles and use the law of cosines.
    (ii) $[\bar{a},\bar{b}]$ leaves comparison situation: we improve the bound on $\bar{\tau}(\bar{a}',\bar{b}')$ by taking a detour through $\bar{p}$.
(C) $\bar{a}' \in [\bar{y}', \bar{z}']$ and $\bar{b}' \in [\bar{x}', \bar{z}']$ (one point on short and long side each)
  (1) $\bar{b}' \in [\bar{p}', \bar{z}']$ (same triangle)
  (2) $\bar{b}' \in [\bar{x}', \bar{p}']$ (different triangle)
    (i) $[\bar{a},\bar{b}]$ stays inside comparison situation
    (ii) $[\bar{a},\bar{b}]$ leaves comparison situation

Finally, observe that switching the labels or assuming $\bar{b}' \ll \bar{a}'$ does not change the proof at all. Hence this is a complete list that covers all possible configurations.
$\underline{\bar{a}' \in [\bar{x}', \bar{y}'] \text{ and } \bar{b}' \in [\bar{y}', \bar{z}'] \text{ (A)}}$: In this first case $\bar{a}' \ll \bar{b}'$ holds anyways. We find comparison points $\bar{a} \in [\bar{x},\bar{y}]$ and $\bar{b} \in [\bar{y},\bar{z}]$, i.e., $\bar{\tau}(\bar{x},\bar{a}) = \bar{\tau}(\bar{x}',\bar{a}')$ and $\bar{\tau}(\bar{y},\bar{b}) = \bar{\tau}(\bar{y}',\bar{b}')$. Consider the triangles $\Delta(\bar{x},\bar{y},\bar{z})$ and $\Delta(\bar{x}',\bar{y}',\bar{z}')$. In these triangles, two sidelengths are the same, namely $|\bar{x}\bar{y}|_\pm = |\bar{x}'\bar{y}'|_\pm$ and $|\bar{y}\bar{z}|_\pm = |\bar{y}'\bar{z}'|_\pm$. For the third side, we have $\bar{\tau}(\bar{x}',\bar{z}') = \bar{\tau}(\bar{x},\bar{p}) + \bar{\tau}(\bar{p},\bar{z}) \leq \bar{\tau}(\bar{x},\bar{z})$, and hence $|\bar{x}\bar{z}|_\pm \leq |\bar{x}'\bar{z}'|_\pm$. Then by the hinge lemma we infer $\measuredangle \bar{x}\bar{y}\bar{z} \geq \measuredangle \bar{x}'\bar{y}'\bar{z}'$, exactly as in Lemma 4.2.1. Now consider the "smaller" triangles $\Delta(\bar{a},\bar{y},\bar{b})$ and $\Delta(\bar{a}',\bar{y}',\bar{b}')$, i.e., instead of the sides $[\bar{x},\bar{y}]$ and $[\bar{y},\bar{z}]$ we consider the (from the point of view of $\bar{y}$) initial segments $[\bar{a},\bar{y}]$ and $[\bar{y},\bar{b}]$ (and the same in the other triangle). Again, two side lengths are pairwise equal. We want to use the hinge lemma in the other direction to obtain estimates on the third side. This is easily possible since $\measuredangle \bar{a}\bar{y}\bar{b}$ and $\measuredangle \bar{a}'\bar{y}'\bar{b}'$ are equal multiples of $\measuredangle \bar{x}\bar{y}\bar{z}$ and $\measuredangle \bar{x}'\bar{y}'\bar{z}'$, respectively. Thus, we obtain $|\bar{a}\bar{b}|_\pm \leq |\bar{a}'\bar{b}'|_\pm$ and hence $\bar{\tau}(\bar{a},\bar{b}) \geq \bar{\tau}(\bar{a}',\bar{b}')$.
We finished the case were both points lie on short sides. We are left with considering pairs of points where one is on the longest side and the other on a short side. These two possibilities clearly are analogous (since they only differ by time orientation), so we will only consider case (B) explicitly. In each of these cases, however, there are subcases which are respectively similar as well. Namely, we have to distinguish whether the point on the long side is chronologically before or after $\bar{p}'$, i.e., if it is in $[\bar{x}',\bar{p}']$ or in $[\bar{p}',\bar{z}']$. In other words, this distinction tells us if the two points

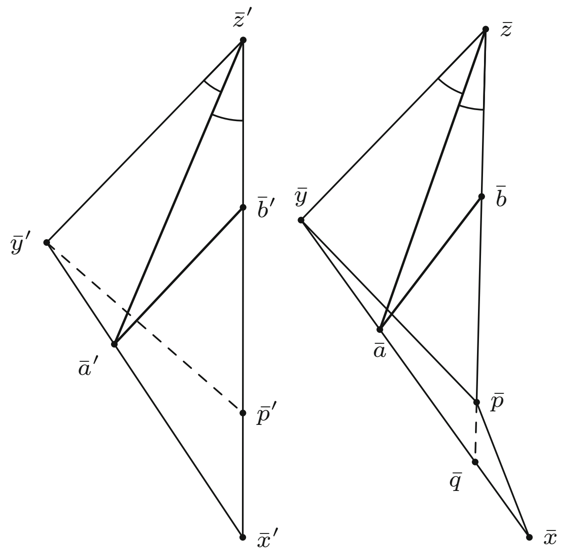

Fig. 10. This configuration requires some additional construction steps

originate from the same triangle or not. Before returning to the proof, observe that by Lemma 4.2.1 we have $\bar{\tau}(\bar{p}, \bar{y}) \geq \bar{\tau}(\bar{p}', \bar{y}')$ as well as $\angle \bar{p}\bar{x}\bar{y} \geq \angle \bar{p}'\bar{x}'\bar{y}'$ and $\angle \bar{p}\bar{z}\bar{y} \geq \angle \bar{p}'\bar{z}'\bar{y}'$.

$\underline{\bar{a}' \in [\bar{x}', \bar{y}'] \text{ and } \bar{b}' \in [\bar{x}', \bar{p}'] \subseteq [\bar{x}', \bar{z}'] \text{ (B.1)}}$: Here, both points are from $\overline{T_1}$. Consider the smaller triangles $\Delta(\bar{x}, \bar{a}, \bar{b})$ and $\Delta(\bar{x}', \bar{a}', \bar{b}')$. We have as in the case (A) above that $\angle \bar{a}\bar{x}\bar{p}$ and $\angle \bar{a}'\bar{x}'\bar{p}'$ are equal multiples of $\angle \bar{y}\bar{x}\bar{p}$ and $\angle \bar{y}'\bar{x}'\bar{p}'$, respectively. Two sides are of equal length by construction. Thus, we infer $|\bar{a}\bar{b}|_\pm \leq |\bar{a}'\bar{b}'|_\pm$ from the hinge lemma and hence $\bar{\tau}(\bar{a}, \bar{b}) \geq \bar{\tau}(\bar{a}', \bar{b}')$. Note that in this case (and also in (C.1)), since both points originate from the same triangles, we immediately obtain $\tau(a, b) \geq \bar{\tau}(\bar{a}, \bar{b})$ as well.

Looking at the list from above, we are now in the case (B.2), where we have to make yet another distinction. The extension of $[\bar{p}, \bar{z}]$ meets $[\bar{x}, \bar{y}]$ in a unique point, denote it by $\bar{q}$. We have to distinguish whether $\bar{a}$ lies chronologically before or after $\bar{q}$. This distinction in particular tells us if $[\bar{a}, \bar{b}]$ lies inside of $\overline{T_1} \cup \overline{T_2}$ or not. We first cover the case where this segment lies inside the triangles (the other case requires even more extra work).

$\underline{\bar{a}' \in [\bar{q}', \bar{y}'] \subseteq [\bar{x}', \bar{y}'] \text{ and } \bar{b}' \in [\bar{p}', \bar{z}'] \subseteq [\bar{x}', \bar{z}'] \text{ (B.2.i)}}$: We try to construct a triangle in $\overline{T_1} \cup \overline{T_2}$ that has both $\bar{a}$ and $\bar{b}$ as vertices and somehow inherits enough properties so that we can deduce the claim, see Fig. 10.

In the end, this will be the timelike triangle $\Delta(\bar{a}, \bar{b}, \bar{z})$. However, before this we need to estimate the hyperbolic angle $\measuredangle_{\bar{z}}(\bar{b}, \bar{a})$. Recall that we have $\angle \bar{y}\bar{z}\bar{p} \geq \angle \bar{y}'\bar{z}'\bar{p}'$. As the two legs have the same time orientation and the same length, we infer $\measuredangle_{\bar{z}}(\bar{y}, \bar{p}) \leq \measuredangle_{\bar{z}'}(\bar{y}', \bar{p}')$, cf. Remark 4.1.3. Now consider the timelike triangles $\Delta(\bar{a}, \bar{y}, \bar{z})$ and $\Delta(\bar{a}', \bar{y}', \bar{z}')$. They have two sides of pairwise equal length. Then $|\bar{a}\bar{z}|_\pm \leq |\bar{a}'\bar{z}'|_\pm$ easily follows from the hinge lemma. As both triangles are timelike, we can change our perspective in the sense that we go from signed lengths and the hinge lemma to (positive) time separation values and the Lorentzian law of cosines, cf. [14, Theorem 3.1.3]. We do this because the adjacent sides of $\angle \bar{a}\bar{z}\bar{y}$ and $\angle \bar{a}'\bar{z}'\bar{y}'$ do not have pairwise equal length. Recall the following consequence of the Lorentzian law of cosines: fixing the two short sides in a timelike triangle and letting the longest side vary, *any* hyperbolic angle is an increasing function in the length (understood as time separation) of the longest side. In our case, $[\bar{a}, \bar{z}]$ and $[\bar{a}', \bar{z}']$ are the longest sides and so we obtain $\measuredangle_{\bar{z}}(\bar{a}, \bar{y}) \geq \measuredangle_{\bar{z}'}(\bar{a}', \bar{y}')$. Also, since $\measuredangle_{\bar{z}}(\bar{y}, \bar{a}) + \measuredangle_{\bar{z}}(\bar{a}, \bar{p}) = \measuredangle_{\bar{z}}(\bar{y}, \bar{p}) \leq \measuredangle_{\bar{z}'}(\bar{y}', \bar{p}') = \measuredangle_{\bar{z}'}(\bar{y}', \bar{a}') + \measuredangle_{\bar{z}'}(\bar{a}', \bar{p}')$, we must have $\measuredangle_{\bar{z}}(\bar{a}, \bar{b}) = \measuredangle_{\bar{z}}(\bar{a}, \bar{p}) \leq \measuredangle_{\bar{z}'}(\bar{a}', \bar{p}') = \measuredangle_{\bar{z}'}(\bar{a}', \bar{b}')$. Finally, consider $\Delta(\bar{a}, \bar{b}, \bar{z})$ and $\Delta(\bar{a}', \bar{b}', \bar{z}')$. We know $B := \bar{\tau}(\bar{b}, \bar{z}) = \bar{\tau}(\bar{b}', \bar{z}')$, $A := \bar{\tau}(\bar{a}, \bar{z}) \geq \bar{\tau}(\bar{a}', \bar{z}') =: A'$ and $\omega := \measuredangle_{\bar{z}}(\bar{a}, \bar{b}) \leq \measuredangle_{\bar{z}'}(\bar{a}', \bar{b}') =: \omega'$ and we want to infer $C := \bar{\tau}(\bar{a}', \bar{b}') \leq \bar{\tau}(\bar{a}', \bar{b}') =: C'$. In two steps, we want to transform one triangle into the other. In $\Delta(\bar{a}', \bar{b}', \bar{z}')$, fix A' and B and decrease ω' to ω. From the law of cosines, it easily follows that C' increases as ω' decreases. In this way we obtain an intermediate triangle with sidelengths A', B and a hyperbolic angle ω opposite of a side $\tilde{C} \geq C'$. Now we want to increase A' to A while keeping B and ω fixed so that $\tilde{C}$ *changes* to C (we do not yet know whether it increases or decreases). But this is precisely the situation in Remark 4.1.7, and so we obtain $C \geq \tilde{C} \geq C'$, as desired.

So we are only left with the case (B.2.ii). In almost all the cases we covered so far, relating the distance in $\overline{T_1} \cup \overline{T_2}$ with the original distances in X is not difficult. Either the points already stem from the same small triangle in X, e.g., $a \in [x, y], b \in [x, p] \subseteq [x, z]$ in (B.1), in which case we have $\tau(a, b) \geq \bar{\tau}(\bar{a}, \bar{b})$ by assumption. Or the points lie in different triangles but the connecting geodesic is entirely contained in the union of $\overline{T_1}$ and $\overline{T_2}$. For instance in (B.2.ii), the very last case we covered, for $\bar{a}' \in [\bar{x}, \bar{q}] \subseteq [\bar{x}', \bar{y}']$ and $\bar{b}' \in [\bar{p}', \bar{z}'] \subseteq [\bar{x}', \bar{z}']$ there exists an intersection point of $[\bar{p}, \bar{y}]$ and $[\bar{a}, \bar{b}]$, call it $\bar{r}$. Then we compute

$$\bar{\tau}(\bar{a}', \bar{b}') \leq \bar{\tau}(\bar{a}, \bar{b}) = \bar{\tau}(\bar{a}, \bar{r}) + \bar{\tau}(\bar{r}, \bar{b}) \leq \tau(a, r) + \tau(r, b) \leq \tau(a, b). \tag{4.3.3}$$

Illustratively, we divided the segment $[\bar{a}, \bar{b}]$ in two parts which both lie inside a single triangle and then apply the respective curvature conditions of T_1 and T_2. The reader may already have suspected that this unfortunately is not always possible. Take $\bar{a} \in [\bar{x}, \bar{y}], \bar{b} \in [\bar{y}, \bar{z}]$ in case (A), where $\bar{a}$ and $\bar{b}$ are close to $\bar{x}$ and $\bar{z}$, respectively. Then the connecting segment leaves the triangles. While $\bar{\tau}(\bar{a}, \bar{b}) \geq \bar{\tau}(\bar{a}', \bar{b}')$ is easy to see, we have no way of relating $\bar{\tau}(\bar{a}, \bar{b})$ with $\tau(a, b)$. Moreover, in the remaining case (B.2.ii) this is by assumption always the case. We rectify this situation by improving the bound on $\bar{\tau}(\bar{a}', \bar{b}')$: in (B.2.ii) and the cases of (A) where the connecting segment leaves the triangle, we want to show

$$\bar{\tau}(\bar{a}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{b}) \geq \bar{\tau}(\bar{a}', \bar{b}'). \tag{4.3.4}$$

To make the following already very lengthy computations a bit easier and also to do both cases kind of simultaneously, we assume for now $\bar{b}' = \bar{z}'$. After showing (4.3.4) in this case, we let $\bar{b}'$ vary. By moving $\bar{b}'$ onto $[\bar{y}', \bar{z}']$ we cover the remaining case of (A) and by moving $\bar{b}'$ onto $[\bar{p}', \bar{z}']$ we cover (B.2.ii). Remember that we denote by $\bar{q}$ the point where the extension of $[\bar{p}, \bar{z}]$ meets $[\bar{x}, \bar{y}]$.

$\bar{a} \in [\bar{x}, \bar{q}] \subseteq [\bar{x}, \bar{y}]$ and $\bar{b} = \bar{z}$ (B.2.ii & A): If we can show (4.3.4), then this implies $\tau(a, z) \geq \tau(a, p) + \tau(p, z) \geq \bar{\tau}(\bar{a}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}) \geq \bar{\tau}(\bar{a}', \bar{z}')$.

In Fig. 11, this effectively means that the detour through $\bar{p}$ in $\overline{T_1} \cup \overline{T_2}$ is still larger than the direct connection in $\overline{T_3}$. We proceed in an elementary yet effective way.

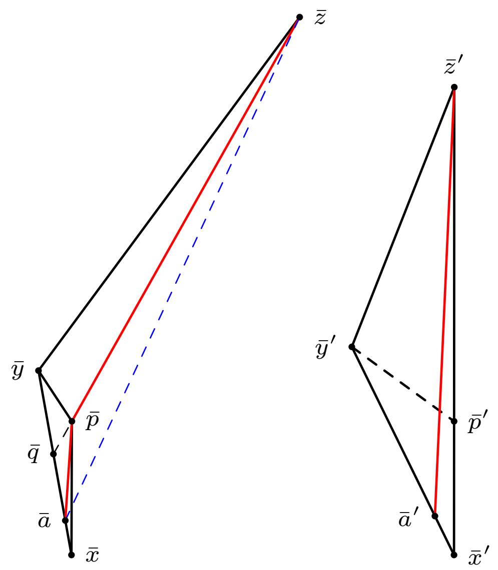

Fig. 11. Improving the bound on $\bar{\tau}(\bar{a}', \bar{z}')$

We define a function that compares these two lengths and show that its sign does not change. Unfortunately, we have to show this separately for the different cases of curvature.

Minkowski space ($K = 0$): Recall that the points $\bar{a}$ and $\bar{a}'$ can be described as $\bar{a} = \gamma_{\bar{x}\bar{y}}(t) = ty + (1-t)x$ and $\bar{a}' = \gamma_{\bar{x}'\bar{y}'}(t) = ty' + (1-t)x'$, respectively. Set $\gamma_{\bar{x}\bar{y}}(m) =: \bar{q} \in [\bar{x}, \bar{y}]$ and similarly $\gamma_{\bar{x}'\bar{y}'}(m) =: \bar{q}' \in [\bar{x}', \bar{y}']$. Then define $f : [0, m] \to \mathbb{R}$ by

$$f(t) := (\bar{\tau}(\bar{a}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}))^2 - \bar{\tau}(\bar{a}', \bar{z}')^2. \tag{4.3.5}$$

We square these values solely for ease of computation. If $t = 0$, then $\bar{a} = \bar{x}$ and $\bar{a}' = \bar{x}'$, and since by construction $\bar{\tau}(\bar{x}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}) = \bar{\tau}(\bar{x}', \bar{z}')$, we have $f(0) = 0$. If $t = m$, then $\bar{a} = \bar{q}$ and since $\bar{q}$ lies on the extension of the segment $[\bar{p}, \bar{z}]$, we have $\bar{\tau}(\bar{q}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}) = \bar{\tau}(\bar{q}, \bar{z}) \geq \bar{\tau}(\bar{q}', \bar{z}')$, hence $f(m) \geq 0$. Thus, if we can show that $f'' \leq 0$, we obtain $f \geq 0$ and so $(\bar{\tau}(\bar{a}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}))^2 \geq \bar{\tau}(\bar{a}', \bar{z}')^2$ which clearly also implies $\bar{\tau}(\bar{a}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}) \geq \bar{\tau}(\bar{a}', \bar{z}')$. Also note that since $\bar{a} \ll \bar{q} \ll \bar{p}$, all distances are timelike. For the calculation, we first observe that

$$\gamma'_{\bar{a}\bar{p}}(0) = \bar{p} - \bar{x} + t(\bar{x} - \bar{y}), \tag{4.3.6}$$

and we get a similar expression for $\gamma'_{\bar{a}'\bar{z}'}(0)$. After calculating the squares one observes that $\langle \gamma'_{\bar{p}\bar{z}}(0), \gamma'_{\bar{p}\bar{z}}(0) \rangle$ does not depend on t at all (since it does not depend on the position of $\bar{a}$), so we can ignore this value for the derivative of f. The following expressions become quite lengthy, so we abbreviate $B := \gamma'_{\bar{a}\bar{p}}(0)$, $A := \frac{d}{dt}\gamma'_{\bar{a}\bar{p}}(0) = \bar{x} - \bar{y}$, $B' := \gamma'_{\bar{a}'\bar{z}'}(0)$ and $A' := \frac{d}{dt}\gamma'_{\bar{a}'\bar{z}'}(0) = \bar{x}' - \bar{y}'$. Then f reads

$$f(t) = -\langle B, B \rangle - 2|\bar{p}\bar{z}|_{\pm}\sqrt{-\langle B, B \rangle} - \langle \gamma'_{\bar{p}\bar{z}}(0), \gamma'_{\bar{p}\bar{z}}(0) \rangle + \langle B', B' \rangle, \tag{4.3.7}$$

and

$$\frac{d}{dt} f(t) = -2\langle A, B \rangle - 2|\bar{p}\bar{z}|_{\pm} \frac{-\langle A, B \rangle}{\sqrt{-\langle B, B \rangle}} + 2\langle A', B' \rangle. \tag{4.3.8}$$

We further compute the second derivative of the middle term:

$$\frac{d}{dt}\frac{-\langle A, B\rangle}{\sqrt{-\langle B, B\rangle}} = \frac{\langle A, A\rangle\langle B, B\rangle - \langle A, B\rangle^2}{\sqrt{(-\langle B, B\rangle)^3}} \leq 0, \tag{4.3.9}$$

where the inequality follows by the reverse Cauchy Schwarz inequality for timelike vectors, cf. [19, Prop. 5.30 (i)]. For the remaining terms of $(\frac{d}{dt})^2 f(t)$, we have

$$\frac{d}{dt}2\langle A, B\rangle = 2\langle A, A\rangle = 2\langle A', A'\rangle = \frac{d}{dt}2\langle A', B'\rangle, \tag{4.3.10}$$

where the middle equality holds since by construction we have $\bar{\tau}(\bar{x}, \bar{y}) = \bar{\tau}(\bar{x}', \bar{y}')$. Thus, these two terms cancel out and the second derivative of f consists only of a nonpositive term. Keep in mind that $-2|\bar{p}\bar{z}|_{\pm}$ is positive as $|\bar{p}\bar{z}|_{\pm}$ is negative.
De Sitter space $(K = 1)$: In de Sitter space, the situation is more involved. Here, we will need Lemma 4.1.5. In this case and the following case of anti-de Sitter space, we choose to omit the bars of points since we are anyways only working in M_K. Let γ_v be the unique geodesic from x to y, i.e., $\gamma_v(t) := \cosh(t)x + \sinh(t)v$, where v is a timelike unit vector perpendicular to x. That is, γ_v is the unit-speed parameterization of γ_{xy}. Let $q := \gamma_v(m)$ be the (unique) point in the intersection of $[\bar{x}, \bar{y}]$ and the geodesic extending $\gamma_{\bar{z}\bar{p}}$. Applying a suitable Lorentz transformation we may assume that $x = x'$, $y = y'$ and so $v = v'$ and $a = a'$. We let a vary on the geodesic γ_v between the parameters 0 and m. The point a is given by $a = \gamma_v(t) = \cosh(t)x + \sinh(t)v = a'$. Then define $f : [0, m] \to \mathbb{R}$,

$$f(t) := \cosh(\tau(a, p) + \tau(p, z)) - \cosh(\tau(a, z'), \tag{4.3.11}$$

where we applied the monotone function cosh to make computation easier. As in the Minkowski case we know $f(0) = 0$ and $f(m) \geq 0$. But we will show $f''(t) - f(t) \leq 0$ instead of $f''(t) \leq 0$, which suffices to infer $f(t) \geq 0$ by Lemma 4.1.5. Note that in de Sitter space we have $\tau(x, y) = \operatorname{arcosh}(\langle x, y\rangle)$ for $x \ll y$, cf. [12]. Then with the help of addition theorems for cosh, the function f simplifies to

$$f(t) = \langle a, p\rangle\langle p, z\rangle + \sinh(\operatorname{arcosh}(\langle a, p\rangle))\sinh(\operatorname{arcosh}(\langle p, z\rangle)) - \langle a, z'\rangle. \tag{4.3.12}$$

Using the relation $\sinh(\operatorname{arcosh}(x)) = \sqrt{x^2 - 1}$, we further obtain

$$f(t) = \langle a, p\rangle\langle p, z\rangle + \sqrt{\langle a, p\rangle^2 - 1}\sqrt{\langle p, z\rangle^2 - 1} - \langle a, z'\rangle. \tag{4.3.13}$$

Now we take the first derivative. Note that of the three points appearing, only a depends on t. For now, we will leave the derivative of a not explicitly calculated. This will be simpler to do later on.

$$\frac{d}{dt}f(t) = \langle\frac{d}{dt}a, p\rangle\langle p, z\rangle + \sqrt{\langle p, z\rangle^2 - 1}\frac{\langle a, p\rangle\langle\frac{d}{dt}a, p\rangle}{\sqrt{\langle a, p\rangle^2 - 1}} - \langle\frac{d}{dt}a, z'\rangle. \tag{4.3.14}$$

For the second derivative, note that $(\frac{d}{dt})^2 a = a$. We calculate and simplify

$$\left(\frac{d}{dt}\right)^2 f(t) = \langle a, p\rangle\langle p, z\rangle + \sqrt{\langle p, z\rangle^2 - 1}$$

$$\left(\frac{\langle \frac{d}{dt}a, p\rangle^2 + \langle a, p\rangle^2}{\sqrt{\langle a, p\rangle^2 - 1}} - \frac{\langle a, p\rangle^2 \langle \frac{d}{dt}a, p\rangle^2}{\sqrt{\langle a, p\rangle^2 - 1}^3}\right) - \langle a, z'\rangle. \tag{4.3.15}$$

Now we observe that this very much resembles the original function f. Using this, we can write

$$\left(\frac{d}{dt}\right)^2 f(t) - f(t) = -\sqrt{\langle a, p\rangle^2 - 1}\sqrt{\langle p, z\rangle^2 - 1} + \sqrt{\langle p, z\rangle^2 - 1}$$
$$\left(\frac{\langle \frac{d}{dt}a, p\rangle^2 + \langle a, p\rangle^2}{\sqrt{\langle a, p\rangle^2 - 1}} - \frac{\langle a, p\rangle^2 \langle \frac{d}{dt}a, p\rangle^2}{\sqrt{\langle a, p\rangle^2 - 1}^3}\right). \tag{4.3.16}$$

Then by Lemma 4.1.5 it suffices to show that the right hand side is lesser than or equal to 0. We further simplify this expression and observe that it is equivalent to

$$\langle a, p\rangle^2 - 1 - \langle \frac{d}{dt}a, p\rangle^2 \leq 0. \tag{4.3.17}$$

Now we insert $a = \cosh(t)x + \sinh(t)v$ and $\frac{d}{dt}a = \sinh(t)x + \cosh(t)v$ into this inequality and at the same time use the bi-linearity of the scalar product. In this way we get

$$(\cosh(t)\langle x, p\rangle + \sinh(t)\langle v, p\rangle)^2 - 1 - (\sinh(t)\langle x, p\rangle + \cosh(t)\langle v, p\rangle)^2 \leq 0. \tag{4.3.18}$$

After computing the big square, we factor appropriately and use $\cosh(t)^2 - \sinh(t)^2 = 1$ to obtain

$$\langle x, p\rangle^2 - 1 \leq \langle v, p\rangle^2. \tag{4.3.19}$$

Since $\tau(x, p) = \operatorname{arcosh}(\langle x, p\rangle)$, we have $\langle x, p\rangle = \cosh(\tau(x, p))$ and so $\cosh(\tau(x, p))^2 - 1 = \sinh(\tau(x, p))^2$. Thus, after taking the root on both sides we finally end up with

$$\sinh(\tau(x, p)) \leq |\langle v, p\rangle|. \tag{4.3.20}$$

We now try to give these reformulations some intuitive value. First note that since v by definition is perpendicular to x, we can write $|\langle v, p\rangle| = |\langle v, p - x\rangle|$. This gives this scalar product a bit more meaning, because v is a (unit) tangent vector at x while p is a point in de Sitter space. Now both $p - x$ and v can be regarded as vectors in the ambient Minkowski space starting at x. In particular, they are timelike (since $x \ll p$ and $x \ll y$) and so we can use the hyperbolic angle:

$$\begin{aligned}|\langle v, p - x\rangle| &= -\langle v, p - x\rangle = \|v\|\|p - x\| \cosh(\measuredangle_x(v, p - x)) \\ &= \|p - x\| \cosh(\measuredangle_x(v, p - x)).\end{aligned} \tag{4.3.21}$$

We now try to minimize this expression while letting v vary among unit tangent vectors. In this way, only $\cosh(\measuredangle_x(v, p - x))$ changes, and it is minimal if and only if $\measuredangle_x(v, p - x)$ is minimal. This angle is minimal precisely if v is the unit vector in the direction of the projection of $p - x$ onto the tangent space. Since p is in a normal neighbourhood of x, we can write $p = \cosh(s)x + \sinh(s)w$ for some appropriate unit vector w and $s = \tau(x, p)$. Then $p - x = (\cosh(s) - 1)x + \sinh(s)w$. We can consider this as an orthogonal decomposition with respect to T_x dS: $(\cosh(s) - 1)x$ is the orthogonal component and $\sinh(s)w$ is the parallel component. Then the angle is minimal if $v = w$. In particular, we then have

$$\begin{aligned}\langle p - x, v\rangle &= \langle(\cosh(s) - 1)x + \sinh(s)v, v\rangle = (\cosh(s) - 1)\langle x, v\rangle\\ &\quad + \sinh(s)\langle v, v\rangle\\ &= -\sinh(s) = -\sinh(\tau(x, p)).\end{aligned}$$

Hence $-\langle p - x, v\rangle = \sinh(\tau(x, p))$ which is precisely the left hand side. So if v is any other (unit) vector than the projection of $x - p$ onto T_x dS then $|\langle v, p\rangle|$ would be even larger. Thus, the inequality holds and Lemma 4.1.5 implies $f(t) \geq 0$, which in turn implies $\cosh(\tau(a, p) + \tau(p, z)) \geq \cosh(\tau(a, z'))$. Since cosh is monotonously increasing, we finally obtain $\tau(a, p) + \tau(p, z) \geq \tau(a, z')$.

Anti-de Sitter space $(K = -1)$: This case is very similar to the case of positive curvature. In fact, one can mostly copy the assumptions and computations from the de Sitter case. This time we define $f : [0, m] \to \mathbb{R}$,

$$f(t) = \cos(\tau(a, p) + \tau(p, z)) - \cos(\tau(a, z')). \tag{4.3.22}$$

We omit the calculations for this case. Keep in mind that as cos is decreasing, we have to reverse some inequalities.
Clearly, the last two cases also work for any other $K \neq 0$, as one only has to respect additional scaling constants.
Letting $\bar{b}$ vary: As announced before, we do not only want to cover the case $\bar{b} = \bar{z}$ but rather let $\bar{b}$ vary. First suppose $\bar{b} \in [\bar{p}, \bar{z}]$. We consider a function $f_b : [0, m] \to \mathbb{R}$,

$$f_b(t) := g(\bar{\tau}(\bar{a}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{b})) - g(\bar{\tau}(\bar{a}', \bar{b}')), \tag{4.3.23}$$

where g is, depending on K, one of the three monotone functions we introduced for easier computations. In the above calculations concerning the second derivative, $\bar{z}$ (and $\bar{z}'$) did not really play an important role, so we could have easily exchanged this point for $\bar{b}$ (or $\bar{b}'$). In particular, we thus have $f_b(t) \geq 0$ if we can show that f_b is nonnegative at the boundary of $[0, m]$. By definition we have $f_b(0) = 0$. For $f_b(m)$, i.e., $\bar{a} = \bar{q}$, note that

$$\bar{\tau}(\bar{q}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{b}) + \bar{\tau}(\bar{b}, \bar{z}) = \bar{\tau}(\bar{q}, \bar{z}) \geq \bar{\tau}(\bar{q}', \bar{z}') \geq \bar{\tau}(\bar{q}', \bar{b}') + \bar{\tau}(\bar{b}', \bar{z}'). \tag{4.3.24}$$

Since we have $\bar{\tau}(\bar{b}, \bar{z}) = \bar{\tau}(\bar{b}', \bar{z}')$ by construction, $\bar{\tau}(\bar{q}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{b}) \geq \bar{\tau}(\bar{q}', \bar{b}')$ follows and so $f_b(m) \geq 0$.
Now assume $\bar{b} \in [\bar{y}, \bar{z}]$. In contrast to $\bar{b} \in [\bar{p}, \bar{z}]$, we have to restrict the position of $\bar{a}$ even further since it may otherwise happen that the connecting geodesic $[\bar{a}, \bar{b}]$ stays inside the triangles. The values of $\bar{a}$ we have to check actually depend on the position of $\bar{b}$ in the following way: denote by $\bar{r} =: \gamma_v(n)$ the point in the intersection of $[\bar{x}, \bar{y}]$ and the extension of $[\bar{p}, \bar{b}]$, where γ_v is the unique geodesic from x to

y. In particular, $n \leq m$ with equality if and only if $\bar{b} = \bar{z}$. Then we consider $f_b : [0, n] \to \mathbb{R}$ as above. As for $f_b(n)$, note that we have

$$\bar{\tau}(\bar{r}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{b}) - \bar{\tau}(\bar{r}', \bar{b}') = \bar{\tau}(\bar{r}, \bar{b}) - \bar{\tau}(\bar{r}', \bar{b}') \geq 0. \qquad (4.3.25)$$

In particular, $\bar{\tau}(\bar{r}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{b}) \geq \bar{\tau}(\bar{r}', \bar{b}')$. By applying g and then bringing everything to one side we infer the claim $f_b(n) \geq 0$. Also, $f_b(0) = g(\bar{\tau}(\bar{x}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{b})) - g(\bar{\tau}(\bar{x}', \bar{b}'))$ is completely analogous to the case of $\bar{a}$ and $\bar{z}$ we explicitly calculated (with reversed time orientation). Thus, we have finally shown all inequalities and so T_3 actually satisfies the same curvature bound from above.
As a last remark, we observe that every argument can be repeated if the timelike relation of p and y is reversed, i.e., if $y \ll p$. This is entirely analogous and basically differs only up to time orientation. □

For applications to a gluing theorem, we of course do not want to restrict the placement of p to only the longest side in a triangle. As the proof of Lemma 4.3.1 basically only uses Lemma 4.2.1, and there the placement of p did not matter, we immediately obtain the following corollary.

Corollary 4.3.2. (Remaining constellations of the gluing lemma) *Let X and U be as in Lemma 4.3.1 and let $\Delta(x, y, z)$ be a timelike triangle in U. Let p be a point on $[x, y]$ (or $[y, z]$) and consider the two resulting subtriangles that share the (timelike) segment $[p, z]$ (or $[x, p]$). If the subtriangles satisfy a timelike curvature bound from above, then so does the original triangle.*

Proof. In this case we mainly use Lemma 4.2.2 and then observe that everything works out similarly as in the situation of the gluing lemma we just proved. □

Finally, we need a version of the gluing lemma that applies even when the shared segment in a Lorentzian pre-length space is not timelike. In general, this would not make any sense in an ordinary Lorentzian pre-length space as we have no concept of spacelike distance in such spaces. However, the main application of these lemmas is when the two spaces in the amalgamation are manifolds. In this case it does make sense to talk about a shared segment which is not timelike.

Lemma 4.3.3. (Gluing lemma for manifolds) *Let X and U be as in Lemma 4.3.1. Let X_1 and X_2 be two strongly causal spacetimes and $U_i \subseteq X_i$, $(i = 1, 2)$ with the same assumptions as in Lemma 4.3.1 as well. Let $T_3 := \Delta(x, y, z)$ be a timelike triangle in $U \subseteq X$ and let $p \in [x, z]$. Assume that there exist geodesic triangles $T_1 := \Delta(x^1, p^1, y^1)$ and $T_2 := \Delta(p^2, y^2, z^2)$ in U_1 and U_2, respectively, such that $\tau(x^1, p^1) = \tau(x, p)$, $\tau(x^1, y^1) = \tau(x, y)$, $\tau(p^2, z^2) = \tau(p, z)$, $\tau(y^2, z^2) = \tau(y, z)$ and $|p^1 y^1|_\pm = |p^2 y^2|_\pm$. If T_1 and T_2 satisfy curvature bounds above by K in the sense of* [2],[6] *then T_3 has timelike curvature bounded above by K in the sense of Definition 2.1.3.*

[6] In [2], a Lorentzian manifold is said to have (sectional) curvature bounded above by $K \in \mathbb{R}$ if spacelike sectional curvatures are $\leq K$ and timelike sectional curvatures are $\geq K$. We say a geodesic triangle satisfies such a curvature bound if $|ab|_\pm \leq |\bar{a}\bar{b}|_\pm$ for all points a, b in the triangle and corresponding comparison points $\bar{a}, \bar{b}$ in the comparison triangle in M_K.

Proof. The proof of Lemma 4.3.1 can be adapted entirely into this setting. One easily observes that the fact that the shared segment is not timelike does not impact the comparison calculations. □

Remark 4.3.4. (Gluing with reverse curvature bound) In the metric case, the gluing theorem only works when considering spaces with an upper curvature bound. Roughly speaking, this is because in Alexandrov's lemma one angle always increases independently of whether one starts with a concave or convex quadrilateral. This in turn prevents the gluing lemma from working in the other direction. In our case we encounter very similar behaviour: in Lemma 4.2.1, the inequality $\angle xyz \geq \angle x'y'z'$ is independent of the general shape of the quadrilateral. Note that in the Lorentzian setting we want this inequality to point in the other direction compared to the two other angle estimates because of the sign of the nonnormalized angle.

Remark 4.3.5. (Hyperbolic angles in Lorentzian pre-length spaces) The first author is currently working together with C. Sämann on adapting the concept of hyperbolic angles to the synthetic setting. The approach is similar to the Alexandrov angle in metric geometry, cf. [Definition I.1.12] [6]. This may simplify some of the calculations in this section while also allowing nicer formulations of some statements.

5. The gluing theorem

This section covers the proof of an analogue of the gluing theorem of Reshetnyak for CAT(k) spaces. As mentioned before, a gluing theorem for Lorentzian pre-length spaces at present does not seem to be possible, at least not in a reasonably general formulation. Indeed, one would require that (locally) any two points in the identified set are timelike related, which basically forces each A_i to be a timelike geodesic (or a "discrete" union of timelike geodesics). The problem is essentially the missing concept of spacelike distance. Our proof idea very much follows the metric case, where one subdivides a triangle into smaller triangles and then applies Lemma 4.3.1. In the Lorentzian setting, however, it may happen that such a division into *timelike* triangles is not always possible: consider two copies of $\mathbb{R}^3_1$ glued along a vertical plane. Then one can construct an arbitrarily small timelike triangle whose intersection with the plane consists of two points which are spacelike related.
Nevertheless, we can formulate the gluing theorem for Lorentzian manifolds viewed as Lorentzian pre-length spaces. In this setting we have access to spacelike distances and the triangle subdivision problem can be overcome.

5.1. Constructing comparison neighbourhoods

One main obstacle in the proof of the gluing theorem is the construction of suitable comparison neighbourhoods. In particular, the comparison neighbourhoods of points in the glued set additionally have to be geodesically convex rather than just ensure the existence of maximizers between timelike related points. This is because

when applying the gluing lemma in the spirit of Reshetnyak it might happen that the triangle subdivision occurs along a spacelike or null geodesic. Moreover, we also need them to be causally convex to ensure the existence of geodesics between the two spaces (obtained as a limit of converging curves).
Fortunately, in recent work by E. Minguzzi such neighbourhoods are constructed, cf. [18]. The basic idea is the following: let M be a spacetime, $p \in M$ and $U \subseteq M$ a normal neighbourhood of p. Define the function $D_p^2 : U \to \mathbb{R}$ by

$$D_p^2(q) := g_p(\exp_p^{-1}(q), \exp_p^{-1}(q)). \tag{5.1.1}$$

That is, the function D_p^2 is the (squared) signed distance from the point p (but applying the sign after the squaring). In the terminology of [2], $D_p^2(q)$ is called the energy $E_q(p) = E_p(q)$. Let γ be a timelike geodesic through p. Then from [18, Lemma 4] we infer the existence of q_1, q_2 in the image of γ, $q_1 \ll p \ll q_2$, and a normal neighbourhood O of p such that, setting $c_1 := D_{q_1}^2(p)$, $c_2 := D_{q_2}^2(p)$ (both of which are negative), for sufficiently close $c_1' > c_1$ and $c_2' > c_2$ the set

$$(D_{q_1}^2)^{-1}(-\infty, c_1') \cap (D_{q_2}^2)^{-1}(-\infty, c_2') \cap O \tag{5.1.2}$$

is geodesically convex and globally hyperbolic. Moreover, in this proof the set O is chosen in such a way that $\overline{O} \subseteq I(q_1, q_2)$ and the (closure of the) desired connected component of $(D_{q_1}^2)^{-1}(-\infty, c_1') \cap (D_{q_2}^2)^{-1}(-\infty, c_2')$ is contained in O. Hence the set in (5.1.2) is causally convex as well. This set might be visually described as a “lens”, i.e., the (bounded) region obtained by intersecting two hyperboloids in Minkowski space. In our setup, this region can more easily be described by using the time separation function τ instead of D_p^2. We will go into a bit more detail on how we adapt this construction in Remark 5.1.2.

Remark 5.1.1. (Assumptions on the glued set) When viewing manifolds as Lorentzian length spaces, one has to be especially careful when dealing with submanifolds. This is because on a Lorentzian submanifold $A \subseteq M$ there are a priori two – potentially strongly differing – structures when viewing them as Lorentzian pre-length spaces. One the one hand, there is the induced substructure on A as a Lorentzian submanifold which is very common in Lorentzian geometry. This leads to concepts such as relative causality relations, usually denoted by $\ll_A$ and $\leq_A$. In particular, a submanifold equipped with this induced structure is always a Lorentzian length space. On the other hand, the restricted structure as a Lorentzian pre-length space just considers the restriction of the causality relations and the time separation function to the submanifold. In this case, A is usually just a Lorentzian pre-length space and one will lose the original description of τ, $\ll$ and $\leq$. For example, if $p, q \in A$, $p \leq q$, then clearly $p \leq|_A q$. But if there is no causal curve from p to q that stays inside A, we do not get $p \leq_A q$. This is why the restricted Lorentzian structure finds little application from a relativistic or differential geometric viewpoint. A subspace of a metric space is a common example where one uses the restricted structure. But also in this case, if one starts with a length space then an arbitrary subspace will only be a metric space and not a length space in general.

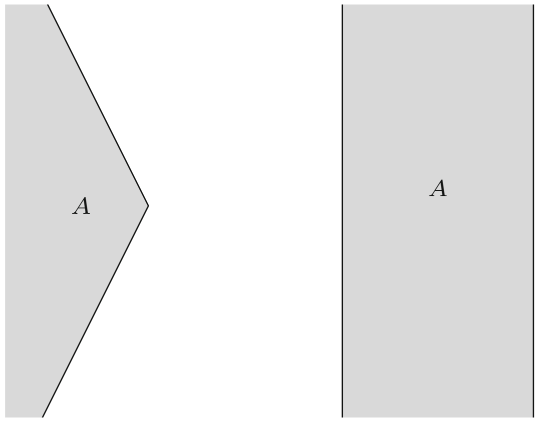

Fig. 12. Two possible choices for A in the Minkowski plane: a "half-space" with cornered boundary and a vertical strip

This observation is one of the reasons we decided to not require the identified sets to be submanifolds (the other of course is a pursuit of generality). We now summarize the properties we require of the identified sets. Let A_1 and A_2 be two closed subsets of strongly causal spacetimes M_1 and M_2, respectively. Let $f : A_1 \to A_2$ be a locally bi-Lipschitz homeomorphism (we will formulate the gluing theorem without an artificial space A). Then we require A_1, A_2 and f to be compatible in the following way:

(i) A_1 and A_2 are non-timelike locally isolating.
(ii) f is τ-preserving and $\leq$-preserving.
(iii) A_1 is "convex" in the following sense: for all $p \in A_1$ there exists a normal neighbourhood $V_1 \subseteq M_1$ of p such that the following holds: whenever $x, y \in U_1 := V_1 \cap A_1$ then the unique geodesic connecting them in V_1 is contained in U_1. The same holds for A_2.
(iv) f locally preserves the signed distance: for all $p \in A_1$ there exists a normal neighbourhood $V_1 \subseteq M_1$ of p such that the following holds: whenever $x, y \in U_1 := V_1 \cap A_1$ then $|xy|_\pm = |f(x)f(y)|_\pm$.

Observe that both (iii) and (iv) also hold for any smaller (convex) normal neighbourhood contained in V. (iii) suggests that the set A_i is at least contained in a totally geodesic submanifold. However, we still gain the possibility of admitting boundaries and even corners, see Fig. 12.

This is also the reason why we do not exactly copy the construction mentioned in [18, Lemma 4], because there the sets in (5.1.2) have both "control points" on a geodesic through p. And if p is a boundary point of A there might not be a geodesic *through* p that stays inside A.

Remark 5.1.2. (Modified lenses) Our modification of the sets in (5.1.2) is very minor. Let $p \in A \subseteq M$ be as in Remark 5.1.1. Choose normal coordinates (φ, U) centered at p. In the chart neighbourhood $\varphi(U)$, let $B_S(0)$ be a Euclidean ball around $0 = \varphi(p)$ such that $\varphi^{-1}(B_S(0))$ is convex (this is equivalent to saying that $B_S(0)$ is convex with respect to the push forward metric induced by g) and moreover such that for all smaller balls entirely contained in $B_S(0)$, their (inverse) image is convex as well. By the non-local timelike isolation of A we find $b_-, b_+ \in B_S(0) \cap \varphi(A)$ which are τ^η-equidistant[7] (and hence, as p is the center of the normal coordinate chart, also τ^g-equidistant) from 0 such that $b_- \ll 0 \ll b_+$. Consider the timelike straight line segments $[b_-, 0]$, $[0, b_+]$ in $B_S(0)$ (which are actually contained in $\varphi(A)$ as well by the convexity of A and moreover correspond to the geodesic segments in the manifold). These segments intersect at 0 in an ordinary Minkowski hyperbolic angle of $\omega < \infty$ (which is also the hyperbolic angle measured in the manifold). Apply a Lorentz transformation if necessary to position these segments in the plane spanned by ∂_0, ∂_1 and such that they are symmetric (from a Euclidean point of view) with respect to the ∂_1 direction. Then the segment $[b_-, b_+]$ is parallel to the ∂_0-axis. If necessary, choose S even smaller such that the ball is still convex after the Lorentz transformation. Introduce a flat metric in the chart neighbourhood via

$$\eta^+ := -(2dx_0)^2 + \sum_{i=1}^{n}(dx_i)^2. \tag{5.1.3}$$

Clearly, $g_p = \eta < \eta^+$[8] and by choosing S smaller if necessary we can assume that $g < \eta^+$ holds on all of $B_S(0)$. Set $R := \tau^{\eta^+}(b_-, b_+)$. Consider the following set:

$$L := \left\{x \in B_S(0) \mid \tau^{\eta^+}(b_-, x) > \frac{R}{3}, \tau^{\eta^+}(x, b_+) > \frac{R}{3}\right\}, \tag{5.1.4}$$

which will be our replacement for the set in (5.1.2). Visually, this set may be described as a “wide lens”. Using the law of cosines, an elementary calculation yields $0 \in L$ independent of ω and R, see Fig. 13.

Say $b_- := \gamma_1(t)$, $b_+ = \gamma_2(t)$, where γ_1 and γ_2 are the (unit speed) geodesics corresponding to the extensions of the straight line segments from above. By choosing t smaller, i.e., by moving b_- and b_+ closer to 0 along these segments (but still keeping them equidistant to 0), it easily follows that R decreases as well. Moreover, L has a “width” which is only dependent on R and can be easily calculated. Denote by q the midpoint of $[b_-, b_+]$, which due to the above applied Lorentz transformation is situated on the ∂_1-axis. Let r be half the width of L. By choosing t small enough, we can achieve that $B_r(q) \subseteq B_S(0)$. In particular, $L \subseteq B_r(q)$ anyways holds and $\varphi^{-1}(B_r(q))$ is geodesically convex. With arguments very similar to [18, Theorem 10, Theorem 11 & Lemma 4] one can then show that $\varphi^{-1}(L)$ is geodesically convex, causally convex and globally hyperbolic.

[7] As the time separation is defined as the supremum of lengths of causal curves, it clearly is dependent on which metric is considered. We denote by τ^g the time separation measured with respect to the metric g. In the same spirit, we can define the signed distance with respect to a certain metric, denoted by $|\cdot\cdot|_{\pm}^{g}$, which is the g-length of the unique g-geodesic connecting two points (of course the two points have to be in a normal neighbourhood with respect to g).

[8] For two Lorentzian metrics g and h on a manifold we say h has wider lightcones than g, and write $g < h$, if for all points the timelike past/future with respect to h contains the causal past/future with respect to g, cf. [18]. In other sources, one may find the notation $g \prec h$ and the formulation $g(v, v) \leq 0 \Rightarrow h(v, v) < 0$ for all $v \in TM$.

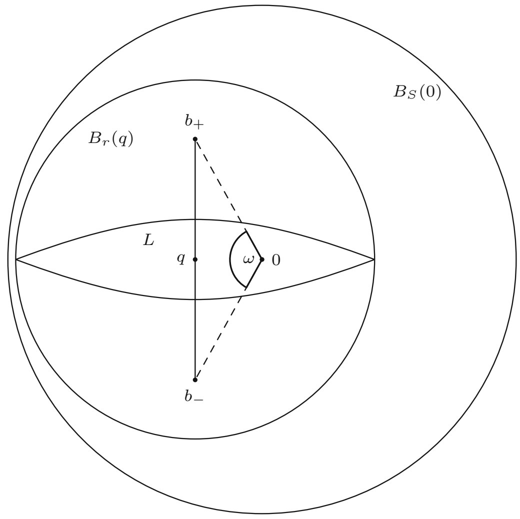

Fig. 13. Constructing nice neighbourhoods in the spirit of [18]

5.2. Formulation and proof of the gluing theorem

Now we can finally state and prove the Lorentzian analogue of the gluing theorem.

Theorem 5.2.1. (Reshetnyak's gluing theorem, Lorentzian version) *Let* (X_1, g_1) *and* (X_2, g_2) *be two smooth and strongly causal spacetimes with* $\dim(X_1) =: n \geq m := \dim(X_2)$. *Let* A_1 *and* A_2 *be two closed non-timelike locally isolating subsets of* X_1 *and* X_2*, respectively. Let* $f : A_1 \to A_2$ *be a* τ*-preserving and* $\leq$*-preserving locally bi-Lipschitz homeomorphism which locally preserves the signed distance. Suppose* A_1 *and* A_2 *are convex in the sense of Remark* 5.1.1*(iii). Suppose* X_1 *and* X_2 *have (sectional) curvature bounded above by* $K \in \mathbb{R}$ *in the sense of* [2]. *Then the Lorentzian amalgamation* $X := X_1 \sqcup_A X_2$[9] *is a Lorentzian pre-length space with timelike curvature bounded above by* K.

Proof. By [16, Proposition 3.5, Example 3.24, Theorem 3.26 & Example 4.9], X_1 and X_2 are strongly causal regular (SR)-localizable Lorentzian length spaces with timelike curvature bounded above by K in the sense of [16]. Let $[p] \in X$ and assume first $[p] = \{p^1\}$. We can choose comparison neighbourhoods in X_1 (and in X_2) to be (small enough) timelike diamonds, cf. [5, Remark 2.2.12]. Since A_1 is closed, we find a neighbourhood $U_1 \subseteq X_1$ of p^1 which does not meet A_1. Since X_1 is strongly causal, there exists a timelike diamond $I_1(x^1, y^1)$ in U_1 containing p^1. In particular, this timelike diamond does not meet A_1 and hence $I_X([x], [y]) = \pi(I_1(x^1, y^1))$ by Lemma 3.3.8. Since $\widetilde{\tau}$ restricts to τ_1 on X_1 by Proposition 3.3.7, it easily follows that $I_X([x], [y])$ is a comparison neighbourhood for $[p] \in X$. Thus, we only have to further investigate points in A.
Let $p^1 \in A_1$. Choose normal coordinates (φ_1, V_1) around p_1 in X_1 such that V_1 is globally hyperbolic, convex in the sense of Remark 5.1.1(iii) and f preserves

[9] Technically, in this formulation the space A does not exist. As it is still convenient in the proof to say when an equivalence class consists of two elements, we keep this notation. One can simply define $A := \pi(A_1 \sqcup A_2)$.

the signed distance on $V_1 \cap A_1 =: U_1$. We can further assume that V_1 is contained in a comparison neighbourhood and a causally closed neighbourhood of p^1 in X_1 and all small enough balls inside V_1 are geodesically convex as well. Moreover, by choosing V_1 smaller if necessary, we can also assume that $f(U_1) := U_2$ is of the form $U_2 = V_2 \cap A_2$ where V_2 has all the properties we imposed on V_1. In particular, (φ_2, V_2) are normal coordinates around p^2 in X_2. As A_1 is non-timelike locally isolating, we find $b^1_-, b^1_+ \in U_1$ which we choose τ^{g_1}-equidistant from p^1 such that $b^1_- \ll_1 p^1 \ll_1 b^1_+$. Since f is τ-preserving and hence $\ll$-preserving, we have $b^2_-, b^2_+ \in U_2$ as well as $b^2_- \ll_2 p^2 \ll_2 b^2_+$ (and they are τ^{g_2}-equidistant from p^2). By applying a Lorentz transformation, we can assume $\varphi_1(p^1) = 0 = \varphi_2(p^2)$, $\varphi_1(b^1_-) = \varphi_2(b^2_-) =: b_-$, $\varphi_1(b^1_+) = \varphi_2(b^2_+) =: b_+$ and the straight line segments $[b_-, 0]$ and $[0, b_+]$ lie in the plane spanned by ∂_0, ∂_1 and b_- and b_+ are symmetric with respect to the ∂_1 direction (as in Remark 5.1.2). Now consider a set as the one described in (5.1.4) and recall the metric defined in (5.1.3). Set $R := \tau^{\eta^+}(b_-, b_+)$. That is, define the set

$$\widetilde{L} := \left\{ x \in \mathbb{R}^n \mid \tau^{\eta^+}(b_-, x) > \frac{R}{3}, \tau^{\eta^+}(x, b_+) > \frac{R}{3} \right\} \tag{5.2.1}$$

and set $L_i := \varphi_i^{-1}(\widetilde{L})$[10]. By the observations in Remark 5.1.2, we have $0 \in \widetilde{L}$ and hence $p^i \in L_i$. In particular, by moving b_- and b_+ closer to 0 along these segments, we can achieve that both L_1 and L_2 have the desired properties of geodesic convexity, causal convexity and global hyperbolicity. While we can make no statement about the equality of L_1 and L_2 as they belong to different manifolds which we cannot relate as a whole, we indeed can say something about their intersections with A_i! Namely, we claim $f(L_1 \cap A_1) = L_2 \cap A_2$. This fact is absolutely essential so that the neighbourhoods of $[p]$ coming from X_1 and X_2, respectively, are compatible. To this end observe that as we are in normal coordinates around p^i and f is signed distance preserving, we have

$$|0\varphi^1(a^1)|^\eta_\pm = |p^1a^1|^{g_1}_\pm = |p^2a^2|^{g_2}_\pm = |0\varphi^2(a^2)|^\eta_\pm. \tag{5.2.2}$$

This in turn yields an equality on the nonnormalized angles:

$$\measuredangle b_- 0 \varphi_1(a^1) = \measuredangle b_- 0 \varphi_2(a^2). \tag{5.2.3}$$

Thus, the η-triangles $\Delta(b_-, 0, \varphi_i(a^i))$, $i = 1, 2$ have two sides of equal length and an equal angle between these sides. Then the law of cosines implies that the opposite side is equal as well, i.e., $|b_-\varphi^1(a^1)|^\eta_\pm = |b_-\varphi^2(a^2)|^\eta_\pm$. An analogous argument gives $|\varphi^1(a^1)b_+|^\eta_\pm = |\varphi^2(a^2)b_+|^\eta_\pm$. Thus, $\varphi_1(a^1)$ and $\varphi_2(a^2)$ both lie on the intersection of two lightcones and/or hyperboloids with respect to η. In any case this intersection is a Euclidean sphere of two dimensions less than the manifold (if the manifolds do not have the same dimension, it is two less than the lower dimensional manifold). Moreover, the center of this sphere is located on $[b_-, b_+]$ and thus, as $[b_-, b_+]$ is parallel to ∂_0, all points on this sphere have the

[10] If the dimension of X_2 is lower, one may want to view φ_2 as embedding $\mathbb{R}^m$ into $\mathbb{R}^n$, i.e., $\varphi_2 : V_2 \to \mathbb{R}^m \times \{0\} \subseteq \mathbb{R}^n$.

same time coordinate. Because of the equality on the signed distance we obtain $\eta(b_- - a^1, b_- - a^1) = \eta(b_- - a^2, b_- - a^2)$. This, together with the fact that the first components of a^1 and a^2 are equal, implies $\eta^+(b_- - a^1, b_- - a^1) = \eta^+(b_- - a^2, b_- - a^2)$ and hence $\tau^{\eta^+}(b_-, a^1) = \tau^{\eta^+}(b_-, a^2)$. In particular, we have $\varphi_1(a^1) \in \widetilde{L}$ if and only if $\varphi_2(a^2) \in \widetilde{L}$ and so $a^1 \in L_1 \cap A_1$ if and only if $a^2 \in L_2 \cap A_2$. We claim that $L := \pi(L_1 \sqcup L_2)$ is a comparison neighbourhood for $[p] \in A$.

We first show that L is timelike geodesic, i.e., for all $[x], [y] \in L$ with $[x] \mathrel{\widetilde{\ll}} [y]$ there exists a $\widetilde{\tau}$-realizing causal curve from $[x]$ to $[y]$ contained in L. To this end we claim the following: Let $\{x^1\} = [x] \in \pi(L_1) \subseteq L$ and $\{y^2\} = [y] \in \pi(L_2) \subseteq L$. Then $J_X([x], [y]) \cap A \subseteq L \cap A$. To see this, let $[q] \in J_X([x], [y]) \cap A$. Then $[q] = \{q^1, q^2\}$. Since $[x] \mathrel{\widetilde{\leq}} [q]$, we find a chain of the form $x^1 \leq_1 a_1^1 \sim a_1^2 \leq_2 a_2^2 \sim a_2^1 \leq_1 \ldots \sim a_n^1 \leq_1 q^1$. Since f is $\leq$-preserving, we obtain $x^1 \leq_1 q^1$ by the transitivity of $\leq_1$. Since $g_i < \eta^+$, we get $\varphi_1(x^1) \ll^{\eta^+} \varphi_1(q^1)$. Clearly, $b_- \ll^{\eta^+} \varphi_1(x^1)$ as $x^1 \in L_1$. Then by the reverse triangle inequality we obtain $\tau^{\eta^+}(b_-, \varphi_1(q^1)) \geq \tau^{\eta^+}(b_-, x) > \frac{R}{3}$. A similar argument implies $\tau^{\eta^+}(\varphi_2(q^2), b_+) \geq \tau^{\eta^+}(\varphi_2(y^2), b_+) > \frac{R}{3}$. By the above arguments, we obtain both inequalities for $\varphi_1(q^1)$ and $\varphi_2(q^2)$. Thus, these points are contained in $\widetilde{L}$ and so $q^i \in L_i \cap A_i$ and $[q] \in L \cap A$ follows.

Now let $[x], [y] \in L$ with $[x] \mathrel{\widetilde{\ll}} [y]$. If both belong to one space, we take the (projection of the) original geodesic connecting them. More precisely, if $x^i \in [x]$, $y^i \in [y]$, $i \in \{1, 2\}$, then $\pi \circ \gamma_{x^i y^i}$ is a $\widetilde{\tau}$-realizing curve connecting $[x]$ and $[y]$ (since $\widetilde{\tau}$ restricts to τ_i on X_i). It is even contained in $\pi(L_i)$ as L_i is causally convex. So the only relevant case is (up to symmetry) $\{x^1\} = [x]$ and $\{y^2\} = [y]$ with $[x] \mathrel{\widetilde{\ll}} [y]$. By Proposition 3.3.7 we find a sequence $([a_n])_{n\in\mathbb{N}}$ such that

$$\lim_{n\to\infty} \tau_1(x^1, a_n^1) + \tau_2(a_n^2, y^2) = \widetilde{\tau}([x], [y]). \tag{5.2.4}$$

By definition, we have $[a_n] \in J_X([x], [y]) \cap A$ for all n. By the above considerations, it then follows that $[a_n] \in L \cap A$ for all n. In particular, we find corresponding sequences $(a_n^1)_{n\in\mathbb{N}}$ in $L_1 \cap A_1$ and $(a_n^2)_{n\in\mathbb{N}}$ in $L_2 \cap A_2$. By construction, we have $L_1 \subseteq J_1(b_-^1, b_+^1)$ and $L_2 \subseteq J_2(b_-^2, b_+^2)$. Since V_1 and V_2 are globally hyperbolic and A_1 and A_2 are closed, we find that these sequences converge to some $a^1 \in A_1$ and $a^2 \in A_2$. In particular, $[a_n] \to [a] \in \overline{J_X([x], [y])} \cap A$. As $[a_n] \in J_X([x], [y]) \cap A$, we have $x^1 \leq_1 a_n^1$ and $a_n^2 \leq_2 y^2$ for all n. Since V_1 and V_2 are contained in a causally closed neighbourhood, it follows that $x^1 \leq_1 a^1$ and $a^2 \leq_2 y^2$ and hence $[a] \in J_X([x], [y]) \cap A$. In summary, we have

$$\begin{aligned}\widetilde{\tau}([x], [y]) &= \lim_{n\to\infty} \tau_1(x^1, a_n^1) + \tau_2(a_n^2, y^2) = \lim_{n\to\infty} \tau_1(x^1, a_n^1) + \lim_{n\to\infty} \tau_2(a_n^2, y^2) \\ &= \tau_1(x^1, a^1) + \tau_2(a^2, y^2),\end{aligned}$$

where the last equality follows since τ_1 and τ_2 are continuous on L_1 and L_2, respectively. Consider the two original geodesics from x^1 to a^1 and from a^2 to y^2 which are contained in L_1 and L_2, respectively, as these sets are causally convex. Then the concatenation of their projections is a $\widetilde{\tau}$-realizing curve from $[x]$ to $[y]$ contained in L.

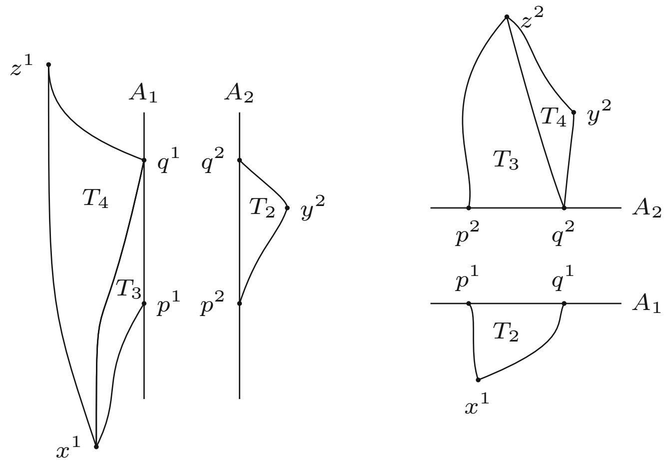

Fig. 14. The two cases of different triangle configurations

The fact that $\widetilde{\tau}|_{L\times L}$ is finite and continuous easily follows from L being timelike geodesic: if $x^i \in [x]$, $y^i \in [y]$, then $\widetilde{\tau}([x],[y]) = \tau_i(x^i, y^i)$ by Proposition 3.3.7. By assumption, this is finite and τ_i is continuous. So again, the case we have to investigate further is (up to symmetry) $\{x^1\} = [x]$ and $\{y^2\} = [y]$. If $\widetilde{\tau}([x],[y]) > 0$, then by the above there exists $[a] \in A$ such that $\tau_1(x^1, a^1) + \tau_2(a^2, y^2) = \widetilde{\tau}([x],[y])$. The left hand side is finite by assumption.

As for the continuity of $\widetilde{\tau}$, note that $\widetilde{\tau}$ is lower semi-continuous by definition. To see that $\widetilde{\tau}$ is upper semi-continuous in $([x],[y])$, let $[x_n] \to [x]$ and $[y_n] \to [y]$ be two sequences in L. We want to show $\widetilde{\tau}([x],[y]) \geq \limsup \widetilde{\tau}([x_n],[y_n])$. If $\limsup \widetilde{\tau}([x_n],[y_n]) = 0$, there is nothing to show. Otherwise (at least for a subsequence converging to the lim sup) we have $[x_n] \mathrel{\widetilde{\ll}} [y_n]$ for large enough n. In this case let $[a_n] \in A$ be such that $\tau_1(x_n^1, a_n^1) + \tau_2(a_n^2, y_n^2) = \widetilde{\tau}([x_n],[y_n])$, which exists since L is timelike geodesic. Similar to the above arguments, it follows that $a_n^1 \in L_1 \cap A_1$ and $a_n^2 \in L_2 \cap A_2$ (here we have $[a_n] \in J_X([x_n],[y_n])$, otherwise the argument is the same). By global hyperbolicity of V_1 and V_2 we infer the existence of a convergent subsequence. Without loss of generality the whole sequence converges, say $[a_n] \to [a]$. Then we compute

$$\begin{aligned}\widetilde{\tau}([x],[y]) &\geq \tau_1(x^1, a^1) + \tau_2(a^2, y^2) = \lim_{n\to\infty} \tau_1(x_n^1, a_n^1) + \lim_{n\to\infty} \tau_2(a_n^2, y_n^2) \\ &= \lim_{n\to\infty} \tau_1(x_n^1, a_n^1) + \tau_2(a_n^2, y_n^2) = \lim_{n\to\infty} \widetilde{\tau}([x_n],[y_n]),\end{aligned}$$

where the first equality holds since τ_1 and τ_2 are (locally) continuous.

So it is only left to show the triangle comparison condition. To this end consider a timelike triangle $T_1 := \Delta([x],[y],[z])$ in L. Clearly, if all three points lie in a single space, then the triangle satisfies the curvature bound by assumption. So we only need to consider triangles passing through A, for which there are two possibilities. Either an endpoint ($[x]$ or $[z]$) of the triangle is isolated, i.e., one endpoint is in one space while the other two points are in the other space. Or the intermediate detour-point ($[y]$) is isolated, see Fig. 14.

The case where $[y]$ is isolated, say $[x],[z] \in X_1$, $[y] \in X_2$, is easily finished: indeed, choose arbitrary points $[p] \in [[x],[y]] \cap A$, $[q] \in [[y],[z]] \cap A$. Then

$[p], [q] \in L \cap A$. Since $[x] \mathrel{\tilde{\ll}} [p] \mathrel{\tilde{\ll}} [y] \mathrel{\tilde{\ll}} [q] \mathrel{\tilde{\ll}} [z]$ and L is timelike geodesic, we find a $\tilde{\tau}$-realizing (causal) curve from $[p]$ to $[q]$ entirely contained in L. We can divide T_1 into three smaller timelike triangles. Consider $T_2 := \Delta([p], [y], [q])$, $T_3 := \Delta([x], [p], [q])$ and $T_4 := \Delta([x], [q], [z])$. By the convexity of A in the sense of Remark 5.1.1(iii), we have $[[p], [q]] \subseteq L \cap A$. Thus, $T_2 \subseteq X_2$ and $T_3 \subseteq X_1$ and so both timelike triangles satisfy the curvature bound. Then by Lemma 4.3.1 also the bigger triangle $T_5 := \Delta([x], [y], [q])$ formed by T_2 and T_3 satisfies the same curvature bound. Applying the gluing lemma once more to the triangles T_4 and T_5 we obtain the desired curvature bound for the whole triangle T_1 and this case is finished.

Finally, we consider the case where an endpoint is isolated, say $[x] \in X_1$, $[y], [z] \in X_2$ (the case of $[z]$ being isolated is clearly symmetric). Choose points $[p] \in [[x], [z]] \cap A$, $[q] \in [[x], [y]] \cap A$. Then $[p], [q] \in L \cap A$. Since L_1 and L_2 are convex in the sense of Remark 5.1.1(iii), there is a unique geodesic from p^1 to q^1 in $L_1 \cap A_1$ and a unique geodesic from p^2 to q^2 in $L_2 \cap A_2$, respectively. Moreover, these geodesics correspond to each other under f and have the same signed length since f preserves the signed distance. We denote this image in X by $[[p], [q]]$ although it may not be a causal curve if $[p]$ and $[q]$ are not causally related. By construction, we can identify the "subtriangles" $T_2 := \Delta([x], [p], [q])$ with $\Delta(x^1, p^1, q^1)$ and $T_3 := \Delta([p], [q], [z])$ with $\Delta(p^2, q^2, z^2)$. In this way we can view T_2 and T_3 as triangles in the manifolds X_1 and X_2, respectively, independently of the causal character of $[[p], [q]]$. Thus, we can apply the manifold version of the gluing lemma, cf. Lemma 4.3.3, to T_2 and T_3. In this way, we obtain that $T_5 := \Delta([x], [q], [z])$ satisfies the desired curvature bound. In particular, this is a valid timelike triangle in X. At last, we apply Lemma 4.3.1 to the triangles $T_4 := \Delta([q], [y], [z])$ and T_5 and we finally obtain that the original triangle T_1 satisfies the curvature bound as well. Thus, the proof is completed. □

Remark 5.2.2. (On regularity) Clearly, one can formulate the above theorem for Lorentzian metrics of regularity C^2. As the techniques in [18] work even in $C^{1,1}$, it is expected that this holds for the gluing theorem as well.

Acknowledgements We want to thank Michael Kunzinger and Clemens Sämann for helpful discussions and comments. This work was supported by research grant P33594 of the Austrian Science Fund FWF.

Declarations

Data Availibility Statement Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Funding Open access funding provided by Austrian Science Fund (FWF).

References

[1] Aké Hau, L., Cabrera Pacheco, A.J., Solis, D.A.: On the causal hierarchy of Lorentzian length spaces. Classical Quantum Gravity **37**(21), 215013 (2020)
[2] Alexander, S.B., Bishop, R.L.: Lorentz and semi-Riemannian spaces with Alexandrov curvature bounds. Commun. Anal. Geom. **16**(2), 251–282 (2008)
[3] Alexander, S. B., Graf, M., Kunzinger, M., Sämann, C.: Generalized cones as Lorentzian length spaces: causality, curvature, and singularity theorems. Commun. Anal. Geom. (2019). to appear, Preprint: arXiv:1909.09575
[4] Alexander, S.B., Kapovitch, V., Petrunin, A.: An Invitation to Alexandrov Geometry: CAT(0) spaces. SpringerBriefs in Mathematics. Springer, Berlin (2019)
[5] Beran, T.: Lorentzian length spaces. Master's thesis, University of Vienna (2020). https://phaidra.univie.ac.at/open/o:1363059
[6] Bridson, M.R., Haeflinger, A.: Metric Spaces of Non-positive Curvature. Comprehensive Studies in Mathematics, vol. 319. Springer, Berlin (1999)
[7] Burago, D., Burago, Y., Ivanov, S.: A Course in Metric Geometry. Graduate Studies in Mathematics, vol. 33. American Mathematical Society, Providence (2001)
[8] Burago, D., Ferleger, S., Kononenko, A.: A geometric approach to semi-dispersing billiards. Ergod. Theory Dyn. Syst. **18**(2), 303–319 (1998)
[9] Burago, D., Ferleger, S., Kononenko, A.: Uniform estimates on the number of collisions in semi-dispersing billiards. Ann. Math. (2) **147**(3), 695–708 (1998)
[10] : Burtscher, A., García-Heveling, L.: Time functions on Lorentzian length spaces. arXiv:2108.02693, (2021)
[11] Cavalletti, F., Mondino, A.: Optimal transport in Lorentzian synthetic spaces, synthetic timelike Ricci curvature lower bounds and applications. arXiv:2004.08934, (2020)
[12] Cunningham, W.J., Krioukov, D., Halverson, J., Rideout, D.: Exact geodesic distances in FLRW spacetimes. Phys. Rrev. D **96**, 103538 (2017)
[13] Grant, J.D.E., Kunzinger, M., Sämann, C.: Inextendibility of spacetimes and Lorentzian length spaces. Ann. Glob. Anal. Geom. **55**(1), 133–147 (2019)
[14] Kirchberger, M.: Lorentzian comparison geometry. Master's thesis, University of Vienna (2018). https://phaidra.univie.ac.at/open/o:1351312
[15] Kunzinger, M., Steinbauer, R.: Null distance and convergence of Lorentzian length spaces. arXiv:2106.05393, (2021)
[16] Kunzinger, M., Sämann, C.: Lorentzian length spaces. Ann. Glob. Anal. Geom. **54**(3), 399–447 (2018)
[17] McCann, R.J., Sämann, C.: A Lorentzian analog for Hausdorff dimension and measure. arXiv:2110.04386 (2021)
[18] Minguzzi, E.: Convex neighbourhoods for Lipschitz connections and sprays. Mon. Math. **177**, 569–625 (2015)

[19] O'Neill, B.: Semi-Riemannian Geometry with Applications to Relativity, volume 103 of Pure and Applied Mathematics. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York (1983)
[20] Sormani, C., Vega, C.: Null distance on a spacetime. Class. Quantum Gravity **33**(8), 085001 (2016)

OPEN ACCESS
IOP Publishing

Classical and Quantum Gravity

Class. Quantum Grav. **40** (2023) 175002 (28pp)

https://doi.org/10.1088/1361-6382/ace585

Gluing of Lorentzian length spaces and the causal ladder

Felix Rott

Faculty of Mathematics, University of Vienna, Vienna, Austria

E-mail: felix.rott@univie.ac.at

Received 2 November 2022; revised 12 June 2023
Accepted for publication 7 July 2023
Published 25 July 2023

Abstract

We investigate the compatibility of Lorentzian amalgamation with various properties of Lorentzian pre-length spaces. In particular, we give conditions under which gluing of Lorentzian length spaces yields again a Lorentzian length space and we give criteria which preserve many steps of the causal ladder. We conclude with some thoughts on the causal properties which seem not so easily transferable.

Keywords: Lorentzian length spaces, gluing constructions, quotient spaces, metric geometry, causality theory

(Some figures may appear in colour only in the online journal)

1. Introduction

Lorentzian length spaces are a comparatively new approach to a synthetic description of Lorentzian geometry, introduced in [19]. This is in spirit very similar to metric geometry and the theory of length spaces, which was key to developing a metric and synthetic point of view of Riemannian geometry, without relying on any differential structure and machinery.

With the paper birthing the theory being only a couple of years old, this theory is, by mathematical standards, still in its infancy. There is, however, rapid progress being made in various directions. The advancements in the theory of Lorentzian length spaces can roughly be sorted into two directions (which are not disjoint, of course): on the one hand, there is the translation of fundamental and almost elementary concepts originally developed for (metric) length spaces aiming to improve the robustness of Lorentzian (pre-)length spaces and bring it to a level of sophistication close to that of the metric counterpart. These range from the introduction of (hyperbolic) angles to the description of Ricci curvature bounds, Hausdorff measure and gluing. Works in this direction include:

(i) [11] introduces optimal transport methods in Lorentzian length spaces, defines timelike Ricci curvature bounds via suitable entropy conditions and gives applications to general relativity (synthetic singularity theorems).
(ii) [18] examines the null distance in Lorentzian length spaces (which was first introduced in [33] for spacetimes) and in turn studies Gromov–Hausdorff convergence, establishing first compatibility results with respect to curvature bounds.
(iii) [23] defines an analogue to Hausdorff measure and dimension on Lorentzian length spaces.
(iv) [6] introduces gluing techniques for Lorentzian pre-length spaces and gives an analogue to the gluing theorem of Reshetnyak for CAT(k) spaces.
(v) [7] introduces the concept of (hyperbolic) angles, a notion very similar to the classical Alexandrov angle in metric geometry, cf [8, definition I.1.12], as well as timelike tangent cones and an exponential map, which mimic the corresponding notions in a metric (length) space, cf [8, definitions II.3.18 and II.4.4]. Via hyperbolic angles, the authors also give a formulation of curvature bounds in terms of angle monotonicity.
(vi) [3] also developed, parallel to the above work [7], hyperbolic angles on Lorentzian length spaces. In addition to a monotonicity formulation of curvature bounds, the authors also introduce a formulation relating the size of angles and their comparison angles, similar to the classical angle condition in metric geometry, cf [9, definition 4.1.5].
(vii) [5] formulates and proves a splitting theorem for (globally hyperbolic) Lorentzian length spaces.

On the other hand, there are attempts to give synthetic versions of various ideas from classical Lorentzian geometry and causality theory. Works in this direction include:

(i) [13] develops a notion of (in)extendibility for Lorentzian length spaces.
(ii) [2] defines an analogue to warped products for Lorentzian length spaces and relates the timelike curvature of the product to the (metric) curvature of the base space.
(iii) [1] expands the causal ladder for Lorentzian length spaces which was originally introduced in [19].
(iv) [10] introduces time functions on Lorentzian length spaces and shows their existence to being equivalent to K-causality.

Simply put, one is interested in 'Lorentzifying' known synthetic concepts from metric geometry as well as synthesizing relativistic concepts from Lorentzian geometry. This article should foremost be regarded as a follow-up to [6], so at first glance appears to be placed more into the first category. However, we are mostly investigating how notions from causality theory behave with respect to gluing, so in essence it is a mix of the two.

There are two main applications which immediately come to mind when thinking about gluing constructions in the context of synthetic Lorentzian geometry. On the one hand, the extendibility of spacetimes, in particular in a low-regularity setting, has gathered increased attention, especially since [32] has shown that the Schwarzschild solution is inextendible even as a spacetime with a continuous Lorentzian metric. As mentioned above, extensions of spacetimes via Lorentzian length spaces have been studied in [13]. In this work, it is shown that a (strongly causal) Lorentzian length space satisfying a synthetic analogue to timelike geodesic completeness, the so-called(TC)-property, is inextendible as a regularly localizable Lorentzian length space, cf [13, theorem 5.3]. Thus, any extension in this setting comes with the price of exhibiting curvature singularities, i.e. curvature unbounded from above in the sense of triangle comparison. The gluing constructions introduced in [6] and further studied in this work not only provide a framework for extensions in general, but in particular might prove useful in establishing extensions which evade curvature singularities. On the other hand, gluing may also find applications in the theory of impulsive gravitational wave geometries, cf [14, chapter 20]. One way of constructing impulsive waves is via the matching of two spacetimes, originally introduced in [28] and further developed in a more systematic way in [22] with recent applications [21]. A particular relevant case of this procedure is to match two parts of Minkowski space along a null hypersurface. With some very minor modifications, this can be regarded as a special case of our construction. Moreover, matching along timelike hypersurfaces anyways directly fits into our framework. Recently, the original construction by Penrose has also been studied in other constant curvature spacetimes, see e.g. [29, 30]. Finally, matching also finds applications in alternative theories of gravity, cf [31].

We begin by recalling some basic results and definitions about Lorentzian pre-length spaces, gluing and the causal ladder. We continue with establishing under which assumptions the additional properties of a Lorentzian length space are retained. Then we discuss the inheritance of most of the steps of the causal ladder. We will usually denote by X the amalgamation of two Lorentzian pre-length spaces X_1 and X_2 and assume that the identifying map preserves the Lorentzian structure (more details on this below).

Finally, we collect the main results of this work in the following two theorems (see below for definitions of the required notions):

Theorem 1.1 (Causal inheritance). *Let X_1 and X_2 be two Lorentzian pre-length spaces and $X := X_1 \sqcup_A X_2$ their Lorentzian amalgamation. Then we have the following preservation of causality conditions.*

(i) If X_1 and X_2 are chronological, then so is X.
(ii) If X_1 and X_2 are causal, then so is X.
(iii) If X_1 and X_2 are extrinsically non-totally imprisoning, then X is non-totally imprisoning.
(iv) If X_1 and X_2 are strongly causal, then so is X.
(iv) If X_1 and X_2 are strongly causal, non-timelike locally isolating and distinguishing, then X is distinguishing.
(v) If X_1 and X_2 are causally path-connected, locally causally closed, d-compatible and globally hyperbolic with A_1 and A_2 time observing, then X is globally hyperbolic.

Theorem 1.2 (Amalgamation of length spaces). *Let X_1 and X_2 be two strongly causal and locally compact Lorentzian length spaces. Then $X := X_1 \sqcup_A X_2$ is a (strongly localizable) Lorentzian length space.*

2. Preliminaries

Here, we briefly collect some essentials regarding the general theory of Lorentzian pre-length spaces as well as Lorentzian amalgamation and the causal ladder. For more details, see [1, 6, 19]. As in [6], we set $\inf \emptyset = \infty$ and $\sup \emptyset = 0$ for convenience, as time separation and distance metric take values only in $[0, \infty]$. Moreover, if $\alpha : [a,b] \to X, \beta : [b,c] \to X$ are two curves, we will denote by $\alpha * \beta : [a,c] \to X$ their concatenation.

Definition 2.1 (Lorentzian pre-length space). A tuple $(X,d,\ll,\leqslant,\tau)$ is called a Lorentzian pre-length space if the following holds:

(i) $(X,\ll,\leqslant)$ is a causal space, i.e. $\leqslant$ is a reflexive and transitive relation on X and $\ll$ is a transitive relation contained in $\leqslant$.
(ii) $\tau : X\times X\to[0,\infty]$ is lower semi-continuous with respect to the metric d (in particular, (X,d) is a metric space).
(iii) If $x\leqslant y\leqslant z$, then $\tau(x,z)\geqslant\tau(x,y)+\tau(y,z)$, and $\tau(x,y)>0 \iff x\ll y$.

By [19, example 2.11], any smooth spacetime is a Lorentzian pre-length space (where the distance metric d is given by a (complete) Riemannian metric). In fact, any continuous and causally plain spacetime is a Lorentzian pre-length space as well, see [19, proposition 5.8].

We will usually abbreviate a Lorentzian pre-length space $(X,d,\ll,\leqslant,\tau)$ by X. By $x<y$ we mean $x\leqslant y$ and $x\neq y$ and we will use common notation for sets described via causality relations, e.g. $J^+(x) := \{y\in X \mid x\leqslant y\}$ and $I(x,z) := \{y\in X \mid x\ll y\ll z\} = I^+(x)\cap I^-(z)$.

In contrast to a metric length space, the correct notion of an 'intrinsic Lorentzian space' is a bit more complicated. To get from a Lorentzian pre-length space to a Lorentzian length space, one first needs to introduce the concept of causal and timelike curves.

Definition 2.2 (Causal/timelike curves). Let $(X,d,\ll,\leqslant,\tau)$ be a Lorentzian pre-length space.

(i) A locally Lipschitz curve $\gamma:[a,b]\to X$ is called future-directed causal (respectively timelike), if $\gamma(s)\leqslant\gamma(t)$ (respectively $\gamma(s)\ll\gamma(t)$) for all $s,t\in[a,b], s<t$. Past-directed curves are defined analogously. Unless explicitly stated otherwise, we assume all causal curves to be future-directed.
(ii) The τ-length of a causal curve γ is given as

$$L_\tau(\gamma) := \inf\left\{\sum_{i=0}^{n}\tau(\gamma(t_i),\gamma(t_{i+1})) \mid a=t_0<t_1<\ldots<t_n=b, n\in\mathbb{N}\right\}. \quad (2.1)$$

If $\gamma(a)=x,\gamma(b)=y$ and $L_\tau(\gamma)=\tau(x,y)$ we say that γ is τ-realizing and we call (the image of) such a curve a geodesic (segment).

The next ingredient is localizability, a concept which is similar to the existence of small convex neighborhoods in Lorentzian manifolds. To this end, recall that the length of a curve $\gamma:[a,b]\to X$ in a metric space (X,d) is given as

$$L^d(\gamma) := \sup\left\{\sum_{i=0}^{n-1}d(\gamma(t_i),\gamma(t_{i+1})) \mid a=t_0<t_1<\ldots<t_n=b, n\in\mathbb{N}\right\}. \quad (2.2)$$

Definition 2.3 (Localizability). Let X be a Lorentzian pre-length space. X is called localizable if every point in X has a neighborhood U such that the following holds:

(i) There is a constant $C>0$ such that $L^d(\gamma)\leqslant C$ for all causal curves γ contained in U.
(ii) There is a continuous map $\omega=\omega_U : U\times U\to[0,\infty)$ such that $(U,d|_{U\times U},\ll|_{U\times U},\leqslant|_{U\times U},\omega)$ is a Lorentzian pre-length space with the following non-triviality condition: $I^\pm(p)\cap U\neq\emptyset$ for all $p\in U$.
(iii) For all $p,q\in U$ with $p\leqslant q$ there is a causal curve γ from p to q contained in U which is maximal in U, i.e. $L_\tau(\gamma)\geqslant L_\tau(\alpha)$ for any causal curve α from p to q in U, and satisfies $L_\tau(\gamma)=\omega(p,q)$.

Note that ω has necessarily the explicit form

$$\omega(p,q) = \max\{L_\tau(\gamma) \mid \gamma \, is\, a\, causal\, curve\, from\, p\, to\, q\, contained\, in\, U\}. \tag{2.3}$$

U is called a localizable neighborhood. If every point has a neighborhood basis of localizable neighborhoods, X is called strongly localizable.

The following definition contains some elementary properties of Lorentzian pre-length spaces, some of which are required for a Lorentzian length space.

Definition 2.4 (Further properties of Lorentzian pre-length spaces). Let X be a Lorentzian pre-length space.

(i) X is called causally path-connected if $x \ll y$ implies that there exists a timelike curve from x to y and $x \leqslant y$ implies there is a causal curve from x to y.
(ii) Let $x \in X$ and let $U \subseteq X$ be a neighborhood of x. U is called causally closed if $\leqslant|_{\overline{U}\times\overline{U}}$ is closed, i.e. for all sequences $(p_n)_{n\in\mathbb{N}}, (q_n)_{n\in\mathbb{N}}$ in U with $p_n \leqslant q_n$ for all $n \in \mathbb{N}$ and $p_n \to p \in \overline{U}, q_n \to q \in \overline{U}$ we have $p \leqslant q$.
(iii) X is called locally causally closed if every point has a causally closed neighborhood.
(iv) Define the function $\mathcal{T} : X \times X \to [0,\infty]$,

$$\mathcal{T}(x,y) := \sup\{L_\tau(\gamma) \mid \gamma \, is\, a\, future-directed\, causal\, curve\, from\, x\, to\, y\}. \tag{2.4}$$

We say X is intrinsic if $\tau = \mathcal{T}$.
(v) X is called d-compatible if for all $x \in X$ there is a neighborhood $U \subseteq X$ and a positive constant C such that $L^d(\gamma) \leqslant C$ for all causal curves γ contained in U.
(vi) A subset A of X is called causally convex if $J(p,q) \subseteq A$ for all $p,q \in A$.[1]
(vii) X is called interpolative if for all $x,z \in X$ with $x \ll z$ there exists $y \in X$ such that $x \ll y \ll z$.

Definition 2.5 (Lorentzian length space). A locally causally closed, causally path-connected, localizable and intrinsic Lorentzian pre-length space is called a Lorentzian length space.

When dealing with metric spaces or even topological spaces, the notion of a subspace is of great importance and hence also desired in our setting.

Definition 2.6 (Lorentzian subspace). Let A be a subset of a Lorentzian pre-length space $(X,d,\ll,\leqslant,\tau)$. There is a natural way to view A as a subspace, namely restricting the original structure of X to A, i.e. the subspace A is the Lorentzian pre-length space $(A, d|_{A\times A}, \ll|_{A\times A}, \leqslant|_{A\times A}, \tau|_{A\times A})$.

As in the metric case, a subspace of a Lorentzian length space is in general not a Lorentzian length space anymore. The next definition introduces the key property required of subsets along which one wants to glue.

Definition 2.7 (Timelike isolation). Let X be a Lorentzian pre-length space.

(i) X is said to contain no $\ll$-isolated points if $I^\pm(x) \neq \emptyset$ for all $x \in X$.

[1] By definition, $I^\pm(x), J^\pm(x)$ are causally convex for all $x \in X$. Moreover, intersections of causally convex sets are causally convex and hence in particular causal and timelike diamonds are causally convex. Note that if X is causally path-connected, then causal convexity can be formulated as in the case of spacetimes, namely all causal curves between points in A are entirely contained in A.

(ii) A subset $A \subseteq X$ is said to be non-future locally isolating if for all $a \in A$ with $I^+(a) \neq \emptyset$ and for all neighborhoods U of a there exists $b_+ \in U \cap A$ such that $a \ll b_+$. Similarly, we define a non-past locally isolating set. We say A is non-timelike locally isolating if it satisfies both properties.

Note that X (as a subset of itself) is non-timelike locally isolating if and only if no open subset of X (viewed as a subspace) contains any $\ll$-isolated points.

Amalgamation in the metric case consists of two steps: first forming the disjoint union and then forming the quotient with respect to an equivalence relation which results from the identification of distinguished subsets. The concept of a disjoint union construction can be adapted naturally. For more details on amalgamation in the metric case see [8, 9].

Definition 2.8 (Lorentzian disjoint union). Let $(X_1, d_1, \ll_1, \leqslant_1, \tau_1)$ and $(X_2, d_2, \ll_2, \leqslant_2, \tau_2)$ be two Lorentzian pre-length spaces and set $X := X_1 \sqcup X_2$. Define $\leqslant := \leqslant_1 \sqcup \leqslant_2$, i.e. '$\leqslant \subseteq X \times X$' and $x \leqslant y :\iff \exists i \in \{1,2\} : x, y \in X_i \wedge x \leqslant_i y$. Similarly, define $\ll := \ll_1 \sqcup \ll_2$. Let d be the disjoint union metric on X, i.e.

$$d(x,y) = \begin{cases} d_i(x,y) & x, y \in X_i \\ \infty & \text{else.} \end{cases} \tag{2.5}$$

Define $\tau : X \times X \to [0, \infty]$ by

$$\tau(x,y) := \begin{cases} \tau_i(x,y) & x, y \in X_i \\ 0 & \text{else.} \end{cases} \tag{2.6}$$

We call $(X, d, \ll, \leqslant, \tau)$ the Lorentzian disjoint union of X_1 and X_2.

It is easily seen that this always results in a Lorentzian pre-length space, cf [6, proposition 3.2.2]. This is in some contrast to the following quotient construction. While the map introduced below is well-defined, it may happen that it is not lower semi-continuous. For a counterexample see [6, example 3.1.8].

Definition 2.9 (Quotient Lorentzian structure). Let X be a Lorentzian pre-length space and let $\sim$ be an equivalence relation on X. The quotient time separation $\widetilde{\tau} : X \times X \to [0, \infty]$ is defined as

$$\widetilde{\tau}([x],[y]) := \sup \left\{ \sum_{i=1}^{n} \tau(x_i, y_i) \mid x \sim x_1 \leqslant y_1 \sim x_2 \leqslant y_2 \sim \ldots \sim x_n \leqslant y_n \sim y, n \in \mathbb{N} \right\}. \tag{2.7}$$

We call a sequence $(x_1, y_1, \ldots, x_n, y_n)$ as above an n-chain from $[x]$ to $[y]$, or simply a chain. We define $[x] \,\widetilde{\ll}\, [y] :\iff \widetilde{\tau}([x],[y]) > 0$ and $[x] \,\widetilde{\leqslant}\, [y] :\iff$ there exists a chain $(x_1, y_1, \ldots, x_n, y_n)$ such that $x \sim x_1 \leqslant y_1 \sim x_2 \leqslant y_2 \sim \ldots \sim x_n \leqslant y_n \sim y$. In words, we have $[x] \,\widetilde{\leqslant}\, [y]$ if and only if there exists a causal chain from x to y and $[x] \,\widetilde{\ll}\, [y]$ if and only if there exists a causal chain of positive τ-length between x and y. Moreover, recall the definition of the quotient semi-metric,

$$\widetilde{d}([x],[y]) := \inf \left\{ \sum_{i=1}^{n} d(x_i, y_i) \mid x \sim x_1, x_{i+1} \sim y_i, y_n \sim y, n \in \mathbb{N} \right\}, \tag{2.8}$$

cf [9, definition 3.1.12].

Under the assumption that the identified subsets are closed and non-timelike locally isolating, the Lorentzian amalgamation is always a Lorentzian pre-length space. Morally, we need the glued subsets to be (topologically) closed so that the background quotient semi-metric is definite, and the non-timelike local isolation, which can be though of as sime kind of openness in timelike direction, to ensure that the lower semi-continuity of $\widetilde{\tau}$ is guaranteed.

Definition 2.10 (Lorentzian amalgamation). Let $(X_1, d_1, \ll_1, \leqslant_1, \tau_1)$ and $(X_2, d_2, \ll_2, \leqslant_2, \tau_2)$ be two Lorentzian pre-length spaces. Let A_1 and A_2 be closed and non-timelike locally isolating subspaces of X_1 and X_2, respectively. Let $f : A_1 \to A_2$ be a locally bi-Lipschitz homeomorphism[2] and suppose that the causality of A_1 and A_2 are compatible in the following sense: for all $a \in A_1$ we have $I_1^{\pm}(a) \neq \emptyset \iff I_2^{\pm}(f(a)) \neq \emptyset$. Let $(X_1 \sqcup X_2, d, \ll, \leqslant, \tau)$ be the Lorentzian disjoint union of X_1 and X_2 and consider the equivalence relation $\sim$ on $X_1 \sqcup X_2$ generated by $a \sim f(a)$ for all $a \in A_1$. Then $((X_1 \sqcup X_2)/\sim, \widetilde{d}, \widetilde{\ll}, \widetilde{\leqslant}, \widetilde{\tau})$ is called the Lorentzian amalgamation of X_1 and X_2 and is denoted[3] by $X_1 \sqcup_A X_2$.

Note that Lorentzian amalgamation can be defined with surprisingly little compatibility of the structure in the two different subsets. However, without such assumptions this construction is usually poorly behaved, see [6, example 3.3.4] for an extreme counterexample. In particular, there is no hope for any compatibility results. Thus, unless explicitly stated otherwise, we will always assume that f is τ-preserving and $\leqslant$-preserving, i.e. $\tau(a,b) = \tau(f(a), f(b))$ and $a \leqslant b \iff f(a) \leqslant f(b)$ for all $a, b \in A_1$. In similar fashion, when dealing with Lorentzian amalgamation, X will always denote the amalgamation of two Lorentzian pre-length spaces X_1 and X_2 along the subsets A_1 and A_2, and we will identify X_i with its image $\pi(X_i)$. For an easy to visualize example of this quite complicated construction, one might best think of gluing two copies of the Minkowski plane along the t-axis.

We also want to explain and justify the notation we will be using from now on. Dealing with causality in three different spaces at once can be overwhelming. Therefore, we decided to highlight in the notation of causal and timelike pasts and futures in which space the respective set is considered. Moreover, we prefer to not use the identifying map f when denoting points in A_2. Rather, we feel it is advantageous to highlight with an index from which space a point comes originally, in addition to using bracket-notation for equivalence classes for points in X. For example, if $[x] \in X_1 \setminus A_1$, then $[x] = \{x^1\}$. If $[a] \in A$, then we write $[a] = \{a^1, a^2\}$ instead of $\{a, f(a)\}$. Further, the (timelike) future of $x^1 \in X_1$ is denoted by $I_1^+(x^1)$, while the future of its equivalence class in X is denoted by $I_X^+([x])$. Similar notation will be used for sets related to the causal relation, e.g. $J_2(x^2, y^2)$ and $J_X([x],[y])$ for causal diamonds in X_2 and X, respectively. Finally, the usage of A is technically a bit misleading since this set is not yet introduced. On the one hand, we feel it is quite clear what is meant by A, namely the shared set in the glued space. On the other hand, one can set $A := \pi(A_1 \sqcup A_2)$ for a rigorous definition. Compared to [6], one tiny detail in notation will appear different to an attentive reader, however: it is essentially never necessary to mention the original relations $\ll_1, \leqslant_1, \ll_2, \leqslant_2$ or the original time separation functions τ_1, τ_2. Rather, it is enough to deal with the corresponding notions in the disjoint union. That is, in the language of definition 2.10 we have $x^1 \leqslant y^1$ if and only if $x^1 \leqslant_1 y^1$, and thus omitting the index in $\leqslant$ leads to easier readability.

For a more detailed investigation of Lorentzian amalgamation and its elementary properties, we refer to [6]. We will nevertheless collect some of the most useful facts for the current work here. We emphasize one final time that we only consider gluing constructions where the identifying map is well-behaved (otherwise, the following properties are of course not valid).

[2] This means that every $a \in A_1$ has a neighborhood $U \subseteq A_1$ such that $f|_U : U \to f(U)$ and its inverse are Lipschitz.

[3] The notation $X_1 \sqcup_A X_2$ is debatable as arguably the function f also plays a defining role in how the gluing behaves. The main reason for our convention is to be in alignment with the major sources from metric geometry. The use of A is justified in the paragraph below.

Proposition 2.11 (Useful facts about Lorentzian amalgamation). *Let $X := X_1 \sqcup_A X_2$ be the Lorentzian amalgamation of two Lorentzian pre-length spaces X_1 and X_2.*

(i) $\widetilde{\tau}$ has the following simplified form:

$$\widetilde{\tau}([x],[y]) = \begin{cases} \tau(x^i, y^i) & x^i, y^i \in X_i, i \in \{1,2\}, \\ \sup\limits_{[a] \in J_X([x],[y]) \cap A} \{\tau(x^i, a^i) + \tau(a^j, y^j)\} & x^i \in X_i, y^j \in X_j, \{i,j\} = \{1,2\}. \end{cases} \tag{2.9}$$

Moreover, if $\widetilde{\tau}([x],[y]) > 0$, the supremum can be approximated by only considering timelike chains.

(ii) $\widetilde{\tau}([x],[y]) \geqslant \tau(x^i, y^j), \widetilde{d}([x],[y]) \leqslant d(x^i, y^j), x^i \ll y^j \Rightarrow [x] \mathrel{\widetilde{\ll}} [y]$ and $x^i \leqslant y^j \Rightarrow [x] \mathrel{\widetilde{\leqslant}} [y], i, j \in \{1,2\}$.

Proof. See [6, remark 3.1.3, proposition 3.3.7, lemma 3.3.3, corollary 3.1.7]. □

Remark 2.12 (Convergence in amalgamation). Since $\widetilde{d} \leqslant d$, for any sequence $p_n^i \to p^i$ it follows that $[p_n] \to [p]$. Conversely, if $[p_n] \to [p]$ and $[p] \in X_i$, then for all subsequences in X_i we have $p_{n_k}^i \to p^i$. In particular, there always exists a converging subsequence in at least one of the original spaces, cf [6, remark 3.3.5]. Indeed, if $[p] \in X_i \setminus A_i$, i.e. $[p] = \{p^i\}$, then $[p_n] = \{p_n^i\}$ for large enough n as well. Let U be an open neighborhood of p^i in X_i that does not meet A_i (in fact, p^i has a neighborhood basis of such sets). Then $\pi(U)$ is a neighborhood of $[p]$ since $\pi^{-1}(\pi(U)) = U$ is open, and as such contains all but finitely many sequence members $[p_n]$. Thus, all but finitely many members p_n^i are contained in U and so $p_n^i \to p^i$ follows. Let $[p] \in A$. For any $[p_n] \in X$, there exists $i \in \{1,2\}$ such that $p_n^i \in X_i$, thus there exists a subsequence in at least one of the original spaces. Say $(p_{n_k}^1)_{k \in \mathbb{N}}$ is a subsequence in X_1. Let U_1 be an open neighborhood of p^1 in X_1. Then we find an open neighborhood U_2 of p^2 in X_2 such that $f(U_1 \cap A_1) = U_2 \cap A_2$. Then $U := \pi(U_1 \sqcup U_2)$ is an open neighborhood of $[p]$ in X and as such $[p_{n_k}] \in U$, i.e. $p_{n_k}^1 \in U_1$ for almost all k. Hence $p_{n_k}^1 \to p^1$.

This remark remains valid for gluing finitely many spaces. Moreover, all statements in this work concerning gluing easily generalize to finitely many spaces, it is just more convenient to formulate everything with only two spaces at hand.

At last, we will briefly present the causal ladder for Lorentzian pre-length spaces. As noted in the introduction, this was first introduced in [19] and further developed in [1]. Both works formulated the steps for Lorentzian length spaces, but in the spirit of generality, we try to do everything in the setting of Lorentzian pre-length spaces and rather specify the precise additional assumptions that are needed.

On a somewhat related note, there is another notion of local causal closure introduced in [1, definition 2.19] and also used in [10] which is more general than definition 2.4(iii) and appears more advantageous in intrinsic and/or causally path-connected settings. We stick to the original definition for two reasons: on the one hand, we prefer to not use intrinsic notions explicitly whenever possible, and on the other hand we use local causal closedness anyways only when we also assume strong causality, which is why we do not really lose any generality.

Definition 2.13 (Preparing definitions for causal ladder). Let X be a Lorentzian pre-length space.

(i) X is future-distinguishing if $I^+(x) = I^+(y) \Rightarrow x = y$ for all $x, y \in X$ and past-distinguishing if $I^-(x) = I^-(y) \Rightarrow x = y$ for all $x, y \in X$. X is distinguishing if it is both future-distinguishing and past-distinguishing.
(ii) X is future-reflective if $I^-(x) \subseteq I^-(y) \Rightarrow I^+(y) \subseteq I^+(x)$ for all $x, y \in X$ and past-reflective if $I^+(x) \subseteq I^+(y) \Rightarrow I^-(y) \subseteq I^-(x)$ for all $x, y \in X$. X is reflective if it is both future-reflective and past-reflective.
(iii) The relation K is defined as the smallest transitive and closed relation that contains $\leqslant$.

Definition 2.14 (The causal ladder). Let X be a Lorentzian pre-length space.

(i) X is chronological if $\ll$ is irreflexive, i.e. $x \not\ll x$ for all $x \in X$.
(ii) X is causal if $\leqslant$ is antisymmetric, i.e. $x \leqslant y$ and $y \leqslant x$ imply $x = y$.
(iii) X is non-totally imprisoning if for every compact set $B \subseteq X$ there exists a constant $C > 0$ such that $L^d(\gamma) \leqslant C$ for all causal curves γ contained in B.
(iv) X is strongly causal if $\mathcal{I} := \{I(x, y) \mid x, y \in X\}$ forms a sub-basis for the (metric) topology $\mathcal{D}$ on X.
(v) X is K-causal[4] if the relation K is antisymmetric.
(vi) X is causally continuous if it is distinguishing and reflective.
(vii) X is causally simple if it is causal and $J^\pm(x)$ is closed for all $x \in X$.
(viii) X is globally hyperbolic if it is non-totally imprisoning and $J(x, y)$ is compact for all $x, y \in X$.

Theorem 2.15 (Implications of the causal ladder). *Let X be a Lorentzian length space. With the additional assumption of local compactness in $(v) \Rightarrow (iv)$, each step of the causal ladder implies the previous one.*

Proof. These implications have been established in [19, theorem 3.26] and [1, theorem 3.16]. □

3. Basics and preparations

Proposition 3.1 (Past and future representation). *Let X be a Lorentzian pre-length space.*

(i) $J^\pm(x) = \cup_{y \in J^\pm(x)} J^\pm(y)$ *for all* $x \in X$.
(ii) If X is interpolative, then $I^\pm(x) = \cup_{y \in I^\pm(x)} I^\pm(y)$ *for all* $x \in X$.

Proof.

(i) '$\subseteq$' holds anyways since $x \in J^\pm(x)$. Conversely, if $z \in J^\pm(y)$ for some $y \in J^\pm(x)$, then $z \in J^\pm(x)$ by the transitivity of $\leqslant$.
(ii) '$\supseteq$' follows as above via the transitivity of $\ll$. Conversely, if, say, $z \in I^+(x)$, then since X is interpolative we find $y \in X$ such that $x \ll y \ll z$, hence $z \in I^+(y)$ and $y \in I^+(x)$. □

Example 3.2 (Being interpolative is a necessary condition). Without requiring that X is interpolative, the statement in Proposition 3.1(ii) is wrong in general. Consider, e.g. Minkowski space with an open rectangle removed, depicted as in figure 1.

[4] In [1], this is called stable causality, which is equivalent in the smooth case.

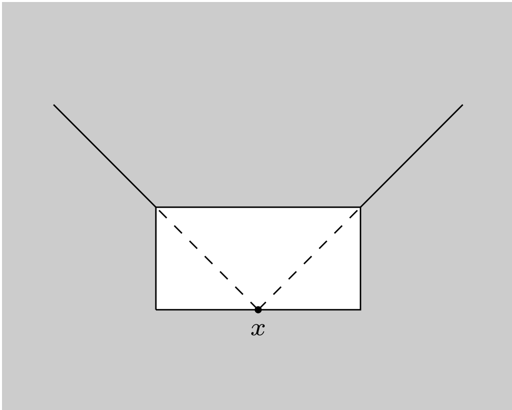

Figure 1. In this case, $I^+(x) = \cup_{y\in I^+(x)} I^+(y)$ does not hold.

Then for a point x on the lower boundary of the rectangle equality does not hold in the future case, since points on the upper boundary are not captured in any future of points in the future of x.

Lemma 3.3 (Quotient relation reformulation). *In an amalgamation $X := X_1 \sqcup_A X_2$, we have $[x] \mathrel{\tilde{\leqslant}} [y]$ if and only if $x^i \leqslant y^i$ for $i \in \{1,2\}$ or there exists $[a] \in A$ such that $x^i \leqslant a^i \sim a^j \leqslant y^j$ for $\{i,j\} = \{1,2\}$. Similarly for $\mathrel{\tilde{\ll}}$.*

Proof. '$\Leftarrow$' of the statement follows immediately by proposition 2.11(ii). Conversely, $[x] \mathrel{\tilde{\leqslant}} [y]$ states by definition that we find a causal chain of the form[5] $x^i \leqslant a^i_1 \sim a^j_1 \leqslant a^j_2 \sim a^i_2 \leqslant \ldots \sim a^j_n \leqslant y^j$ (whether the chain starts and ends in the same space does not matter for the proof). By the causal preservation of the identifying map $f : A_1 \to A_2$, we have $a^1_k \leqslant a^1_{k+1} \iff a^2_k \leqslant a^2_{k+1}$. Thus, one can 'shorten' $a^i_k \leqslant a^i_{k+1} \sim a^j_{k+1} \leqslant a^j_{k+2}$ to $a^i_k \leqslant a^i_{k+2}$. If the two endpoints in the chain are from the same space, then the chain reduces to a relation in the original space. If the two endpoints are from different spaces, then one single point in A is needed to cross to the other side.

With proposition 2.11(i), the timelike case follows in complete analogy. □

Proposition 3.4 (Glued past and future representation). *Let $X := X_1 \sqcup_A X_2$ be the amalgamation of two Lorentzian pre-length spaces X_1 and X_2. Then we have the following decomposition of (pasts and) futures in X:*

(i) If $[x] \in A$, then $I^{\pm}_X([x]) = \pi(I^{\pm}_1(x^1) \sqcup I^{\pm}_2(x^2))$.
(ii) If, say, $[x] = \{x^1\}$, then $I^+_X([x]) = \pi(I^+_1(x^1) \sqcup (\cup_{x^1 \ll a^1} I^+_2(a^2)))$,[6] and similarly for the past case.

The same holds for causal pasts and futures.

[5] This follows by [6, remark 3.1.7]: points outside of A form their own equivalence class, so these 'trivial' equivalences can be immediately cut out via the transitivity of the relations.

[6] By $\cup_{x^1 \ll a^1} I^+_2(a^2)$ we mean the union of $I^+_2(a^2)$ over all $a^1 \in I^+_1(x^1) \cap A_1$. This should be clear as the letter a is usually reserved for points from the identified sets and the expression anyways does not make sense for points outside of A.

Proof. (i) This is essentially one half of [6, lemma 3.3.8]: '$\supseteq$' follows immediately by proposition 2.11(ii). Conversely, suppose $[y] \in I^+_X([x])$. Since $[x] = \{x^1, x^2\}$ we can find a chain as in lemma 3.3 that starts in x^i and ends in y^i, which then reduces to an original relation.

(ii) Here, '$\supseteq$' follows immediately as well. Let $[y] \in I^+_X([x])$. Then by lemma 3.3, either $x^1 \ll y^1$ or $x^1 \ll a^1 \sim a^2 \ll y^2$. These conditions precisely describe the right hand side of the equation.

□

At present, it seems that strong causality is the 'correct threshold' at which Lorentzian pre-length spaces behave reasonably. This is because at the stage of strong causality, the topology and causality genuinely interact with each other in a desirable way. In some cases, it is very convenient that $\mathcal{I}$ is not just a sub-basis but a basis for the topology. As it turns out, the deciding property is that of non-timelike local isolation. This result was first covered in [4, theorem 1.6.33].

Lemma 3.5 (Basis for Alexandrov topology). *Let X be a strongly causal and non-timelike locally isolating Lorentzian pre-length space. Then $\mathcal{I}$ forms a basis for the topology.*

Proof. Let U be a $\mathcal{D}$-neighborhood of a point $p \in X$, where $\mathcal{D}$ denotes the metric topology on X induced by d. Then by strong causality of X we find points $x_1, y_1, \ldots, x_n, y_n$ such that $p \in \cap_{i=1}^n I(x_i, y_i) \subseteq U$. Clearly, this intersection is an open causally convex neighborhood of p. Hence, by the non-timelike local isolation of X we find points $q_-, q_+ \in \cap_{i=1}^n I(x_i, y_i)$ such that $q_- \ll p \ll q_+$. Then $p \in I(q_-, q_+) \subseteq J(q_-, q_+) \subseteq \cap_{i=1}^n I(x_i, y_i) \subseteq U$, showing that any neighborhood of any point contains a timelike diamond which contains that point. Indeed, it is even possible to choose the 'endpoints' of the diamond to lie in the original neighborhood as well. □

Corollary 3.6 (Strongly causal Lorentzian length spaces). *Let X be a strongly causal Lorentzian length space. Then $\mathcal{I}$ forms a basis for the topology.*

Proof. This follows immediately by the definition of localizing neighborhoods. Indeed, let $p \in V := \cap_{i=1}^n I(x_i, y_i) \subseteq U$ be as above, where U is a localizable neighborhood of p. Since $I^\pm(p) \cap U \neq \emptyset$ and X is causally path-connected, we find a timelike curve $\gamma : [a,b] \to X$ through p, which by continuity is contained for some time in V, i.e. $p \in \gamma((-\varepsilon, \varepsilon)) \subseteq V$. Since V is causally convex, any timelike diamond formed by points on $\gamma((-\varepsilon, \varepsilon))$ is contained in V and hence in U. □

In particular, causal path-connectedness and the absence of $\ll$-isolated points imply that a Lorentzian pre-length space X is non-timelike locally isolating. The following three statements are slight generalizations of some of the points in [19, theorem 3.26].

Lemma 3.7 (Reformulating strong causality). *In a causally path connected Lorentzian pre-length X space with no $\ll$-isolated points, strong causality is equivalent to the non-existence of almost closed causal curves.*

Proof. Strong causality implying the formulation with curves is already shown in [19, lemma 2.38]. Conversely, let $x \in X$ and let U be a neighborhood of x. Then we find a neighborhood V of x contained in U such that $J(p,q) \subseteq U$ for all $p, q \in V$. Since $I^\pm(x) \neq \emptyset$, there is a timelike curve $\gamma : [a,b] \to X$ through x. In particular, setting $\gamma(0) =: x$, there is $\varepsilon > 0$ such that $\gamma([-\varepsilon, \varepsilon]) \subseteq V$. Then by assumption $J(p,q) \subseteq U$ for $p, q \in \gamma([-\varepsilon, \varepsilon])$. Since $I(p,q) \subseteq J(p,q)$, we have found an element of the sub-basis of the Alexandrov topology inside an arbitrary neighborhood of x. □

Lemma 3.8 (Global hyperbolicity implies strong causality). *A causally path-connected, locally causally closed and globally hyperbolic Lorentzian pre-length space X with no $\ll$-isolated points is strongly causal. In particular, such a space is locally compact.*

Proof. If X were not strongly causal, by lemma 3.7 there exists $x\in X$ and a neighborhood U of x such that for all neighborhoods $V\subseteq U$ of x there are points $p,q\in V$ and a causal curve between them leaving U. Since $I^{\pm}(x)\neq\emptyset$ and X is causally path-connected, there is a timelike curve $\gamma:[a,b]\to X$ through x. By choosing points on this curve close enough to x, we find $p,q\in\gamma([a,b])\cap U, p\ll x\ll q$, where U is the supposed neighborhood where strong causality fails. In particular, $x\in I(p,q)$. Since this set is open, there exists $r>0$ such that $B_r(x)\subseteq I(p,q)\subseteq J(p,q)$, where $J(p,q)$ is compact by assumption. Since X is supposed to be not strongly causal, for each $r'<r$ there is a causal curve with endpoints in $B_{r'}(x)$ that leaves U. Consider a sequence $(\gamma_n)_{n\in\mathbb{N}}$ of such curves with endpoints $p_n,q_n\in B_{\frac{1}{n}}(x)$ for large enough n. The image of each γ_n is contained in the compact and causally convex set $J(p,q)$, and since X is non-totally imprisoning, there exists a constant $C>0$ such that $L_d(\gamma_n)\leqslant C$ for all n. Parameterized by d-arclength, each curve γ_n has $[0,l_n]$ as domain, with $l_n\leqslant C$ and Lipschitz constant 1. Denote by $\tilde{\gamma}:[0,C]\to X$ the reparametrization of γ_n given by $\tilde{\gamma}_n(t):=\gamma_n(\frac{l_n}{C}t)$. Then

$$d(\tilde{\gamma}_n(s),\tilde{\gamma}_n(t))=d\left(\gamma_n\left(\frac{l_n}{C}s\right),\gamma_n\left(\frac{l_n}{C}t\right)\right)=\frac{l_n}{C}|t-s|\leqslant|t-s|. \tag{3.1}$$

In summary, $\tilde{\gamma}_n:[0,C]\to X$ and $\mathrm{Lip}(\tilde{\gamma}_n)\leqslant 1$ for all n. Thus, all requirements of the limit curve theorem [19, theorem 3.7] are satisfied. Since $p_n,q_n\to x$, we infer the existence of a causal loop through x. Thus, X is not causal, a contradiction to X being non-totally imprisoning, cf [19, theorem 3.26(ii)] (the cited theorem is of course valid in a causally path-connected Lorentzian pre-length space).

Concerning local compactness, let $p\in X$ and let $U\subseteq X$ be a neighborhood of p. By strong causality we find points $x_1,y_1,\ldots,x_n,y_n$ such that $p\in\cap_{i=1}^n I(x_i,y_i)\subseteq U$. Thus, $\cap_{i=1}^n J(x_i,y_i)$ is a compact neighborhood of p as each $J(x_i,y_i)$ is compact. □

Lemma 3.9 (Strong causality implies non-total imprisonment). *A strongly causal, locally causally closed and d-compatible Lorentzian pre-length space is non-totally imprisoning.*

Proof. This is just extracting the right properties of a Lorentzian length space needed in [19, lemmas 3.12, 3.15 and theorem 3.26]. □

Corollary 3.10 (Conditions for global causal closedness). *A globally hyperbolic, causally path-connected and locally causally closed Lorentzian pre-length space X with no $\ll$-isolated points is globally causally closed, i.e. for any $p_n\to p, q_n\to q$ with $p_n\leqslant q_n$ for all n we have $p\leqslant q$.*

Proof. This is a consequence of the Limit curve theorem [19, Theorem 3.7]: let $p_n\to p, q_n\to q$ with $p_n\leqslant q_n$ for all n. If $p=q$ we are done so assume $p\neq q$. By lemma 3.8, X is strongly causal, so there exist timelike diamonds $I(x_1,y_1)$ and $I(x_2,y_2)$ containing p and q as well as infinitely many sequence members p_n and q_n, respectively. Then $p_n,q_n\in I(x_1,y_2)\subseteq J(x_1,y_2)$ by push-up, cf [19, lemma 2.10]. Moreover, $J(x_1,y_2)$ is compact by assumption and hence closed, so $p,q\in J(x_1,y_2)$ as well. By the causal path-connectedness we infer the existence of causal curves γ_n connecting p_n and q_n, which are all contained in $J(x_1,y_2)$. By the non-total imprisonment of X, we find $C>0$ such that $L^d(\gamma_n)\leqslant C$ for all n. With a reparametrization argument as in lemma 3.8, we can apply the limit curve theorem to obtain a causal curve from p to q, so in particular $p\leqslant q$. □

Lemma 3.11 (Recovering subsets). *Let X_1 and X_2 be topological spaces, $A_i \subseteq X_i$, and $f : A_1 \to A_2$ bijective. Consider the quotient space of $X_1 \sqcup X_2$ generated by the equivalence relation $a \sim f(a)$ for all $a \in A_1$. Let $Y_i \subseteq X_i$ be subsets such that $\pi(Y_1 \cap A_1) = \pi(Y_2 \cap A_2)$. Then $\pi^{-1}(\pi(Y_1 \sqcup Y_2)) = Y_1 \sqcup Y_2$.*

Proof. We first show that $\pi^{-1}(\pi((Y_1 \cap A_1) \sqcup (Y_2 \cap A_2))) = (Y_1 \cap A_2) \sqcup (Y_2 \cap A_2)$. The inclusion '$\supseteq$' is always true. So let $x \in \pi^{-1}(\pi((Y_1 \cap A_1) \sqcup (Y_2 \cap A_2)))$, then $\pi(x) = [x] \in \pi((Y_1 \cap A_1) \sqcup (Y_2 \cap A_2))$. Clearly, $[x] = \{x^1, x^2\}$, so there exists $i \in \{1, 2\}$ such that $x^i \in Y_i \cap A_i$, say without loss of generality $i = 1$. Since $\pi(Y_1 \cap A_1) = \pi(Y_2 \cap A_2)$ by assumption, it then follows that $[x] \in \pi(Y_2 \cap A_2)$ as well. So we have $x^1, x^2 \in Y_1 \cap A_2 \sqcup Y_2 \cap A_2$.

The full statement easily follows from the following calculation, since, roughly speaking, for points outside of A there is a unique pre-image:

$$\begin{aligned}
\pi^{-1}(\pi(Y_1 \sqcup Y_2)) &= \pi^{-1}(\pi((Y_1 \cap A_1) \cup (Y_1 \setminus A_1) \sqcup (Y_2 \cap A_2) \cup (Y_2 \setminus A_2))) \\
&= \pi^{-1}(\pi((Y_1 \cap A_1 \sqcup Y_2 \cap A_2) \cup (Y_1 \setminus A_1 \sqcup Y_2 \setminus A_2))) \\
&= \pi^{-1}(\pi(Y_1 \cap A_1 \sqcup Y_2 \cap A_2) \cup \pi(Y_1 \setminus A_1 \sqcup Y_2 \setminus A_2)) \\
&= \pi^{-1}(\pi(Y_1 \cap A_1 \sqcup Y_2 \cap A_2)) \cup \pi^{-1}(\pi(Y_1 \setminus A_1 \sqcup Y_2 \setminus A_2)) \\
&= (Y_1 \cap A_1 \sqcup Y_2 \cap A_2) \cup (Y_1 \setminus A_1 \sqcup Y_2 \setminus A_2) = Y_1 \sqcup Y_2.
\end{aligned}$$

□

4. Gluing Lorentzian length spaces

Next, we investigate the question of whether the amalgamation of Lorentzian length spaces is again a Lorentzian length space. We cover each defining property separately and obtain the final result in the end.

Example 4.1 (Being intrinsic vs. causal path-connectedness). Note that an intrinsic Lorentzian pre-length space is not necessarily causally path-connected, as there might be points which are null related but do not have a null curve joining them. Indeed, consider the subspace (in the sense of definition 2.6) of the Minkowski plane $J_+((0,0)) \setminus \{(x,y) \mid 1 < x = y < 2\}$, i.e. a causal future with a segment of a light ray removed. Then $(1,1) \leqslant (2,2), \tau((1,1),(2,2)) = 0$, but there is no causal curve connecting the two points, see figure 2 on the left. Moreover, the two properties are independent, as the open subset of the Minkowski plane viewed as a subspace depicted on the right in figure 2 is causally path-connected but not intrinsic.

Proposition 4.2 (Causal path-connectedness). *Let X_1 and X_2 be two causally path-connected Lorentzian pre-length spaces. Then $X := X_1 \sqcup_A X_2$ is causally path-connected.*

Proof. Let $[x], [y] \in X$ with $[x] \mathrel{\tilde{\leqslant}} [y]$. If $[x], [y] \in X_i$, then $x^i \leqslant y^i$ and the projection of a causal curve in X_i joining x^i and y^i is a causal curve in X joining $[x]$ and $[y]$. Thus, the only case left to check is, up to symmetry, $[x] = \{x^1\}$ and $[y] = \{y^2\}$. In this case, $[x] \mathrel{\tilde{\leqslant}} [y]$ implies that there exists $[a] \in A$ such that $x^1 \leqslant a^1 \sim a^2 \leqslant y^2$, cf lemma 3.3. Since X_1 and X_2 are causally path-connected, there exist causal curves from x^1 to a^1 in X_1 and from a^2 to y^2 in X_2, respectively. The concatenation of their projections is a causal curve from $[x]$ to $[y]$ in X.

Using proposition 2.11(i), the timelike case is entirely analogous. □

Before dealing with the property of being intrinsic, we state the following lemma.

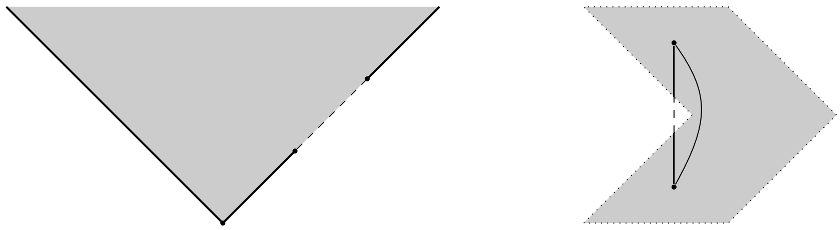

Figure 2. The space on the left is intrinsic but not causally path-connected. The space on the right is causally path-connected but not intrinsic.

Lemma 4.3 (Quotient length of causal curves). *Let $X := X_1 \sqcup_A X_2$ be the Lorentzian amalgamation of two Lorentzian pre-length spaces X_1 and X_2.*

(i) Let $\gamma : [a,b] \to X_1$ be a causal curve in X_1. Assume that f is not necessarily τ-preserving. Then $L_{\widetilde{\tau}}(\pi \circ \gamma) \geqslant L_\tau(\gamma)$. If f is τ-preserving, then $L_{\widetilde{\tau}}(\pi \circ \gamma) = L_\tau(\gamma)$. The same is true for causal curves in X_2.
(ii) Let $\gamma : [a,c] \to X$ be a causal curve resulting from the concatenation of projections of two causal curves $\alpha : [a,b] \to X_1$ and $\beta : [b,c] \to X_2$, i.e. $\gamma = (\pi \circ \alpha) * (\pi \circ \beta)$. Then $L_{\widetilde{\tau}}(\gamma) = L_\tau(\alpha) + L_\tau(\beta)$.

Proof.

(i) This follows immediately by proposition 2.11(ii). Let $a = t_0 < t_1 < \ldots < t_n = b$ be a partition of $[a,b]$. Then $\sum_{i=1}^{n-1} \tau(\gamma(t_i)^1, \gamma(t_{i+1})^1) \leqslant \sum_{i=1}^{n-1} \widetilde{\tau}([\gamma(t_i)], [\gamma(t_{i+1})])$. Since the length of a causal curve is defined as the supremum over such expressions, the claim follows. The second claim in (i) follows with the same arguments and proposition 2.11(i), since in this case we have equality for each partition.
(ii) Note that we return here to the familiar setting of a structure preserving identification map f. This claim would be wrong in general without this assumption. Via (i) and [19, lemma 2.25] we compute $L_{\widetilde{\tau}}(\gamma) = L_{\widetilde{\tau}}(\gamma|_{[a,b]}) + L_{\widetilde{\tau}}(\gamma|_{[b,c]}) = L_{\widetilde{\tau}}(\pi \circ \alpha) + L_{\widetilde{\tau}}(\pi \circ \beta) = L_\tau(\alpha) + L_\tau(\beta)$. □

Proposition 4.4 (Intrinsicality). *Let X_1 and X_2 be two intrinsic Lorentzian pre-length spaces. Then $X := X_1 \sqcup_A X_2$ is intrinsic.*

Proof. Let $[x],[y] \in X$. If $\widetilde{\tau}([x],[y]) = 0$, there is nothing to show, so suppose $\widetilde{\tau}([x],[y]) > 0$. If $[x],[y] \in X_i$, then $\widetilde{\tau}([x],[y]) = \tau(x^i, y^i)$. As X_i is intrinsic, we find a causal curve γ from x^i to y^i such that $L_\tau(\gamma) > \tau(x^i,y^i) - \varepsilon$. Then $L_{\widetilde{\tau}}(\pi \circ \gamma) = L_\tau(\gamma) > \tau(x^i,y^i) - \varepsilon = \widetilde{\tau}([x],[y]) - \varepsilon$. In other words, if the two points are from the same space, the causality is unaffected by the gluing process and hence a well-approximating causal curve in the original space projects to a well-approximating causal curve in the quotient.

Again, the only case left to check is, up to symmetry, $[x] = \{x^1\}$ and $[y] = \{y^2\}$. By proposition 2.11(i), we have $\widetilde{\tau}([x],[y]) = \sup\{\tau(x^1,a^1) + \tau(a^2,y^2) \mid [a] \in J_X([x],[y]) \cap A\}$. Let $[a]$ be such that $\tau(x^1,a^1) + \tau(a^2,y^2) > \widetilde{\tau}([x],[y]) - \varepsilon$. As X_1 and X_2 are intrinsic (and in fact, $X_1 \sqcup X_2$ is intrinsic as well), we find causal curves γ_1 from x^1 to a^1 in X_1 and γ_2 from a^2 to y^2 in X_2 such that, respectively, $L_\tau(\gamma_1) > \tau(x^1,a^1) - \varepsilon$ and $L_\tau(\gamma_2) > \tau(a^2,y^2) - \varepsilon$. Denote by γ the concatenation of the projections of these curves. Then by lemma 4.3(ii), $L_{\widetilde{\tau}}(\gamma) = L_\tau(\gamma_1) + L_\tau(\gamma_2) > \tau(x^1,a^1) - \varepsilon + \tau(x^2,a^2) - \varepsilon > \widetilde{\tau}([x],[y]) - 3\varepsilon$. □

Proposition 4.5 (Local causal closedness). *Let X_1 and X_2 be two locally causally closed, strongly causal and locally compact Lorentzian pre-length spaces. Then $X := X_1 \sqcup_A X_2$ is locally causally closed.*

Proof. At first, note that any subset of a causally closed neighborhood is again causally closed. Let $[x] \in X \setminus A$. Then there is an open neighborhood $U_i \subseteq X_i$ of x^i which does not meet A_i, and we can choose U_i small enough to be contained in a causally closed neighborhood of x^i in X_i. Then $\pi(U_i)$ is a causally closed neighborhood of $[x]$ in X. Indeed, $\pi(U_i)$ is open as $\pi^{-1}(\pi(U_i)) = U_i$ is open. Then for any $[p] \mathrel{\widetilde{\leqslant}} [q]$ in $\pi(U_i)$ it follows that $p^i \leqslant q^i$.

Let $[x] \in A$. Let V_1 and V_2 be open and causally closed neighborhoods of x^1 and x^2 in X_1 and X_2, respectively. Let V_1' be such that $V_1' \cap A_1 = f^{-1}(V_2 \cap A_2)$ and let V_2' be such that $V_2' \cap A_2 = f(V_1 \cap A_1)$. Then $x^1 \in U_1 := V_1 \cap V_1'$ and $x^2 \in U_2 := V_2 \cap V_2'$ are open as well as causally closed and satisfy $f(U_1 \cap A_1) = U_2 \cap A_2$[7]. It follows that $U := \pi(U_1 \sqcup U_2)$ is an open neighborhood of $[x]$ in X, cf lemma 3.11. The next step is very similar to the construction in lemma 3.5: by strong causality we find points $p_1^1, q_1^1, \ldots, p_n^1, q_n^1$ in X_1 and $y_1^2, z_1^2, \ldots, y_m^2, z_m^2$ in X_2 such that $x^1 \in S_1 := \cap_{i=1}^n I_1(p_i^1, q_i^1) \subseteq U_1$ and $x^2 \in S_2 := \cap_{j=1}^m I_2(y_j^2, z_j^2) \subseteq U_2$. Then we find a neighborhood O_i of x^i which is contained in S_i such that $f(O_1 \cap A_1) = O_2 \cap A_2$. As A_1 and A_2 are non-timelike locally isolating, we find $b_-^i, b_+^i \in O_i \cap A_i$ such that $b_-^i \ll x^i \ll b_+^i, i = 1, 2$. In particular, $x^i \in I_i(b_-^i, b_+^i) \subseteq J_i(b_-^i, b_+^i) \subseteq S_i \subseteq U_i$ since S_i is causally convex. Additionally, by the local compactness of X_i, we can assume that both these timelike diamonds are contained in a compact neighborhood as well. We claim that (the closure of) $D := I_X([b_-], [b_+]) = \pi(I_1(b_-^1, b_+^1) \sqcup I_2(b_-^2, b_+^2)) =: \pi(D_1 \sqcup D_2)$ is a causally closed neighborhood of $[x]$ in X.

Let $([p_n])_{n\in\mathbb{N}}, [q_n])_{n\in\mathbb{N}}$ be two sequences in D converging to $[p]$ and $[q]$, respectively, and suppose $[p_n] \mathrel{\widetilde{\leqslant}} [q_n]$ for all $n \in \mathbb{N}$. This means that for all n we either have $p_n^i \leqslant q_n^i$ or there exists $[a_n] \in A$ such that, say, $p_n^1 \leqslant a_n^1 \sim a_n^2 \leqslant q_n^2$. Recall remark 2.12 and suppose first that there exist subsequences such that $p_{n_k}^i \leqslant q_{n_k}^i$. Then $p^i \in [p], q^i \in [q]$ as X_i is closed in X. Then $p^i \leqslant q^i$ follows since $\bar{D}_i$ is causally closed. Otherwise, there exist subsequences that are up to symmetry of the form $p_{n_k}^1 \leqslant a_n^1 \sim a_{n_k}^2 \leqslant q_{n_k}^2$. By definition, $[a_n] \in J_X([p_n], [q_n]) \subseteq D$. In particular, $p_{n_k}^1, a_{n_k}^1 \in D_1$ as well as $a_{n_k}^2, q_{n_k}^2 \in D_2$ for all k. As D_i is relatively compact, we can extract from $(a_{n_k})_{k\in\mathbb{N}}$ a converging subsequence, say without loss of generality the whole sequence already converges to some $a^i \in \bar{D}_i$. Since $\bar{D}_i$ is causally closed, we infer $p^1 \leqslant a^1$ and $a^2 \leqslant q^2$. In particular, we have $[p] \mathrel{\widetilde{\leqslant}} [a] \mathrel{\widetilde{\leqslant}} [q]$ and hence $[p] \mathrel{\widetilde{\leqslant}} [q]$, showing that $\bar{D}$ is a causally closed neighborhood of $[x]$ in X. □

The next two lemmas will be helpful for showing the preservation of localizability, but they are also interesting on their own.

Lemma 4.6 (Localizing neighborhoods and strong causality). *In a strongly causal, non-timelike locally isolating and localizable Lorentzian pre-length space X, each point has a neighborhood basis of localizable timelike diamonds, hence X is strongly localizable.*

Proof. Let U be a localizing neighborhood of $x \in X$. By strong causality and non-timelike local isolation we find, as in lemma 3.5, $p, q \in U$ such that $x \in I(p,q) \subseteq U$ (note that X has no $\ll$-isolated points as it is localizable). By the causal convexity of $I(p,q)$, every causal curve with endpoints in $I(p,q)$ is entirely contained in $I(p,q)$ and hence also in the original neighborhood U. Thus, we can take the same d-compatibility constant $C > 0$ for $I(p,q)$ that we used for U. If $y \in I(p,q)$, then as $I(p,q)$ is a neighborhood of y and X is non-timelike

[7] We will use this construction of neighborhoods which agree on A via f and share some property several times throughout this work. We will not give a detailed description in future cases.

locally isolating, we find $y_-, y_+ \in I(p,q)$ such that $y_- \ll y \ll y_+$, which immediately yields $I^\pm(y) \cap I(p,q) \neq \emptyset$, so the non-triviality condition is satisfied. Finally, recall that $\omega_U(x,y) = \max\{L(\gamma) \mid \gamma \text{is a causal curve from } x \text{ to } y \text{ in } U\}$. We claim that $\omega_U|_{I(p,q)\times I(p,q)} = \omega_{I(p,q)}$, i.e. the maximal curve between points in $I(p,q)$ that is contained in U is already contained in $I(p,q)$. This follows immediately by the causal convexity of $I(p,q)$.

Timelike diamonds form a basis for the topology, cf lemma 3.5, so X is strongly localizable. □

Via lemma 4.6, we obtain the following very slight generalization of [13, lemma 4.3].

Lemma 4.7 (τ determines ω). *Let X be a strongly causal, non-timelike locally isolating, intrinsic and localizable Lorentzian pre-length space. Then each point in X has a neighborhood basis of localizable neighborhoods where ω agrees with τ. In other words, X is locally geodesic.*

Proof. Lemma 4.6 establishes the existence of a causally convex and localizable neighborhood basis. Being intrinsic is the only additional property of a strongly causal and localizable Lorentzian pre-length space that is required in the proof of [13, lemma 4.3]. □

Proposition 4.8 (Localizability). *Let X_1 and X_2 be two strongly causal and locally compact Lorentzian length spaces. Then $X := X_1 \sqcup_A X_2$ is localizable.*

Proof. Let $[x] \in X \setminus A$. Then we find a timelike diamond $I_i(p^i, q^i)$ in X_i which is a localizing neighborhood of x^i and does not meet A_i. Then $I_X([p],[q]) = \pi(I_i(p^i,q^i))$ is a localizable neighborhood of $[x]$.

Let $[x] \in A$. Similar to the construction in proposition 4.5, we find a timelike diamond $I_X([p],[q]) = \pi(I_1(p^1,q^1) \sqcup I_2(p^2,q^2))$ such that both original diamonds are localizing neighborhoods in X_1 and X_2, respectively. We claim that $I_X([p],[q])$ is a localizing neighborhood of $[x]$ in X. Let $\gamma : [a,b] \to X$ be a causal curve in $I_X([p],[q])$. Then $\pi^{-1}(\gamma([a,b])$ is either a causal curve in X_i or it consists of pieces of causal curves in X_1 and X_2. In the first case one can take the original constant $C_i > 0$ for γ by proposition 2.11(ii), so assume we are in the second case. Say γ starts out in in X_1, leaves at $[y] = \gamma(s) \in A$ and enters again at $[z] = \gamma(t) \in A$. Then $\pi^{-1}(\gamma|_{[a,s]})$ is a causal curve in X_1 ending in y^1 and $y^1 \leqslant z^1$. As X_1 is causally path-connected, there is a causal curve between y^1 and z^1, any of which by definition is contained in $J_1(y^1,z^1) \subseteq I_1(p^1,q^1)$. Denote by α such a curve. Then the concatenation of $\pi^{-1}(\gamma|_{[a,s]})$ and α is a causal curve in X_1. In this way, we can piece together all parts of γ that lie in X_1 to obtain one single causal curve, denoted by γ_1, which is still contained in $I_1(p^1,q^1)$. Doing the same procedure in X_2 and denoting the corresponding curve by γ_2, we infer that $L_{\tilde{d}}(\gamma) \leqslant L_{\tilde{d}}(\pi \circ \gamma_1) + L_{\tilde{d}}(\pi \circ \gamma_2) \leqslant L_d(\gamma_1) + L_d(\gamma_2) \leqslant C_1 + C_2$.

Concerning the non-triviality condition, let $[y] \in I_X([p],[q])$. Then $[p] \mathrel{\widetilde{\ll}} [y] \mathrel{\widetilde{\ll}} [q]$ and a timelike curve through $[y]$ must lie partially in this neighborhood, hence $I_X^\pm([y]) \cap I_X([p], [q]) \neq \emptyset$.

Finally, we need to show the existence of (continuously varying) maximal causal curves in $I_X([p],[q])$. We essentially need to show that the map $\tilde{\omega} : I_X([p],[q]) \times I_X([p],[q]) \to [0,\infty]$, defined via

$$\tilde{\omega}([y],[z]) := \sup\{L_{\tilde{\tau}}(\gamma) \mid \gamma \text{is a causal curve from } [y] \text{ to } [z] \text{ contained in } I_X([p],[q])\}, \tag{4.1}$$

is continuous and there exists a curve which realizes the supremum. Note that $\tilde{\tau}$ is intrinsic by proposition 4.4, and $\tilde{\tau} = \tilde{\omega}$ on $I_X([p],[q])$ since the neighborhood is causally convex. Let $[y],[z] \in I_X([p],[q])$ with $[y] \mathrel{\widetilde{\leq}} [z]$. If $\tilde{\tau}([y],[z]) = 0$, then either $\tau(y^i,z^i) = 0$ or for all $[a] \in J_X([y],[z]) \cap A$ we have $\tau(y^1,a^1) = \tau(a^2,z^2) = 0$. In both cases, we find a null curve between the points in X, which by definition is contained in $J_X([y],[z]) \subseteq I_X([p],[q])$. In the first case this is immediate with the projection and in the second case we concatenate two null curves which exist by the causal path-connectedness of X_i. So we are left with the case $[y] \widetilde{\ll} [z]$. If $[y],[z] \in X_i$, then by the τ-preservation of $f: A_1 \to A_2$ and the fact that $\widetilde{\tau}$ is intrinsic, a maximal curve in $I_i(p^i,q^i)$ is still maximal in $I_X([p],[q])$. Indeed, $\tilde{\omega}([y],[z]) = \widetilde{\tau}([y],[z]) = \tau(y^i,z^i) = \omega_i(y^i,z^i)$, cf lemma 4.7, where $\omega_i : I_i(p^i,q^i) \times I_i(p^i,q^i) \to [0,\infty)$ is the continuous map returning the maximal length of causal curves contained in $I_i(p^i,q^i)$. So suppose without loss of generality $[y] = \{y^1\}, [z] = \{z^2\}$. Let $(\gamma_n)_{n\in\mathbb{N}}$ be a sequence of causal curves from $[y]$ to $[z]$ in $I_X([p],[q])$ such that their lengths converge to the supremum. By proposition 2.11(i), each of these causal curves can be assumed to consist of (the projections of) two maximal causal curves from y^1 to a^1_n in $I_1(p^1,q^1)$ and from a^2_n to z^2 in $I_2(p^2,q^2)$, respectively, for some $[a_n] \in A \cap J_X([y],[z]) \subseteq I_X([p],[q])$. This restriction does not decrease the length of curves and we are considering a sequence whose lengths converge to the supremum. By the local compactness of X_i we can assume that the neighborhoods were chosen small enough such that even $J_i(p^i,q^i)$ is contained in a compact subset. Thus, we infer the existence of a converging subsequence of $(a^i_n)_{n\in\mathbb{N}}$, say without loss of generality $a^i_n \to a^i \in A_i \cap \overline{J_i(p^i,q^i)}$. By construction we have $x^1 \leqslant a^1_n \sim a^2_n \leqslant y^2$ for all n. Since $J_i(p^i,q^i)$ can be assumed to be causally closed as well (just shrink everything), we infer $x^1 \leqslant a^1 \sim a^2 \leqslant y^2$ and hence $[a] \in A \cap J_X([y],[z]) \subseteq I_X([p],[q])$. By hypothesis, we find maximal causal curves from x^1 to a^1 in $J_1(p^1,q^1)$ and from a^2 to y^2 in $J_2(p^2,q^2)$. Denote the (projections of the) two pieces by γ^1 and γ^2, respectively, and their concatenation by γ. This is a maximal causal curve from $[y]$ to $[z]$ in $I_X([p],[q])$: using lemma 4.3, we clearly have $L_{\widetilde{\tau}}(\gamma_n) = L_\tau(\gamma^1_n) + L_\tau(\gamma^2_n)$, where γ^1_n and γ^2_n are the (maximal) parts of γ_n in $I_1(p^1,q^1)$ and $I_2(p^2,q^2)$, respectively. Then we compute

$$\begin{aligned}\tilde{\omega}([y],[z]) &= \lim_{n\to\infty} L_{\widetilde{\tau}}(\gamma_n) = \lim_{n\to\infty}(L_\tau(\gamma^1_n) + L_\tau(\gamma^2_n)) \\ &= \lim_{n\to\infty} L_\tau(\gamma^1_n) + \lim_{n\to\infty} L_\tau(\gamma^2_n) = \lim_{n\to\infty}\omega_1(y^1,a^1_n) + \lim_{n\to\infty}\omega_2(a^2_n,z^2) \\ &= \omega_1(y^1,a^1) + \omega_2(a^2,z^2) = L_\tau(\gamma^1) + L_\tau(\gamma^2) = L_{\widetilde{\tau}}(\gamma).\end{aligned}$$

Let $[y_m] \to [y]$ be a sequence, then $[y_m] = \{y^1_m\}$ for large m. This yields a sequence of jump points $[b_m]$ and corresponding maximal causal curves β_m from $[y_m]$ to $[z]$. In particular, $([b_m])_{m\in\mathbb{N}}$ converges as well and so we obtain a causal curve β from $[y]$ to $[z]$ through, say, $[b]$. By the above calculation we know that γ is a maximal causal curve from $[y]$ to $[z]$, so we have $\lim_{m\to\infty}\tilde{\omega}([y_m],[z]) = L_{\widetilde{\tau}}(\beta) \leqslant L_{\widetilde{\tau}}(\gamma) = \tilde{\omega}([y],[z])$. Conversely, recall that $L_{\widetilde{\tau}}(\gamma) = L_\tau(\gamma_1) + L_\tau(\gamma_2)$. Say γ_1 is parameterized by $[0,1]$ and consider the map $t \mapsto L_\tau(\gamma_1|_{[t,1]})$. This map is continuous by [19, lemma 3.33] (local continuity of τ is clearly sufficient in the proof, i.e. the curve is contained in a neighborhood where τ is continuous). For appropriately small $\varepsilon > 0$, let t be such that $L_\tau(\gamma_1|_{[t,1]}) = L_\tau(\gamma_1) - \varepsilon$. As $y^1 \ll \gamma_1(t)^1$, we have $y^1_m \ll \gamma_1(t)^1$ by the openness of $\ll$. Since X_1 is causally path-connected, we find a causal curve from y^1_m to $\gamma_1(t)^1$, denote it by α_m. Then we compute

$$\begin{aligned}\lim_{m\to\infty}\tilde{\omega}([y_m],[z]) &= \lim_{m\to\infty} L_{\widetilde{\tau}}(\beta_m) \geqslant \lim_{m\to\infty} L_{\widetilde{\tau}}((\pi\circ(\alpha_m * \gamma_1|_{[t,1]})) * (\pi\circ\gamma_2)) \\ &\geqslant L_{\widetilde{\tau}}(\gamma) - \varepsilon = \tilde{\omega}([y],[z]) - \varepsilon.\end{aligned}$$

This establishes that $\tilde{\omega}$ is continuous on $(I_X([p],[q]) \times I_X([p],[q])) \cap ((X_1 \setminus A_1 \times X_2 \setminus A_2) \cup (X_2 \setminus A_2 \times X_1 \setminus A_1))$. We also know that $\tilde{\omega}$ is continuous on $(I_X([p],[q]) \times I_X([p],[q])) \cap ((X_1 \times X_1) \cup (X_2 \times X_2))$, which together cover $I_X([p],[q]) \times I_X([p],[q])$. The only thing left to check is that the two cases fit together in the following scenario: let $[y] = \{y^1\}, [z] \in A$ and let $[z_k] \to [z]$ be a sequence such that $[z_k] = \{z^2_k\}$ for all k. We have to show $\tilde{\omega}([y],[z_k]) \to \tilde{\omega}([y],[z])$. By above calculations, we have $\tilde{\omega}([y],[z_k]) = \tau(y^1,c^1_k) + \tau(c^2_k,z^2)$ and $\tilde{\omega}([y],[z]) = \tau(y^1,z^1)$ for some sequence $([c_k])_{k\in\mathbb{N}}$ in A. The convergence of $([c_k])_{k\in\mathbb{N}}$ gives a limit jump point $[c]$ for $[z]$. On the one hand, we have $\lim_k \tilde{\omega}([y],[z_k]) = \lim_k(\tau(y^1,c_k^1)+\tau(c_k^2,z_k^2)) = \tau(y^1,c^1)+\tau(c^2,z^2) = \tau(y^1,c^1)+\tau(c^1,z^1) \leqslant \tau(y^1,z^1) = \widetilde{\tau}([y],[z]) = \tilde{\omega}([y],[z])$ by the maximality of $\tilde{\omega}$ or by the reverse triangle inequality for τ. On the other hand, we know $\widetilde{\tau}$ is lower semi-continuous and $\widetilde{\tau} = \tilde{\omega}$, so $\lim_k \tilde{\omega}([y],[z_k]) \geqslant \tilde{\omega}([y],[z])$. This establishes that $\tilde{\omega}$ is continuous and hence $I_X([p],[q])$ is a localizable neighborhood of $[x]$. □

Collecting the previous propositions, we obtain the following theorem:

Theorem 4.9 (Amalgamation of length spaces). *Let X_1 and X_2 be two strongly causal and locally compact Lorentzian length spaces. Then $X := X_1 \sqcup_A X_2$ is a (strongly localizable) Lorentzian length space.*

Proof. All assumptions in propositions 4.2, 4.4, 4.5 and 4.8 are satisfied, so X is a Lorentzian length space. The claim on strong localizability follows by corollary 3.6 and lemma 4.6.[8] □

5. The causal inheritance

In this chapter, we take a look at the compatibility of the causal ladder with Lorentzian amalgamation. We consider various steps of the causal ladder in separate statements and then sum everything up in a final theorem. As it turns out, many causality conditions are able to be inherited by Lorentzian amalgamation if one is willing to admit some additional properties.

Proposition 5.1 (Chronology and causality). *Let X_1 and X_2 be two Lorentzian pre-length spaces and $X := X_1 \sqcup_A X_2$ their amalgamation.*

(i) If X_1 and X_2 are chronological, then so is X.
(ii) If X_1 and X_2 are causal, then so is X.

Proof.

(i) This is immediate from proposition 2.11(i), as for any $[x] \in X_i$ we have $\widetilde{\tau}([x],[x]) = \tau(x^i,x^i) = 0$ since X_i is chronological.
(ii) If $[x]$ and $[y]$ are from the same space, say X_1, then this follows from the causality of X_1. So assume $[x] \in X_1 \setminus A_1, [y] \in X_2 \setminus A_2$ with $[x] \,\widetilde{\leqslant}\, [y]$ and $[y] \,\widetilde{\leqslant}\, [x]$. Then we find $[a],[b] \in A$ such that

$$x^1 \leqslant a^1 \sim a^2 \leqslant y^2 \text{and } y^2 \leqslant b^2 \sim b^1 \leqslant x^1. \tag{5.1}$$

By the transitivity of $\leqslant$ it follows that $b^1 \leqslant a^1$ and $a^2 \leqslant b^2$. Since X_1 and X_2 are causal and f is $\leqslant$-preserving, we obtain $a^1 = b^1$ and $a^2 = b^2$, i.e. $[a] = [b]$. Then (5.1) implies $[x] = [a] = [b] = [y]$, a contradiction. □

Proposition 5.2 (Strong causality). *Let X_1 and X_2 be two strongly causal Lorentzian pre-length spaces. Then $X := X_1 \sqcup_A X_2$ is strongly causal.*

Proof. If $[z] \in X \setminus A$, let U be a neighborhood of $[z]$ that does not meet A. Then $\pi^{-1}(U)$ is a neighborhood of, say, z^1 in X_1 that does not meet A_1. By strong causality of X_1, we find points

[8] Technically, for the strong localizability one needs to either slightly adapt the construction of neighborhoods in proposition 4.5 or apply proposition 5.2.

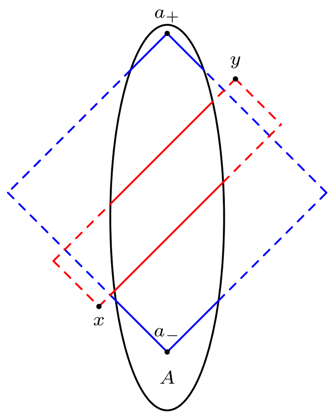

Figure 3. Showcasing the property of being time observing.

$x_1^1, y_1^1, x_2^1, y_2^1, \ldots, x_n^1, y_n^1$ in X_1 such that $z^1 \in \cap_{i=1}^n I_1(x_i^1, y_i^1) \subseteq \pi^{-1}(U)$. As $\pi|_{X_1}$ is injective, we compute $[z] \in \pi(\cap_{i=1}^n I_1(x_i^1, y_i^1)) = \cap_{i=1}^n \pi(I_1(x_i^1, y_i^1)) = \cap_{i=1}^n I_X([x_i], [y_i]) \subseteq \pi(\pi^{-1}(U)) = U$.

Let $[z] \in A$ and let $U \subseteq X$ be a neighborhood of $[z]$. Then $\pi^{-1}(U)$ is open in $X_1 \sqcup X_2$, i.e. $\pi^{-1}(U) = U_1 \sqcup U_2$, where U_i is open in $X_i, i = 1, 2$. Moreover, by definition we have $f(U_1 \cap A_1) = U_2 \cap A_2$. By strong causality, we find points $x_1^1, y_1^1, x_2^2, y_2^2, \ldots, x_n^1, y_n^1$ in X_1 and $p_1^2, q_1^2, p_2^2, q_2^2, \ldots, p_m^2, q_m^2$ in X_2 such that $z^1 \in D_1 := \cap_{i=1}^n I_1(x_i^1, y_i^1) \subseteq U_1$ and $z^2 \in D_2 := \cap_{j=1}^m I_1(p_j^2, q_j^2) \subseteq U_2$. Then we find open neighborhoods $V_i \subseteq D_i$ such that $f(V_1 \cap A_1) = V_2 \cap A_2$. As A_i is non-timelike locally isolating, there exist $b_+^i, b_-^i \in V_i \cap A_i$, such that $b_-^i \ll z^i \ll b_+^i$. Since D_i is causally convex, we have $I(b_-^i, b_+^i) \subseteq D_i \subseteq U_i$. Finally, $I(b_-^1, b_+^1) \sqcup I(b_-^2, b_+^2) \subseteq U_1 \sqcup U_2 = \pi^{-1}(U)$ and hence

$$U = \pi(\pi^{-1}(U)) = \pi(U_1 \sqcup U_2) \supseteq \pi(I(b_-^1, b_+^1) \sqcup I(b_-^2, b_+^2)) = I_X([b_-], [b_+]), \tag{5.2}$$

where the first equality holds since π is surjective and the last equality is due to proposition 3.4. To summarize, in an arbitrary neighborhood of $[z]$ we found an intersection of elements of the sub-basis of the Alexandrov topology, showing that X is strongly causal. □

Definition 5.3 (Time observation). A subset A of a Lorentzian pre-length space A is called future observing if for every $x \in X$ there exists $a_- \in A$ such that $J^+(x) \cap A \subseteq J^+(a_-) \cap A$. Similarly, we define a past observing set. We say A is time observing if it is both future and past observing. In particular, we then infer for all $x, y \in X$ the existence of $a_-, a_+ \in A$ such that $J(x, y) \cap A \subseteq J(a_-, a_+) \cap A$, see figure 3.

Note that $J^-(A) \cap J^+(A) = X$ is a sufficient condition for A being time observing. Indeed, in this case for all $x \in X$ we find a_- and a_+ in A such that $a_- \leqslant x \leqslant a_+$ and hence $J^-(x) \subseteq J^-(a_+)$ as well as $J^+(x) \subseteq J^+(a_-)$.

Proposition 5.4 (Global hyperbolicity). *Let X_1 and X_2 be two globally hyperbolic, causally path connected, locally causally closed and d-compatible Lorentzian pre-length spaces. Let A_1 and A_2 be time observing. Then $X := X_1 \sqcup_A X_2$ is globally hyperbolic.*

Proof. By lemma 3.8, X_1 and X_2 are strongly causal, so by proposition 5.2, X is strongly causal. It is easy to infer via proposition 4.8 that X is d-compatible as well. Then X fulfills all the assumptions to apply lemma 3.9, showing that X is non-totally imprisoning. It is left to show that diamonds in X are compact. Let $[x], [y] \in X$ with $[x] \widetilde{\leqslant} [y]$. If $[x], [y] \in X_i \setminus A_i$ and $J_i(x^i, y^i) \cap A_i = \emptyset$, then $J_X([x], [y]) = \pi(J_i(x^i, y^i))$ inherits the compactness by assumption. If

$[x],[y]\in A$, the compactness of $J_X([x],[y])=\pi(J_1(x^1,y^1)\sqcup J_2(x^2,y^2))$ follows by assumption as well. The two cases left to show are $[x],[y]\in X_i$ and their causal diamond meets A or, say, $[x]=\{x^1\},[y]=\{y^2\}$.

Assume we are in the first case, and both points are contained in, say, $X_1\setminus A_1$. By proposition 3.4, we then have

$$J_X([x],[y])=\pi(J_1(x^1,y^1)\sqcup(\cup_{x^1\leqslant a^1}J_2^+(a^2))\cap(\cup_{b^1\leqslant y^1}J_2^-(b^2)))\,. \tag{5.3}$$

Let $([p_n])_{n\in\mathbb{N}}$ be a sequence in $J_X([x],[y])$. If there exists a subsequence $(p^1_{n_k})_{k\in\mathbb{N}}$ in $J_1(x^1,y^1)$, then by the compactness of $J_1(x^1,y^1)$ we can extract a convergent subsequence which yields a convergent subsequence of $([p_n])_{n\in\mathbb{N}}$. So suppose all (but finitely many) sequence members are from the second space originally. By definition, this means that for all large enough $n\in\mathbb{N}$ we find $a^1_n,b^1_n\in A_1$ such that $a^2_n\leqslant p^2_n\leqslant b^2_n$. In particular, $x^1\leqslant a^1_n\leqslant b^1_n\leqslant y^1$, so $a^1_n,b^1_n\in J_1(x^1,y^1)\cap A_1$. Since A_1 is time observing, we find $q^1_-,q^1_+\in A_1$ such that $J_1(x^1,y^1)\cap A_1\subseteq J_1(q^1_-,q^1_+)\cap A_1$. Then also $a^1_n,b^1_n\in J_1(q^1_-,q^1_+)\cap A_1$ and since f is causality preserving, we conclude $a^2_n,b^2_n\in J_2(q^2_-,q^2_+)\cap A_2$ as well. Thus, $(p^2_n)_{n\in\mathbb{N}}$ is contained in a compact set and has a convergent subsequence, say without loss of generality the whole sequence already converges to some $p^2\in\overline{(\cup_{x^1\leqslant a^1}J_2^+(a^2))\cap(\cup_{b^1\leqslant y^1}J_2^-(b^2))}=:\overline{D}$.

It is left to show that p^2 is not only in the closure but inside the set D itself. To this end consider the resulting sequences $(b^2_n)_{n\in\mathbb{N}}$ and $(a^2_n)_{n\in\mathbb{N}}$. Clearly, $x^1\leqslant a^1_n\leqslant b^1_n\leqslant y^1$ for all n by assumption. The two sequences $(b^1_n)_{n\in\mathbb{N}}$ and $(a^1_n)_{n\in\mathbb{N}}$ also possess convergent subsequences since they are contained in the compact set $J_1(x^1,y^1)$. Say $a^1_n\to a^1,b^1_n\to b^1$, then $x^1\leqslant a^1$ and $b^1\leqslant y^1$ by the (global) causal closedness of X_1. Thus, $J_2^+(a^2)$ and $J_2^-(b^2)$ are also part of the unions in D. By the causal closedness of X_2, we infer from $a^2_n\leqslant p^2_n\leqslant b^2_n$ that $a^2\leqslant p^2\leqslant b^2$. So p^2 is an element of D and not just of $\overline{D}$. In particular, $[p]\in J_X([x],[y])$. This shows that $J_X([x],[y])$ is compact. The case of $[x]\in X\setminus A,[y]\in A$ is similar and less complicated than the one we just proved.

Finally, consider the case $[x]=\{x^1\},[y]=\{y^2\}$. Then

$$J_X([x],[y])=\pi(J_1^+(x^1)\cap(\cup_{y^2\geqslant a^2}J_1^-(a^1))\sqcup(\cup_{x^1\leqslant b^1}J_2^+(b^2))\cap J_2^-(y^2))\,. \tag{5.4}$$

Let $([p_n])_{n\in\mathbb{N}}$ be a sequence in $J_X([x],[y])$, then there is a subsequence in one of the two original spaces. Say without loss of generality $p^1_n\in J_1^+(x^1)\cap(\cup_{y^2\geqslant a^2}J_1^-(a^1))$ for all n. Then for all n we have that there exists $[a_n]\in A:x^1\leqslant p^1_n\leqslant a^1_n\sim a^2_n\leqslant y^2$. As A_2 is time observing, we find $c^2\in A_2$ such that $J_2^-(c^2)\cap A_2\supseteq J_2^-(y^2)\cap A_2$. Thus, we have $c^2\geqslant a^2_n$ for all n, and hence $c^1\geqslant a^1_n$ as well. In particular, $p^1_n,a^1_n\in J_1(x^1,c^1)$ for all n, so we infer the existence of converging subsequences, say $p^1_n\to p^1$ and $a^1_n\to a^1$. By the global causal closedness, it follows that $p^1\leqslant a^1$. Moreover, by $y^2\geqslant a^2_n$ for all n we infer $y^2\geqslant a^2$, so $p^1\in J_1^+(x^1)\cap(\cup_{y^2\geqslant a^2}J_1^-(a^1))$ and hence $[p]\in J_X([x],[y])$, establishing that diamonds are compact. □

Proposition 5.5 (Distinction). *Let X_1 and X_2 be two strongly causal, non-timelike locally isolating and distinguishing Lorentzian pre-length spaces. Then $X:=X_1\sqcup_A X_2$ is distinguishing.*

Proof. We need to show that $I_X^\pm([x])=I_X^\pm([y])$ implies $[x]=[y]$ for all $[x],[y]\in X$. We distinguish several cases depending on where the points originally come from. We will only consider the future case, as the past case is completely analogous. Suppose first $[x],[y]\in A$ with $I^+([x])=I_X^+([y])$. By proposition 3.4, we then have $\pi(I_1^+(x^1)\sqcup I_2^+(x^2))=\pi(I_1^+(y^1)\sqcup I_2^+(y^2))$. As f is $\ll$-preserving, we clearly have $\pi(I_1^+(x^1)\cap A_1)=\pi(I_2^+(x^2)\cap A_2)$ as well as $\pi(I_1^+(y^1)\cap A_1)=\pi(I_2^+(y^2)\cap A_2)$. By lemma 3.11, we then infer $I_1^+(x^1)\sqcup I_2^+(x^2)=\pi^{-1}(\pi(I_1^+(x^1)\sqcup I_2^+(x^2)))=\pi^{-1}(\pi(I_1^+(y^1)\sqcup I_2^+(y^2)))=I_1^+(y^1)\sqcup I_2^+(y^2)$. So in particular $I_1^+(x^1)=I_1^+(y^1)$ and $I_2^+(x^2)=I_2^+(y^2)$, hence $[x]=[y]$ by X_1 and X_2 being distinguishing.

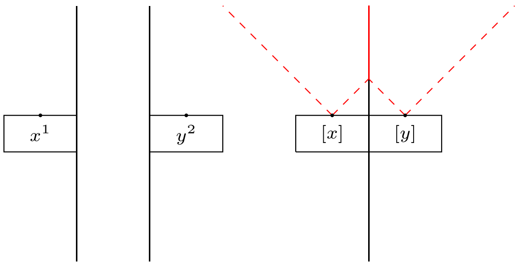

Figure 4. In this example, we have $I^+_X([x]) = I^+_X([y])$ but $[x] \neq [y]$.

Now suppose, say, $[x] \in X_1$ and $[y] \in X_1 \setminus A_1$, with $I^+_X([x]) = I^+_X([y])$. If $[x] \in X_1 \setminus A_1$, then we have $\pi(I^+_1(x^1) \sqcup (\cup_{x^1 \ll a^1} I^+_2(a^2)) = \pi(I^+_1(y^1) \sqcup (\cup_{y^1 \ll b^1} I^+_2(b^2))$ by proposition 3.4. Similarly as above, we have $\pi(I^+_1(x^1) \cap A_1) = \pi(\cup_{x^1 \ll a^1} I^+_2(a^2) \cap A_2)$ by the $\ll$-preservation of f. Indeed, let $b^1 \in I^+_1(x^1) \cap A_1$, then since $\ll$ is open, there is a neighborhood U of b^1 such that $U \subseteq I^+_1(x^1)$. As A_1 is non-timelike locally isolating, we find $b^1_- \in U \cap A_1$ such that $b^1_- \ll b^1$. Then also $b^2_- \ll b^2$ and since $x^1 \ll b^1_-$, it follows that $b^2 \in \cup_{x^1 \ll a^1} I^+_2(a^2) \cap A_2$. Conversely, if $b^2 \in \cup_{x^1 \ll a^1} I^+_2(a^2) \cap A_2$, then we obtain the existence of $c^1 \in I^+_1(x^1) \cap A_1$ such that $c^2 \ll b^2$. Then also $c^1 \ll b^1$ and hence $x^1 \ll b^1$, establishing $b^1 \in I^+_1(x^1) \cap A_1$. In summary, we have shown $b^1 \in I_1(x^1) \cap A_1$ whenever $b^2 \in \cup_{x^1 \ll a^1} I^+_2(a^2) \cap A_2$, i.e. $\pi(I^+_1(x^1) \cap A_1) = \pi(\cup_{x^1 \ll a^1} I^+_2(a^2) \cap A_2)$. Then we have $I^+_1(x^1) \sqcup (\cup_{x^1 \ll a^1} I^+_2(a^2) = I^+_1(y^1) \sqcup (\cup_{y^1 \ll b^1} I^+_2(b^2)$ by lemma 3.11. So $I^+_1(x^1) = I^+_1(y^1)$ and hence $x^1 = y^1$ and $[x] = [y]$ by X_1 being distinguishing.

In the case of $[y] \in X_1 \setminus A_1$ and $[x] \in A$, we can proceed analogously using $\pi(I^+_1(x^1) \sqcup I^+_2(x^2)) = \pi(I^+_1(y^1) \sqcup (\cup_{y^1 \ll b^1} I^+_2(b^2))$.

It is left to cover the case of $[x] \in X_1 \setminus A_1$ and $[y] \in X_2 \setminus A_2$. We then have $I^+_X([x]) = \pi(I^+_1(x^1) \sqcup \cup_{x^1 \ll a^1} I^+_2(a^2)) = \pi(\cup_{y^2 \ll a^2} I^+_1(a^1) \sqcup I^+_2(y^2)) = I^+_X([y])$. By similar arguments as before we conclude $I^+_1(x^1) = \cup_{y^2 \ll a^2} I^+_1(a^1)$ and $\cup_{x^1 \ll a^1} I^+_2(a^2) = I^+_2(y^2)$. Note that using the same letter in both unions is in fact not sloppy but accurate: if, say, $y^2 \ll a^2$, then by the above equality we infer the existence of some $b^1 \in I^+_1(x^1)$ such that $b^2 \ll a^2$. Then $b^1 \ll a^1$ and since $x^1 \ll b^1$, we have $x^1 \ll a^1$. We can apply the same argument vice versa. To emphasize, we have $x^1 \ll a^1$ if and only if $y^2 \ll a^2$ for all $[a] \in A$. Now we use similar arguments as in proposition 5.2. Since $x^1 \in X_1 \setminus A_1$, we find a neighborhood U_1 containing x^1 that does not meet A_1. By lemma 3.5 we find a timelike diamond D_1 inside U_1 which contains x^1 and whose endpoints lie in U_1. Then for all $q^1 \in D_1 \cap I^+_1(x^1)$ (which is nonempty by non-timelike local isolation), there cannot exist $a^1 \in I^+_1(x^1) \cap A_1$ such that $a^1 \ll q^1$, since otherwise $a^1 \in D_1$. In particular, $q^1 \notin \cup_{y^2 \ll a^2} I^+_1(a^1)$ while $q^1 \in I^+_1(x^1)$, a contradiction. So under our assumptions, two points from two different spaces cannot have the same future. In summary, we conclude that X is distinguishing. □

Example 5.6 (Double flagpole). The following example demonstrates that without the assumptions of strong causality and non-timelike local isolation, proposition 5.5 fails. Consider two 'flagpoles' in the Minkowski plane, i.e. the union of (part of) the t-axis with a rectangle, see figure 4. Both of these spaces are distinguishing. Gluing them along the t-axis results in a space which is not distinguishing anymore.

Until now, most of the properties were able to be transferred either as is or with relatively harmless additional assumptions. This changes now, as the remaining steps of the causal ladder appear to be less suitable for preservation via gluing constructions.

Let us begin with non-total imprisonment. It is easily shown that a compact set in X is built from two compact sets in X_1 and X_2, i.e. for any compact $K \subseteq X$ we have $K = \pi(K_1 \sqcup K_2)$ with $K_i \subseteq X_i, i = 1,2$ compact. Further, any causal curve in K can thus be decomposed into causal curves in K_1 and K_2. Then the first problem is that the causal curve may consist of infinitely many pieces. At first glance one might eliminate this problem by imposing causal path-connectedness on X_1 and X_2, which would enable us to connect all pieces in one space into a single curve. Then of course the problem is that these connecting pieces cannot be guaranteed to stay inside K_1 or K_2. This is yet another indicator that below strong causality, there is no way of taming the causality and relating it to the topology.

The following definition can be regarded as a 'non-intrinsic' formulation of non-total imprisonment. We will show that under this stronger condition, the property of non-total imprisonment is inherited.

Definition 5.7 (Extrinsic non-total imprisonment). Let X be a Lorentzian pre-length space. X is called extrinsically non-totally imprisoning if for every compact set $K \subseteq X$ there is a constant $C > 0$ such that for all causal sequences $(x_1, x_2, \ldots, x_n)$ in K we have $\sum_{i=1}^{n} d(x_i, x_{i+1}) \leqslant C$, where a causal sequence satisfies $x_i \leqslant x_{i+1}$ for all i. Since L_d is given as the supremum of causal sequences, we immediately get that this property implies non-total imprisonment. Moreover, we also have the same bound for infinite sequences, since if there exists an infinite causal sequence whose sum of distances is larger than C, then this must already occur after summing up finitely many distances.

Remark 5.8 (Placing the extrinsic version on the causal ladder). Recall the classical example of a spacetime which is causal but not strongly causal, cf [27, example 14.12]: the Lorentz cylinder $M := \mathbb{S}^1 \times \mathbb{R}$ with two horizontal strips removed. This spacetime is also non-totally imprisoning, as in fact *any* causal curve has a (uniformly) bounded d-length since wrapping around is forbidden via the removed pieces. Thus, also any causal chain has uniformly bounded d-length. So M is an example of a extrinsically non-totally imprisoning space which is not strongly causal. This suggests at least that this property is not stronger than strong causality.

Lemma 5.9 (Decomposition of compact subsets in amalgamation). *Let X_1 and X_2 be metric spaces, $A_i \subseteq X_i$ closed and $f : A_1 \to A_2$ a locally bi-Lipschitz homeomorphism. Consider the quotient space X of $X_1 \sqcup X_2$ generated by the equivalence relation $a \sim f(a)$ for all $a \in A_1$. Let $K \subseteq X$ be compact. Then there exist compact $K_i \subseteq X_i, i = 1,2$ such that $\pi^{-1}(K) = K_1 \sqcup K_2$.*

Proof. Let $K \subseteq X$ be compact. Then by definition $\pi^{-1}(K) = K_1 \sqcup K_2 \subseteq X_1 \sqcup X_2$ for some $K_i \subseteq K$. Let, say, $(x_n^1)_{n\in\mathbb{N}}$ be a sequence in K_1. Then $([x_n])_{n\in\mathbb{N}}$ is a sequence in K, which contains a convergent subsequence as K is compact, say $([x_n])_{n\in\mathbb{N}} \to [x] \in K$. Note that $\pi(X_1)$ is closed in X, hence $[x] \in K \cap \pi(X_1) = \pi(K_1)$. Thus, $x_n^1 \to x^1$ follows easily, cf remark 2.12. This shows that K_1 is compact and similarly for K_2. □

Proposition 5.10 (Non-total imprisonment). *Let X_1 and X_2 be two extrinsically non-totally imprisoning Lorentzian pre-length spaces. Then X is non-totally imprisoning.*

Proof. Let K be a compact subset of X. Then $\pi^{-1}(K) = K_1 \sqcup K_2$ for compact $K_1 \subseteq X_1$ and $K_2 \subseteq X_2$.

Let $\gamma : [a,b] \to K$ be a causal curve. Then either $\pi^{-1}(\gamma([a,b]))$ is a causal curve in K_i or it consists of pieces of causal curves in K_1 and K_2. In the first case we can take the original constant C_i of K_i by the extrinsic non-total imprisonment of X_i, so assume we are in the second case. Without loss of generality, assume γ starts out in X_1, leaves X_1 through $[p] = \gamma(t_1) \in A$ and enters X_1 again at $[q] = \gamma(t_2) \in A$. By the $\leqslant$-preservation of f, we obtain $p^1 \leqslant q^1$. By the extrinsic non-total imprisonment of X_1, we have $L_d(\gamma|_{[a,t_1]}) + d(p^1, q^1) \leqslant C_1$ (remember that

L_d is given as the supremum of lengths of causal chains, each of which satisfies the given bound by assumption). Doing this iteratively for all jump points, we end up with two causal sequences (each of which may have infinitely[9] many members, as γ may jump infinitely often) in K_1 and K_2, respectively. In the end, this implies $L_{\tilde{d}}(\gamma) \leqslant C_1 + C_2$. □

Finally, we collect all compatibility results regarding the causal ladder and amalgamation into the following theorem.

Theorem 5.11 (Causal inheritance). *Let X_1 and X_2 be two Lorentzian pre-length spaces and $X := X_1 \sqcup_A X_2$ their Lorentzian amalgamation. Then we have the following preservation of causality conditions.*

(i) If X_1 and X_2 are chronological, then so is X.
(ii) If X_1 and X_2 are causal, then so is X.
(iii) If X_1 and X_2 are extrinsically non-totally imprisoning, then X is non-totally imprisoning.
(iv) If X_1 and X_2 are strongly causal, then so is X.
(v) If X_1 and X_2 are strongly causal, non-timelike locally isolating and distinguishing, then X is distinguishing.
(vi) If X_1 and X_2 are causally path-connected, locally causally closed, d-compatible and globally hyperbolic with A_1 and A_2 time observing, then X is globally hyperbolic.

Proof. This is a summary of propositions 5.1, 5.10, 5.2, 5.5 and 5.4. □

6. Outlook on missing properties

The remaining properties of the causal ladder, for now, are not inherited by gluing under 'reasonable' additional assumptions. These are K-causality, reflectivity and causal simplicity. Nevertheless, we will give our (unsuccessful) thoughts and discuss these properties in some detail.

6.1. K-causality

Let us begin with K-causality. In [10], it is shown that for certain Lorentzian pre-length spaces this is equivalent to the existence of a time function, which seems a bit more convenient to work with. The most obvious approach for constructing a time function on a glued space X under the assumption that X_1 and X_2 do possess time functions is the following: start with a time function $T_1 : A_1 \to \mathbb{R}$ and consider the corresponding time function $T_2 := T_1 \circ f^{-1} : A_2 \to \mathbb{R}$. Suppose we could extend T_1 and T_2 to time functions defined on the entirety of the spaces X_1 and X_2 (and denote the extensions again by T_1 and T_2, respectively). Then consider a function $T : X \to \mathbb{R}$ defined by $T([x]) := T_i(x^i)$. This function is well defined since for $[x] \in A$ we have $T_1(x^1) = T_2(x^2)$. If $[x] \widetilde{<} [y]$ in X, then either $x^i < y^i$ or there exists $[a] \in A$ such that $x^i < a^i \sim a^j < y^j$. In the first case we have $T([x]) = T_i(x^i) < T_i(y^i) = T([y])$, since T_i is a time function on X_i. In the second case, we know $T([x]) = T_i(x^i) < T_i(a^i) = T_j(a^j) < T_j(y^j) = T([y])$. Thus, T is a time function (it is continuous since T_1 and T_2 are).

Of course, the difficulty lies in extending a time function from a subset to the whole space in such a way that it stays a time function. Historically, techniques and results from utility theory

[9] Notice there are at most countably infinitely many jump points. Indeed, the domain $[a, b]$ can only be covered by countably many non-trivial subintervals, as, e.g. a measure theoretic argument shows.

have enabled an entirely different approach to showing the existence of a strictly increasing function with respect to the K-relation compared to original results [12, 15, 16] (see [24, page 2] for a detailed discussion). Most importantly, [24] noted that the existence of time functions on K-causal spacetimes is an immediate consequence of Levin's theorem:

Theorem 6.1 (Levin's theorem, [20]). *Let X be a second countable and locally compact topological Hausdorff space and $\leqslant$ a closed, reflexive and transitive relation on X. Then there exists a continuous function $T : X \to \mathbb{R}$ such that $x \leqslant y \Rightarrow T(x) \leqslant T(y)$ with equality if and only if $y \leqslant x$ as well.*

Let for now X be a topological space equipped with a pre-order, i.e. a reflexive and transitive relation, denoted by $\leqslant$. For the task of extending a merely increasing function[10] from a subset to the whole space X, [25] gives the following necessary and sufficient criterion, also found in [26, theorem 2]. To this end, recall that a subset A is called decreasing if $x \in A, y \leqslant x \Rightarrow y \in A$, and an increasing set is defined analogously. Then $D(A)$ denotes the smallest closed decreasing subset that contains A. Similarly, $I(A)$ denotes the smallest closed increasing subset containing A. Finally, a normally pre-ordered space is a topological space equipped with a pre-order such that whenever A is closed and decreasing and B is closed and increasing and $A \cap B = \emptyset$, then there exist open sets $A' \supseteq A, B' \supseteq B$ such that $A^{'}$ is decreasing and $B^{'}$ is increasing and $A' \cap B' = \emptyset$.

Theorem 6.2 (Extension criterion, [25, theorem 3.1]). *Let X be a normally pre-ordered space and let A be a subspace. Let $T : A \to [0,1]$ be a continuous isotone function. Then T possesses an extension to a continuous isotone function on X if and only if the following holds for all $s,t \in [0,1]$:*

$$s < t \Rightarrow D(T^{-1}([0,s])) \cap I(T^{-1}([t,1])) = \emptyset. \tag{6.1}$$

Two problems become apparent rather quickly: on the one hand, denoting the K-relation by $\leqslant^K$, we do not know whether composing a $\leqslant^K$-increasing function on A_1 with f^{-1} yields a $\leqslant^K$-increasing function on A_2. For this one really needs to take into consideration the topology and causality of the whole spaces and not just of the homeomorphic subsets. This is why it is expected that requiring f to be $\leqslant^K$-preserving is strictly necessary. On the other hand, time functions are not just isotone, rather they are strictly increasing with respect to $\leqslant$, i.e. if $x < y$[11] then $f(x) < f(y)$.

In this regard there are also some fundamental topological results: to extend a strictly increasing function in the same fashion as above, [17, theorem 2.1][12] demands, among other things, a partial order $\leqslant$ to be order-separable, which includes that there exists a countable subset C of X such that for all $x,z \in X$ with $x < z$ there is $y \in C$ such that $x < y < z$. By considering the disjoint and uncountable collection of lightlike segments from $(0,a)$ to $(1,a+1)$ for $a \in [0,1]$ in the Minkowski plane, one sees that the causal relation $\leqslant$ (and hence also the K-relation) is not order-separable even in very nice spaces. Thus, this is an unreasonable requirement. Next, one could try to apply this theory to the timelike relation, which is order-separable if X is separable: indeed, $I(x,y)$ is open for all $x,y \in X$, and separability yields the existence of a countable subset $\{p_i\}_{i\in\mathbb{N}}$ such that each nonempty open set in X contains some p_i. This

[10] Here, $f : X \to \mathbb{R}$ is increasing if $x \leqslant y \Rightarrow f(x) \leqslant f(y)$ holds for all $x,y \in X$. In the context of order topology and/or utility theory, such functions are called isotone.

[11] Notice that in this setting, $x < y$ is defined as $x \leqslant y$ and $y \nleqslant x$. If the relation is antisymmetric, this is equivalent to $x \leqslant y$ and $x \neq y$.

[12] To not distract too much with new notation which is then not really used anyways, we decided to only give a reference in this case.

would at least lead to the existence of a semi-time function[13] on a glued space. However, $\ll$ is neither closed nor a pre-order. In summary, it seems that there is currently no machinery available to guarantee the extension of a (semi-)time function.

As a small consolation, one can establish the existence of a continuous isotone function on an amalgamation. This is done by imposing the topological conditions required in Levin's theorem on the individual spaces and observe that these are inherited. Notice that we now switch back to our usual notation of Lorentzian pre-length spaces X_1 and X_2 and their amalgamation X.

Proposition 6.3 (Existence of isotone functions). *Let X_1 and X_2 be two locally compact and second countable Lorentzian pre-length spaces. Then X admits a continuous $\leqslant$-isotone function.*

Proof. We only need to show that X is second countable and locally compact, to then apply Levin's theorem to the K-relation on X. To this end, recall that being second countable and Lindelöf is equivalent for metric spaces. So let $\{U^i\}_{i\in I}$ be an open cover of X. Then $\pi^{-1}(U^i) = U_1^i \sqcup U_2^i$ is open in $X_1 \sqcup X_2$, i.e. U_1^i and U_2^i are open in X_1 and X_2, respectively, for all $i \in I$. Moreover, $\{U_1^i\}_{i\in I}$ and $\{U_2^i\}_{i\in I}$ are open covers of X_1 and X_2, respectively. Thus, there exist countable subcovers $\{U_1^{i_j}\}_{j\in\mathbb{N}}$ and $\{U_2^{h_j}\}_{j\in\mathbb{N}}$ of X_1 and X_2, respectively. In particular, $\{U_1^{i_j} \sqcup U_2^{i_j}\}_{j\in\mathbb{N}} \cup \{U_1^{h_j} \sqcup U_2^{h_j}\}_{j\in\mathbb{N}}$ is a countable subcover of $X_1 \sqcup X_2$ such that the projection of each set is open in X. This yields a countable subcover of X, establishing that X is Lindelöf and hence second countable.

Concerning local compactness, let first, say, $[x] \in X_1 \setminus A_1$. Let U_1 be a compact neighborhood of x^1 which does not meet A_1 and let $V_1 \subseteq U_1$ be an open set containing x^1. Then $\pi(U_1)$ is a compact neighborhood of $[x]$. It is compact since U_1 is compact and π is continuous, and $[x] \in \pi(V_1)$ is open since $\pi^{-1}(\pi(V_1)) = V_1$ is open. The case of $[x] \in A$ is similar: we then find compact neighborhoods U_i of x^i in $X_i, i = 1,2$ such that $f(U_1 \cap A_1) = U_2 \cap A_2$. Let $V_i \subseteq U_i$ be open sets containing x^i such that $f(V_1 \cap A_1) = V_2 \cap A_2$. Then $\pi(U_1 \sqcup U_2)$ is a compact neighborhood of $[x]$. It is compact since π is continuous and $[x] \in \pi(V_1 \sqcup V_2)$ is open since $\pi^{-1}(\pi(V_1 \sqcup V_2)) = V_1 \sqcup V_2$ is open, cf lemma 3.11. Thus, X satisfies the assumptions for Levin's theorem for the K-relation. $\square$

6.2. Reflectivity

The next step on the ladder is causal continuity, or more precisely, the reflectivity part of this property as we already showed that being distinguishing is inherited. Our by now standard approach of distinguishing cases of where the points originally come from starts out very promising, as all cases except both points being isolated in different spaces follow very naturally. Indeed, in all these other cases the inclusions $I_X^+([x]) \subseteq I_X^+([y])$ lead back to, say, $I_1^+(x^1) \subseteq I_1^+(y^1)$, and from this the other, more complicated side can be resolved as well. For instance, consider the case $[x], [y] \in X_1 \setminus A_1$, then via proposition 3.4 and lemma 3.11, $I_X^+([x]) \subseteq I_X^+([y])$ yields $I_1^+(x^1) \subseteq I_1^+(y^1)$ and $\cup_{x^1 \ll a^1} I_2^+(a^2) \subseteq \cup_{y^1 \ll b^1} I_2^+(b^2)$. By the reflectivity of X_1 we obtain $I_1^-(y^1) \subseteq I_1^-(x^1)$. In particular, for any $b^1 \ll y^1$ it then follows that $b^1 \ll x^1$, so $\cup_{b^1 \ll y^1} I_2^-(b^2) \subseteq \cup_{a^1 \ll x^1} I_2^-(a^2)$. This implies $I_X^-([y]) \subseteq I_X^-([x])$. Thus, the only case left to consider is $[x] = \{x^1\}, [y] = \{y^2\}$, for which $I_X^+([x]) \subseteq I_X^+([y])$ yields $I_1^+(x^1) \subseteq \cup_{y^2 \ll b^2} I_1^+(b^1)$ and $\cup_{x^1 \ll a^1} I_2^+(a^2) \subseteq I_2^+(y^2)$. The seemingly easier half, in this case in X_1, can be achieved

[13] A semi-time function is a continuous function $T : X \to \mathbb{R}$ satisfying $x \ll y \Rightarrow T(x) < T(y)$.

with the additional assumption that τ is continuous and X_1 and X_2 are non-timelike locally isolating. That is, under these assumptions one can infer $\cup_{b^2 \ll y^2} I_1^-(b^1) \subseteq I_1^-(x^1)$. Indeed, given $p^1 \in \cup_{b^2 \ll y^2} I_1^-(b^1)$ we find $[b] \in A$ such that $p^1 \ll b^1 \sim b^2 \ll y^2$. Consider a sequence $x_n^1 \to x^1, x_n^1 \in I_1^+(x^1)$.[14] Since $I_1^+(x^1) \subseteq \cup_{y^2 \ll b^2} I_1^+(b^1)$ by assumption, for all x_n^1 we find $[a_n] \in A$ such that $x_n^1 \gg a_n^1 \sim a_n^2 \gg y^2$. In particular, $\tau(p^1, x_n^1) \geqslant \tau(p^1, b^1) + \tau(b^1, a_n^1) + \tau(a_n^1, x_n^1) > \tau(p^1, b^1) =: K > 0$.

As τ is continuous, we get $\tau(p^1, x^1) \geqslant K$, i.e. $p^1 \in I_1^-(x^1)$. Continuity of the time separation function implies that a space is reflecting, see [1, proposition 3.17], so this is already a comparatively strong assumption. However, if one would follow the same approach, the other desired inclusion, $I_2^-(y^2) \subseteq \cup_{a^1 \ll x^1} I_2^-(a^2)$, suggests that one would have to be able to enclose a sequence of jump points into a compact subset to infer the existence of a limit which serves as a jump point for x^1. More precisely, given a sequence $x_n^1 \to x^1, x_n^1 \in I_1^+(x^1)$ and some $p^2 \in I_2^-(y^2)$, we find $[b_n]$ such that $p^2 \ll b_n^2 \sim b_n^1 \ll x_n^1$ for all n, but there is no lower bound on $\tau(b_n^1, x_n^1)$ and we do not know whether $([b_n])_{n\in\mathbb{N}}$ converges. To guarantee this, one would essentially be already on the level of global hyperbolicity, at which point the inheritance of being reflective would be redundant. Finally, note that example 5.6 also serves as a counterexample for the inheritance of reflectivity, only in this case the omitted assumptions do not appear to be useful.

6.3. *Causal simplicity*

At last, let us briefly discuss causal simplicity. This is similar to reflectivity in the sense that some compactness argument is required, which then requires the original spaces to be not far removed from global hyperbolicity. Indeed, showing the closedness of, say, $J_X^+([x]) = \pi(J_1^+(x^1) \sqcup \cup_{x^1 \leqslant a^1} J_2^+(a^2))$, requires a subsequence of jump points a_n^i to converge. More precisely, let $([p_n])_{n\in\mathbb{N}}$ be a sequence in $J_X^+([x])$ converging to $[p]$. If there is a subsequence $p_{n_k}^1 \to p^1$ in X_1, then the desired relation follows by the closedness of $J_1^+(x^1)$. So suppose that for all (large enough) n there exists $[a_n]$ such that $x^1 \leqslant a_n^1 \sim a_n^2 \leqslant p_n^2$. Convergence of $([a_n])_{n\in\mathbb{N}}$ appears to be the only way of constructing $[a]$ such that $x^1 \leqslant a^1 \sim a^2 \leqslant p^2$. The most obvious assumptions that guarantee this are precisely the ones used in proposition 5.4.

At this point, it might also be important to point out that these issues about convergence mentioned in the subsections about causal simplicity and reflectivity appear at first glance to be solvable by assuming that A itself is compact. However, compactness and non-timelike local isolation are in fact a very restrictive combination: indeed, A will most likely have a boundary, and on the, so to say, spacelike part of the boundary, points in A will not have timelike related points from A nearby. To put it in another way, a compact A might work out if the original spaces X_1 and X_2 had spacelike boundary to begin with.

Data availability statement

No new data were created or analysed in this study.

[14] This exists by non-timelike local isolation: choose $x_n^1 \in I^+(x^1) \cap B_{\frac{1}{n}}(x^1)$.

Acknowledgments

I want to thank Tobias Beran for valuable input throughout various stages of this project. I also want to acknowledge the kind hospitality of the Erwin Schrödinger International Institute for Mathematics and Physics (ESI), where this research was in part concluded.

This work was supported by research Grant P33594 of the Austrian Science Fund FWF.

ORCID iD

Felix Rott https://orcid.org/0000-0001-7314-0889

References

[1] Aké Hau L, Cabrera Pacheco A and Solis D 2020 On the causal hierarchy of Lorentzian length spaces *Class. Quantum Grav.* **37** 215013

[2] Alexander S B, Graf M, Kunzinger M and Sämann C 2019 Generalized cones as Lorentzian length spaces: causality, curvature and singularity theorems *Commun. Anal. Geom.* (https://doi.org/10.48550/arXiv.1909.09575) accepted

[3] Barrera W, Montes de Oca L and Solis D 2022 Comparison theorems for Lorentzian length spaces with lower timelike curvature bounds (arXiv:2204.09612)

[4] Beran T 2020 Lorentzian length spaces *Master's Thesis* University of Vienna (available at: https://phaidra.univie.ac.at/open/o:1363059)

[5] Beran T, Ohanyan A, Rott F and Solis D A 2023 The splitting theorem for globally hyperbolic Lorentzian length spaces with non-negative timelike curvature *Lett. Math. Phys.* **113** 48

[6] Beran T and Rott F 2022 Gluing constructions for Lorentzian length spaces (arXiv:2201.09695)

[7] Beran T and Sämann C 2022 Hyperbolic angles in Lorentzian length spaces and timelike curvature bounds (arXiv:2204.09491)

[8] Bridson M R and Haefliger A 1999 *Metric Spaces of non-Positive Curvature* (*Comprehensive Studies in Mathematics* vol 319) (Springer)

[9] Burago D, Burago Y and Ivanov S 2001 *A Course in Metric Geometry* (*Graduate Studies in Mathematics* vol 33) (American Mathematical Society)

[10] Burtscher A and García-Heveling L 2021 Time functions on Lorentzian length spaces (arXiv:2108.02693)

[11] Cavalletti F and Mondino A 2020 Optimal transport in Lorentzian synthetic spaces, synthetic timelike Ricci curvature lower bounds and applications (arXiv:2004.08934)

[12] Geroch R 1970 Domain of dependence *J. Math. Phys.* **11** 437–49

[13] Grant J D E, Kunzinger M and Sämann C 2019 Inextendibility of spacetimes and Lorentzian length spaces *Ann. Glob. Anal. Geom.* **55** 133–47

[14] Griffiths J B and Podolský J 2009 *Exact Space-Times in Einstein's General Relativity* (*Cambridge Monographs on Mathematical Physics*) (Cambridge University Press)

[15] Hawking S and Bondi H 1969 The existence of cosmic time functions *Proc. R. Soc.* A **308** 433–5

[16] Hawking S and Ellis G F R 1973 *The Large Scale Structure of Space-Time* (*Cambridge Monographs on Mathematical Physics*) (Cambridge University Press)

[17] Husseinov F 2018 Extension of strictly monotonic functions in order-separable spaces (arXiv:ssrn.3260586)

[18] Kunzinger M and Steinbauer R 2021 Null distance and convergence of Lorentzian length spaces (arXiv:2106.05393)

[19] Kunzinger M and Sämann C 2018 Lorentzian length spaces *Ann. Glob. Anal. Geom.* **54** 399–447

[20] Levin V 1983 Continuous utility theorem for closed preorders on a metrizable σ-compact space *Dokl. Akad. Nauk SSSR* **273** 800–4 (in Russian) (available at: www.mathnet.ru/links/df6e362ddd9e48885ecdd043663315d0/dan9854.pdf)

[21] Manzano M and Mars M 2021 Null shells: general matching across null boundaries and connection with cut-and-paste formalism *Class. Quantum Grav.* **38** 155008, 34

[22] Mars M 2020 Hypersurface data: general properties and Birkhoff theorem in spherical symmetry *Mediterr. J. Math.* **17** 206, 45

[23] McCann R and Sämann C 2021 A Lorentzian analog for Hausdorff dimension and measure (arXiv:2110.04386)
[24] Minguzzi E 2010 Time functions as utilities *Commun. Math. Phys.* **298** 855–68
[25] Minguzzi E 2013 Normally preordered spaces and utilities *Order* **30** 137–50
[26] Nachbin L 1965 *Topology and Order* (*Vand Nostrand Mathematical Studies* vol 4) (Van Nostrand Co., Inc)
[27] O'Neill B 1983 *Semi-Riemannian Geometry With Applications to Relativity* (*Pure and Applied Mathematics* vol 103) (Academic, Inc. (Harcourt Brace Jovanovich, Publishers))
[28] Penrose R 1972 The geometry of impulsive gravitational waves *General Relativity (Papers in Honour of J. L. Synge)* (Clarendon) pp 101–15
[29] Podolský J, Sämann C, Steinbauer R and Švarc R 2019 Cut-and-paste for impulsive gravitational waves with Λ: the geometric picture *Phys. Rev.* D **100** 024040, 8
[30] Podolský J and Steinbauer R 2022 Penrose junction conditions with Λ: geometric insights into low-regularity metrics for impulsive gravitational waves *Gen. Relativ. Gravit.* **54** 96, 24
[31] Reina B, Senovilla Jé M M and Vera R 2016 Junction conditions in quadratic gravity: thin shells and double layers *Class. Quantum Grav.* **33** 105008, 41
[32] Sbierski J 2018 The c^0-inextendibility of the Schwarzschild spacetime and the spacelike diameter in Lorentzian geometry *J. Differ. Geom.* **108** 319–78
[33] Sormani C and Vega C 2016 Null distance on a spacetime *Class. Quantum Grav.* **33** 085001

Letters in Mathematical Physics (2023) 113:48
https://doi.org/10.1007/s11005-023-01668-w

The splitting theorem for globally hyperbolic Lorentzian length spaces with non-negative timelike curvature

Tobias Beran[1] · **Argam Ohanyan**[1] · **Felix Rott**[1] · **Didier A. Solis**[2]

Received: 10 November 2022 / Revised: 20 March 2023 / Accepted: 28 March 2023 /
Published online: 18 April 2023

Abstract

In this work, we prove a synthetic splitting theorem for globally hyperbolic Lorentzian length spaces with global non-negative timelike curvature containing a complete timelike line. Inspired by the proof for smooth spacetimes (Beem et al. in Global differential geometry and global analysis 1984, Springer, pp. 1–13, 1985), we construct complete, timelike asymptotes which, via triangle comparison, can be shown to fit together to give timelike lines. To get a control on their behaviour, we introduce the notion of parallelity of timelike lines in the spirit of the splitting theorem for Alexandrov spaces as proven in Burago et al. (A course in metric geometry, vol 33, American Mathematical Society, Providence, 2001) and show that asymptotic lines are all parallel. This helps to establish a splitting of a neighbourhood of the given line. We then show that this neighbourhood has the *timelike completeness* property and is hence inextendible by a result in Grant et al. (Ann Glob Anal Geom 55(1):133–147, 2019), which globalises the local result.

Keywords Lorentzian length spaces · Synthetic curvature bounds · Splitting theorems

✉ Didier A. Solis
didier.solis@correo.uady.mx

Tobias Beran
tobias.beran@univie.ac.at

Argam Ohanyan
argam.ohanyan@univie.ac.at

Felix Rott
felix.rott@univie.ac.at

1 Department of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Wien, Austria

2 Facultad de Matemáticas, Universidad Autónoma de Yucatán, Anillo Periférico, Tablaje 13615, Mérida, Mexico

Mathematics Subject Classification 53C50 · 53C23 · 53B30

Contents

1 Introduction

As originally formulated in the context of Riemannian geometry, splitting theorems are a class of results which establish that a Riemannian manifold (M, g) subject to certain hypotheses on its curvature and its geodesic structure actually *splits*, i.e. it is isometric to a product. In most cases, the curvature assumption involves a weak inequality (usually on sectional or Ricci curvature) while the geodesic assumption is related to the existence of a *line*, that is, a globally distance realising complete geodesic. Since lines are conjugate point free and positive curvature promotes the appearance of conjugate points, both features are expected to hold only under very special circumstances.

The very first attempt to establish the metric product structure of a Riemannian manifold with non-negative curvature and having a line is attributed to Cohn–Vossen [27], who proved that a non-compact complete manifold of dimension 2 with sectional curvature $K \geq 0$ is either diffeomorphic to $\mathbb{R}^2$ or flat. However, the proof of this result relied on the Gauss-Bonnet Theorem and therefore was not suitable for generalisation to arbitrary dimensions. The breakthrough that allowed such a generalisation was due to Toponogov, who in the mid 1960s proved his celebrated Triangle Comparison Theorem, which asserts that "the angles of an arbitrary triangle in a Riemannian space made up of minimising arcs are no greater than the corresponding angles of the plane Euclidean triangle with sides of the same length" [62]. Armed with this tool, Toponogov established the archetype of a splitting theorem [61, 62]:

Theorem 1.1 *Let (M, g) be a complete Riemannian manifold with $K \geq 0$ containing a line, then (M, g) is isometric to $(\mathbb{R}^k \times N, g_0 \oplus h)$ where $(\mathbb{R}^k, g_0)$ is the standard Euclidean space and (N, h) is complete and does not have any lines.*

Shortly after Toponogov's splitting theorem was published, Cheeger and Gromoll made substantial progress by generalising it under the less stringent condition $\mathrm{Ric} \geq 0$

[24]. Although their original motivation was rooted in the investigation of topological obstructions on manifolds of non-negative curvature [25], their result sparked the interest in the study of splitting theorems in many different contexts, like in the theory of orbifolds [14] or in the study of curvature inequalities relating other tensors, (like the curvature operator [56] or Bakry–Emery tensor [29, 64]). Furthermore, some more general versions of the Cheeger–Gromoll splitting theorem have been proven. For example, the curvature condition can be replaced by averaged Ricci inequalities [31], or by almost positivity of the Ricci tensor [21, 57]. There are also some local versions in which the splitting occurs either on tubular neighbourhoods or the complement of a compact set disjoint from a line [20, 22].

Remarkably, Yau posed in 1982 the problem of establishing a Lorentzian analogue of the Cheeger–Gromoll splitting theorem [65], thus inaugurating a very active field in which different versions of the Lorentzian splitting theorem were established [10, 11, 28, 30, 32, 55],[1] and some important problems have remained unsolved even to this day, like the famous Bartnik Conjecture [8, 33]. In analogy to its Riemannian counterparts, the first Lorentzian splitting theorem was proven for globally hyperbolic spacetimes with non-positive sectional curvature on timelike planes. In Beem et al. [10, 11] the authors proved what can be thought as the Lorentzian analogue of Toponogov's splitting theorem:

Theorem 1.2 *Let (M, g) be a spacetime of dimension $n \geq 2$ that satisfies the following conditions:*

(i) *(M, g) is globally hyperbolic.*
(ii) *$K(\Pi) \leq 0$ for all timelike planes Π.*
(iii) *M has a timelike line.*

Then (M, g) splits isometrically as $(M, g) \simeq (\mathbb{R} \times N, -dt^2 \oplus h)$, where (N, h) is a complete Riemannian manifold.

Notice that global hyperbolicity can be seen as a more suitable analogue of Riemannian completeness than timelike geodesic completeness, since the latter fails to guarantee connectability by distance realising geodesic segments. The original proof of this result not only uses techniques related to Toponogov's theorem—a Lorentzian triangle comparison due to Harris [43]—, but also tools used in Cheeger-Gromoll's theorem, like Busemann functions and Hessian estimates.

Some of the key techniques used in Toponogov's original proof can be abstracted into an axiomatic setting rather than derived, thus expanding its range of applications to non-smooth contexts, particularly, to the realm of metric length spaces. Recall that a length space (X, d) is a metric space whose distance is intrinsic, in other words, the distance $d(p, q)$ can be recovered as the infimum of the lengths of curves joining p to q. This is the setting for the so-called synthetic methods in geometry, that have proven instrumental in the development of recent mathematical landmarks such as the study of geometry in the large [41, 42], precompactness theorems [17, 18], geometric flows [6, 39] and optimal transport [60, 63]. An *Alexandrov space with curvature bounded from below by k*—or $\mathrm{Alex}(k)$ for short—is a locally compact, complete and path connected

[1] Refer to chapter 14 in [9] for a detailed account.

(metric) length space (X, d) on which the triangle comparison theorem holds [17]. For Alex(0) spaces an analogue of Theorem 1.1 was first proved by Milka [53]. In this result, instead of the existence of a line, the existence of an m-affine function, (i.e. a function $g: M \to \mathbb{R}$ that when restricted to any unit speed geodesics satisfies the differential equation $g'' + mg = 0$) is assumed. A warped product $I \times_g F$ with g an m-affine function is called a *cone* [18]. Cones naturally have m-affine functions, and conversely, an Alex(k) space with a non-constant m-affine function splits as a cone, provided that a boundary condition is met [3, 51]. In the particular case of a complete Riemannian manifold, Innami [46] showed that a 0-affine function exists if and only if M is isometric to a product $\mathbb{R} \times N$. Moreover, under the hypothesis of Theorem 1.1 the Busemann functions associated to a line are affine and thus Innami's theorem implies Toponogov's splitting theorem. Finally, let us note that in [17] there is a proof of Milka's splitting theorem for Alexandrov spaces that resembles more closely the original works of Toponogov. The precise statement is as follows:

Theorem 1.3 *Let (X, d) be an* Alex(0) *space containing a line. Then X splits as a metric product $\mathbb{R} \times Y$, where Y is an* Alex(0) *space.*

In their seminal work [48] Kunzinger and Sämann set the foundations for a synthetic approach to Lorentzian geometry. Their novel notion of Lorentzian (pre-)length space is suited to accommodate several different non-smooth scenarios such as cones [5, 54], spacetimes with C^0 metrics [26, 50], contact structures [45] or causal boundaries [1]. Moreover, in [48] the authors also introduced a notion of timelike curvature bounds in the same spirit of Alex(k) spaces. Among the recent developments in this fast growing field we have detailed analyses of the causal structure [2, 36], extendibility [35, 40], convergence [49], gluing techniques [12, 58] and the basis for a comparison theory [7, 13].

In this work, we aim at proving a splitting theorem for Lorentzian length spaces in the spirit of Theorems 1.1, 1.2 and 1.3. In precise terms, we establish the following splitting result for Lorentzian length spaces:

Theorem 1.4 (Splitting) *Let $(X, d, \ll, \leq, \tau)$ be a connected, regularly localisable, globally hyperbolic Lorentzian length space with proper metric d and global non-negative timelike curvature satisfying timelike geodesic prolongation and containing a complete timelike line $\gamma : \mathbb{R} \to X$. Then there is a τ and causality preserving homeomorphism $f : \mathbb{R} \times S \to X$, where S is a proper, strictly intrinsic metric space of Alexandrov curvature ≥ 0.*

The paper is organised as follows: In Sect. 2, we collect some well-known facts about Lorentzian (pre-)length spaces while focusing on (versions of) results best suited for our needs. We review the basic theory of Lorentzian (pre-)length spaces and discuss concepts like timelike curvature bounds, angles and extensions. Section 3 is dedicated to the study of product Lorentzian pre-length spaces of the form $\mathbb{R} \times X$, where (X, d) is a metric space. Section 4 is the heart of the proof and deals with the construction of co-rays and asymptotes, where we follow [11]. In order to have a control on the behaviour of lines formed as asymptotes to a given line, we introduce the notion of parallelity of timelike lines which is motivated by the splitting theorem for Alexandrov spaces as treated in [17] and show that asymptotic lines are always timelike, of infinite τ-length, and parallel to each other. In Sect. 5 we first establish a local splitting result by endowing a cross section S of the parallel asymptotes spanning $I(\gamma)$ with a natural metric and then showing that $I(\gamma)$ splits as $\mathbb{R} \times S$. Then we show that $I(\gamma)$ has the timelike completeness (TC) property, from which, via an inextendibility argument, $I(\gamma) = X$ follows. Finally, in Sect. 6 we note some classes of spacetimes which naturally satisfy the technical assumptions in our results, and hence split synthetically. We then give an outlook on open problems in the context of synthetic splitting theorems that may be addressed in future projects.

1.1 Notation and conventions

Let us collect some notation and conventions that will be used throughout the paper.

$A \subset B$ means A is a subset of B (not necessarily a proper one). $\mathbb{R}^{1,1}$ is two dimensional Minkowski space with the Lorentzian metric $-dt^2+dx^2$ and coordinates (t, x). $\overline{\tau}$ denotes the time separation on $\mathbb{R}^{1,1}$. A *proper* metric space (X, d) is a metric space such that all closed balls are compact. A metric space is called *length space* or *intrinsic* if the distance between two points is the infimum of lengths of curves running between them, and *strictly intrinsic* if that infimum is a minimum, i.e. between any two points there exists a distance-realising curve.

We denote an open ball with radius R around a point x in a metric space by $B_R(x)$, and the corresponding closed ball by $\bar{B}_R(x)$. Although in general this ball is not the closure $\overline{B_R(x)}$ of the open ball, notice that for length spaces and proper metric spaces we do have the equality $\bar{B}_R(x) = \overline{B_R(x)}$.

2 Basic theory of Lorentzian (pre-)length spaces

2.1 Lorentzian (pre-)length spaces

We give a brief review of the theory of Lorentzian (pre-)length spaces. We focus on the specific results that we need and give proofs if they cannot be found in the literature, otherwise we give precise references. In particular, we refer to [48] for a detailed treatment. We start with a few basic definitions.

Definition 2.1 (*Lorentzian pre-length spaces*) A *Lorentzian pre-length space* $(X, d, \ll, \leq, \tau)$ consists of a metric space (X, d), a reflexive and transitive relation $\leq$ (the *causal relation*), a transitive relation $\ll$ (the *timelike relation*) contained in $\leq$, and a lower semi-continuous map (*time separation*) $\tau : X \times X \to [0, \infty]$ with the following properties: $\tau(x, y) = 0$ if $x \not\leq y$, and $\tau(x, y) > 0$ if and only if $x \ll y$. Moreover, if $x \leq y \leq z$, then the following *reverse triangle inequality* holds:

$$\tau(x, z) \geq \tau(x, y) + \tau(y, z).$$

In the synthetic theory, we employ the usual nomenclature and well-known notation from Lorentzian geometry, such as $I^{\pm}(x)$, $J^{\pm}(x)$ for timelike and causal pasts and futures, as well as $I(x, y) := I^+(x) \cap I^-(y)$ and $J(x, y) := J^+(x) \cap J^-(y)$ for timelike and causal diamonds, respectively.

Lemma 2.2 (Openness of timelike futures and push-up) *Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space and $x \ll y \leq z$ or $x \leq y \ll z$. Then $x \ll z$. In particular, for any $x \in X$, the sets $I^{\pm}(x)$ are open. Moreover, the relation $\ll$ is open in $X \times X$.*

Proof See [48, Lem. 2.10, Lem. 2.12, Prop. 2.13]. □

Definition 2.3 (*Causal curves*) Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space. A non-constant curve $\gamma : I \to X$, I an interval, is called *future directed causal* (resp. *future directed timelike*) if it is locally Lipschitz and for all $t_1, t_2 \in I$ with $t_1 < t_2$ we have that $\gamma(t_1) \leq \gamma(t_2)$ (resp. $\gamma(t_1) \ll \gamma(t_2)$). It is called *future directed null* if it is future directed causal and no two points on the curve are $\ll$-related. *Past directed causal/timelike/null* curves are defined dually. From now on, unless explicitly stated otherwise, we assume all causal curves to be future directed.

We will sometimes deal with causal curves which are not locally Lipschitz. For these curves, some important results (e.g. the limit curve theorem) will not hold in general, however, many important properties (especially those that are purely causal-theoretic in nature) will continue to be true.

Definition 2.4 (*Continuous causal curves*) Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space. A non-constant curve $\gamma : I \to X$, I an interval, is called a *future directed continuous causal curve* (resp. *future directed continuous timelike curve*) if it is continuous and for all $t_1, t_2 \in I$ with $t_1 < t_2$ we have that $\gamma(t_1) \leq \gamma(t_2)$ (resp. $\gamma(t_1) \ll \gamma(t_2)$). It is called a *future directed continuous null curve* if it is a future directed continuous causal curve and no two points on the curve are $\ll$-related. Past versions of these notions are defined dually. From now on, unless explicitly stated otherwise, we assume all continuous causal curves to be future directed.

We emphasise that a *causal/timelike/null curve* is always understood to be locally Lipschitz. If *continuous* causal/timelike/null curves are meant, we will state that explicitly. Wherever possible, we will state definitions and results for continuous causal curves instead of (locally Lipschitz) causal curves for greater generality. In most cases, the proofs given in the literature for these results apply word for word to the continuous case.

Definition 2.5 (*Extensions of causal curves*) Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space and let $\gamma : I \to X$ be a continuous causal curve, where I is any interval. We say that γ is *(continuously) extendible* if there exists a (continuous) causal curve $\tilde{\gamma} : J \to X$, J an interval with $J \supsetneq I$, such that $\tilde{\gamma}|_I = \gamma$. Otherwise, we say γ is *(continuously) inextendible*. We refer to (in)extendibility to points in the future resp. past directions of γ as *future/past (in)extendibility*.[2]

[2] Geodesics defined on intervals $[a, b]$ with $a = -\infty$ or $b = \infty$ should be thought of as properly reparametrised.

Definition 2.6 (*τ-length*) Let $\gamma : [a, b] \to X$ be a future directed continuous causal curve. Then its *τ-length* is defined by

$$L_\tau(\gamma) := \inf \left\{ \sum_{i=0}^{N-1} \tau(\gamma(t_i), \gamma(t_{i+1})) \mid a = t_0 < t_1 < \cdots < t_N = b \right\}.$$

If γ is past directed causal, then $L_\tau(\gamma)$ is defined analogously. For continuous causal curves defined on half-open intervals, e.g. $[a, b)$, one takes the limit of $L_\tau(\gamma|_{[a,c]})$ as $c \to b$, and similarly in the other cases.

The τ-length of a continuous causal curve is invariant under reparametrisations. We will have more to say on these topics later.

Definition 2.7 (*Maximising causal curves*) A future directed (continuous) causal curve $\gamma : [a, b] \to X$ is called *maximising* (or *τ-maximising*) if $\tau(\gamma(a), \gamma(b)) = L_\tau(\gamma)$, and analogously for past directed (continuous) causal curves. We also refer to such curves as *(continuous) distance realisers*.

Lorentzian pre-length spaces where between each pair of causally related points there is a maximising (locally Lipschitz) causal curve are referred to as *strictly intrinsic* or *geodesic*.

Next, we define some of the steps on the causal ladder for Lorentzian pre-length spaces.

Definition 2.8 (*Causality conditions*) A Lorentzian pre-length space $(X, d, \ll, \leq, \tau)$ is called

(i) *non-totally imprisoning* if for every compact set $K \subset X$ there is $C > 0$ such that the d-lengths of all causal curves in K is bounded by C,

(ii) *strongly causal* if the Alexandrov topology generated by the subbasis of timelike diamonds $\{I(x, y) \mid x, y \in X\}$ coincides with the metric topology,

(iii) *globally hyperbolic* if it is non-totally imprisoning and all *causal diamonds* $J(x, y)$ are compact.

Lemma 2.9 *Let $(X, d, \ll, \leq, \tau)$ be a strongly causal Lorentzian pre-length space. Then each $x \in X$ has a neighbourhood basis consisting of causally convex*[3] *neighbourhoods. Even more is true: X has a topological basis of causally convex neighbourhoods.*

Proof This is trivial since by definition finite intersections of timelike diamonds are a basis for the topology, but these sets are causally convex. □

Definition 2.10 (*Causal path-connectedness*) A Lorentzian pre-length space X is called *causally path-connected* if for all $x < y$ there is a causal curve from x to y and for all $x \ll y$ there is a timelike curve from x to y.

[3] A set $V \subset X$ is called *causally convex* if for all $p, q \in V$ we have $J(p, q) \subset V$.

For the next definition, we use the version given in [2, Def. 2.16] rather than [48, Def. 3.4] as it more closely resembles the case of smooth spacetimes. Before doing so, note that on causally path-connected Lorentzian pre-length spaces X, we can introduce a *local causal relation*: Let $U \subset X$ be open and $x, y \in U$, then we say $x \leq_U y$ if there exists a future causal curve from x to y entirely contained in U.

Definition 2.11 (*Local causal closedness*) Let X be causally path-connected and $U \subset X$ be open. We say that U is *causally closed* if for all sequences p_n, q_n in U converging to $p, q \in U$ and $p_n \leq_U q_n$, we have that $p \leq_U q$. We say that X is *locally causally closed* if every point in X has a causally closed neighbourhood.

Definition 2.12 (*Global causal closedness*) A Lorentzian pre-length space X is called *globally causally closed* if $\leq$ is a closed relation.

Recall that in the theory of smooth spacetimes, a globally causally closed and causal spacetime is *causally simple*. Such spacetimes are one step below global hyperbolicity in the classical causal ladder.

Definition 2.13 (*d-compatibility*) A Lorentzian pre-length space $(X, d, \ll, \leq, \tau)$ is called *d-compatible* if for every $x \in X$ there is a neighbourhood U of x and a constant $C > 0$ such that the d-lengths of all causal curves contained in U are bounded above by C.

Lemma 2.14 *Let $(X, d, \ll, \leq, \tau)$ be a locally causally closed, d-compatible, globally hyperbolic Lorentzian pre-length space. Then no compact set in X contains an inextendible causal curve.*

Proof See [48, Cor. 3.15]. □

Definition 2.15 (*Localisability*) A Lorentzian pre-length space $(X, d, \ll, \leq, \tau)$ is called *localisable* if for each $x \in X$ there exists a neighbourhood (called *localising neighbourhood*) Ω_x containing x with the following properties:

(i) Causal curves in Ω_x have uniformly bounded d-length, i.e., Ω_x is d-compatible.
(ii) For each $y \in \Omega_x$, $I^{\pm}(y) \cap \Omega_x \neq \emptyset$.
(iii) There is a continuous map (*local time separation*) $\omega_x : \Omega_x \times \Omega_x \to [0, \infty)$ such that Ω_x is a Lorentzian pre-length space upon restricting $d, \ll, \leq$.
(iv) For all $p, q \in \Omega_x$ with $p < q$ there exists a causal curve γ_{pq} from p to q entirely in Ω_x with maximal τ-length among all causal curves from p to q contained in Ω_x, as well as $L_\tau(\gamma_{pq}) = \omega_x(p, q)$.

Note that this necessarily means that ω_x is of the form

$$\omega_x(p, q) := \max\{L_\tau(\gamma) \mid \gamma \text{ is a causal curve from } p \text{ to } q \text{ in } \Omega_x.\}. \tag{2.1}$$

If in addition, for $p, q \in \Omega_x$ with $p \ll q$ the curve γ_{pq} is timelike and strictly longer than any causal curve from p to q in Ω_x containing a null segment, then Ω_x is called a *regular localising neighbourhood*, and X *regularly localisable* if any point has a regular localising neighbourhood. Moreover, X is called *strongly localisable* if each point has a neighbourhood base of localising neighbourhoods and *SR-localisable* if each point has a neighbourhood base of regular localising neighbourhoods.

Lemma 2.16 *Let $(X, d, \ll, \leq, \tau)$ be a strongly causal Lorentzian pre-length space.*

(i) *If X is localisable, then it is strongly localisable.*
(ii) *If X is regularly localisable, then it is SR-localisable.*

Proof By Lemma 2.9, each point in X has a neighbourhood basis of causally convex neighbourhoods. Now suppose X is localisable, take $x \in X$ and let Ω_x be a localising neighbourhood of x and let $U \subset \Omega_x$ be any neighbourhood of x. Then there is a causally convex neighbourhood V of x such that $V \subset U$. We are done if we show that V is a localising neighbourhood. But this is easy to see since any causal curve starting and ending in V is entirely contained in there. If X is regularly localisable, then V is even a regular localising neighbourhood. □

Regularly localisable spaces have the following important property:

Theorem 2.17 (Causal character of maximising causal curves) *In a regularly localisable Lorentzian pre-length space, maximising causal curves are either timelike or null.*

Proof See [48, Thm. 3.18]. □

Note that in localisable Lorentzian pre-length spaces, the sets $I^{\pm}(y)$ are never empty for any $y \in X$.

Definition 2.18 (*Geodesic*) Let X be a localisable Lorentzian pre-length space and $\gamma : I \to X$ a (continuous) causal curve. Then γ is a (future-directed causal) *(continuous) geodesic* if for each $t_0 \in I$ there exists a localising neighbourhood Ω of $\gamma(t_0)$ and a neighbourhood $[c, d] \subset I$ of t_0 such that $\gamma([c, d]) \subset \Omega$ and $\gamma|_{[c,d]}$ is maximising in Ω with respect to the local time separation.

We will sometimes need the notion of *geodesic (in)extendibility* which is defined in terms of the existence of a geodesic extending a given geodesic. Note that inextendibility implies geodesic inextendibility but the converse need not hold in general.

For many arguments in the proof of the splitting theorem, the property that geodesics can always be extended to open domains will be essential. While this is a consequence of the ODE theory of the geodesic equation in the case of smooth spacetimes, we will need to assume it here:

Definition 2.19 (*Timelike geodesic prolongation*) A localisable Lorentzian pre-length space X is said to have the *timelike geodesic prolongation* property if any maximising timelike segment $\gamma : [a, b] \to X$ can be extended to a timelike geodesic with an open domain, i.e. there is an $\varepsilon > 0$ and a timelike geodesic $\bar{\gamma} : (a - \varepsilon, b + \varepsilon) \to X$ with $\bar{\gamma}|_{[a,b]} = \gamma$.

It is clear that this is satisfied in smooth strongly causal spacetimes considered as a Lorentzian pre-length space. One easily sees that spacetimes with timelike boundary do not satisfy this.

Definition 2.20 (*Locally quasiuniformly maximising*) A localisable Lorentzian pre-length space X is called *locally quasiuniformly maximising* if for each $x \in X$ there exists a localisable neighbourhood U containing x such that each causal geodesic entirely contained in U is maximising in U with respect to the local time separation.

It is unclear to us how the property of being locally quasiuniformly maximizing is related to other notions for Lorentzian pre-length spaces, such as being intrinsic, strongly causal or (locally) causally closed.

Proposition 2.21 (Upper semicontinuity of Lorentzian arclength) *Let $(X, d, \ll, \leq, \tau)$ be a strongly causal and localisable Lorentzian pre-length space. Then L_τ is upper semi-continuous, i.e. if $\gamma_n : [a, b] \to X$ are causal curves converging d-uniformly to a causal curve $\gamma : [a, b] \to X$, then*

$$L_\tau(\gamma) \geq \limsup_{n\to\infty} L_\tau(\gamma_n).$$

Proof See [48, Prop. 3.17]. □

Proposition 2.22 *Let X be a strongly causal, localisable, locally quasiuniformly maximising Lorentzian pre-length space and let $\gamma : [a, b) \to X$ be a causal geodesic. Then γ is geodesically extendible to $[a, b]$ if and only if it is extendible as a continuous curve to $[a, b]$.*

Proof See [40, Prop. 4.6] and note that the proof really requires the space to be locally quasiuniformly maximising. □

We now state the limit curve theorem in the case of strongly causal Lorentzian pre-length spaces, which will be sufficient for our needs. Note that in this case, by [2, Prop. 2.21], Definition 2.11 is equivalent to [48, Def. 3.4].

Theorem 2.23 (Limit curve theorem) *Let $(X, d, \ll, \leq, \tau)$ be a locally causally closed, localisable, strongly causal, d-compatible Lorentzian pre-length space with proper metric d. Let $\gamma_n : [0, L_n] \to X$ be a sequence of causal curves parametrised by d-arclength, and suppose that $L_n = L^d(\gamma_n) \to \infty$. If the sequence $\gamma_n(0)$ has an accumulation point $x \in X$, then some subsequence of the γ_n converges locally uniformly to a future inextendible causal curve $\gamma : [0, \infty) \to X$. Moreover, if the γ_n are maximising, then so is γ. In this case, for any $T > 0$, $L_\tau(\gamma|_{[0,T]}) = \lim_n L_\tau(\gamma_n|_{[0,T]})$.*

Proof The statement about the existence of the limit curve is shown in [48, Thm. 3.14]. Now suppose that γ_n are maximising and w.l.o.g. let γ_n itself converge locally uniformly to γ. Consider any $T > 0$. Then by upper semicontinuity of L_τ and lower semicontinuity of τ,

$$\begin{aligned}\limsup_n L_\tau(\gamma_n|_{[0,T]}) &\leq L_\tau(\gamma|_{[0,T]}) \leq \tau(\gamma(0), \gamma(T)) \\ &\leq \liminf_n \tau(\gamma_n(0), \gamma_n(T)) = \liminf_n L_\tau(\gamma_n|_{[0,T]}).\end{aligned}$$

Thus, equality holds everywhere. In particular, γ is maximising and $\lim_n L_\tau(\gamma_n|_{[0,T]}) = L_\tau(\gamma|_{[0,T]})$. □

In the context of the splitting theorem, we are interested in Lorentzian pre-length spaces with significantly more structure, which we will discuss now.

Definition 2.24 (*Lorentzian length spaces*) A locally causally closed, causally path-connected and localisable Lorentzian pre-length space $(X, d, \ll, \leq, \tau)$ is called *Lorentzian length space* if for all $x \leq y$, $x \neq y$,

$$\tau(x, y) = \sup\{L_\tau(\gamma) \mid \gamma \text{ future directed causal from } x \text{ to } y\}.$$

Note that since by definition $L_\tau(\gamma) \leq \tau(p, q)$ for any (even continuous) causal curve γ from p to q, one can equivalently take the supremum over continuous causal curves to obtain the time separation in the case of Lorentzian length spaces.

Proposition 2.25 (Causality of Lorentzian length spaces)

(i) *A Lorentzian length space is strongly causal if and only if each point has a neighbourhood base of causally convex neighbourhoods.*
(ii) *Strongly causal Lorentzian length spaces are non-totally imprisoning.*
(iii) *Globally hyperbolic Lorentzian length spaces are strongly causal. Moreover, τ is continuous and finite and for any two causally related points there exists a maximising causal curve connecting them.*

Proof See [48, Thm. 3.26, Thm. 3.28, Thm. 3.30]. □

Lemma 2.26 *Let $(X, d, \ll, \leq, \tau)$ be a strongly causal Lorentzian pre-length space such that for each $x \in X$ there exists a timelike curve $\gamma : [-\varepsilon, \varepsilon] \to X$ with $\gamma(0) = x$. Then $\{I(x, y) \mid x, y \in X\}$ is a base for the topology on X. Moreover, for each $x \in X$, each timelike curve $\gamma : [-\varepsilon, \varepsilon] \to X$ with $\gamma(0) = x$ and parameters $0 < s_n, t_n \to 0$, $\{I(\gamma(-s_n), \gamma(t_n)) \mid n \in \mathbb{N}\}$ is a neighbourhood base at x.*

Proof Let $x \in X$ with $x \in U$, $U \subset X$ open. By the definition of the Alexandrov topology, finite intersections of the sets $I(y, z)$ with $y, z \in X$ are a base. Thus, there exist $x_1, \ldots, x_n, y_1, \ldots, y_n \in X$ such that

$$x \in V := I(x_1, y_1) \cap \cdots \cap I(x_n, y_n) \subset U.$$

Now let $\gamma : [-\varepsilon, \varepsilon] \to X$ be a timelike curve with $\gamma(0) = x$. By continuity, maybe after making ε smaller, we may assume that $\gamma([-\varepsilon, \varepsilon]) \subset V$. Then $W := I(\gamma(-\varepsilon), \gamma(\varepsilon))$ is an open set containing x and we claim that $W \subset V$. To see this, let $z \in W$, then for any $i = 1, \ldots, n$ we have $x_i \ll \gamma(-\varepsilon) \ll z$, hence $z \in I^+(x_i)$. Similarly, $z \ll \gamma(\varepsilon) \ll y_i$, hence $z \in I^-(y_i)$. Thus $z \in V$. This proves that $\{I(x, y) \mid x, y \in X\}$ is a base for the topology of X, and the second claim then easily follows from the arguments above. □

Lorentzian pre-length spaces where each point lies in the interior of a timelike curve are future and past approximating in the sense of [19, Def. 2.17] (see also [19, Lem. 2.18])

Lemma 2.27 (Geodesics in strongly causal Lorentzian length spaces) *Let $(X, d, \ll, \leq, \tau)$ be a strongly causal Lorentzian length space. A (continuous) causal curve $\gamma : I \to X$ is a (continuous) geodesic if and only if for each $t_0 \in I$ there exists a subinterval $[c, d] \subset I$ with $t_0 \in [c, d]$ such that $\gamma|_{[c,d]}$ is τ-maximising.*

Proof This follows from [40, Lem. 4.3]. □

To conclude this subsection, we discuss d-arclength and τ-arclength reparametrisations of causal curves.

Remark 2.28 (On d-arclength parametrisations) Let $(X, d, \ll, \leq, \tau)$ be a globally hyperbolic Lorentzian length space with proper metric d and $\gamma : [a, b) \to X$ a causal curve that is inextendible at b. Then γ has infinite d-length: By Lemma 2.14, γ has to leave every compact set, in particular it leaves each compact ball $\overline{B}_n(\gamma(a))$, which shows that $L^d(\gamma) = \infty$. If we reparametrise γ by d-arclength (note that this is in general not possible for continuous causal curves), it is thus defined on $[0, \infty)$. Similarly, if we have a doubly inextendible causal curve $\tilde{\gamma} : (a, b) \to X$, then upon reparametrising it by d-arclength, it is defined on $\mathbb{R}$, so that $L^d(\tilde{\gamma}|_{[c,d]}) = |d - c|$ (see also [19, Rem. 2.5]). Note that d-arclength parametrisations are globally 1-Lipschitz.

Lemma 2.29 (τ-arclength parametrisations for inextendible curves) *Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space with continuous time separation τ which satisfies $\tau(x, x) = 0$ for all $x \in X$. Let $\gamma : [a, b) \to X$ be a future inextendible continuous timelike curve defined on a half-open interval with $L_\tau(\gamma|_{[a,c]}) < \infty$ for each $c \in [a, b)$ (notice we might have $L_\tau(\gamma) = \infty$). Then the map $\phi : [a, b) \to [0, L_\tau(\gamma))$, $t \mapsto L_\tau(\gamma|_{[a,t]})$ is continuous and strictly increasing. Moreover, $\tilde{\gamma} := \gamma \circ \phi^{-1} : [0, \infty) \to X$ is a weak reparametrisation of γ with respect to τ-arclength. Similarly, a doubly inextendible continuous timelike curve $\gamma : (a, b) \to X$ can be weakly parametrised by τ-arclength and is then defined on $\big(-L_\tau(\gamma|_{(a,t_0]}), L_\tau(\gamma|_{[t_0,b)})\big)$, where $t_0 \in (a, b)$ is arbitrary.*

Proof This can be seen by applying [48, Lem. 3.33, Lem. 3.34] to each $\gamma|_{[a,\tilde{b}]}$ with $\tilde{b} < \infty$. □

Note that, in general, the τ-arclength parametrisation results in a continuous timelike curve, even if the original curve is locally Lipschitz.

In the following, whenever τ-arclength parametrised curves are mentioned, it is always implicitly understood that those curves have finite τ-length on compact subintervals. In localisable spaces this is anyway the case.

2.2 Timelike curvature bounds

Establishing an analogue to sectional curvature bounds in semi-Riemannian manifolds á la [4] was one of the main goals in the first work on Lorentzian length spaces [48]. Historically, splitting theorems always used some type of curvature condition, either Ricci or sectional curvature. By M_k we denote the two-dimensional Lorentzian model space of constant sectional curvature k and, in this subsection only, $\bar{\tau}$ denotes the time separation in M_k (in general $\bar{\tau}$ will only be used as the time separation in $M_0 = \mathbb{R}^{1,1}$). For more details on this, see [48, Definition 4.5]. We will always assume that all mentioned triangles satisfy size bounds for M_k, cf. [48, Lemma 4.6]. However, since we are mainly working with $k = 0$, we do not have to worry about this.

We call three points $x_1, x_2, x_3 \in X$ that are timelike related together with maximisers between them a *timelike triangle*. We denote such triangles by $\Delta(x_1, x_2, x_3)$.

Definition 2.30 Let X be a Lorentzian pre-length space and $\Delta(x_1, x_2, x_3)$ be a timelike triangle. Then a timelike triangle $\Delta(\bar{x}_1, \bar{x}_2, \bar{x}_3)$ in M_k such that $\tau(x_i, x_j) = \bar{\tau}(\bar{x}_i, \bar{x}_j)$ is called a *comparison triangle* for $\Delta(x_1, x_2, x_3)$. It always exists (if the size bounds are satisfied) and is unique up to isometries of M_k.

Definition 2.31 (*Timelike curvature bounds*) A Lorentzian pre-length space X is said to satisfy a *timelike curvature bound from below by k*, if each point in X has a so-called comparison neighbourhood U, satisfying the following:

(i) $\tau|_{U\times U}$ is finite and continuous.
(ii) For all $x, y \in U$ with $x \ll y$ there exists a causal maximiser γ from x to y which is entirely contained in U.
(iii) Let $\Delta(x, y, z)$ be a triangle in U and let $\Delta(\bar{x}, \bar{y}, \bar{z})$ be its comparison triangle in M_k. Let $p, q \in \Delta(x, y, z)$ and let $\bar{p}, \bar{q}$ be the corresponding comparison points in $\Delta(\bar{x}, \bar{y}, \bar{z})$. Then $\tau(p, q) \leq \bar{\tau}(\bar{p}, \bar{q})$.

For curvature bounds from above, the inequality in (iii) is reversed, but we will not need these types of curvature bounds. If $k = 0$, we also say that X has *non-negative timelike curvature* resp. *non-positive timelike curvature*.

Evidently, curvature bounds are only formulated locally as of yet. There is, however, a very natural way of globalising this concept.

Definition 2.32 (*Global curvature bound*) A Lorentzian pre-length space X is said to have *global curvature bounded below by k* if X itself is a comparison neighbourhood in the sense of Definition 2.31.

Remark 2.33 (Continuous triangles vs. Lipschitz triangles) If X is a globally hyperbolic Lorentzian length space with global nonnegative timelike curvature, then in fact curvature comparison even holds for timelike triangles where the maximisers are continuous: Indeed, suppose $\Delta := \Delta_{C^0}(x, y, z)$ is a continuous timelike triangle, and let $p, q \in \Delta$. Due to the properties of X, we find Lipschitz maximisers from the endpoints of Δ to p, q, respectively. The concatenations at p and q, respectively, of those maximisers are again maximisers because the sides on Δ are (continuous) maximisers. Hence we have realised p, q on a Lipschitz triangle. If X is in addition regularly localisable, then in fact every continuous timelike maximiser is Lipschitz reparametrisable [52], so the above discussion is superfluous in this case.

One of the most commonly used implications of lower (timelike) curvature bounds is the prohibition of branching of distance-realisers. A formulation of this result for Lorentzian pre-length spaces was first given in [48, Theorem 4.12]. However, with the introduction of hyperbolic angles in [13] it was possible to generalise this result by omitting some of the additional assumptions:

Theorem 2.34 (Timelike non-branching) *Let X be a strongly causal Lorentzian pre-length space with timelike curvature bounded below by $k \in \mathbb{R}$. Then timelike distance realisers cannot branch, i.e., if $\alpha, \beta : [-\varepsilon, \varepsilon] \to X$ are timelike distance realisers such that there exists $t_0 \in (-\varepsilon, \varepsilon)$ with $\alpha|_{[-\varepsilon,t_0]} = \beta|_{[-\varepsilon,t_0]}$, then $\alpha = \beta$.*

Proof See [13, Theorem 4.7]. □

The non-branching of timelike distance realisers is a key property of spaces with lower curvature bounds and will appear in various forms in our proof of the splitting theorem.

2.3 Angles and comparison angles

Hyperbolic angles in Lorentzian pre-length spaces were introduced in [7, 13], where the former puts a bigger focus on comparison results. We will follow the conventions of the latter reference.

Lemma 2.35 (The law of cosines ($k = 0$)) *Let $X = \mathbb{R}^{1,1}$ be 2D-Minkowski space and x_1, x_2, x_3 be three points which are timelike related (an (unordered) timelike triangle). Let $a_{ij} = \max(\tau(x_i, x_j), \tau(x_j, x_i))$ (note one of these is zero anyway). Let ω be the hyperbolic angle between the straight lines x_1x_2 and x_2x_3 at x_2. Set $\sigma = 1$ if x_2 is not a time endpoint of the triangle (i.e., $x_1 \ll x_2 \ll x_3$ or $x_3 \ll x_2 \ll x_1$) and $\sigma = -1$ if x_2 is a time endpoint of the triangle (i.e., $x_2 \ll x_1, x_3$ or $x_1, x_3 \ll x_2$). Then we have:*

$$a_{13}^2 = a_{12}^2 + a_{23}^2 + \sigma a_{12}a_{23} \cosh(\omega).$$

In particular, when only changing one side-length, the angle ω is a monotonously increasing function of the longest side-length and monotonously decreasing in the other side-lengths.

Proof See [13, Appendix A]. □

Note that for $a_{ij} > 0$ satisfying a reverse triangle inequality and choosing an appropriate $\sigma = \pm 1$, we can always solve this equation for ω.

Definition 2.36 (*Comparison angles*) Let X be a Lorentzian pre-length space and x_1, x_2, x_3 three timelike related points. Let $\bar{x}_1, \bar{x}_2, \bar{x}_3 \in \mathbb{R}^{1,1}$ be a comparison triangle for x_1, x_2, x_3. We define the *comparison angle* $\tilde{\measuredangle}_{x_2}(x_1, x_3)$ as the hyperbolic angle between the straight lines $\bar{x}_1\bar{x}_2$ and $\bar{x}_2\bar{x}_3$ at $\bar{x}_2$. It can be calculated with the law of cosines using $a_{ij} = \max(\tau(x_i, x_j), \tau(x_j, x_i))$ and σ, where we set $\sigma = 1$ if x_2 is not a time endpoint of the triangle and $\sigma = -1$ if x_2 is a time endpoint of the triangle. σ is called the *sign* of the comparison angle (even though we always have $\tilde{\measuredangle}_{x_2}(x_1, x_3) > 0$). For a reduction of the number of case distinctions we also define the *signed comparison angle* $\tilde{\measuredangle}^{\mathrm{S}}_{x_2}(x_1, x_3) = \sigma \tilde{\measuredangle}_{x_2}(x_1, x_3)$.

Definition 2.37 (*Angles*) Let X be a Lorentzian pre-length space and $\alpha, \beta : [0, \varepsilon) \to X$ be two timelike curves (future or past directed or one of each) with $x := \alpha(0) =$

$\beta(0)$. Then we define the *upper angle*

$$\measuredangle_x(\alpha, \beta) = \limsup_{(s,t)\in D, s,t\to 0} \tilde{\measuredangle}_x(\alpha(s), \beta(t)),$$

where $D = \{(s,t) : s, t > 0,\ \alpha(s), \beta(t)$ timelike related$\}$. If the limes superior is a limit and finite, we say *the angle exists* and call it an *angle*.

Note that the sign of the comparison angle is independent of $(s,t) \in D$. We define the *sign* of the (upper) angle σ to be that sign, and define the *signed (upper) angle* to be $\measuredangle_x^{\mathrm{S}}(\alpha, \beta) = \sigma \measuredangle_x(\alpha, \beta)$.

If maximisers between any two timelike related points are unique (as e.g. in $\mathbb{R}^{1,1}$), then we simply write $\measuredangle_p(x, y)$ for the angle at p between the maximisers from p to x and p to y.

We give an alternative definition of timelike curvature bounds in the case $k = 0$ following Definition 4.9 in [13]:

Definition 2.38 Let X be a Lorentzian pre-length space. We say that X has *timelike curvature bounded below by $k = 0$ (bounded above by $k = 0$) in the sense of monotonicity comparison* if every point in X has a neighbourhood U such that

- (i) $\tau|_{U\times U}$ is finite and continuous.
- (ii) U is *timelike geodesically connected*, i.e. whenever $x, y \in U$ with $x \ll y$, there exists a future-directed timelike distance realiser α in U from x to y, and any distance realiser from x to y contained in U is timelike.
- (iii) Whenever $\alpha : [0, a] \to U$, $\beta : [0, b] \to U$ are timelike distance realisers in U with $x := \alpha(0) = \beta(0)$, we define the function $\theta : (0, a] \times (0, b] \supset D \to \mathbb{R}$ by $\theta(s,t) := \tilde{\measuredangle}_x^{\mathrm{S}}(\alpha(s), \beta(t))$, where $(s,t) \in D$ precisely when $\alpha(s), \beta(t)$ are timelike related. We require this to be monotonically increasing (decreasing) in s and t. Additionally, in the case of non-positive timelike curvature, we consider $0 < s' \leq s$ and $0 < t' \leq t$ and require that if $\theta(s', t')$ is not defined (i.e. $\alpha(s'), \beta(t')$ are not timelike related) but $\theta(s,t)$ is, then also the comparison points for $\alpha(s'), \beta(t')$ in a comparison triangle for $\Delta(x, \alpha(s), \beta(t))$ are not timelike related.

We call the curvature bound *global* if X is a comparison neighbourhood in the above sense.

Theorem 2.39 (Timelike curvature: Equivalence of definitions) *Let X be a locally timelike geodesically connected Lorentzian pre-length space. Then X has non-negative (non-positive) timelike curvature in the sense of Definition 2.31 if and only if it has non-negative (non-positive) timelike curvature in the sense of monotonicity comparison (see Definition 2.38).*

Proof See [13, Theorem 4.12]. □

Remark 2.40 Note that if X is *globally timelike geodesically connected*, i.e. between any two timelike related points there is a timelike distance realiser connecting them, and any distance realiser connecting them is timelike (which is certainly true if X is a regularly localisable and globally hyperbolic Lorentzian length space), then X has global non-negative (non-positive) timelike curvature in the sense of Definition 2.32 if and only if it has global non-negative (non-positive) timelike curvature in the sense of monotonicity comparison.

The timelike geodesic prolongation property plays a technical role in our proof of the splitting theorems mainly because of the following result.

Proposition 2.41 (Continuity of angles in spaces with timelike curvature bounded below) *Let X be a strongly causal, localisable, timelike geodesically connected and locally causally closed Lorentzian pre-length space with global non-negative timelike curvature and which satisfies timelike geodesic prolongation. Let $\alpha_n, \alpha, \beta_n, \beta$ be future or past directed timelike geodesics all starting at $\alpha_n(0) = \alpha(0) = \beta_n(0) = \beta(0) =: x$ and with $\alpha_n \to \alpha$ and $\beta_n \to \beta$ pointwise (in particular, α_n is future directed if and only if α is, and similarly for β_n and β). Then*

$$\measuredangle_x(\alpha, \beta) = \lim_n \measuredangle_x(\alpha_n, \beta_n)$$

Proof See [13, Proposition 4.14]. □

The following two propositions will be proven together.

Proposition 2.42 (Alexandrov lemma: across version) *Let X be a Lorentzian pre-length space. Let $\Delta := \Delta(x, y, z)$ be a timelike triangle (in particular the distance realisers between the endpoints exist). Let p be a point on the side xz with $p \ll y$, such that the distance realiser between p and y exists. Then we can consider the smaller triangles $\Delta_1 := \Delta(x, p, y)$ and $\Delta_2 := \Delta(p, y, z)$. We construct a comparison situation consisting of a comparison triangle $\bar{\Delta}_1$ for Δ_1 and $\bar{\Delta}_2$ for Δ_2, with $\bar{x}$ and $\bar{z}$ on different sides of the line through $\bar{p}\bar{y}$ and a comparison triangle $\tilde{\Delta}$ for Δ with a comparison point $\tilde{p}$ for p on the side xz. This contains the subtriangles $\tilde{\Delta}_1 := \Delta(\tilde{x}, \tilde{y}, \tilde{p})$ and $\tilde{\Delta}_2 := \Delta(\tilde{p}, \tilde{y}, \tilde{z})$, see Fig. 1.*

Then the situation $\bar{\Delta}_1$, $\bar{\Delta}_2$ is convex (concave) at p (i.e. $\tilde{\measuredangle}_p(x, y) = \measuredangle_{\bar{p}}(\bar{x}, \bar{y}) \geq \measuredangle_{\bar{p}}(\bar{y}, \bar{z}) = \tilde{\measuredangle}_p(y, z)$ (or $\leq$)) if and only if $\tau(p, y) \leq \bar{\tau}(\bar{p}, \bar{y})$ (or $\geq$). If this is the case, we have that

- *each angle in the triangle $\bar{\Delta}_1$ is $\geq$ (or $\leq$) than the corresponding angle in the triangle $\tilde{\Delta}_1$,*
- *each angle in the triangle $\bar{\Delta}_2$ is $\geq$ (or $\leq$) than the corresponding angle in the triangle $\tilde{\Delta}_2$.*

In any case, we have that

- $\measuredangle_{\bar{y}}(\bar{x}, \bar{z}) \geq \measuredangle_{\tilde{x}}(\tilde{x}, \tilde{z}) = \tilde{\measuredangle}_y(x, z)$.

The same is true if p is a point on the side xz such that $y \ll p$. Note that if X has non-negative (non-positive) timelike curvature and Δ is within a comparison neighbourhood, the condition is satisfied, i.e. $\tau(p, y) \leq \bar{\tau}(\tilde{p}, \tilde{y})$ (or $\geq$).

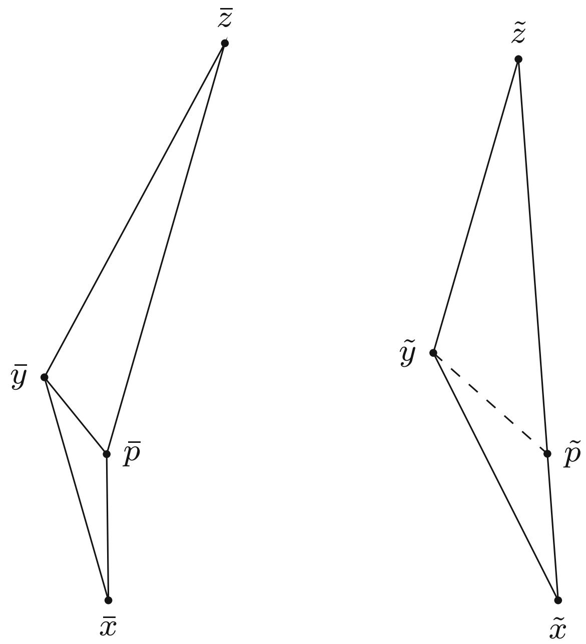

Fig. 1 A concave situation in the across version

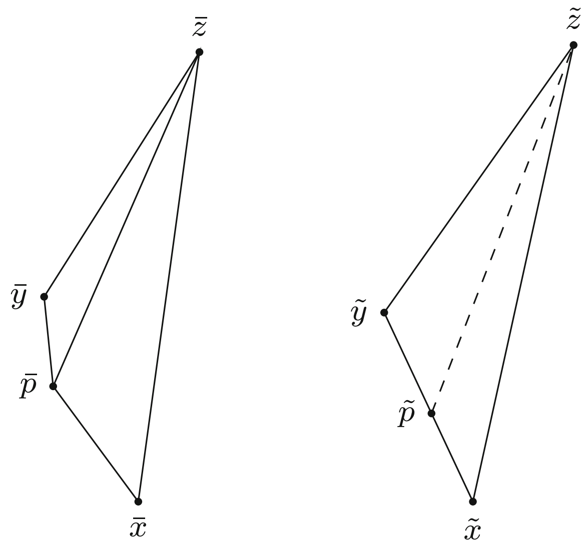

Fig. 2 A convex situation in the future version

Proposition 2.43 (Alexandrov lemma: future version) *Let X be a Lorentzian pre-length space. Let $\Delta := \Delta(x, y, z)$ be a timelike triangle (in particular the distance realisers between the endpoints exist). Let p be a point on the side xy, such that the distance realiser between p and z exists. Then we can consider the smaller triangles $\Delta_1 := \Delta(x, p, z)$ and $\Delta_2 := \Delta(p, y, z)$. We construct a comparison situation consisting of a comparison triangle $\bar{\Delta}_1$ for Δ_1 and $\bar{\Delta}_2$ for Δ_2, with $\bar{x}$ and $\bar{y}$ on different sides of the line through $\bar{p}\bar{z}$ and a comparison triangle $\tilde{\Delta}$ for Δ with a comparison point $\tilde{p}$ for p on the side xy. This contains the subtriangles $\tilde{\Delta}_1 := \Delta(\tilde{x}, \tilde{p}, \tilde{z})$ and $\tilde{\Delta}_2 := \Delta(\tilde{p}, \tilde{y}, \tilde{z})$, see Fig. 2.*

Then the situation $\bar{\Delta}_1, \bar{\Delta}_2$ is convex (concave) at p (i.e. $\measuredangle_{\bar{p}}(\bar{y}, \bar{z}) \geq \measuredangle_{\bar{p}}(\bar{x}, \bar{z})$ (or $\leq$)) if and only if $\tau(p, z) \leq \bar{\tau}(\bar{p}, \bar{z})$ (or $\geq$). If this is the case, we have that

- *each angle in the triangle $\bar{\Delta}_1$ is $\geq$ (or $\leq$) than the corresponding angle in the triangle $\tilde{\Delta}_1$,*

- *each angle in the triangle $\bar{\Delta}_2$ is $\leq$ (or $\geq$) than the corresponding angle in the triangle, $\tilde{\Delta}_2$.*

In any case, we have that

- $\measuredangle_{\bar{z}}(\bar{x}, \bar{y}) \leq \measuredangle_{\tilde{z}}(\tilde{x}, \tilde{y}) = \tilde{\measuredangle}_z(x, y)$.

Note that if X has non-negative (non-positive) timelike curvature and Δ is within a comparison neighbourhood, the condition is satisfied, i.e. $\tau(p, z) \leq \bar{\tau}(\tilde{p}, \tilde{z})$ (or $\geq$).

Proof (Proof for both Alexandrov lemmas) It is sufficient to show the "if" part of the statement. Indeed, under the assumption that the triangles Δ_1 and Δ_2 are non-degenerate, the "if" part holds with strict inequalities, which then also implies the "only if" part of the statement. The degenerate cases are trivial.

For the triangles $\bar{\Delta}_1$ vs. $\tilde{\Delta}_1$, only one side-length changes, and it is not the longest side of the triangle in both versions. Thus, the monotonicity statement in the law of cosines, cf. Lemma 2.35, implies that in the across version, the relation between the desired angle $\measuredangle_{\bar{p}}(\bar{x}, \bar{y})$ in $\bar{\Delta}_1$ to the corresponding angle $\measuredangle_{\tilde{p}}(\tilde{x}, \tilde{y})$ in $\tilde{\Delta}_1$ is pointing in the opposite direction than the relation of the changing side-length $\tau(p, y) = \bar{\tau}(\bar{p}, \bar{y})$ to the side-length $\bar{\tau}(\tilde{p}, \tilde{y})$. In the future version, we obtain that the relation between the desired angle $\measuredangle_{\bar{p}}(\bar{x}, \bar{z})$ in $\bar{\Delta}_1$ to the corresponding angle $\measuredangle_{\tilde{p}}(\tilde{x}, \tilde{z})$ in $\tilde{\Delta}_1$ is pointing in the other direction than the relation of the changing side-length $\tau(p, z) = \bar{\tau}(\bar{p}, \bar{z})$ to the side-length $\bar{\tau}(\tilde{p}, \tilde{z})$.

Similarly, for the triangles $\bar{\Delta}_2$ vs. $\tilde{\Delta}_2$, only one side-length changes. In the across version, it is not the longest side, and for the future version, it is. Thus, in the future version, the monotonicity statement in the law of cosines yields that the relation between the angle $\measuredangle_{\bar{p}}(\bar{y}, \bar{z})$ in $\bar{\Delta}_1$ to the corresponding angle $\measuredangle_{\tilde{p}}(\tilde{y}, \tilde{z})$ in $\tilde{\Delta}_1$ is pointing in the same direction as the relation of the changing side-length $\bar{\tau}(\bar{p}, \bar{y})$ to the side-length $\bar{\tau}(\tilde{p}, \tilde{y})$. Similarly, in the across version, we get that the relations between the angles $\measuredangle_{\bar{p}}(\bar{y}, \bar{z})$ in $\bar{\Delta}_1$ and $\measuredangle_{\tilde{p}}(\tilde{y}, \tilde{z})$ in $\tilde{\Delta}_1$ points in the opposite direction as the relation of the changing side-length $\bar{\tau}(\bar{p}, \bar{y})$ to the side-length $\bar{\tau}(\tilde{p}, \tilde{y})$. Note that $\tilde{\Delta}_1$ and $\tilde{\Delta}_2$ together form the triangle $\tilde{\Delta}$, so we have $\measuredangle_{\tilde{p}}(\tilde{x}, \tilde{y}) = \measuredangle_{\tilde{p}}(\tilde{y}, \tilde{z})$ in the across version and $\measuredangle_{\tilde{p}}(\tilde{x}, \tilde{z}) = \measuredangle_{\tilde{p}}(\tilde{y}, \tilde{z})$ in the future version. From this, the desired inequalities follow.

For the "split" angle, we use the triangle equality along lines and the reverse triangle inequality on the broken side, giving $\bar{\tau}(\bar{x}, \bar{z}) \geq \bar{\tau}(\bar{x}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{z}) = \tau(x, p) + \tau(p, z) = \tau(x, z) = \bar{\tau}(\tilde{x}, \tilde{z})$ in the across statement and $\bar{\tau}(\bar{x}, \bar{y}) \geq \bar{\tau}(\bar{x}, \bar{p}) + \bar{\tau}(\bar{p}, \bar{y}) = \tau(x, p) + \tau(p, y) = \tau(x, y) = \bar{\tau}(\tilde{x}, \tilde{y})$ in the future statement. For the triangles $\bar{\Delta} = \Delta(\bar{x}, \bar{y}, \bar{z})$ and $\tilde{\Delta} = \Delta(\tilde{x}, \tilde{y}, \tilde{z})$, only one side-length changes. In the across version it is the longest side of the triangle, and in the future version, it is not. Thus, in the future version, the monotonicity statement in the law of cosines implies that the relation between the angle $\measuredangle_{\bar{z}}(\bar{x}, \bar{y})$ in $\bar{\Delta}$ to the corresponding angle $\measuredangle_{\tilde{z}}(\tilde{x}, \tilde{y})$ in $\tilde{\Delta}$ is pointing in the opposite direction than the relation of the changing side-lengths $\tau(x, y) = \bar{\tau}(\bar{x}, \bar{y})$ and $\bar{\tau}(\tilde{x}, \tilde{y})$. Similarly, in the across version, we obtain that the relation of the angle $\measuredangle_{\bar{y}}(\bar{x}, \bar{z})$ in $\bar{\Delta}$ and $\measuredangle_{\tilde{y}}(\tilde{x}, \tilde{z})$ in $\tilde{\Delta}$ points in the same direction as the inequality between $\tau(x, z) = \bar{\tau}(\bar{x}, \bar{z})$ and $\bar{\tau}(\tilde{x}, \tilde{z})$.

Finally, note that all monotonicity arguments work the same if the timelike relation between y and p is reversed in the across version. $\square$

2.4 Extensions of Lorentzian length spaces

As in [11], in the proof of the splitting theorem we will first show that $I(\gamma) := I^+(\gamma) \cap I^-(\gamma)$ splits, and then we will use an inextendibility argument to show that $I(\gamma) = X$. To this end, let us recall the notion of extensions of Lorentzian (pre-)length spaces as introduced in [40, Def. 3.1].

Definition 2.44 (*Extensions*) Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space. A Lorentzian pre-length space $(\tilde{X}, \tilde{d}, \tilde{\ll}, \tilde{\leq}, \tilde{\tau})$ is called an *extension* of $(X, d, \ll, \leq, \tau)$ if

(i) $(\tilde{X}, \tilde{d})$ is connected,
(ii) there exists an isometric embedding $\iota : (X, d) \to (\tilde{X}, \tilde{d})$,
(iii) $\iota(X) \subsetneq \tilde{X}$ is open,
(iv) ι preserves $\leq$ and $\ll$,
(v) ι preserves τ-lengths of causal curves.

Remark 2.45 (Convex extensions) An important class of examples (and the only one we will need) is the case of open, causally convex subsets of a suitable Lorentzian length space. In this case, ι even preserves τ.

Next, we recall the definition of *timelike completeness* for (localisable) Lorentzian pre-length spaces, cf. [40, Def. 5.1]. However, since continuous extendibility, extendibility as a causal curve and geodesic extendibility are in general inequivalent concepts, we will define the timelike completeness property precisely in such a way that it is sufficient for the subsequent inextendibility result.

Definition 2.46 (*Timelike completeness (TC)*) Let $(X, d, \ll, \leq, \tau)$ be a localisable Lorentzian pre-length space. We say X satisfies the *timelike completeness (TC) property* if each (future or past directed) timelike geodesic $\alpha : [a, b) \to X$ that is inextendible to $[a, b]$ as a continuous curve has infinite τ-length.

Theorem 2.47 (Inextendibility of (TC) spaces) *Let $(X, d, \ll, \leq, \tau)$ be a strongly causal Lorentzian length space satisfying the (TC) property. Then X is inextendible as a regularly localisable Lorentzian length space.*

Proof See [40, Thm. 5.3]. □

3 Lorentzian products

This section is dedicated to the treatment of Lorentzian product spaces. In the context of Lorentzian pre-length spaces, (warped) products were first introduced in [5]. On the one hand, this investigation is too general for our purposes, we only need products instead of warped products. On the other hand, the formulation we use also works for non-intrinsic spaces.

Definition 3.1 (*Lorentzian product*) Let (X, d) be a metric space. Define the space $Y := \mathbb{R} \times X$. Equip it with the product metric

$$D : Y \times Y \to \mathbb{R},\ D((s, x), (t, y)) := \sqrt{|t - s|^2 + d(x, y)^2}.$$

Define the timelike relation as $(s, x) \ll (t, y) :\iff t - s > d(x, y)$ and the causal relation as $(s, x) \leq (t, y) :\iff t - s \geq d(x, y)$. Define the product time separation $\tau : Y \times Y \to \mathbb{R}$ via

$$\tau((s, x), (t, y)) := \sqrt{(t - s)^2 - d(x, y)^2}$$

if $(s, x) \leq (t, y)$, and 0 otherwise.

Then $(Y, D, \ll, \leq, \tau)$ is called the Lorentzian product of X with $\mathbb{R}$. If it is clear that the Lorentzian product is meant, we simply denote it by $\mathbb{R} \times X$.

The following is more or less immediate from the definition:

Proposition 3.2 (Properties of Lorentzian products) *Let (X, d) be a metric space. Then the Lorentzian product $(Y, D, \ll, \leq, \tau)$ is a Lorentzian pre-length space. Moreover, τ is continuous and $\leq$ is closed on $Y \times Y$.*

Proof τ is continuous by definition, so in particular, lower semi-continuous. The inclusion of $\ll$ in $\leq$ is clear from the definition as well. The transitivity of both relations as well as the reverse triangle inequality of τ for $\leq$-related triples follow (via an elementary calculation) from the triangle inequality for d. The reflexivity of $\leq$ is clear from the definition. Finally, the definition of $\ll$ is a reformulation of $\tau > 0$.

Note that for converging real sequences $n_i \to n, m_i \to m$ such that $n_i \leq m_i$ for all i, we have $n \leq m$. The closedness of $\leq$ on $Y \times Y$ follows immediately from this fact. □

It is easily seen that generalised cones in the sense of [5] with warping function $f \equiv 1$ are Lorentzian products in the sense of Definition 3.1. Indeed, in this case we have, in the notation of [5, Def. 3.9], that $m_{s,t} := \min_{r \in [s,t]} f(r) = 1$ for all $s, t \in \mathbb{R}$. Then [5, Lem. 3.10] shows that a causal curve in the sense of [5, Def. 3.2] is a causal curve in the usual sense in our setting, with timelikeness being inherited as well. Moreover, the variational length of [5, Def. 3.9] is precisely the τ-length of a causal curve in the sense of Definition 2.6. It is easily seen that if X is strictly intrinsic, then the product $\mathbb{R} \times X$ in our sense can be canonically identified with the warped product $\mathbb{R} \times_f X$ with $f \equiv 1$.

We now show that Lorentzian products are always strongly causal and non-totally imprisoning.

Proposition 3.3 (Diamonds form basis) *Let $Y := \mathbb{R} \times X$ be a Lorentzian product. Then $\{I(p, q) \mid p, q \in Y\}$ forms a basis for the topology. In particular, Y is strongly causal.*

Proof Consider the open set $O := (a, c) \times B_R(x)$ for $a, c \in \mathbb{R}, R > 0, x \in X$ and let $(b, y) \in O$. Choose $\varepsilon = \min(b - a, c - b, R - d(x, y)) > 0$, then $p := (b - \varepsilon, y), q := (b + \varepsilon, y) \in \bar{O}$. Then we have that $I(p, q) \subseteq O$: If $(s, z) \in I(p, q)$, then $s \in (b - \varepsilon, b + \varepsilon) \subseteq (a, c)$. In particular, at least one of $|s - (b - \varepsilon)| \leq \varepsilon$ and $|s - (b + \varepsilon)| \leq \varepsilon$ holds. By definition of $\ll$, we then also have $d(x, z) \leq d(x, y) + d(y, z) < R$, so $y \in B_R(x)$. Thus, $(s, z) \in O$. □

Proposition 3.4 (Non-total imprisonment of Lorentzian products) *Any Lorentzian product $Y = \mathbb{R} \times X$, where (X, d) is a metric space, is non-totally imprisoning.*

Proof For any two points $(s, x) \leq (t, y) \in Y$ we have $d(x, y) \leq t - s$, so $D((s, x), (t, y)) = \sqrt{(t - s)^2 + d(x, y)^2} \leq \sqrt{2}(t - s)$. In particular, for any future directed causal curve $\gamma : [a, b] \to Y$ starting at $\gamma(a) = (s, x)$ and ending at $\gamma(b) = (t, y)$ we have $L_D(\gamma) \leq \sqrt{2}(t - s)$: In any partition of $[a, b]$, the above bound in any of the subintervals forms a telescopic sum, making $\sqrt{2}(t - s)$ an upper bound on the length of the polygon approximation of γ corresponding to the partition. Let now K be compact, then we can enclose it in a bounded set: $K \subseteq [s, t] \times \bar{B}_R(x)$ for some $s < t$, $R > 0$ and $x \in X$. We know that the D-length of future directed causal curves contained in $[s, t] \times \bar{B}_R(x)$ is bounded by $\sqrt{2}(t - s)$, so we have that Y is non-totally imprisoning. □

The next result characterises distance realisers in $Y = \mathbb{R} \times X$. Note that we make use of the following: A distance minimiser in a metric space can always be parametrised by unit speed, which is evidently a Lipschitz parametrisation, cf. [17, Proposition 2.5.9].

Proposition 3.5 (Characterisation of distance realisers) *Let (X, d) be a metric space and let $Y := \mathbb{R} \times X$ the Lorentzian product. Then a continuous causal curve $\gamma = (\alpha, \beta) : [a, b] \to Y$ is a distance realiser if and only if either $\beta = const$ or $\beta : [a, b] \to X$ is a (metric) distance realiser and when reparametrising γ such that β is unit speed parametrised, $\alpha : [a, b] \to \mathbb{R}$ is affine, i.e. of the form $\alpha(t) = ct + d$ with $c = 1$ (which implies that γ future directed null) or $c > 1$ (which implies that γ is future directed timelike). In particular, γ has causal character and a reparametrisation as a (Lipschitz) causal curve. If in addition Y is localisable, then the same is true for continuous geodesics.*

Proof γ is a distance realiser if and only if for all $r < s < t$ we have

$$\tau(\gamma(r), \gamma(t)) = \tau(\gamma(r), \gamma(s)) + \tau(\gamma(s), \gamma(t)). \tag{$\star$}$$

We denote $\gamma(r) = (t_1, x_1)$, $\gamma(s) = (t_2, x_2)$ and $\gamma(t) = (t_3, x_3)$, then ($\star$) reads:

$$\sqrt{(t_3 - t_1)^2 - d(x_1, x_3)^2} = \sqrt{(t_2 - t_1)^2 - d(x_1, x_2)^2} + \sqrt{(t_3 - t_2)^2 - d(x_2, x_3)^2}$$

We now define $\lambda = \frac{t_2 - t_1}{t_3 - t_1}$ and the function $f(x) = \sqrt{1 - x^2}$. Then ($\star$) simplifies to

$$f\left(\frac{d(x_1, x_3)}{t_3 - t_1}\right) = \lambda f\left(\frac{d(x_1, x_2)}{\lambda(t_3 - t_1)}\right) + (1 - \lambda) f\left(\frac{d(x_2, x_3)}{(1 - \lambda)(t_3 - t_1)}\right). \tag{$\star\star$}$$

We use that f is concave and monotonously decreasing on $[0, 1]$ and $d(x_1, x_2) + d(x_2, x_3) \geq d(x_1, x_3)$ on the right hand side of $(\star\star)$:

$$\lambda f\left(\frac{d(x_1, x_2)}{\lambda(t_3 - t_1)}\right) + (1-\lambda) f\left(\frac{d(x_2, x_3)}{(1-\lambda)(t_3 - t_1)}\right) \\ \leq f\left(\frac{d(x_1, x_2) + d(x_2, x_3)}{(t_3 - t_1)}\right) \leq f\left(\frac{d(x_1, x_3)}{(t_3 - t_1)}\right).$$

So we get that $(\star\star)$ is equivalent to having equality here. As f is strictly concave, the first inequality has equality if and only if

$$\frac{d(x_1, x_2)}{t_2 - t_1} = \frac{d(x_2, x_3)}{t_3 - t_2}, \tag{3.1}$$

where we got rid of λ again for generality. As f is strictly monotonously decreasing, the second inequality has equality if and only if

$$d(x_1, x_2) + d(x_2, x_3) = d(x_1, x_3). \tag{3.2}$$

We translate this back to what happens for general $r < s < t$: Eq. 3.2 reads $d(\beta(r), \beta(s)) + d(\beta(s), \beta(t)) = d(\beta(r), \beta(t))$, i.e., it holds if and only if β is a distance realiser. Equation 3.1 reads $\frac{d(\beta(r),\beta(s))}{\alpha(s)-\alpha(r)} = \frac{d(\beta(s),\beta(t))}{\alpha(t)-\alpha(s)}$, i.e. that quantity is a constant $\frac{1}{c}$ (if $\frac{1}{c}$ is 0 we get that β is constant, in this case the result is trivial). In particular, if we reparametrise γ such that β is unit speed (i.e. $d(\beta(s), \beta(t)) = t - s$), $\alpha(t) - \alpha(s) = c(t - s)$, i.e. α is affine.

Now we look at what happens for different parameters c: If $c > 1$, we easily see that $f(c) > 0$, so $\tau(\gamma(r), \gamma(t)) = (\alpha(t) - \alpha(r)) f(c) > 0$ and the curve is timelike. For $c = 1$, $f(c) = f(1) = 0$, so $\alpha(t) - \alpha(r) = d(\beta(r), \beta(t))$ and the curve is null.

Now if Y is localisable and γ is only a geodesic, we can cover the domain with intervals J_i that are open in $[a, b]$ where it is distance realising. Then on each of these J_i we apply the above. As the J_i overlap, the constant c agrees, and we get that $\beta|_{J_i}$ is distance realising, thus β is a geodesic. The converse follows equivalently. We even get that the subintervals of the domain where the curves are distance realising agree. □

Corollary 3.6 (τ-arclength parametrisations of distance realisers) *Any continuous timelike maximiser $\gamma : [a, b] \to \mathbb{R} \times X$ has a Lipschitz τ-arclength parametrisation, which is of the form (α, β), with $\alpha : [0, L_\tau(\gamma)] \to \mathbb{R}$, $\alpha(t) = ct + d$, and $\beta : [0, L_\tau(\gamma)] \to X$ a constant speed minimiser of speed $\sqrt{c^2 - 1}$. If $\mathbb{R} \times X$ is localisable, then the same is true for timelike geodesics (with β constant speed geodesic).*

Note that the results above continue to hold for maximisers resp. geodesics defined on any interval I, not just a closed one.

To finish this section, we show that global hyperbolicity of the Lorentzian product $Y = \mathbb{R} \times X$ is equivalent to the metric properness of X.

Proposition 3.7 (Global hyperbolicity and properness) *A Lorentzian product $Y := \mathbb{R} \times X$ is globally hyperbolic if and only if X is a proper metric space.*

Proof Since products are always non-totally imprisoning by Proposition 3.4, we only need to check that causal diamonds in Y are compact if and only if (X, d) is a proper metric space. First, suppose that Y is globally hyperbolic. Let $(t, x) \in Y$ and fix $R > 0$. Consider the points $p := (t - R, x), q := (t + R, x)$. Then $(t, x) \in J(p, q)$. By definition of $\leq$, any point (t, y) with the same $\mathbb{R}$-coordinate as (t, x) is in $J(p, q)$ if and only if $d(x, y) \leq R$. Thus, the set $J(p, q) \cap \{(t, x) \mid x \in X\}$, which is compact by assumption, can be identified with the closed ball $\bar{B}_R(x)$ in X of radius R around x, establishing the fact that closed balls in X are compact.

Conversely, suppose that X is proper and let $p = (r, x), q = (t, z) \in Y$ with $p \leq q$. By definition of $\leq$, if $(s, y) \in J(p, q)$, then $r \leq s \leq t$. Moreover, we have $|s| \leq |r| + |t|$. Set $R := 2|r| + 2|t|$. Then for $(s, y) \in J(p, q)$, we also have $d(x, y) \leq |s - r| \leq |s| + |r| \leq R$. Thus, $y \in \bar{B}_R(x)$, which is compact by assumption. In total, we conclude that $J(p, q) \subseteq [r, t] \times \bar{B}_R(x)$, which is a compact set as the Cartesian product of two compact sets. The fact that $J(p, q)$ is closed follows immediately from the closedness of $\leq$, so $J(p, q)$ is compact as well. □

4 Rays, lines, co-rays and asymptotes

4.1 Rays, lines, co-rays, asymptotes, timelike co-ray condition

In this subsection, we study causal rays and lines and show how to obtain them as limits of causal maximisers. Moreover, we analyse triangles where one side is a segment on a timelike line and show that the angles adjacent to the line are equal to their comparison angles. This in fact follows from the more general principle that one can stack triangle comparisons of nested triangles with two endpoints on a timelike line. The latter situation arises in the construction of asymptotes, and via the stacking principle, one can show any future directed and any past directed asymptote from a common point to a given timelike line fit together to give a (timelike) asymptotic line. In constructing co-rays and asymptotes we follow [11], whereas the stacking principle and equality of angles can be viewed as Lorentzian analogues of [17, Lem. 10.5.4].

Definition 4.1 (*Rays, Lines*) Let $(X, d, \ll, \leq, \tau)$ be a Lorentzian pre-length space. A *future directed causal ray* is a future inextendible, future directed causal curve $c : I \to X$ that maximises the time separation between any of its points, where I is either a closed interval $[a, b]$ or a half-open interval $[a, b)$. A *future directed causal line* is a (doubly) inextendible, future directed causal curve $\gamma : I \to X$ that maximises the time separation between any of its points. Here I can in general be open, closed, or half-open. More generally, for $S \subset X$, a *future directed causal S-ray* is a future directed, future inextendible causal curve starting in S and satisfying $\tau(S, c(t)) := \sup_{p \in S} \tau(p, c(t)) = L_\tau(c|_{[0,t]})$. Past directed rays and lines are defined analogously. A future directed causal ray $c : I \to X$ is called *complete* if $L_\tau(c) = \infty$. Similarly, a future directed causal line γ is called *complete* if there is $t_0 \in I$ such that the past and future directed rays obtained from dividing γ at $\gamma(t_0)$ are complete.

Unless explicitly stated otherwise, all causal rays and lines are understood to be future directed.

Remark 4.2 (i) Clearly, any ray c is a $\{c(0)\}$-ray. Conversely, if $c : [a, b) \to X$ is a future directed S-ray for some $S \subset X$, then $\tau(c(a), c(t)) \leq \tau(S, c(t)) = L_\tau(c|_{[a,t]}) \leq \tau(c(a), c(t))$, so c is a ray. A similar statement is true for past directed rays.

(ii) If X is regularly localisable, then any causal ray/line c is either timelike or null by Theorem 2.17.

(iii) If X is localisable and causally path-connected, then any ray c is defined on a half-open interval and any line γ is defined on an open interval.

In the following, unless we specify the assumptions on X, we always take X to be as in the splitting theorem.

Proposition 4.3 *Let $z_n \to z$ in X. Let $p_n \in I^+(z_n)$ and let $c_n : [0, a_n] \to X$ be a sequence of future directed, maximising causal curves from z_n to p_n in d-arclength parametrisation. If $\tau(z_n, p_n) \to \infty$, then there is a limit causal ray $c : [0, \infty) \to X$ from z.*

Proof We only have to show $a_n \to \infty$, the rest follows from the limit curve theorem. Suppose $a_n \to a < \infty$. This implies that all p_n are contained in a common compact set K: Indeed, if we assume this to be false, then for each $n > 0$ there would be a p_k, $k = k(n)$, such that $p_k \notin \overline{B}_n(z) = \{x \in X : d(x, z) \leq n\}$ (recall that (X, d) is proper), which in particular means that

$$a_{k(n)} = L^d(c_{k(n)}) \geq d(z_{k(n)}, p_{k(n)}) \geq d(p_{k(n)}, z) - d(z_{k(n)}, z) \geq n - d(z_{k(n)}, z),$$

which is a contradiction since the right hand side tends to ∞. So, up to a choice of subsequence, we may assume that $\gamma_n(a_n) = p_n \to p$. But then, by continuity of τ,

$$\tau(z, p) = \lim_n \tau(z_n, p_n) = \infty,$$

a contradiction to the finiteness of τ on all of $X \times X$. □

Let now $c : [0, \infty) \to X$ be a complete timelike S-ray, and fix $z \in I^-(c) \cap I^+(S)$. Let $z_n \to z$ and $p_n := c(r_n)$ for some sequence of parameters $r_n \to \infty$. Then it is easily seen that $\tau(z_n, p_n) \to \infty$: Indeed, let $r > 0$ be large and assume w.l.o.g. that all z_n are in $I^-(c(r))$, then

$$\tau(z_n, p_n) \geq \tau(z_n, c(r)) + \tau(c(r), c(r_n)) \to \infty.$$

Hence, by the above result, we may construct causal rays as follows: Let μ_n be maximising timelike curves from z_n to p_n and let μ be a limit causal ray from z whose existence we just proved.

Definition 4.4 (*Co-rays, asymptotes*) Any causal ray μ constructed by the above method is called a *co-ray* to the ray c at z. If $z_n = z$ for all n, then μ is called an *asymptote* to c at z.

Fig. 3 The domain where the y_i can lie in is in the double line

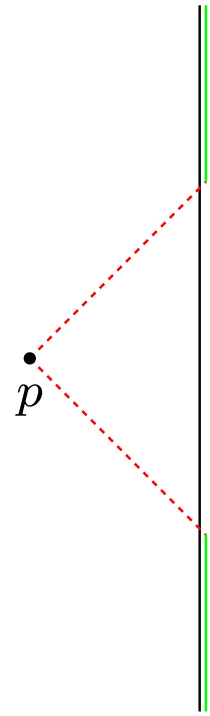

Since maximising causal curves have causal character, a co-ray is either timelike or null.

Definition 4.5 (*Timelike co-ray condition*) Let $c : [0, \infty) \to X$ be a complete timelike S-ray and let $p \in I^-(c) \cap I^+(S)$. We say the *timelike co-ray condition (TCRC)* holds at p if every co-ray to c at p is timelike.

We turn to the treatment of triangles adjacent to lines and begin with the aforementioned stacking principle. It will be an essential technical tool for controlling the behaviour of asymptotes to a line. The following two results are true for rather general Lorentzian pre-length spaces, and the exact conditions are specified.

Proposition 4.6 (Comparison situations stack along a geodesic) *Let X be a timelike geodesically connected Lorentzian pre-length space with global timelike curvature bounded below by* 0 *and $\gamma : \mathbb{R} \to X$ be a complete timelike line. Let $p \in X$ be a point not on γ. Let $t_1 < t_2 < t_3$ such that all $y_i := \gamma(t_i)$ are timelike related to p, see Fig. 3. Let $\bar{\Delta}_{12} := \Delta(\bar{p}, \bar{y}_1, \bar{y}_2)$ be a comparison triangle for $\Delta_{12} := \Delta(p, y_1, y_2)$ and extend the side $\bar{p}, \bar{y}_2$ to a comparison triangle $\bar{\Delta}_{23} := \Delta(\bar{p}, \bar{y}_2, \bar{y}_3)$ for $\Delta_{23} := \Delta(p, y_2, y_3)$. We choose it in such a way that $\bar{y}_1$ and $\bar{y}_3$ lie on opposite sides of the line through $\bar{y}_1, \bar{y}_2$. Then $\bar{y}_1, \bar{y}_2, \bar{y}_3$ are collinear. That makes $\bar{\Delta}_{13} := \Delta(\bar{p}, \bar{y}_1, \bar{y}_3)$ a comparison triangle for $\Delta_{13} := \Delta(p, y_1, y_3)$.*

Proof We set $s_- = \sup(\gamma^{-1}(I^-(p)))$ and $s_+ = \inf(\gamma^{-1}(I^+(p)))$. Then the set of s where $\gamma(s)$ is timelike related to p is $(-\infty, s_-) \cup (s_+, +\infty)$. We assume $p \ll y_2$, the other case $y_2 \ll p$ can be reduced to this one by flipping the time orientation of the space.

We create comparison situations for an Alexandrov situation: We take the triangles $\bar{\Delta}_{12} = \Delta(\bar{p}, \bar{y}_1, \bar{y}_2)$ and $\bar{\Delta}_{23} = \Delta(\bar{p}, \bar{y}_2, \bar{y}_3)$ as in the statement. We also create a comparison triangle $\tilde{\Delta}_{13} = \Delta(\tilde{p}, \tilde{y}_1, \tilde{y}_3)$ for (p, y_1, y_3) and get a comparison point $\tilde{y}_2$ for y_2 on the side $y_1 y_3$. Then we can apply curvature comparison to get $\bar{\tau}(\tilde{p}, \tilde{y}_2) \geq \tau(p, y_2) = \bar{\tau}(\bar{p}, \bar{y}_2)$. By Lemmas 2.42 and 2.43,[4] this means the situation is convex, i.e.

$$\tilde{\measuredangle}_{y_2}(p, y_1) \leq \tilde{\measuredangle}_{y_2}(p, y_3). \tag{4.1}$$

If $p \ll y_1$ we also get

$$\tilde{\measuredangle}_{y_1}(p, y_2) \leq \tilde{\measuredangle}_{y_1}(p, y_3) \tag{4.2}$$

and if $y_1 \ll p$, this inequality flips.

Now we apply Eq. (4.2) to the situation where we fix y_2 and move y_3 and y_1: We get $\tilde{\measuredangle}_{y_2}(p, \gamma(t_3))$ is monotonously increasing in t_3. Similarly, the time-reversed situation gives that if $p \ll \gamma(t_1)$, $\tilde{\measuredangle}_{y_2}(p, \gamma(t_1))$ is monotonously increasing in t_1 and if $\gamma(t_1) \ll p$ it is also monotonously increasing in t_1 (note it is also increasing when switching from one case to the other). Then remember $\tilde{\measuredangle}_{y_2}(p, \gamma(t_1)) \leq \tilde{\measuredangle}_{y_2}(p, \gamma(t_3))$ by the convexity of the Alexandrov situation (Eq. (4.1)). Note that the left hand side is decreasing with decreasing t_1 and the right hand side is increasing with increasing t_3.

<u>Claim</u>: For each t_1 and t_3, Eq. (4.1) has equality.

<u>Then</u>: The comparison situation is straight, i.e. $\bar{y}_1, \bar{y}_2, \bar{y}_3$ are collinear, and by triangle equality along straight lines we have $\tau(y_1, y_3) = \tau(y_1, y_2) + \tau(y_2, y_3) = \bar{\tau}(\bar{y}_1, \bar{y}_2) + \bar{\tau}(\bar{y}_2, \bar{y}_3) = \bar{\tau}(\bar{y}_1, \bar{y}_3)$, i.e., $\bar{p}, \bar{y}_1, \bar{y}_3$ is a comparison triangle for p, y_1, y_3.

<u>Proof</u>: We indirectly assume there is a t_1^0, t_3^0 such that the comparison situation $\bar{\Delta}_{12}, \bar{\Delta}_{23}$ is strictly convex. We realise this situation such that $\bar{p} = 0$, $\bar{y}_2$ is to the right of the t-axis, such that the side $\bar{y}_2\bar{y}_3^0$ slopes to the left and $\bar{y}_1^0\bar{y}_2$ slopes to the right (i.e., $\frac{x(\bar{y}_3^0 - \bar{y}_2)}{t(\bar{y}_3^0 - \bar{y}_2)} < 0$, $\frac{x(\bar{y}_2 - \bar{y}_1^0)}{t(\bar{y}_2 - \bar{y}_1^0)} > 0$). When we vary t_1 and t_3, the comparison situation is chosen such that $\bar{p}$ and $\bar{y}_2$ stay fixed. As the comparison angle $\tilde{\measuredangle}_{y_2}(p, \gamma(t_3))$ is monotonously increasing in t_3, we get that for $t_3 \geq t_3^0$ the slope $\frac{x(\bar{y}_3 - \bar{y}_2)}{t(\bar{y}_3 - \bar{y}_2)}$ is smaller than the slope of $\bar{y}_2\bar{y}_3$. In particular, $\bar{y}_3$ lies to the left of the extension of the side $\bar{y}_2\bar{y}_3^0$. Thus, for large enough t_3, $\bar{y}_3$ lies to the left of the t-axis, and for a certain t_3^* $\bar{y}_3^*$ lies on it. Similarly, we find a t_1^* such that $\bar{y}_1^*$ lies on the t-axis. See Fig. 4for a visualisation of the construction. But then

$$\begin{aligned}\tau(y_1^*, p) + \tau(p, y_3^*) &= \bar{\tau}(\bar{y}_1^*, \bar{p}) + \bar{\tau}(\bar{p}, \bar{y}_3^*) > \bar{\tau}(\bar{y}_1^*, \bar{y}_2) + \bar{\tau}(\bar{y}_2, \bar{y}_3^*) \\ &= \tau(y_1^*, y_2) + \tau(y_2, y_3^*) = \tau(y_1^*, y_3^*)\end{aligned}$$

in contradiction to the reverse triangle inequality. Thus we get the claim. □

Remark 4.7 The proof of the statement can also be used in more general situations: Let X be a locally timelike geodesically connected Lorentzian pre-length space with local timelike curvature bounded below by 0, and let γ be a (possibly extendible) timelike

[4] If $p \ll y_1 \ll y_2 \ll y_3$, we use Lemma 2.43 and if $y_1 \ll p \ll y_2 \ll y_3$ we use Lemma 2.42.

Fig. 4 We assume that the path $\bar{y}_1^0\bar{y}_2\bar{y}_3^0$ is longer than the path $\bar{y}_1^0\bar{p}\bar{y}_3^0$. If the angle at $\bar{y}_2$ is not straight, we extend some lines and as points further from $\bar{p}$ are more on the left, we find some t_1^* and t_3^* such that $\bar{y}_1^*\bar{p}\bar{y}_3^*$ are collinear. But then the path $\bar{y}_1^0\bar{y}_2\bar{y}_3^0$ is shorter than $\bar{y}_1^0\bar{p}\bar{y}_3^0$, which then yields a contradiction to γ being maximising

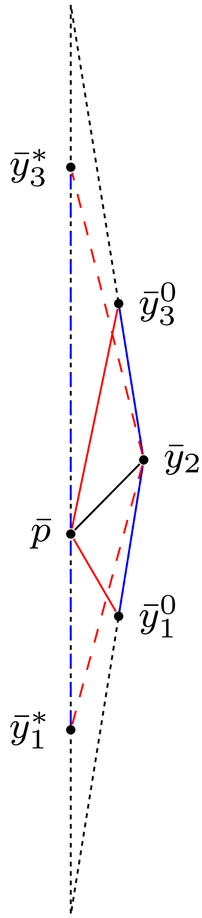

maximiser. Let $p \ll x = \gamma(t_2)$ be points in a comparison neighbourhood. Assume the statement is not true (with certain $t_1 < t_2 < t_3$). Then we get $\omega_1 = \tilde{\measuredangle}_{y_2}(p, y_1)$, $\omega_3 = \tilde{\measuredangle}_{y_2}(p, y_3)$. Let $d_1 + d_2 = \omega_3 - \omega_1$ and $e = \tau(p, x)\sinh(d_1 + \omega_1)$. Then for the parameters $t_1 = -\frac{e}{\sinh(d_1)}$ and $t_3 = \frac{e}{\sinh(d_3)}$ we have one of the following:

- γ is not defined at/inextendible to one of them,
- γ stops being distance realising between t_1^* and t_3^* (and we have found a longer curve),
- there is no curvature comparison neighbourhood containing both $\gamma(t_i^*)$ $(i = 1, 3)$.

For this, realise the above with $\measuredangle_{\bar{y}_2}(\partial_t, \bar{y}_3) = d_3$ and $\measuredangle_{\bar{y}_2}(\partial_t, \bar{y}_1) = d_1$. Then e is the difference of the x-coordinates of p and x.

Note the similarity of this result and classical arguments for conjugate points along geodesics in Riemannian/Lorentzian geometry.

Proposition 4.8 (Angle = comparison angle) *Let X be a timelike geodesically connected Lorentzian pre-length space with global timelike curvature bounded below by 0 and $\gamma : \mathbb{R} \to X$ be a complete timelike line and $x := \gamma(t_0)$ a point on it. We split γ into the future part $\gamma_+ = \gamma|_{[t_0,+\infty)}$ and the past part $\gamma_- = \gamma|_{(-\infty,t_0]}$. Let $p \in X$ be a point not on γ with x and p timelike related and $\alpha : x \rightsquigarrow p$ a connecting distance realiser. Then for all $s \neq t$ such that p and $\gamma(s)$ are timelike related, we have:*

$$\tilde{\measuredangle}_x(p, \gamma(s)) = \measuredangle_x(\alpha, \gamma_+) = \measuredangle_x(\alpha, \gamma_-),$$

Fig. 5 Illustration: These angles are the same, and have the same value as if they are considered as comparison angles

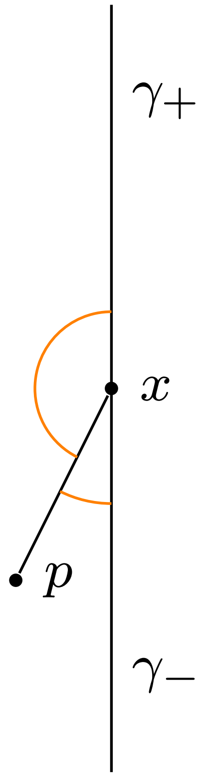

i.e. the comparison angle is equal to the angle, and the same in both directions (see Fig. 5).

Proof First, we check that the comparison angle $\tilde{\measuredangle}_x(p, \gamma(s))$ is constant in s: For s_1 and s_2 for which this is defined, we have three parameters on γ involved: s_1, s_2, t_0. The previous result (Proposition 4.6) tells us we can construct a comparison situation for all three triangles at once. As the comparison situations have the comparison angle $\tilde{\measuredangle}_x(p, \gamma(s_1))$ resp. $\tilde{\measuredangle}_x(p, \gamma(s_2))$ as the angle in $\bar{x}$, they are equal.

Now we look at the angle $\measuredangle_x(\alpha, \gamma_\pm)$: We assume $p \ll x$. We already know $\tilde{\measuredangle}_x(\alpha(s), \gamma(t))$ is constant in t. We now have to look at its dependence on s: By Theorem 2.39, $\tilde{\measuredangle}_x^{\mathrm{S}}(\alpha(s), \gamma(t))$ is monotonously increasing in s. Note now that for $t < t_0$, the sign of this angle is $\sigma = -1$, and for $t > t_0$, the sign of this angle is $\sigma = +1$. So choose some $t_- < t_0$ and $t_+ > t_0$ for which all the necessary angles exist (i.e. $p \ll \gamma(t_-)$ and $t_+ > t_0$) we have that $\tilde{\measuredangle}_x(\alpha(s), \gamma(t_-)) = \tilde{\measuredangle}_x(\alpha(s), \gamma(t_+))$ is both a monotonously decreasing and increasing function in s. Thus it is constant, and $\tilde{\measuredangle}_x(\alpha(s), \gamma(t)) = \measuredangle_x(\alpha, \gamma_-) = \measuredangle_x(\alpha, \gamma_+)$, which includes the desired equalities. □

Corollary 4.9 *Let X be a timelike geodesically connected, globally causally closed Lorentzian pre-length space satisfying $J^\pm(p) = \overline{I^\pm(p)}$ for all $p \in X$.*[5] *and having global timelike curvature bounded below by 0 and let $\gamma : \mathbb{R} \to X$ be a complete timelike line. Then for any point $p \in X$ and two points $x_1 = \gamma(t_1)$ and $x_2 = \gamma(t_2)$ which are timelike related to p, we get a timelike triangle $\Delta = \Delta(x_1, x_2, p)$. Let q_1, q_2 be any points on Δ, one of them lying on the side x_1x_2. We form a comparison triangle $\bar{\Delta}$ for Δ and find comparison points $\bar{q}_1, \bar{q}_2$ for q_1, q_2. Then $q_1 \leq q_2$ if and*

[5] Approximating and globally causally closed Lorentzian pre-length spaces satisfy this, cf. [19, Def. 2.17].

only if $\bar{q}_1 \leq \bar{q}_2$ *and* $\tau(q_1, q_2) = \bar{\tau}(\bar{q}_1, \bar{q}_2)$*, where* $\bar{\tau}$ *denotes the time separation of* $2D$*-Minkowski space.*

Proof First assume $q_1 \leq q_2$. As global curvature is bounded below by 0, we get that $\tau(q_1, q_2) \leq \bar{\tau}(\bar{q}_1, \bar{q}_2)$, and a continuity argument then shows that $\bar{q}_1 \leq \bar{q}_2$.

Conversely, let $\bar{q}_1 \leq \bar{q}_2$. We distinguish which sides the q_i lie on: Note we assumed one of them is on the side x_1x_2. We only prove the case where q_1 is on the side α connecting x_1x_2 (say $q_1 = \alpha(s)$) and q_2 is on the side β connecting x_1, p (say $q_2 = \beta(t)$), the proof of the other cases is easily adapted.

Consider now the hinge $(\alpha|_{[0,s]}, \beta|_{[0,t]})$, which has base-point x_1 and two tips at q_1, q_2. By the previous Proposition 4.8, we get that $\omega := \measuredangle_{x_1}(\alpha, \beta) = \tilde{\omega} := \tilde{\measuredangle}_{x_1}(p, x_2)$. We now have two comparison situations: $\bar{\Delta}$ has angle $\tilde{\measuredangle}_{x_1}(p, x_2)$ at x_1, and a comparison hinge $(\tilde{\alpha}, \tilde{\beta})$ with tips $\tilde{q}_1, \tilde{q}_2$. But note that the angles and distances at $\bar{x}_1$ resp. $\tilde{x}_1$ are the same, so we have that $\bar{\tau}(\bar{q}_1, \bar{q}_2) = \bar{\tau}(\tilde{q}_1, \tilde{q}_2)$. But for the comparison hinge, [13, Cor. 4.11] gives that $\tau(q_1, q_2) \geq \bar{\tau}(\tilde{q}_1, \tilde{q}_2) = \bar{\tau}(\bar{q}_1, \bar{q}_2)$. In total, this gives that $\tau(q_1, q_2) = \bar{\tau}(\bar{q}_1, \bar{q}_2)$, as desired. A similar continuity argument then gives that $q_1 \leq q_2$, thus completing the proof. □

We return to the situation of the splitting theorem. Let $\gamma : \mathbb{R} \to X$ be the complete timelike line whose existence we assume, and let it be in any locally Lipschitz parametrisation (e.g. parametrisation by d-arclength) defined on $\mathbb{R}$. We call any co-ray to any of its past or future directed subrays a co-ray to the line γ, so in particular, we can construct past and future directed co-rays from all points on $I(\gamma) := I^+(\gamma) \cap I^-(\gamma)$. The next result shows that the timelike co-ray condition holds on $I(\gamma)$, i.e. all co-rays from all points in $I(\gamma)$ are timelike. For our purposes it would be sufficient to show this for asymptotes, since we only work with those in the proof.

Proposition 4.10 *The timelike co-ray condition holds on* $I(\gamma)$.

Proof Suppose there is a point $x \in I(\gamma)$ such that the TCRC does not hold at x, so w.l.o.g. there is a sequence $x_n \to x$, $r_n \to \infty$ and maximal timelike curves σ_n from x_n to $\gamma(r_n)$ such that σ_n converge to a future directed null ray σ. Choose some $q \in \gamma \cap I^-(x)$ and let μ_n be maximal timelike curves from q to x_n (assuming n to be large enough, $q \ll x_n$). Suitably pre- and post-composing μ_n and then applying the limit curve theorem, it is easily seen that (up to a choice of subsequence) the μ_n converge locally uniformly to a maximising limit causal curve μ from q to x, which is timelike since $q \ll x$. Denote by γ_n the piece of γ that runs between q and $\gamma(r_n)$. Set $a_n := L_\tau(\mu_n)$, $b_n := L_\tau(\sigma_n)$ and $c_n := L_\tau(\gamma_n)$. Let $\beta_n := \measuredangle_q(\mu_n, \gamma_n)$ and $\theta_n := \measuredangle_{x_n}(\mu_n, \sigma_n)$ (see Fig. 6). Then $a_n \to a := \tau(q, x)$ and by the continuity of angles (cf. Proposition 2.41), $\beta_n \to \beta$, where β is the angle between μ and γ. Consider the comparison triangle $(\overline{\mu}_n, \overline{\sigma}_n, \overline{\gamma}_n)$ in $\mathbb{R}^{1,1}$, then the angle $\overline{\beta}_n$ between $\overline{\mu}_n$ and $\overline{\gamma}_n$ satisfies $\overline{\beta}_n \leq \beta_n$ and similarly $\overline{\theta}_n \geq \theta_n$. Since $\beta_n \to \beta$, there is some $C > 0$ such that $\overline{\beta}_n \leq C$ for all n. The hyperbolic law of cosines in $\mathbb{R}^{1,1}$ (see Lemma 2.35) gives

$$b_n^2 = a_n^2 + c_n^2 - 2a_nc_n \cosh(\overline{\beta}_n),$$
$$c_n^2 = a_n^2 + b_n^2 + 2a_nb_n \cosh(\overline{\theta}_n).$$

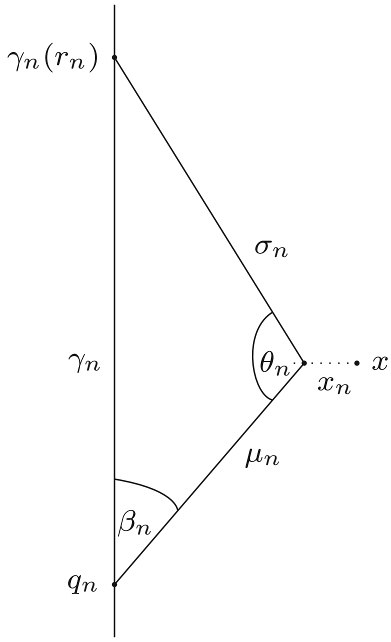

Fig. 6 The construction of maximal curves along $I(\gamma)$

Using these two equations and solving for $\cosh(\overline{\theta}_n)$, we get

$$\cosh(\overline{\theta}_n) = -\frac{a_n}{b_n} + \frac{c_n}{b_n}\cosh(\overline{\beta}_n).$$

By the initial equation for b_n, it is easy to see that $b_n/c_n \to 1$ and $b_n \to \infty$ (using that $a_n \to a$ and $\overline{\beta}_n \leq C$), hence there is some constant $\tilde{C} > 0$ such that $\cosh(\overline{\theta}_n) \leq \tilde{C}$ and thus also $\theta_n \leq C'$ for some constant C'. Using the monotonicity condition, we get that $\theta_n = \measuredangle_{x_n}(\sigma_n, \mu_n) \geq \overline{\measuredangle}_{x_n}(\sigma_n(s), \mu_n(t)))$ for each s and t, so this is bounded too. We will see that this is incompatible with σ_n "getting more and more null": Note that we get the following estimate for some constant C'':

$$C'' \geq \cosh(\overline{\measuredangle}_{x_n}(\sigma_n(s), \mu_n(t))) = \frac{\tau(\mu_n(t), \sigma_n(s))^2 - \tau(\mu_n(t), x_n)^2 - \tau(x_n, \sigma_n(s))^2}{2\tau(x_n, \sigma_n(s))\tau(\mu_n(t), x_n)}.$$

Since the denominator goes to 0 for $n \to \infty$ (as $\tau(x_n, \sigma_n(s)) \to \tau(x, \sigma(s)) = 0$ and $\tau(\mu_n(t), x_n) \to \tau(\mu(t), x) > 0$), the numerator has to go to 0 as well, in particular this means

$$\tau(\mu(t), \sigma(s)) = \tau(\mu(t), x)$$

for all s, t. This implies that running along μ from $\mu(t)$ to x and then along σ to $\sigma(s)$ gives a maximiser, but this curve has a timelike and a null piece, a contradiction to Theorem 2.17. □

Next, we show that any asymptote to γ (which we now know to be timelike) is complete.

Proposition 4.11 *Let $x \in I(\gamma)$ and let η be a timelike asymptote to γ at x. Then $L_\tau(\eta) = \infty$.*

Proof W.l.o.g. let η be future-directed. Suppose $L := L_\tau(\eta) < \infty$. By construction, $\eta : [0,\infty) \to X$ arises as a locally uniform limit of timelike maximisers $\eta_n : [0, a_n] \to X$ from x to $\gamma(t_n)$, where $t_n \to \infty$. By continuity of angles (see Proposition 2.41), $\measuredangle_x(\eta, \eta_n) \to 0$ for $n \to \infty$. Let $\varepsilon > 0$ and let $N \in \mathbb{N}$ be such that for $n \geq N$, $\measuredangle_x(\eta, \eta_n) < \varepsilon$ (in fact, any finite bound would suffice here, we do not necessarily need an arbitrarily small one). Since $\tau(x, \gamma(t_n)) \to \infty$, we may assume that $\tau(x, \gamma(t_n)) \geq 3L\cosh(\varepsilon)$ (upon possibly choosing a larger N). Now note that for any t_n with $n \geq N$, $\partial J^-(\gamma(t_n)) \cap \eta = (J^-(\gamma(t_n)) \setminus I^-(\gamma(t_n))) \cap \eta$ is non-empty: Certainly, $x \in I^-(\gamma(t_n))$, so if this intersection were empty, then η would be imprisoned in the compact set $J^+(x) \cap J^-(\gamma(t_n))$, which cannot happen. So we find a point $y_0 \in \eta$ that is null-related to $\gamma(t_n)$, i.e. $y_0 < \gamma(t_n)$ and $\tau(y_0, \gamma(t_n)) = 0$. By continuity, there is y on η near y_0 such that $0 < \tau(y, \gamma(t_n)) < 3L\cosh(\varepsilon)/2$. Let now ν be the part of η from x to y with length $L_\tau(\nu) =: a$, σ a timelike maximiser from y to $\gamma(t_n)$ with length $b := L_\tau(\sigma) = \tau(y, \gamma(t_n))$. Moreover, we write $c := L_\tau(\eta_n)$. Then (ν, σ, η_n) forms a timelike triangle with angle $\beta := \measuredangle_x(\eta, \eta_n) < \varepsilon$ at x, consider a corresponding comparison triangle $(\overline{\nu}, \overline{\sigma}, \overline{\eta}_n)$ with angle $\overline{\beta}$ at $\overline{x}$. By the law of cosines, we get

$$b^2 = a^2 + c^2 - 2ac\cosh(\overline{\beta}). \tag{4.3}$$

Our curvature assumption gives that $\overline{\beta} \leq \beta < \varepsilon$. Moreover, as we have argued, $c \geq 3L\cosh(\varepsilon)$ and $b < c/2$ and $a < L$. Inserting all of this, we get

$$\begin{aligned} b^2 &\geq a^2 + c^2 - 2ac\cosh(\varepsilon) \geq a^2 + c^2 - 2Lc\cosh(\varepsilon) \geq c^2 - 2Lc\cosh(\varepsilon) \\ &= c^2\left(1 - \frac{2L\cosh(\varepsilon)}{c}\right) \geq c^2/3, \end{aligned}$$

contradicting $b < c/2$. □

To conclude this subsection, we show that any future directed and any past directed (timelike) asymptote to γ from a common point fit together to give a timelike line.

Proposition 4.12 (Asymptotic lines) *Let $p \in I(\gamma)$ and consider any future and past rays $\sigma^+ : [0,\infty) \to X$ and $\sigma^- : (-\infty, 0] \to X$ from p to γ. Then $\sigma := \sigma^-\sigma^+ : \mathbb{R} \to X$ is a complete timelike line.*

Proof Let σ_n^+ and σ_n^- be two sequences of timelike maximisers from p to $\gamma(r_n)$ and $\gamma(-r_n)$, respectively, such that σ^+ and σ^- arise as limits of these sequences as $r_n \to \infty$. To show that σ is a line, it is sufficient to show that for any $t > 0$, $\tau(\sigma(-t), \sigma(t)) = L_\tau(\sigma|_{[-t,t]})$. To see this, let $q_+ := \sigma(t)$ and $q_- := \sigma(-t)$, and $q_+^n := \sigma_n^+(t)$, $q_-^n := \sigma_n^-(-t)$. Then $q_\pm = \lim_n q_\pm^n$. Consider the triangle going from x via σ_n^+ to the endpoint of σ_n^+, then following γ down to the endpoint of σ_n^-, and finally running the latter up to x again. Consider a comparison triangle with points $\overline{q_\pm^n}$ corresponding to $q_\pm^n$. Sending $n \to \infty$, we see that the stacked comparison triangles in $\mathbb{R}^{1,1}$ converge to a vertical line (here we use Proposition 4.6 and the completeness of γ), hence our curvature bound gives

$$\tau(q_-, q_+) = \lim_{n\to\infty} \tau(q_-^n, q_+^n) \leq \lim_{n\to\infty} \overline{\tau}(\overline{q_-^n}, \overline{q_+^n}) = \lim_n L_\tau(\sigma_n^- \sigma_n^+|_{[-t,t]})$$
$$= L_\tau(\sigma|_{[-t,t]}),$$

which is what we wanted to show, as the other inequality is trivial. □

4.2 Parallel lines

This subsection introduces the notion of parallelity for complete timelike lines. This approach is better suited for the synthetic case as it circumvents the analysis of Busemann functions. In the following, we call a map $f : Y_1 \to Y_2$ between Lorentzian pre-length spaces $(Y_1, d_1, \ll_1, \leq_1, \tau_1)$ and $(Y_2, d_2, \ll_2, \leq_2, \tau_2)$ *τ-preserving* if for all $p, q \in Y_1$, $\tau_1(p, q) = \tau_2(f(p), f(q))$, and we call f *causality preserving* if $p \leq_1 q$ if and only if $f(p) \leq_2 f(q)$.

Definition 4.13 (*Parallel lines*) Let α, β be two complete timelike lines in a Lorentzian pre-length space X defined on open intervals, w.l.o.g. on $\mathbb{R}$. They are called *parallel* if there exists a τ- and causality preserving map $f : (\alpha(\mathbb{R}) \cup \beta(\mathbb{R})) \to \mathbb{R}^{1,1}$ such that $f(\alpha(\mathbb{R}))$ and $f(\beta(\mathbb{R}))$ are parallel timelike lines in Minkowski (in the sense of parallel lines in affine spaces). We call such a map f a *parallel realisation* of α and β.

Remark 4.14 Note that by post-composing this by an isometry of Minkowski space, we can always achieve that $f(\alpha(\mathbb{R})) = \{(t, 0) : t \in \mathbb{R}\} \subseteq \mathbb{R}^{1,1}$ and $f(\beta(\mathbb{R})) = \{(t, c) : t \in \mathbb{R}\}$ for some $c \geq 0$. In this form, if α and β are parallel and both parametrized by τ-arclength (this is possible if τ is continuous and $\tau(x, x) = 0$ for all $x \in X$, cf. Lemma 2.29), we get that $f(\alpha(t)) = (t + a, 0)$ and $f(\beta(t)) = (t + b, c)$. By doing a shift in Minkowski, we can make a to be 0 (changing b to $b - a$).

Lemma 4.15 (Properties of parallel realisations) *Let X be a strongly causal Lorentzian pre-length space with continuous time separation, let $\alpha, \beta : \mathbb{R} \to X$ be two complete timelike lines and let $f : \alpha(\mathbb{R}) \cup \beta(\mathbb{R}) \to \mathbb{R}^{1,1}$ be a parallel realisation. Then f is a homeomorphism onto its image.*

Proof We first show that f is injective. Certainly, due to the fact that f preserves time separations, $f|_{\alpha(\mathbb{R})}$ and $f|_{\beta(\mathbb{R})}$ are injective: Indeed suppose that e.g. $f(\alpha(t)) = f(\alpha(s)) = (r, x)$ for some $s < t$, then $0 < \tau(\alpha(s), \alpha(t)) = \overline{\tau}((r, x), (r, x)) = 0$, a contradiction. Now suppose that the lines $f(\alpha(\mathbb{R}))$ and $f(\beta(\mathbb{R}))$ in $\mathbb{R}^{1,1}$ are different, then we are done. Otherwise, by nature of lines in Minkowski space, they have to be equal if they intersect. In that case, each point on that line is reached by one point on α and one point on β via f. So suppose that $(r, x) = f(\alpha(t)) = f(\beta(s))$. Since f is τ-preserving and τ is continuous, for each $\varepsilon > 0$ there are $\delta, \tilde{\delta} > 0$ such that $f(\beta(s + \delta)) = (r + \varepsilon, x)$ and $f(\beta(s - \tilde{\delta})) = (r - \varepsilon, x)$. But this means that $\alpha(t) \in \bigcap_{\delta,\tilde{\delta}\downarrow 0} I(\beta(s-\tilde{\delta}), \beta(s+\delta))$. By strong causality, these sets are a neighbourhood basis at $\beta(s)$, hence $\alpha(t) = \beta(s)$. This concludes the proof that f is injective.

Clearly, $\alpha(\mathbb{R}) \cup \beta(\mathbb{R})$ is a strongly causal Lorentzian pre-length space with the restriction of the structure of X, and similarly for their images in Minkowski space.

Moreover, f maps the timelike diamonds in $\alpha(\mathbb{R}) \cup \beta(\mathbb{R})$ to the timelike diamonds in $f(\alpha(\mathbb{R})) \cup f(\beta(\mathbb{R}))$. As these form topological bases due to strong causality, we conclude that f is a homeomorphism onto its image. □

Note that in the above setting, any parallel realisation f gives us a τ-arclength parametrisation of the parallel lines α, β: Let $\tilde{\alpha}$ be a τ-arclength parametrisation of the line $f(\alpha(\mathbb{R}))$ in Minkowski, then $f^{-1} \circ \tilde{\alpha}$ is τ-arclength parametrised, similarly for β.

Definition 4.16 (*Synchronised parallel lines*) Let X be a Lorentzian pre-length space with continuous time separation τ satisfying $\tau(x, x) = 0$ for all $x \in X$. Let $\alpha : \mathbb{R} \to X, \beta : \mathbb{R} \to X$ be two τ-arclength parametrised, parallel and complete timelike lines. They are called *synchronised parallel*[6] if the parallel realisation $f : (\alpha(\mathbb{R}) \cup \beta(\mathbb{R})) \to \mathbb{R}^{1,1}$ can be chosen to be of the form $f(\alpha(t)) = (t, 0)$ and $f(\beta(t)) = (t, c)$ for some $c \geq 0$. The constant c is a well-defined property of (α, β) and is called the *distance* of the parallel lines. For any two parallel lines α and β parametrised by τ-arclength, one can shift β such that $(\alpha, \beta \circ (t \mapsto t - a))$ is synchronised parallel.

Lemma 4.17 (The c-criterion for parallel lines) *Let X be a Lorentzian pre-length space with continuous time separation τ satisfying $\tau(x, x) = 0$ for all $x \in X$. Let $\alpha : \mathbb{R} \to X$ and $\beta : \mathbb{R} \to X$ be two τ-arclength parametrised complete timelike lines. We define the following partial functions:*

- $c_{\alpha\beta}(s, t) = \sqrt{(t - s)^2 - \tau(\alpha(s), \beta(t))^2} (\in \mathbb{C})$ *if* $\alpha(s) \leq \beta(t)$ *(otherwise undefined),*
- $c_{\beta\alpha}(s, t) = \sqrt{(s - t)^2 - \tau(\beta(t), \alpha(s))^2} (\in \mathbb{C})$ *if* $\beta(t) \leq \alpha(s)$ *(otherwise undefined),*
- $c^N_{\alpha+}(s) = \inf\{t - s : \alpha(s) \leq \beta(t)\}$,
- $c^N_{\beta+}(s) = \inf\{t - s : \beta(s) \leq \alpha(t)\}$.

These are all constant where defined, have the same value c and the infima are minima if and only if α and β are synchronised parallel. This constant c is the distance between α and β.

If X is additionally globally causally closed and $\alpha \cap I^+(\beta) \neq \emptyset$ and $\beta \cap I^+(\alpha) \neq \emptyset$, then the infima in $c^N_{\alpha+}$ and $c^N_{\beta+}$ are automatically minima and the conditions that $c^N_{\alpha+} = c^N_{\beta+} = c$ are automatically satisfied if $c_{\alpha\beta} = c_{\beta\alpha} = c$. In that case, all of the constants agree.

Proof We define f as in the definition of synchronised parallel with distance c: $f(\alpha(s)) := (s, 0)$, $f(\beta(s)) := (s, c)$. We will prove that f is causality preserving if and only if $c^N_{\alpha+}$ and $c^N_{\beta+}$ are constantly c, and under the assumption that this is the case f is τ-preserving if and only if $c_{\alpha\beta}$ and $c_{\beta\alpha}$ are constantly c wherever defined.

As α and β are future directed τ-arclength parametrised timelike lines, it is clear that f is causality and τ-preserving along α and along β, i.e. we only have to check the conditions on $f(\alpha(s)) \leq f(\beta(t))$, their τ-distance and the same for α, β reversed.

[6] This agrees with the usual notion of synchronising clocks in the sense that this is the parametrisation coming from synchronising the clocks on the particles α and β and parametrising α and β by their clock values.

So first, for causality preserving: We know that $f(\alpha(s)) \leq f(\beta(t)) \Leftrightarrow (s, 0) \leq (t, c) \Leftrightarrow (t - s) \geq c$. On the other hand, $\alpha(s) \leq \beta(t) \Leftrightarrow t - s \geq c^N_{\alpha+}(s)$ if this is a minimum and $\alpha(s) \leq \beta(t) \Leftrightarrow t - s > c^N_{\alpha+}(s)$ if it is only an infimum. In particular, these conditions are the same precisely when $c^N_{\alpha+}(s) = c$ and it is a minimum. Thus, $\alpha(s) \leq \beta(t) \Leftrightarrow f(\alpha(s)) \leq f(\beta(t))$ if and only if $c^N_{\alpha+}$ is constantly c and is always a minimum. Analogously, we get that $\beta(s) \leq \alpha(t) \Leftrightarrow f(\beta(s)) \leq f(\alpha(t))$ if and only if $c^N_{\beta+}$ is constantly c and is always a minimum.

We already proved that f is causality preserving. Next, we show that it preserves τ. On the Minkowski side, we have: $\bar{\tau}(f(\alpha(s)), f(\beta(t))) = \sqrt{(t-s)^2 - c^2}$ (if $f(\alpha(s)) \leq f(\beta(t))$). On the X side, we solve the defining equation for $c_{\alpha\beta}$ for $\tau(\alpha(s), \beta(t))$ to get: $\tau(\alpha(s), \beta(t)) = \sqrt{(t-s)^2 - c_{\alpha\beta}(s,t)^2}$ (if $\alpha(s) \leq \beta(t)$). Now note that as f is causality preserving, the side conditions when these equations can be applied match up (if they are not applied, $\tau(\alpha(s), \beta(t)) = 0 = \bar{\tau}(f(\alpha(s)), f(\beta(t)))$ anyway). If the side conditions are satisfied, we note that the equation differ only in one place: $\tau(\alpha(s), \beta(t))$ uses $c_{\alpha\beta}(s,t)$ and $\bar{\tau}(f(\alpha(s)), f(\beta(t)))$ uses c. Thus, we easily see that $\tau(\alpha(s), \beta(t)) = \bar{\tau}(f(\alpha(s)), f(\beta(t)))$ if and only if $c_{\alpha\beta}(s,t) = c$.

Analogously, we get that $\tau(\beta(t), \alpha(s)) = \bar{\tau}(f(\beta(t)), f(\alpha(s)))$ if and only if $c_{\beta\alpha}(s,t) = c$.

For the additional claim, fix s. If X satisfies the additional assumptions, $\{t : \alpha(s) \leq \beta(t)\}$ is closed and bounded (as is easily seen from the fact that $c_{\alpha\beta} \in \mathbb{R}$ when defined due to the reverse triangle inequality), thus it has a minimum automatically. We consider t such that $\alpha(s) \ll \beta(t)$, that is $t \in (t_s, +\infty) = \beta^{-1}(I^+(\alpha(s)))$. We get that $\tau(\alpha(s), \beta(t)) \to 0$ as $t \searrow t_s$. We transform the equation $c = \sqrt{(t-s)^2 - \tau(\alpha(s), \beta(t))^2}$ to $\tau(\alpha(s), \beta(t)) = \sqrt{(t-s)^2 - c^2}$. By continuity of τ, this still holds in the limit $t \searrow t_s$, i.e. $\sqrt{(t_s - s)^2 - c^2} = \tau(\alpha(s), \beta(t_s)) = 0$, so $t_s - s = c$ and $\{t : \alpha(s) \leq \beta(t)\} = [t_s, +\infty)$ as claimed. □

Lemma 4.18 (The strong causality trick) *Let X be a strongly causal Lorentzian pre-length space with continuous τ and $\tau(x,x) = 0$ for all $x \in X$. Let $\alpha, \beta : [0, b) \to \mathbb{R}$ be two τ-arclength parametrised timelike distance realisers with $x := \alpha(0) = \beta(0)$. Assume that for all s, t such that $\alpha(s)$ and $\beta(t)$ are timelike related, the comparison angle is $\tilde{\measuredangle}_x(\alpha(s), \beta(t)) = 0$. Then $\alpha = \beta$.*

Proof We set $f_+(s,t) = \tau(\alpha(s), \beta(t))$ and $f_-(s,t) = \tau(\beta(s), \alpha(t))$. They are both monotonically increasing in t and continuous. We want to describe the set where $f_+ > 0$ resp. $f_- > 0$. We set $t_s^+ = \inf\{t : f_+(s,t) > 0\}$, then $\lim_{t \searrow t_s^+} f_+(s,t) = 0$. Thus in the law of cosines (Lemma 2.35), we get $\lim_{t \searrow t_s^+} \cosh(\tilde{\measuredangle}_x(\alpha(s), \beta(t))) = 1 = \lim_{t \searrow t_s^+} \frac{s^2 + t^2 - f_+(s,t)}{2st} = \frac{s^2 + (t_s^+)^2}{2st_s^+}$, so $s = t_s^+$. Analogously, we get $s = \inf\{t : f_-(s,t) > 0\}$, in total $f_\pm(s,t) > 0$ for $s < t$.

That means that whenever $s_- < t < s_+$ we have that $\alpha(t) \in I(\beta(s_-), \beta(s_+))$. By Lemma 2.26, the $I(\beta(s_-), \beta(s_+))$ form a neighbourhood basis of the point $\beta(t)$, and $\alpha(t)$ is inside of all of these neighbourhoods. As X is Hausdorff, we get that $\alpha(t) = \beta(t)$ for all $t \in (0, b)$, so $\alpha = \beta$. □

Lemma 4.19 (Parallel lines are unique) *Let X be a strongly causal, timelike geodesically connected Lorentzian pre-length space with continuous τ, $\tau(x,x) = 0$ for all $x \in X$. Suppose that X has global timelike curvature bounded below by 0, let $\alpha : \mathbb{R} \to X$ be a complete timelike line and $p \in X$ a point. Then there is (up to reparametrisation) at most one parallel line to α through p.*

Proof We indirectly assume there are two parallel lines to α through p, namely $\beta : \mathbb{R} \to X$ and $\tilde{\beta} : \mathbb{R} \to X$. We assume all three curves are parametrised by τ-arclength, such that $\beta(0) = \tilde{\beta}(0) = p$ and in such a way that α, β are synchronised parallel. Then also $\alpha, \tilde{\beta}$ are parametrised in a way that they are synchronised parallel, and both α, β and $\alpha, \tilde{\beta}$ have the same distance: set $s_+ = \inf \alpha^{-1}(I^+(p))$ and $s_- = \sup \alpha^{-1}(I^-(p))$ the boundary of the points on α timelike related to p. Then the distance of α, β and that of $\alpha, \tilde{\beta}$ can be calculated as $\frac{s_+ - s_-}{2}$, so they are the same, and the necessary shift to make β and $\tilde{\beta}$ synchronised parallel to α can be calculated by $\frac{s_+ + s_-}{2}$, so they are the same too. As α, β are assumed to be synchronised already, this shift is 0.

We construct a comparison situation for all three lines at once: We choose the parallel realisation f for α and β given by $f(\alpha(s)) = (s, 0)$ and $f(\beta(s)) = (s, c)$, and similarly we choose $\tilde{f}$ for α and $\tilde{\beta}$ given by $\tilde{f}(\tilde{\beta}(s)) = (s, -c)$ and $\tilde{f}(\alpha(s)) = (s, 0)$, i.e. we realise β and $\tilde{\beta}$ on opposite sides of α.

We calculate the angle $\omega := \measuredangle_p(\beta|_{[0,+\infty)}, \tilde{\beta}|_{[0,+\infty)})$: By Proposition 4.8, this is equal to any comparison angle $\tilde{\measuredangle}_p(\beta(s), \tilde{\beta}(t))$ as long as $\beta(s) \ll \tilde{\beta}(t)$ or conversely. We know that $\beta(s) \ll \alpha(s + c + \varepsilon) \ll \tilde{\beta}(s + 2c + 2\varepsilon)$ and set $t = s + 2c + 2\varepsilon$. We now look at the law of cosines to estimate the comparison angle $\omega = \tilde{\measuredangle}_p(\beta(s), \tilde{\beta}(t))$: $\cosh(\omega) = \frac{s^2 + t^2 - \tau(\beta(s), \tilde{\beta}(t))^2}{2st} \leq \frac{2t^2}{2s^2}$ which converges to 1 as $s \to +\infty$. In particular, $\omega = 0$. We can now apply the strong causality trick to get $\beta = \tilde{\beta}$ (so they are even equal in this parametrisation), so there is only one line parallel to α through p. $\square$

Let us now show for X as in the splitting theorem that asymptotic lines to γ constructed in Proposition 4.12 are parallel to γ.

Lemma 4.20 *Let $\alpha : \mathbb{R} \to X$ be a complete timelike asymptotic line to γ. Then α, γ are parallel.*

Proof By construction, α arises as the limit of timelike maximising segments α_n^+ from $p := \alpha(0)$ to $\gamma(r_n)$ and α_n^- from p to $\gamma(-r_n)$, where $r_n \to \infty$. We will use the stacking principle (cf. Proposition 4.6) to show that α and γ are parallel. Indeed, the triangles $\Delta(\gamma(-r_n), p, \gamma(r_n))$ stack in $\mathbb{R}^{1,1}$, hence by orienting the side corresponding to the segment on γ vertically (see Fig. 7), we may assume that for $n \to \infty$ the comparison triangles converge to two vertical lines in $\mathbb{R}^{1,1}$.

We now argue that the map sending $\alpha(\mathbb{R})$ and $\gamma(\mathbb{R})$ to the corresponding limit lines $\bar{\alpha}$ and $\bar{\gamma}$ is a parallel realisation. For any t, s, consider $\tau(\alpha(t), \gamma(s)) = \lim_n \tau(\alpha_n(t), \gamma(s))$. For n so large that $-r_n < s < r_n$, $\gamma(s)$ is part of the triangle $\Delta(\gamma(-r_n), p, \gamma(r_n))$. From Corollary 4.9 we conclude that $\alpha_n(t) \ll \gamma(s)$ if and only if $\overline{\alpha}_n(t) \ll \overline{\gamma}(s)$, and $\tau(\alpha_n(t), \gamma(s)) = \overline{\tau}(\overline{\alpha}_n(t), \overline{\gamma}(s))$, where the bars denote the corresponding points on the comparison side. This shows that the realisation of α as $\bar{\alpha}$ and γ as $\bar{\gamma}$ is a parallel realisation. $\square$

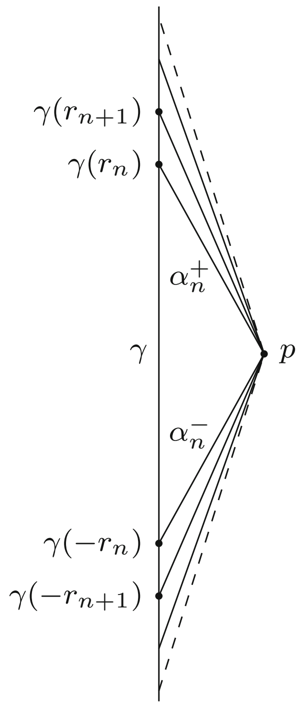

Fig. 7 Stacking comparison triangles

Lemma 4.21 (Shifting asymptotes) *Let $\alpha : \mathbb{R} \to X$ be a complete timelike asymptote to γ through $\alpha(0)$. Then the asymptote to γ through $\alpha(s)$ is $\alpha(\cdot - s)$ after parametrising both by τ-arclength.*

Proof Let $\tilde{\alpha}$ be the asymptotic line through $\alpha(s)$. Then the previous result shows that both γ and α as well as γ and $\tilde{\alpha}$ are parallel and they meet at $\alpha(s) = \tilde{\alpha}(0)$. Thus, as parallel lines are unique (see Lemma 4.19), we have that (after synchronising) $\alpha = \tilde{\alpha}$. To get the correct parameters before shifting, we know $\alpha(s) = \tilde{\alpha}(0)$, which fixes the shift parameter. □

Corollary 4.22 (Asymptotes stay in $I(\gamma)$) *Asymptotic lines to γ from points in $I(\gamma)$ stay in $I(\gamma)$.*

Proof Since asymptotic lines to γ are parallel to γ by Lemma 4.20, this readily follows. □

Now that we have established the parallelity of γ and its asymptotes, we may reparametrise all of them by τ-arclength and fix a shift on each asymptote so that they are synchronised to γ. Note that we fix a τ-arclength parametrisation of γ, namely the one keeping $\gamma(0)$ the same.

Definition 4.23 (*Busemann parametrisation*) Let $\alpha : \mathbb{R} \to X$ be a complete timelike asymptote to γ. Then we call the (unique) τ-arclength parametrisation of α that synchronises it to γ its *Busemann parametrisation*.

Remark 4.24 (Busemann functions) Up to this point, we have avoided the use of Busemann functions in our treatment. Though it will not really be necessary in what follows, let us discuss them briefly here. Let X be as in the splitting theorem and γ the given timelike line in τ-arclength parametrisation. Then we define the (*future*) *Busemann* *function* of γ as $b^+ : I(\gamma) \to \mathbb{R} \cup \{\pm\infty\}$ via $b^+(x) := \lim_{t\to\infty}(t - \tau(x, \gamma(t)))$. If $\alpha : \mathbb{R} \to X$ is any asymptotic line to γ in τ-arclength parametrisation and we choose the parallel realisation between α and γ so that $\gamma(t)$ is mapped to $(t, 0) \in \mathbb{R}^{1,1}$ (which we always do) and $\alpha(s)$ is mapped to $(s + S, c) \in \mathbb{R}^{1,1}$, then

$$b^+(\alpha(s)) = \lim_{t\to\infty}(t - \tau(\alpha(s), \gamma(t))) = \lim_{t\to\infty}(t - \sqrt{(t - (s + S))^2 - c^2}) = s + S.$$

Hence, b^+ indicates the shift in the time parameter along any τ-arclength parametrised asymptote. In particular, b^+ is finite-valued on all of $I(\gamma)$ since any point lies on an asymptote, and in fact it is not hard to show that b^+ is continuous on $I(\gamma)$ in a similar fashion as [11, Lem. 3.3], but we will not need this. Note that if α is in Busemann parametrisation, then $b^+(\alpha(s)) = s$ for all $s \in \mathbb{R}$, hence the name.

Lemma 4.25 (Parallelity is weakly transitive) *Let X be a strongly causal, timelike geodesically connected Lorentzian pre-length space with continuous τ, $\tau(x, x) = 0$ for all x, and global non-negative timelike curvature. Let $\alpha, \beta, \gamma : \mathbb{R} \to X$ be timelike lines such that α, β and β, γ are parallel. Assume that there is a point p on γ such that there is a parallel line to α through p. Then this parallel line is already γ.*

Proof We assume the parallel lines are even synchronised parallel in a suitable τ-arclength parametrisation. Let $\tilde{\gamma}$ be the synchronised parallel line to α through p. Let a be the distance between α, β, b be the distance between β, γ and c be the distance between $\alpha, \tilde{\gamma}$. We have two lines through p: γ and $\tilde{\gamma}$. We assume the parameters for p are $p = \gamma(t_0) = \tilde{\gamma}(\tilde{t}_0)$. By Proposition 4.8, we know that $\tilde{\measuredangle}_p(\gamma(s), \tilde{\gamma}(t))$ is constant in s and t as long as $\gamma(s) \ll \tilde{\gamma}(t)$ (or conversely). We know that $\gamma(s) \ll \beta(s+b+\varepsilon) \ll \alpha(s + a + b + 2\varepsilon) \ll \tilde{\gamma}(s + a + b + c + 3\varepsilon)$ and set $t = s + a + b + c + 3\varepsilon$. We now look at the law of cosines to estimate the comparison angle $\omega = \tilde{\measuredangle}_p(\gamma(s), \tilde{\gamma}(t))$: $\cosh(\omega) = \frac{(s-t_0)^2+(t-\tilde{t}_0)^2-\tau(\gamma(s),\tilde{\gamma}(t))}{2(s-t_0)(t-\tilde{t}_0)} \leq \frac{2(t-\tilde{t}_0)^2}{2(s-t_0)^2}$ which converges to 0 as $s \to +\infty$. In particular, $\omega = 0$. We can now apply the strong causality trick to get that γ and $\tilde{\gamma}$ are just shifts of each other, so α and γ are parallel.

By doing the same argument backwards as well, i.e. $\alpha(s - a - b - 2\varepsilon) \ll \gamma(s) \ll \alpha(s + a + b + 2\varepsilon)$, we get that the synchronised version only differs by a shift of at most $a + b$ (upon taking $\varepsilon \to 0$). □

Lemma 4.26 (Verticality in Minkowski space) *Let $a = (0, 0) \ll b$ be points in Minkowski space $\mathbb{R}^{1,1}$. Let the t-coordinate difference be $\Delta t = t(b) - t(a)$. Let c_n be points with $b \ll c_n$, the t-coordinate $t(c_n) \to +\infty$, c_n above the straight line through a, b and $\tau(a, c_n) - \tau(b, c_n) \to \Delta t$ as $n \to +\infty$. Then $\frac{x(c_n)}{t(c_n)} \to 0$, or equivalently, $\frac{c_n}{\|c_n\|} \to \partial_t$.*[7]

Proof We show that $\forall\varepsilon > 0$ there are constants $T > 0$ and $\delta > 0$ such that for all p with $\tau(a, p) \geq T$ and $|\tau(a, p) - \tau(b, p) - (t_b - t_a)| < \delta$, we have that the angle of ap with the t-axis satisfies $\measuredangle_a(\partial_t, p) < \varepsilon$.

We set $\omega_1 = \measuredangle_a(\partial_t, b)$ the angle between the vertical and ab, and $\omega_2 = \measuredangle_a(p, b)$ the angle between the vertical and ap. We will prove $|\omega_1 - \omega_2| < \varepsilon$. By assumption,

[7] So c_n converges to the point $[\partial_t]$ on the limit sphere.

p and ∂_t lie on the same side of the straight line through ab, so by angle additivity in the plane we also get that $\measuredangle_a(\partial_t, p) < \varepsilon$.

We apply the law of cosines:

$$\tau(b,p)^2 - \tau(a,p)^2 = \tau(a,b)^2 - 2\tau(a,b)\tau(a,p)\cosh(\omega_2)$$

and use that $\cosh(\omega_1)\tau(a,b) = t_b - t_a$ to get:

$$\tau(a,p) - \tau(b,p) = (t_b - t_a)\underbrace{\frac{(2\tau(a,p)\cosh(\omega_2) - 1)}{\cosh(\omega_1)(\tau(b,p) + \tau(a,p))}}_{=:F}.$$

We now have to check that whenever $|\omega_1 - \omega_2| \geq \varepsilon$, the factor F is bounded away from 1.

We have two cases to cover: First, we indirectly assume $\omega_1 > \omega_2 + \varepsilon$ (then $1 - \frac{\cosh(\omega_2)}{\cosh(\omega_1)} > 1 - \cosh(\varepsilon)$). (Assuming b lies to the right of the t-axis, this is the case where p also lies on the right.) We will show $F < 1 - \delta$ (if $\tau(a,p)$ is large enough and $\tau(a,p), \tau(b,p)$ are close enough). This is equivalent to:

$$(2\tau(a,p)\cosh(\omega_2) - 1) < (1-\delta)\cosh(\omega_1)(\tau(b,p) + \tau(a,p))$$

Dividing by $\tau(a,p)$, this is obvious: $\frac{\tau(b,p)}{\tau(a,p)} \to 1$ as $\tau(a,p) \to +\infty$ and the difference is bounded. We can e.g. choose

$$\delta < \cosh(\varepsilon) - 1$$

and T large enough.

Now we indirectly assume $\omega_1 < \omega_2 - \varepsilon$ (then $\frac{\cosh(\omega_2)}{\cosh(\omega_1)} - 1 > \cosh(\varepsilon) - 1$) (Assuming b lies to the right of the t-axis, this is the case where p lies on the left.) We will show $F > 1 + \delta$ (if $\tau(a,p)$ is large enough and $\tau(a,p), \tau(b,p)$ are close enough). We need to check:

$$(2\tau(a,p)\cosh(\omega_2) - 1) > (1+\delta)\cosh(\omega_1)(\tau(b,p) + \tau(a,p))$$

Dividing by $\tau(a,p)$, this is obvious: $\frac{\tau(b,p)}{\tau(a,p)} \to 1$ as $\tau(a,p) \to +\infty$ and the difference is bounded. We can e.g. choose

$$\delta < \cosh(\varepsilon) - 1$$

and T large enough. □

We now run into a subtlety: being parallel is in general only weakly transitive, but being synchronised parallel is probably not. Luckily, for X as in the splitting theorem, the fact that we consider asymptotic lines to γ simplifies the situation:

Lemma 4.27 (Two asymptotes are parallel) *Let $x, y \in I(\gamma)$ and α, β the asymptotes to γ through x resp. y in Busemann parametrisation. Then α, β are synchronised parallel.*

Proof We dive into the construction of asymptotes: We first assume $x \ll y$. We set $s_0 = b^+(x)$ and $t_0 = b^+(y)$, then there are maximisers α_n from x to $\gamma(T_n)$ parametrised in τ-arclength such that $\alpha_n(s_0) = x$ and β_n from y to $\gamma(T_n)$ parametrised in τ-arclength such that $\beta_n(t_0) = y$, which converge (pointwise) to the upper parts of α and β in Busemann parametrisation for $T_n \to \infty$, respectively. Note that $\alpha_n(s_0 + \tau(x, \gamma(T_n))) = \beta_n(t_0 + \tau(y, \gamma(T_n))) = \gamma(T_n)$.

We try to prove the c-criterion (Lemma 4.17): Clearly, due to our assumptions on X, we do not need to consider the null functions. We look at $c_{\alpha\beta}(s_0, t_0)$ and $c_{\alpha\beta}(s', t')$ or $c_{\beta\alpha}(s', t')$ for $s_0 \le s'$, $t_0 \le t'$: We require $a := x = \alpha(s_0) \ll b := y = \beta(t_0)$. We get points converging to these parameters: $a'_n := \alpha_n(s') \to a' := \alpha(s')$ and $b'_n := \beta_n(t') \to b' := \beta(t')$ (note that $\alpha_n(s_0) = a$ and $\beta_n(t_0) = b$ anyway). We get a timelike triangle $\Delta_n = \Delta(a, b, c_n := \gamma(T_n))$ containing the points (a'_n, b'_n). We form a comparison situation for this in Minkowski space $\mathbb{R}^{1,1}$: $\bar{\Delta}_n = \Delta(\bar{a}, \bar{b}, \bar{c}_n)$ with comparison points $\bar{a}'_n, \bar{b}'_n$. As X has global curvature bounded below by 0, we get that $\tau(a'_n, b'_n) \le \bar{\tau}(\bar{a}'_n, \bar{b}'_n)$ and $\tau(b'_n, a'_n) \le \bar{\tau}(\bar{b}'_n, \bar{a}'_n)$.

We would now like to let $T_n \to +\infty$, so we need control over what the comparison triangle $\bar{\Delta}_n$ converges to. For this, we select which way to realise $\bar{\Delta}_n$ in $\mathbb{R}^{1,1}$. We choose:

- $\bar{a} = (s_0, 0)$ constant in n, i.e. at the right t-coordinate and with x-coordinate 0,
- $\bar{b} = (t_0, c_{\alpha\beta}(s_0, t_0))$ constant in n. Note this has the right τ-distance to $\bar{a}$ by definition of $c_{\alpha\beta}$.
- $\bar{c}_n$ above the straight line $\bar{a}$ and $\bar{b}$ lie on.

We claim that in the limit $T_n \to +\infty$, the line $\bar{a}\bar{c}_n$ becomes vertical (the line $\bar{b}\bar{c}_n$ also becomes vertical). For this, we look at what $\bar{\tau}(\bar{a}, \bar{c}_n) - \bar{\tau}(\bar{b}, \bar{c}_n)$ does as $T_n \to +\infty$: Using Landau small-oh notation, we get that $\bar{\tau}(\bar{a}, \bar{c}_n) = T_n - b^+(a) + o(1)$ and $\bar{\tau}(\bar{b}, \bar{c}_n) = T_n - b^+(b) + o(1)$ as $T_n \to +\infty$, so the difference is $b^+(b) - b^+(a) + o(1)$, which approaches the t-coordinate-difference $t_0 - s_0$ of $\bar{a}$ and $\bar{b}$. We can now apply Lemma 4.26 to get that $\frac{x(c_n)}{t(c_n)} \to 0$ and $t(c_n) \to +\infty$, thus these lines become vertical.

We have proven that the comparison points converge, so we have limit comparison points: $\bar{a} = (s_0, 0)$ and $\bar{b} = (t_0, c_{\alpha\beta}(s_0, t_0))$ stay fixed anyway, $\bar{a}'_n \to \bar{a}' = (s', 0)$, $\bar{b}'_n \to \bar{b}' = (t', c_{\alpha\beta}(s_0, t_0))$. We also set the comparison sides $\bar{\alpha}(s) = (s, 0)$ and $\bar{\beta}(t) = (t, c_{\alpha\beta}(s_0, t_0))$. By continuity of τ and $\bar{\tau}$, our curvature assumption gives $\tau(a', b') \le \bar{\tau}(\bar{a}', \bar{b}')$ in the limit (similarly with arguments flipped). Now we compare the definition of $c_{\alpha\beta}(s', t')$ resp. $c_{\beta\alpha}(s', t')$ with the c-functions for $\bar{\alpha}$ and $\bar{\beta}$ at the same parameters:

$$c_{\alpha\beta}(s', t') = \sqrt{(t' - s')^2 - \tau(a', b')^2},$$

$$c_{\bar{\alpha}\bar{\beta}}(s', t') = \sqrt{(t' - s')^2 - \bar{\tau}(\bar{a}', \bar{b}')^2} = x(\bar{b}') - x(\bar{a}') = c_{\alpha\beta}(s_0, t_0),$$

whenever defined. The last equality holds because we know $\bar{\alpha}$ and $\bar{\beta}$ are synchronised parallel with distance $c_{\alpha\beta}(s_0, t_0)$. Note that $a' \le b'$ implies $\bar{a}' \le \bar{b}'$ by the curvature bound and a simple continuity argument, so whenever the first line is defined so is the second. Notice these equations only differ in the τ term, and we know $\tau(a', b') \leq \bar{\tau}(\bar{a}', \bar{b}')$, so we get $c_{\alpha\beta}(s', t') \geq c_{\alpha\beta}(s_0, t_0)$ whenever the former is defined. Similarly, if $b' \leq a'$ we get $c_{\beta\alpha}(s', t') \geq c_{\alpha\beta}(s_0, t_0)$ for all $s' \geq s_0$ and $t' \geq t_0$. We can also do this if $a \gg b$, giving $c_{\alpha\beta}(s', t') \geq c_{\beta\alpha}(s_0, t_0)$ and $c_{\beta\alpha}(s', t') \geq c_{\beta\alpha}(s_0, t_0)$.

Doing the same construction towards the past, we similarly get $c_{\alpha\beta}(s', t') \geq c_{\alpha\beta}(s_0, t_0)$ and $c_{\beta\alpha}(s', t') \geq c_{\alpha\beta}(s_0, t_0)$ resp. $c_{\alpha\beta}(s', t') \geq c_{\beta\alpha}(s_0, t_0)$ and $c_{\beta\alpha}(s', t') \geq c_{\beta\alpha}(s_0, t_0)$ (depending on whether $a \ll b$ or $b \ll a$) for all $s' \leq s_0$ and $t' \leq t_0$.

Now we use Lemma 4.21 to get that the (Busemann parametrised) asymptote to γ through $\alpha(s)$ is α, and similarly the asymptote to γ through $\beta(t)$ is β. In particular, we can use the above argument again for $\alpha(s)$ instead of x and $\beta(t)$ instead of y. We know that $b^+(\alpha(s)) = s$ and $b^+(\beta(t)) = t$. The above (to the future, $a \ll b$) then gives: $c_{\alpha\beta}(s', t') \geq c_{\alpha\beta}(s, t)$ for $s \leq s', t \leq t'$ such that $a \ll b$ and $a' \ll b'$, extending continuously, $c_{\alpha\beta}$ is monotonously increasing where defined. On the other hand, the above (to the past, $a \ll b$) gives: $c_{\alpha\beta}(s', t') \geq c_{\alpha\beta}(s, t)$ for $s \geq s', t \geq t'$ such that $a \ll b$ and $a' \ll b'$, extending continuously, $c_{\alpha\beta}$ is monotonously decreasing where defined. Via a two-step process, we see that $c_{\alpha\beta}$ is constant where defined. Similarly, we get that also $c_{\beta\alpha}$ is constant where defined, and that they have the same value.

Thus, the c-criterion (Lemma 4.17) yields that α and β are synchronised parallel. □

This result establishes the transitivity of synchronised parallel lines: Any two asymptotes to γ in Busemann parametrisation are synchronised parallel.

5 Proof of the main result

Let us summarise what we have shown so far: Let $(X, d, \ll, \leq, \tau)$ be a connected, regularly localisable, globally hyperbolic Lorentzian length space satisfying timelike geodesic prolongation, with proper metric d and global non-negative timelike curvature containing a complete timelike line $\gamma : \mathbb{R} \to X$. Then from each point in $I(\gamma) = I^+(\gamma) \cap I^-(\gamma)$ we can construct (unique) asymptotic rays to γ, all of which are timelike with infinite τ-length. Future directed and past directed rays from a common point fit together to give a timelike line which is parallel to γ. We fix a τ-arclength parametrisation of γ fixing $\gamma(0)$, and always consider it to be mapped to the t-axis in $\mathbb{R}^{1,1}$ in any parallel realisation. Then, by way of the Busemann parametrisation, all asymptotic lines become synchronised parallel to γ and to each other. This synchronisation selects a "spacelike" slice (namely $\{b^+ = 0\}$) in X which will provide the metric part of the splitting.

In this section, we complete the proof of the splitting theorem in two steps: First, we show that $I(\gamma)$ splits, then we show that $I(\gamma)$ has the (TC) property, which implies $I(\gamma) = X$ due to Theorem 2.47.

Definition 5.1 (*Spacelike slice*) We call the set of all $\alpha(0)$, where α is a Busemann parametrised timelike asymptotic line to γ, the *spacelike slice* and denote it by S. For $\alpha(0), \beta(0) \in S$, we define $d_S(\alpha(0), \beta(0))$ to be the distance between α and β in the sense of parallel lines.

Lemma 5.2 (S, d_S) *is a metric space.*

Proof As asymptotes to γ are parallel to γ and parallel lines are unique, d_S is well-defined.

Let $p, q \in S$. It is obvious from the definition of the distance of two parallel lines that $d_S(p, q) \geq 0$. If $d_S(p, q) = 0$, we get that the asymptotes through p and q are the same curve α. As the Busemann parametrisation is unique, we get that $p = q$.

For the triangle inequality, let $p, q, r \in S$. We get asymptotes α through p, β through q and γ through r, and by Lemma 4.27, they are pairwise synchronised parallel. Let $d_1 = d_S(p, q)$ and $d_2 = d_S(q, r)$, then $\alpha(0) \leq \beta(d_1)$ by the $c^N_{\alpha+}(0)$ criterion (see Lemma 4.17) for α, β and $\beta(d_1) \leq \gamma(d_1 + d_2)$ by the $c^N_{\beta+}(d_1)$ criterion for β, γ. Thus, we get that $\alpha(0) \leq \gamma(d_1 + d_2)$, giving $d_S(p, r) \leq d_1 + d_2$ by the $c^N_{\alpha+}(0)$ criterion for α, γ. ☐

Definition 5.3 (*The splitting map*) We define $f : \mathbb{R} \times S \to I(\gamma) \subset X$ by $f(s, p) = \alpha_p(s)$, where α_p is the asymptote to γ through p in Busemann parametrisation, that is $\alpha_p(0) = p$.

Proposition 5.4 (Local splitting) *Let X be a connected, regularly localisable, globally hyperbolic Lorentzian length space with proper metric d and global non-negative timelike curvature satisfying timelike geodesic prolongation and containing a complete timelike line $\gamma : \mathbb{R} \to X$. Then $I(\gamma) \subset X$ is a causally convex open set that is itself a path-connected, regularly localisable, globally hyperbolic Lorentzian length space of global non-negative timelike curvature with the metric, relations and time separation induced from X. Moreover, the spacelike slice S is a proper (hence complete), strictly intrinsic metric space, the Lorentzian product $\mathbb{R} \times S$ is a path-connected, regularly localisable, globally hyperbolic Lorentzian length space and the splitting map $f : \mathbb{R} \times S \to I(\gamma)$ is a τ- and causality preserving homeomorphism.*

Proof First, it is clear that $I(\gamma)$ is path-connected, causally convex in X and has global non-negative timelike curvature. It is hence causally path-connected since X is and it is trivially locally causally closed. Moreover, if $x \in I(\gamma)$ and U is a regular localising neighbourhood of x in X, then $U \cap I(\gamma)$ is a regular localising neighbourhood of x in $I(\gamma)$, hence $I(\gamma)$ is regularly localisable. By causal convexity, the time separation between causally related points in $I(\gamma)$ is achieved as the supremum of lengths of causal curves running between them which have to stay inside $I(\gamma)$. The causal diamonds in $I(\gamma)$ are precisely those in X, since they must be contained in $I(\gamma)$. Finally, $I(\gamma)$ is non-totally imprisoning (it inherits this from X), thus we have shown all the claims on $I(\gamma)$.

Next, we argue that f is τ- and causality preserving: Let $p, q \in S$. Let α be the asymptote to γ through p and let β be the asymptote to γ through q. Then α and β are synchronised parallel with distance $d_S(p, q)$, so there is a parallel realisation $\tilde{f} : \alpha(\mathbb{R}) \cup \beta(\mathbb{R}) \to \mathbb{R}^{1,1}$ defined by $\tilde{f}(\alpha(s)) = (s, 0)$ and $\tilde{f}(\beta(t)) = (t, d_S(p, q))$. In particular, $\alpha(s) \leq_X \beta(t) \Leftrightarrow t - s \geq d_S(p, q)$ and if this is true, $\tau(\alpha(s), \beta(t)) = \sqrt{(t - s)^2 - d_S(p, q)^2}$. But that is just the definition of $(s, p) \leq_{\mathbb{R}\times S} (t, q)$ resp. $\tau_{\mathbb{R}\times S}((s, p), (t, q))$ in $\mathbb{R} \times S$. This is true for all $p, q \in S$ and $s, t \in \mathbb{R}$, so f is τ- and causality preserving. f is injective as asymptotes to γ are parallel to γ and parallel lines are unique, and surjective since any point in $I(\gamma)$ lies on an asymptote.

From the discussion above, it is easy to see that f maps timelike (and causal) diamonds in $I(\gamma)$ to timelike (and causal) diamonds in $\mathbb{R}\times S$. As both sides are strongly causal, this implies that f is a continuous open bijection, hence a homeomorphism. Since $\mathbb{R} \times S$ is always non-totally imprisoning (cf. Proposition 3.4) and its causal diamonds are compact (as continuous images of compact causal diamonds in $I(\gamma)$), we conclude that $\mathbb{R}\times S$ is globally hyperbolic, hence S is proper by Proposition 3.7. To see that S is a strictly intrinsic space, fix $p, q \in S$ and connect any two timelike related points on the corresponding asymptotes by a distance realiser in $I(\gamma)$. The image of that distance realiser under f is a continuous distance realiser in $\mathbb{R} \times S$, hence by Proposition 3.5 the projection onto S gives a distance minimiser in S between p and q.

Finally, we need to prove the remaining claimed properties of $\mathbb{R} \times S$. Path-connectedness is inherited from $I(\gamma)$ via f, and products are always globally causally closed (cf. Proposition 3.2). Note that $I(x, y)$ for $x, y \in I(\gamma)$ are regular localising neighbourhoods in $I(\gamma)$. Since $f(I(x, y)) = I(f(x), f(y))$ and the d-lengths of causal curves in timelike diamonds can always be uniformly bounded in products (cf. the proof of Proposition 3.4), timelike diamonds in $\mathbb{R}\times S$ are in fact (regular) localising neighbourhoods: For the local time separation, take the restriction of $\tau_{\mathbb{R}\times S}$, and note that maximisers in $I(\gamma)$ (or in any $I(x, y) \subset I(\gamma)$) map to continuous maximisers in $\mathbb{R} \times S$, which are always Lipschitz reparametrisable and are hence causal curves (cf. Proposition 3.5). □

Proposition 5.5 (Global splitting) *Let $(X, d, \ll, \leq, \tau)$ satisfy the assumptions in Proposition 5.4. Then $I(\gamma) = X$, i.e. the splitting in Proposition 5.4 is global.*

Proof Suppose $\alpha : [a, b) \to I(\gamma)$ is a future directed timelike geodesic with finite τ-length. We reparametrise α by τ-arclength and denote the reparametrised curve again by α, it is then defined on $[0, L)$, where $L := L_\tau(\alpha) < \infty$. Then $f^{-1} \circ \alpha$ is a continuous timelike geodesic in $\mathbb{R} \times S$, where f is the splitting map. We decompose it into time- and space-component: $f^{-1} \circ \alpha = (\alpha^t, \alpha^x)$. If α^x is constantly p, we see that $\alpha = \alpha_p$, a contradiction since asymptotes have infinite τ-length. Note that since the τ-lengths of $f^{-1} \circ \alpha$ and α agree, $f^{-1} \circ \alpha$ also has finite τ-length in $\mathbb{R} \times S$. By Corollary 3.6, α^x has constant speed and domain $[0, L)$, thus finite d_S-length.

α^t can certainly be extended to L. As S is a complete metric space, we can extend α^x continuously to L. Mapping this extension back to $I(\gamma)$ via f, we get a continuous extension in $I(\gamma)$ of α to b (in its original parametrisation). This shows that $I(\gamma)$ has the (TC) property, thus $I(\gamma) = X$ by Theorem 2.47. This concludes the proof of the splitting in Theorem 1.4. □

Recall the notion of Cauchy sets and (Cauchy) time functions on Lorentzian pre-length spaces from [19, Sec. 5.1]: A Cauchy set is any subset that is met exactly once by doubly inextendible causal curves, and a Cauchy time function is a continuous function $t : X \to \mathbb{R}$ such that $x < y$ implies $t(x) < t(y)$ and the image under t of any doubly inextendible causal curve is all of $\mathbb{R}$.

Corollary 5.6 *Let $(X, d, \ll, \leq, \tau)$ satisfy the assumptions of Theorem 1.4 and let $f : \mathbb{R} \times S \to X$ be the splitting. Then the sets $S_t := f(\{t\} \times S)$ are Cauchy sets in X that* *are all homeomorphic to S. Moreover, the map $pr_1 \circ f^{-1}$ is a Cauchy time function. Moreover, all Cauchy sets in X are homeomorphic to S.*

Proof Let $\alpha : (a, b) \to X$ be any doubly inextendible causal curve in X meeting any S_t twice, then its image under f^{-1} meets $\{t\} \times S$ twice, which cannot happen since $\{t\} \times S$ is acausal in $\mathbb{R} \times S$. Next we argue that any such α meets S_t: For simplicity let $t = 0$, and suppose that α does not meet S, so w.l.o.g. we may assume that $\alpha \subset I^+(S)$ (this is because $X = I^-(S) \cup S \cup I^+(S)$ due to the splitting, and this union is disjoint). Let $t_0 \in (a, b)$, then $\alpha((a, t_0]) \subset J^-(\alpha(t_0)) \cap J^+(S)$ which is easily seen to be a compact set by considering the corresponding situation in $\mathbb{R} \times S$. But this is a contradiction, since X is non-totally imprisoning. Hence S (and any S_t) is a Cauchy set. They are homeomorphic to S since X has the topology of $\mathbb{R} \times S$.

We now argue that $pr_1 \circ f^{-1}$ is a Cauchy time function for X. It is clearly a time function, as it is continuous and the time separation on $\mathbb{R} \times S$ is more or less a reformulation of that. Now let $\alpha : (a, b) \to X$ be a doubly inextendible causal curve, we need to show that $pr_1 \circ f^{-1} \circ \alpha((a, b)) = \mathbb{R}$. Suppose not, so there is some time value t_0 that is not attained. W.l.o.g. suppose $t_0 \geq 0$, so the image is contained in $(-\infty, t_0]$. Similarly to before, this would imply that $\alpha|_{[t_0, b)}$ is contained in the compact set $J^+(\alpha(t_0)) \cap J^-(S_{t_0})$, a contradiction.

Finally, let C be any Cauchy set in X. Then the projection $C \to S$ (via the splitting) is continuous. Its inverse is given by sending each $p \in S$ to the unique point on α_p meeting C. This is a continuous map since C is achronal. This shows that C is homeomorphic to S. □

Note that in general, Cauchy sets in globally hyperbolic Lorentzian pre-length spaces need not be homeomorphic, as [19, Ex. 5.7, Ex. 5.8] show.

Corollary 5.7 *Let $(X, d, \ll, \leq, \tau)$ satisfy the assumptions of Theorem 1.4. Then (S, d_S) has Alexandrov curvature ≥ 0.*

Proof Due to the splitting, the space $\mathbb{R} \times S$ has global nonnegative timelike curvature (cf. Remark 2.33; the splitting maps Lipschitz triangles in $\mathbb{R} \times S$ to continuous triangles in X, but this does not lead to any issues), and then it follows from [5, Thm. 5.7] that S has Alexandrov curvature ≥ 0. □

6 Applications and outlook

Let us now give several reformulations of our splitting result for spacetimes. We start with smooth spacetimes, where we get a weaker version of the smooth splitting result proven in [11].

Corollary 6.1 (Smooth spacetimes) *Let (M, g) be a smooth, globally hyperbolic spacetime with non-positive timelike sectional curvature containing a complete timelike line $\gamma : \mathbb{R} \to M$. Then the Lorentzian length space induced by (M, g) splits in the sense of Theorem 1.4.*

Proof Fix any complete Riemannian metric h on M and let d_h be the induced distance, then $(M, d_h, \ll, \leq, \tau)$ is a connected, regularly localisable (via convex neighbourhoods) Lorentzian length space and d_h is a proper metric (note that the τ-length of causal curves agrees with the usual notion of length in (M, g) due to [48, Prop. 2.32]). The fact that (M, g) has global non-negative timelike curvature in the sense of Definition 2.32 is a consequence of the Lorentzian Toponogov theorem (see [43, p. 303]). The remaining technical assumptions are easily seen to hold. Hence the Lorentzian length space induced by (M, g) can be split according to Theorem 1.4. □

Remark 6.2 In the setting of Corollary 6.1, once one has shown that the spacelike slice S is a spacelike hypersurface in M and the normal asymptotic distance d_S is the Riemannian distance on S (e.g. via an analysis of the gradient of the Busemann function as in [11]), the Hawking–King–McCarthy theorem [44], which states that any τ-preserving homeomorphism between strongly causal spacetimes is a smooth isometry, provides a somewhat alternative proof of the smooth splitting result proven in [11].

Next, we turn to $C^{1,1}$-spacetimes. First note that the Lorentzian pre-length space induced by any strongly causal Lipschitz spacetime is in fact a Lorentzian length space by [48, Thm. 5.12]. For $C^{1,1}$-metrics, the notion of sectional curvature is not available any more, so we have to assume synthetic curvature bounds. However, most of the other technical assumptions are easily seen to be satisfied due to the existence of convex neighbourhoods. For details on low regularity spacetimes, see e.g. the recent review article [59] and the references therein.

Corollary 6.3 ($C^{1,1}$-spacetimes) *Let (M, g) be a globally hyperbolic $C^{1,1}$-spacetime with global non-negative timelike curvature containing a complete timelike line $\gamma : \mathbb{R} \to M$. Then the Lorentzian length space induced by (M, g) splits in the sense of Theorem 1.4.*

Finally let us mention the case of C^1-spacetimes, which are substantially different. Notably, the Lorentzian length space induced by a globally hyperbolic C^1-spacetime need not satisfy the timelike geodesic prolongation property. This is because solutions of the g-geodesic equation need not maximise locally, nor does a maximising segment necessarily extend as a local maximiser beyond its endpoints. However, points in C^1-spacetimes still have neighbourhood bases consisting of globally hyperbolic, causally convex neighbourhoods, so regular localisability is satisfied. Assuming the necessary synthetic properties as well, we have the following (note that geodesics in the synthetic sense are locally maximising curves, which are necessarily C^2-solutions of the geodesic equation even if the metric is just C^1):

Corollary 6.4 (C^1-spacetimes) *Let (M, g) be a globally hyperbolic C^1-spacetime with global non-negative timelike curvature containing a complete timelike line $\gamma : \mathbb{R} \to M$. Suppose that each maximising timelike geodesic segment extends to a locally maximising timelike geodesic on an open domain. Then the Lorentzian length space induced by (M, g) splits in the sense of Theorem 1.4.*

In this work, we have established a synthetic splitting result for Lorentzian length spaces using triangle comparison. Ideas for future projects include the generalisation to Lorentzian length spaces lower on the causal ladder (e.g. (TC) spaces instead of globally hyperbolic ones, an approach to this that seems to be adaptable to the synthetic situation is [34]), an analysis of the aforementioned splittings of low regularity spacetimes (regarding regularity of the splitting, level set, etc.), as well as an investigation into synthetic timelike Ricci curvature bounds introduced in [23] and expanded in [15]. In the case of low regularity spacetimes, it is not yet clear how these different notions of curvature bounds relate to each other (some recent progress in this direction includes [16, 47]). For metric measure spaces, it is known that the CD-condition is not sufficient and the full RCD-condition is needed for a splitting result (see [37, 38]). An analogue of the RCD-condition is currently unavailable for Lorentzian length spaces and will probably have to be developed before a splitting theorem with synthetic timelike Ricci curvature bounds can be considered.

Acknowledgements Argam Ohanyan and Felix Rott were supported by project P 33594 of the Austrian Science Fund FWF. Argam Ohanyan was also supported by the ÖAW-DOC scholarship of the Austrian Academy of Sciences. Didier Solis acknowledges the support of Conacyt under grant SNI 38368 and UADY under program FMAT-PTA2022, as well as the hospitality of the University of Vienna, where parts of this work were conducted. Tobias Beran acknowledges the support of University of Vienna. The authors would like to thank the Erwin Schrödinger Institute for support and hospitality during the programme *Non-regular Spacetime Geometry* where part of this work was undertaken. The authors would like to thank Matteo Calisti, Gregory Galloway, Melanie Graf, Michael Kunzinger, Clemens Sämann, Benedict Schinnerl and Roland Steinbauer for helpful discussions and valuable feedback. The authors would like to acknowledge the thorough work of the reviewer, whose suggestions helped improving this article.

Availability of data and materials No data was collected nor produced in this work.

Declarations

Conflict of interest/Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

1. Ake Hau, L., Burgos, S., Solis, D.A.: Causal completions as Lorentzian pre-length spaces. Gen. Relativ. Gravit. **54**(9), 1–20 (2022)
2. Ake Hau, L., Cabrera Pacheco, A.J., Solis, D.A.: On the causal hierarchy of Lorentzian length spaces. Class. Quantum Gravity **37**(21), 215013 (2020)
3. Alexander, S., Bishop, R.: A cone splitting theorem for Alexandrov spaces. Pac. J. Math. **218**, 1–15 (2005)
4. Alexander, S.B., Bishop, R.L.: Lorentz and semi-Riemannian spaces with Alexandrov curvature bounds. Commun. Anal. Geom. **16**(2), 251–282 (2008)
5. Alexander, S. B., Graf, M., Kunzinger, M., Sämann, C.: Generalized cones as Lorentzian length spaces: Causality, curvature, and singularity theorems. arXiv:1909.09575 (2019)
6. Ambrosio, L., Gigli, N., Savare, G.: Gradient flows in metric spaces and in the space of probability measures. Birkhauser (2008)
7. Barrera, W., Montes de Oca, L., Solis, D. A.: Comparison theorems for Lorentzian length spaces with lower timelike curvature bounds. arXiv:2204.09612 (2022)
8. Bartnik, R.: Remarks on cosmological spacetimes and constant mean curvature surfaces. Commun. Math. Phys. **117**(4), 615–624 (1988)
9. Beem, J., Ehrlich, P., Easley, K.: Global Lorentzian Geometry. Marcel Dekker (1996)

10. Beem, J., Ehrlich, P., Markvorsen, S., Galloway, G.: Decomposition theorems for Lorentzian manifolds with nonpositive curvature. J. Differ. Geom. **22**, 29–42 (1985)
11. Beem, J., Ehrlich, P., Markvorsen, S., Galloway, G.: A Toponogov splitting theorem for Lorentzian manifolds. In: Global Differential Geometry and Global Analysis 1984, pp. 1–13. Springer (1985)
12. Beran, T., Rott, F.: Gluing constructions for Lorentzian length spaces. arXiv:2201.09695 (2022)
13. Beran, T., Sämann, C.: Hyperbolic angles in Lorentzian length spaces and timelike curvature bounds. J. Lond. Math. Soc. **2**, 1–58 (2023)
14. Borzellino, J., Zhu, S.: The splitting theorem for orbifolds. Ill. J. Math. **38**, 679–691 (1994)
15. Braun, M.: Rényi's entropy on Lorentzian spaces. Timelike curvature-dimension conditions. arXiv:2206.13005 (2022)
16. Braun, M., Calisti, M.: Timelike Ricci bounds for low regularity spacetimes by optimal transport. arXiv:2209.03802 (2022)
17. Burago, D., Burago, Y., Ivanov, S.: A Course in Metric Geometry, vol. 33. American Mathematical Society, Providence, RI (2001)
18. Burago, Y., Gromov, M., Perelman, G.: AD Alexandrov spaces with curvatures bounded below I. Russian Math. Surv. **47**, 1–58 (1992)
19. Burtscher, A., García-Heveling, L.: Time functions on Lorentzian length spaces. arXiv:2108.02693 (2021)
20. Cai, M.: Ends of Riemannian manifolds with nonnegative Ricci curvature outside of a compact set. Bull. Am. Math. Soc. **24**, 371–377 (1991)
21. Cai, M.: A splitting theorem for manifolds of almost nonnegative Ricci curvature. Ann. Glob. Anal. Geom. **11**, 373–385 (1993)
22. Cai, M., Galloway, G., Liu, Z.: Local splitting theorems for Riemannian manifolds. Proc. Am. Math. Soc. **120**, 1231–1239 (1994)
23. Cavalletti, F., Mondino, A.: Optimal transport in Lorentzian synthetic spaces, synthetic timelike Ricci curvature lower bounds and applications. arXiv:2004.08934 (2020)
24. Cheeger, J., Gromoll, D.: The splitting theorem for manifolds of nonnegative Ricci curvature. J. Differ. Geom. **6**(1), 119–128 (1971)
25. Cheeger, J., Gromoll, D.: On the structure of complete manifolds of nonnegative curvature. Ann. Math. **96**, 413–443 (1972)
26. Chruściel, P.T., Grant, J.D.E.: On Lorentzian causality with continuous metrics. Class. Quantum Gravity **29**(14), 145001 (2012)
27. Cohn-Vossen, S.: Totalkrümmung und geodätische Linien auf einfach zusammenhängenden, offenen, vollstandingen Flächenstucken. Math. Sb. **43**, 139–163 (1936)
28. Eschenburg, J.-H.: The splitting theorem for space–times with strong energy condition. J. Differ. Geom. **27**(3), 477–491 (1988)
29. Fang, F., Li, X., Zhang, Z.: Two generalizations of Cheeger–Gromoll splitting theorem via Bakry–Emery Ricci curvature. Ann. Inst. Fourier **59**, 563–573 (2009)
30. Galloway, G.: Splitting theorems for spatially closed space–times. Commun. Math. Phys. **96**(4), 423–429 (1984)
31. Galloway, G.: A generalization of the Cheeger–Gromoll splitting theorem. Arch. Math. **47**, 372–375 (1986)
32. Galloway, G.: The Lorentzian splitting theorem without the completeness assumption. J. Differ. Geom. **29**, 373–387 (1989)
33. Galloway, G.: Existence of CMC Cauchy surfaces and spacetime splitting. Pure Appl. Math. Q. **15**, 667–682 (2019)
34. Galloway, G., Horta, A.: Regularity of Lorentzian Busemann functions. Trans. Am. Math. Soc. **348**(5), 2063–2084 (1996)
35. Galloway, G., Ling, E., Sbierski, J.: Timelike completeness as an obstruction to C^0-extensions. Commun. Math. Phys. **359**(3), 937–949 (2018)
36. García-Heveling, L.: Causality theory of spacetimes with continuous Lorentzian metrics revisited. Class. Quantum Gravity **38**(14), 145028 (2021)
37. Gigli, N.: The splitting theorem in non-smooth context. arXiv:1302.5555 (2013)
38. Gigli, N.: An overview of the proof of the splitting theorem in spaces with non-negative Ricci curvature. Anal. Geom. Metric Spaces **2**(1) (2014)
39. Gigli, N., Pasqualetto, E.: Lectures on Non-smooth Differential geometry. Springer (2018)

40. Grant, J.D.E., Kunzinger, M., Sämann, C.: Inextendibility of spacetimes and Lorentzian length spaces. Ann. Glob. Anal. Geom. **55**(1), 133–147 (2019)
41. Gromov, M.: Synthetic geometry in Riemannian manifolds. Proc. Int. Congress Math. **1**, 415–419 (1978)
42. Gromov, M.: Metric Structures for Riemannian and Non-Riemannian Spaces. Birkhauser (1999)
43. Harris, S.: A triangle comparison theorem for Lorentz manifolds. Indiana Univ. Math. J. **31**(3), 289–308 (1982)
44. Hawking, S.W., King, A.R., McCarthy, P.J.: A new topology for curved space-time which incorporates the causal, differential, and conformal structures. J. Math. Phys. **17**(2), 174–181 (1976)
45. Hedicke, J.: Lorentzian distance functions in contact geometry. J. Topol. Anal. 1–21 (2022)
46. Innami, N.: Splitting theorems of Riemannian manifolds. Comput. Math. **47**, 237–247 (1982)
47. Kunzinger, M., Oberguggenberger, M., Vickers, J. A.: Synthetic versus distributional lower Ricci curvature bounds. arXiv:2207.03715 (2022)
48. Kunzinger, M., Sämann, C.: Lorentzian length spaces. Ann. Glob. Anal. Geom. **54**(3), 399–447 (2018)
49. Kunzinger, M., Steinbauer, R.: Null distance and convergence of Lorentzian length spaces. Ann. Henri Poincare 1424 (2022)
50. Ling, E.: Aspects of C^0 causal theory. Gen. Relativ. Gravit. **52**(57) (2020)
51. Mashiko, Y.: A splitting theorem for Alexandrov spaces. Pac. J. Math. **204**, 445–458 (2002)
52. McCann, R. J.: Private communication (2023)
53. Milka, A.D.: Metric structure of a certain class of spaces that contain straight lines. Ukrain. Geometr. Sb. Vyp. **4**, 43–48 (1967)
54. Minguzzi, E.: Causality theory for closed cone structures with applications. Rev. Math. Phys. **31**(05), 1930001 (2019)
55. Newman, R. P.: A proof of the splitting conjecture of S.-T. Yau. J. Differ. Geom. **31**(1):163–184 (1990)
56. Noronha, M.: A splitting theorem for complete manifolds with non-negative curvature operator. Proc. Am. Math. Soc. **105**, 979–985 (1989)
57. Paeng, S.-H.: Manifolds with non-negative Ricci curvature almost everywhere. J. Korean Math. Soc. **36**(1), 125–137 (1999)
58. Rott, F.: Gluing of Lorentzian length spaces and the causal ladder. arXiv:2209.06894 (2022)
59. Steinbauer, R.: The singularity theorems of General Relativity and their low regularity extensions. arXiv:2206.05939 (2022)
60. Sturm, K.T.: On the geometry of measure metric spaces. Acta Math. **196**, 65–131 (2006)
61. Toponogov, V.: Riemannian spaces with straight lines. AMS Transl. **37**, 287–290 (1964)
62. Toponogov, V.: The metric structure of Riemannian spaces with non-negative curvature which contain straight lines. Am. Math. Soc. Transl. **70**, 225–239 (1968)
63. Villani, C.: Optimal Transport, Old and New. Springer (2009)
64. Wang, L.: A splitting theorem for the weighted measure. Ann. Glob. Anal. Geom. **42**, 79–89 (2012)
65. Yau, S.T.: Problem section. In: Annals of Mathematics Studies. Princeton University Press (1982)

Zusammenfassung

Diese Arbeit beschäftigt sich mit fundamentalen Konstruktionen und Konzepten der Theorie der Lorentz-Längenräume, welche als Abstraktion der Lorentz-Geometrie im Sinne der Theorie der metrischen Längenräume und der metrischen Geometrie beschrieben werden kann. Einerseits wird eine Verklebungs-Konstruktion für Lorentz-Längenräume entwickelt. Auf dieser aufbauend wird ein Analog zum Verklebungs-Theorem von Reshetnyak sowie die Erhaltung von diversen kausalen Eigenschafen etabliert. Andererseits wird ein Spaltungs-Theorem für nicht-negative Krümmungsschranken im Sinne von Dreiecksvergleichen präsentiert.