Non-Hausdorff parallelized manifolds over geometric models of conservative programs

Emmanuel Haucourt

Mathematical Structures in Computer (202x), 1–00
doi:10.1017/xxxxx

CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

Non-Hausdorff parallelized manifolds over geometric models of conservative programs

Emmanuel Haucourt

École polytechnique, LIX - UMR 7161, Palaiseau, France
emmanuel.haucourt@polytechnique.edu

(Submitted draft)

Abstract

Every graph G induces a locally ordered metric space $\mathcal{X}_G$ together with a local order $\tilde{\mathcal{X}}_G$ that is locally dihomeomorphic to the standard pospace $\mathbb{R}$; both are related by a morphism $\beta_G : \tilde{\mathcal{X}}_G \to \mathcal{X}_G$ satisfying a universal property. The underlying set of $\tilde{\mathcal{X}}_G$ admits a non-Hausdorff atlas $\mathcal{A}_G$ equipped with a non-vanishing vector field f_G; the latter is associated to $\tilde{\mathcal{X}}_G$ through the correspondence between local orders and cone fields on manifolds. The above constructions are compatible with cartesian products, so the geometric model of a conservative program is lifted through $\beta_{G_1} \times \cdots \times \beta_{G_n}$ to a subset M of the parallelized manifold $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$. By assigning the suitable norm to each tangent space of $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ the length of every directed smooth path γ on M, i.e. $\int |\gamma'(t)|_{\gamma(t)} dt$, corresponds to the execution time of the sequence of multi-instructions associated to γ. This induces a pseudometric $d_{\mathcal{A}}$ whose restrictions to sufficiently small open sets of $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ (we refer to the manifold topology, which is strictly finer than the pseudometric topology) are isometric to open subspaces of $\mathbb{R}^n$ with the α-norm for some $\alpha \in [1, \infty]$. The transition maps of $\mathcal{A}_G$ are vector translations so the representation of a tangent vector does not depend on the chart of $\mathcal{A}_G$ in which it is represented; consequently, differentiable maps between open subsets of $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ are handled as if they were maps between open subsets of $\mathbb{R}^n$. For every directed path γ on M (possibly the representation of a sequence σ of multi-insructions) there is a shorter directed smooth path on M that is arbitrarily close to γ, and that can replace γ as a representation of σ.

Keywords: Concurrency; local order; conal field; directed path; execution time; multi-instruction;

1. Introduction

The concurrent programs we consider are made of finitely many sequential processes running in parallel. The idea that such programs could be interpreted geometrically was already lurking in the work of Dijkstra (1968); which was quickly followed by publications containing pictures of models that are subsets of $\mathbb{R}^n$ (Coffman et. al. (1971); Kung et. al. (1979); Lipski (1981)). The «geometry of concurrency» was formalized a bit later for programs P whose processes neither have branchings nor loops: their representations are subsets of $\mathbb{R}^n$ called *progress graphs* with n being the number of processes of P (Carson & Reynolds, 1987, §4). We recall two prototypical examples, both made of two processes P_1 and P_2. In the first one, both try to take (P(m)) and release (V(m)) the mutex m concurrently (whereas mutexes cannot be held by more than one process at the time). In the second one, both have to wait for each other (W(b)) behind the synchronization barrier b. The corresponding progress graphs are $\mathbb{R}^2 \setminus [1, 2]^2$ (up to subtleties at the boundary of

the square) and $\{(x, y) \in \mathbb{R}^2 \mid x = 1 \Leftrightarrow y = 1\}$, they are pictured below:

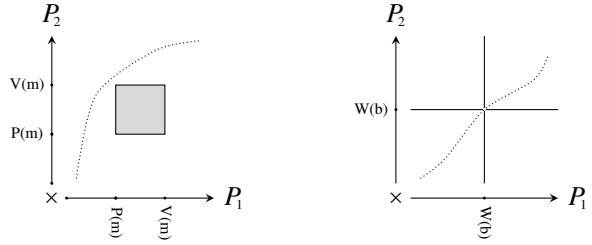

The motivation for progress graphs is the relation between their directed paths (namely the componentwise order-preserving ones) and the execution traces of the programs they represent. As an illustration, the dotted paths on the preceding figures induce the following execution traces:

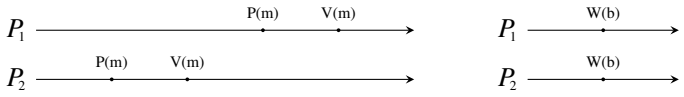

Conversely, any execution trace is induced by a directed path on the progress graph. Following this principle, and according to the postulates a progress graph X is submitted to (Carson & Reynolds, 1987, p.28), the by-product of X we are really interested in is its collection dX of directed paths.

The «geometry of concurrency» described in Carson & Reynolds (1987) only applies to programs without branching nor loop. At the price of some extra mathematical machinery, this limitation can be overcome (Haucourt (2018)). The key feature of a progress graph of dimension n is that its set of directed paths derives from the order and the metric inherited from $\mathbb{R}^n$. Let P be a program made of the processes $P_1, \ldots, P_n$:

i) In accordance with the standard representation of sequential programs used in compilers (Aho et. al., 2007, §9.2), each P_i is given as an automaton which we associate with its underlying graph G_i.
ii) Each graph G_i induces a locally order metric graph $|G_i|$. An execution trace of P_i is seen as a directed path on $|G_i|$, so the set of directed paths on $|G_1| \times \cdots \times |G_n|$ is an overapproximation of the set of execution traces of P.

The process P_i is said to be *conservative* when for any directed path γ on $|G_i|$ the amount of resources held by P_i after the execution of γ only depends on the point where γ stops (this property is decided by a breadth first traversal of G_i).

iii) If all the processes of P are conservative, we define its *geometric model* as the complement in $|G_1| \times \cdots \times |G_n|$ of the *forbidden region* of P (Carson & Reynolds, 1987, §4), (Haucourt, 2018, Definition 4.2).

Every directed path on a geometric model induces a sequence of multi-instructions that respect the constraints imposed by the synchronization mechanisms (mutexes, semaphores, and barriers) (Haucourt, 2018, Theorem 4.1). Moreover, directed paths that are metrically close to each other represent sequences of multi-instructions having the same effect on the state of the abstract machine (Haucourt, 2018, Theorem 6.1). As progress graphs, geometric models are helpful abstractions of their directed paths. In this perspective, we aim at proving that the local order and the metric of the model of P are by-products of a possibly non-Hausdorff parallelized manifold canonically defined from P (it is Hausdorff if, and only, if the program under consideration has no branching).

Motivation. For now let us suppose that the geometric model X is a subset of $\mathbb{R}^n$ (we may as well suppose that X is one of the previous prototypical models). We give some facts about X that will be proven later in the general setting of geometric models:

i) For any directed path γ on X, there is a shorter directed *smooth* path δ on X that is arbitrarily close to γ.
ii) If δ is close enough to γ, then the sequence of multi-instructions they induce have the same effect of the system.
iii) The product metric on $\mathbb{R}^n$ can be chosen so that the lengths of γ and δ correspond to execution times, so «shorter» actually means «faster».
iv) A smooth path on X is directed (i.e. order-preserving in all coordinates) if, and only if, all its tangent vectors belong to $\mathbb{R}^n_+$.

These properties suggest that the differential structure we expect is the standard manifold $\mathbb{R}^n$ equipped with its standard parallelization $(f_1, \dots, f_n)$ (Example 3.28).

As an insight of what happens beyond the case of progress graphs, assume that G represents a basic «if-then-else» construction (see Example 2.16):

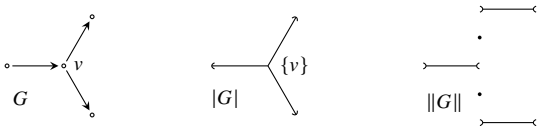

The underlying set of the locally ordered metric graph $|G|$ is $V \cup A\times]0,1[$ with A and V denoting the sets of arrows and vertices of the graph G. Except around the four vertices, which form a 'neglectable' subset of $|G|$, the local order is locally isomorphic to $\mathbb{R}$. In order to get rid of these 'singularities', every vertex is 'blown up' into as many points as ways of going through it in the respect of the local order. Consequently, the vertex v is duplicated and the extremities of G are removed. The resulting set $\|G\|$ equipped with the adhoc (non-Hausdorff) topology often appears in textbooks as an undesirable 1-dimensional smooth manifold. Such oddities are precisely what we need to represent branchings and loops.

Overview. To every program $P = P_1 | \cdots | P_n$ made of n processes running in parallel we associate the n-tuple $(G_1, \dots, G_n)$ of P_i's graphs; the discrete and the continuous models of P are deduced from it (§2). In particular, every graph G comes with a set theoretic map $\beta_G : \|G\| \to |G|$ which 'blows up' every singularity of $|G|$ (Definition 2.11). The manifold and the local order to which the title refers are the *standard atlas* $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ with its *standard parallelization* $(f_1, \dots, f_n)$ (Definitions 3.14 and 3.29), and the *standard local order* $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$ (Definition 5.6); their underlying sets are $\|G_1\| \times \cdots \times \|G_n\|$ and $|G_1| \times \cdots \times |G_n|$ respectively. From a practical point of view, the transition maps between standard charts are vector translations (Lemma 3.13), which makes the derivatives of smooth maps between standard atlases easy to represent (Lemma 3.31). The parallelization $(f_1, \dots, f_n)$ induces a local order $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$ on $\|G_1\| \times \cdots \times \|G_n\|$, and reciprocally (Definition 5.8 and Theorem 5.4); this result is based on an equivalence of category due to J.D. Lawson (1989) (§5.4). The product $\beta_{G_1} \times \cdots \times \beta_{G_n}$ satisfies a universal lifting property relating the local orders $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$ and $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$ (Theorem 5.3). The relation between $(f_1, \dots, f_n)$ and $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$ is thus established through $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$.

From every map ℓ : {arrows of G} $\to \mathbb{R}_+$ admitting a positive lower bound, one defines a metric $d_{|G|}$ on the set $|G|$; metric spaces of this form are called *metric graphs*. In this article (as well as in Haucourt (2018)) they are meant to replace intervals of $\mathbb{R}$ in view of modeling programs with branchings and loops. Due to their prominent role (and also because we have not found

a presentation that fits our needs), we devote §4.2 to a thorough description of these very well-known spaces (Bridson & Haefliger (1999); Papadopoulos (2013)). For our purpose we assume that $|G_1| \times \cdots \times |G_n|$ is equipped with the α-product metric $d^{(\alpha)}$ for some $\alpha \in [1, \infty]$ (Definition 4.10). If we wish α to be in accordance with the *parallel execution time principle* (§6 (25)) we should take $\alpha = \infty$. We transport the α-norm on $\mathbb{R}^n$ to every tangent space of $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ through $(f_1, \ldots, f_n)$ (Definition 4.46 (11)). For every piecewise smooth path γ, we define the *smoothed α-length* $\mathcal{L}^\alpha(\gamma)$ as the sum along γ of the lengths of its tangent vectors (Definition 4.47). Then, assuming that p and q are the images of p' and q' under $\beta_{G_1} \times \cdots \times \beta_{G_n}$, the distance $d^{(\alpha)}(p, q)$ is the infimum of $\mathcal{L}^\alpha(\gamma)$ for γ piecewise smooth path from p' to q' (Theorem 4.3). So for every $\alpha \in [1, \infty]$, the metric $d^{(\alpha)}$ derives from $\mathcal{L}^\alpha$ which derives from $(f_1, \ldots, f_n)$.

From the computer science point of view, for $\varepsilon > 0$ sufficiently small, the sequences of mutli-instructions induced by a directed path and its ε-approximations (Definition 6.1) lead to the same result (Haucourt, 2018, Theorem 6.1). Also, every directed path γ on a tile compatible subset M of $|G_1| \times \cdots \times |G_n|$ (resp. $\|G_1\| \times \cdots \times \|G_n\|$) admits a piecewise affine ε-approximation on M (Definition 2.8, Theorem 6.2, and Corollary 6.10). Since the geometric and the smooth models of a conservative program are tile compatible (Definitions 2.9 and 2.17), any execution trace of the program to model can be represented by a piecewise affine path.

Organization. Every section of the article comes with its own introduction, so we just briefly describe the (straightforward) organization of the paper. Sections 2, 3, 4, 5, and 5.4 respectively deal with the set maps $\beta_{G_1} \times \cdots \times \beta_{G_n}$, the atlases $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ and their parallelizations $(f_1, \ldots, f_n)$, the α-product (psuedo)metrics $d^{(\alpha)}$ and $\tilde{d}^{(\alpha)}$, the local orders $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$ and $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$, and the Lawson correspondance.

2. Underlying sets of models

As we shall see in sections 3, 4, and 5, the continuous models of conservative programs inherit their structures from atlases, metrics, and local orders. We explicitly describe their underlying sets.

A *graph* is a set G, whose elements are called *points*, together with three maps, namely $\dim : G \to \{0, 1\}$, and $\mathrm{src}, \mathrm{tgt} : G^{(1)} \to G^{(0)}$ with $G^{(\varepsilon)} = \{x \in G \mid \dim x = \varepsilon\}$ for $\varepsilon \in \{0, 1\}$. These maps are respectively called *dimension*, *source*, and *target*. A *vertex* (resp. an *arrow*) is a point of G of dimension 0 (resp. 1). Given a vertex v, an arrow α is said to be *ingoing* (resp. *outgoing*) when $\mathrm{tgt}\,\alpha = v$ (resp. $\mathrm{src}\,\alpha = v$). A vertex with no ingoing or no outgoing arrow is called an *endpoint*. We use the same denotation for the graph and its underlying set. A map $f : G \to G'$ is a *morphism of graphs* when it preserves dimensions, sources, and targets. The category of graphs is denoted by **Grph**.

Discrete models of conservative programs. The notion of a conservative program (Haucourt (2018)) is based on a virtual machine allowing parallel execution of *multi-instructions* i.e. partial maps μ on $\{1, \ldots, n\}$ assigning to the i^{th} process of the machine the 'atomic' instruction it has to execute (the i^{th} process is inactive during the execution of μ if the latter is not defined at i).

Definition 2.1. A *sequential process* is a graph G whose arrows are labeled with atomic instructions. The vertices of G are possible positions of the instruction pointer p. At each step p moves from the source v to the target v' of the arrow a carrying the next instruction to execute. However, our setting does not allow p to jump from v to v' without going through an intermediate stage, namely the arrow a itself. The arrows of G are thus legit positions of the instruction pointer. The *execution* of a directed path on G is the execution of the corresponding sequence of instructions.

Without loss of generality, we suppose that every vertex of G with more than one outgoing (resp. ingoing) arrow has at least one ingoing (resp. outgoing) arrow; this property is used exactly once, but crucially, in the proof of Lemma 4.51.

A program P is a tuple of sequential processes running concurrently. The virtual machine on which it is executed comes with a pool of resources shared by the processes of P. These resources are *renewable* in the sense that they can be restored after use: the processes *take* and *release* them by means of atomic instructions. A sequential process over the graph G is said to be *conservative* when the amount of resources mobilized after the execution of any directed path starting at the initial vertex only depends on where it stops (this property is decided by a breadth first traversal of G). A program is said to be *conservative* when so are all its processes. Each process of the program P is identified by an element of $\{1, \dots, n\}$, its underlying graph is denoted by G_i.

Definition 2.2. For the purposes of our study, we only need to know that the *discrete model* of P (Haucourt, 2018, §4.2), which is a subset $[\![P]\!]_d$ of the cartesian product of sets $G_1 \times \cdots \times G_n$, has the following properties:

(1) If the virtual machine fulfills the resource requirements specified in P, then any *directed path* on $[\![P]\!]_d$ (i.e. any finite sequence of points in $[\![P]\!]_d$ satisfying certain conditions (Haucourt, 2018, §3, p.1729)) induces a sequence of multi-instructions that the virtual machine can execute without exhausting its pool of resources.
(2) Any execution trace of P is induced by a directed path on $[\![P]\!]_d$.

In other words, the amount of resources required to execute a conservative program P is known at compile time, and its discrete model induces an overapproximation of its set of execution traces.

Continuum of states. We have already mentioned that the arrows of the graph of a sequential process are intermediate positions of the instruction pointer (Definition 2.1). Pushing this principle further, we specify *where* the pointer stands on a given arrow of G. Suppose that the map

$$\ell \,:\, G^{(1)} \,\longrightarrow\, [R, \infty[$$

assigns to each arrow a *length* which is meant to be the execution time of the atomic instruction it carries (§6). The real number R is supposed to be *positive* because execution times should not be arbitrarily small. This constraint have important consequences on the mathematical properties of the models that will be defined later. We end up with a continuum of states:

Definition 2.3. The *continuous support* of G is the set

$$|G| \quad = \quad G^{(0)} \;\cup\; \bigcup \Big\{\; \{a\} \times]0, \ell(a)[\;\;\Big|\;\; a \in G^{(1)} \;\Big\} \quad .$$

The quotient map $\pi_G : |G| \to G$ is defined by $\pi_G(v) = v$ on $G^{(0)}$ and by $\pi_G(a,t) = a$ elsewhere. The local order on $|G|$ (Definition 5.6) encodes the constraints that force the pointer p to move continuously with respect to the direction imposed by the arrows. The source and the target of an arrow a play the role of $(a, 0)$ and $(a, \ell(a))$ which are intentionally excluded from $|G|$. An *arrow* of $|G|$ is a subset of the form $\{a\} \times]0, \ell(a)[$ for some $a \in G^{(1)}$; by extension its source and its target are those of a.

Definition 2.4. Subsets of the form $\{a\} \times J$ are said to be *initial* (resp. *final*) when J is an *initial* (resp. *final*) subinterval of $]0, \ell(a)[$. A subset S of $|G|$ is said to be a *star* when it is of the form

$$S(v,i,f) \quad = \quad \bigcup_{\mathrm{tgt}\, a = v} \{a\} \times]\ell(a) - f(a), \ell(a)[\quad \cup \quad \{v\} \quad \cup \quad \bigcup_{\mathrm{src}\, a = v} \{a\} \times]0, i(a)[$$

with the functions $i : \mathrm{src}^{-1}\{v\} \to \mathbb{R}_+\backslash\{0\}$, $f : \mathrm{tgt}^{-1}\{v\} \to \mathbb{R}_+\backslash\{0\}$ satisfying $i(a), f(a) \leqslant \ell(a)$ whenever it makes sense. We just write $S(v)$ when $i(a) = \ell(a)$ and $f(a) = \ell(a)$ whenever it makes sense. If both i and f are constant and equal to $r < \frac{R}{2}$, we write $S(v,r)$ instead of $S(v,i,f)$, and call

$S(v,r)$ the star centered at v of radius r. The *canonical star* centered at v is the set $S(v,\frac{\ell}{2},\frac{\ell}{2})$, i.e. $i(a)=\ell(a)/2$ for every $a\in\mathrm{src}^{-1}\{v\}$ and $f(a)=\ell(a)/2$ for every $a\in\mathrm{tgt}^{-1}\{v\}$. In every case, the vertex v is called the *centre* of the star.

Definition 2.5. A *segment* of G centered at (a,t) with $a\in G^{(1)}$ and $t\in\,]0,\ell(a)[$, is a subset of $|G|$ of the form $\{a\}\times\,]t-\varepsilon,t+\varepsilon[$ with $\varepsilon\leqslant\min(t,\ell(a)-t)$.

Remark 2.6. The intersection of two stars is either a star or a (possibly empty) disjoint union of segments. The intersection of a segment with a star (resp. a segment) is either a segment, or the union of two disjoint segments, or empty. Stars and segments thus form a base of topology. The stars centered at v (resp. the segments centered at (a,t)) form a base of neighborhoods of v (resp. (a,t)). We give some immediate consequences: the topology of $|G|$ is Hausdorff; the boundary of $\{a\}\times\,]0,\ell(a)[$ is $\{\mathrm{src}\,a,\mathrm{tgt}\,a\}$; the boundary of $S(v)$ is the set $\{\mathrm{src}(a),\mathrm{tgt}(a)\mid a\in G^{(1)},\ v\in\{\mathrm{src}(a),\mathrm{tgt}(a)\}\}\setminus\{v\}$ (for any vertex v); and the connected components of $|G|\setminus\{\text{vertices}\}$ are the segments $\{a\}\times\,]0,\ell(a)[$ for a arrow of G.

Definition 2.7. The *topology of* $|G|$ is the one described in Remark 2.6.

Definition 2.8. The *tile* over $p\in G_1\times\cdots\times G_n$ is the set $(\pi_{G_1}\times\cdots\times\pi_{G_n})^{-1}\{p\}$, i.e. $\tau_1\times\cdots\times\tau_n$ with $\tau_i=\{p_i\}$ if p_i is a vertex, and $\tau_i=\{p_i\}\times\,]0,\ell(p_i)[$ if p_i is an arrow. A subset of $|G_1|\times\cdots\times|G_n|$ is said to be *tile compatible* when it contains any tile it meets.

Definition 2.9. (Haucourt, 2018, §6.2). The support of the geometric model of P is the set

$$|P| \quad=\quad (\pi_{G_1}\times\cdots\times\pi_{G_n})^{-1}[\![P]\!]_d$$

with $[\![P]\!]_d$ denoting the discrete model (Definition 2.2); it is tile compatible.

Blowing up vertices. For every arrow a of G the canonical bijection $\{a\}\times\,]0,\ell(a)[\to\,]0,\ell(a)[$ induces a *dihomeomorphism* (i.e. a homeomorphism that is also an order isomorphism).

Definition 2.10. A *traversal* at v is an ordered pair of arrows (a,b) such that $\mathrm{tgt}\,a=v=\mathrm{src}\,b$, we say that v is *regular* when there is *exactly* one traversal at v in G; otherwise v is said to be *singular*.

From the topological and order theoretic points of view, the set $|G|$ is 'almost everywhere' like $\mathbb{R}$. The only possible exceptions are the singular vertices of G. As for curves in Algebraic Geometry[(1)] we 'resolve singularities' by 'blowing them up': we replace each vertex v by the traversals at v :

Definition 2.11. The *support* of the *blowup* of G, which we denote by $\|G\|$, is the set $(|G|\setminus G^{(0)})\cup\{\text{traversals of }G\}$, that is to say

$$\|G\| \quad=\quad \bigcup\big\{\{a\}\times\,]0,\ell(a)[\;\big|\;a\in G^{(1)}\big\}\ \cup\ \big\{(a,b)\in G^{(1)}\times G^{(1)}\;\big|\;\mathrm{tgt}\,a=\mathrm{src}\,b\big\}.$$

The *desingularizator* of G is the map $\beta_G:\|G\|\to|G|$ defined by $\beta_G(a,b)=\mathrm{tgt}\,a$ (or $\mathrm{src}\,b$) for every traversal (a,b) of G, and by $\beta_G(a,t)=(a,t)$ for every $a\in G^{(1)}$ and $t\in\,]0,\ell(a)[$. A subset of $\|G_1\|\times\cdots\times\|G_n\|$ is said to be *tile compatible* when it is the inverse image under $\beta_{G_1}\times\cdots\times\beta_{G_n}$ of some tile compatible set (Definition 2.8). From now on we write β_i instead of β_{G_i}.

The image of β_G is precisely $|G|\setminus\{\text{endpoints of }G\}$.

---

(1) « The effect of blowing up is thus to separate out branches of curves passing through O according to their slopes » (Hartshorne, 1977, p.30). In our context, the traversals at v play the role of the slopes of branches at O.

Remark 2.12. The constructions $|\_|$ and $\|\_\|$ readily extend to functors from **Grph** to **Set**, and the collection of desingularizators $\beta_G : \|G\| \to |G|$, for G running through the class of graphs, forms a natural transformation from $\|\_\|$ to $|\_|$.

Remark 2.13. It is natural to define the *topology of* $\|G\|$ as the initial topology of the desingularizator β_G: a base of this topology is given by the inverse images under β_G of the stars and segments (Remark 2.6). By observing that ${\beta_G}^{-1}(S(v,i,f))$ is the set

$$\bigcup_{\text{tgt } a=v=\text{src } b} \{a\}\times]\ell(a)-f(a),\ell(a)[\;\cup\; \{(a,b)\} \;\cup\; \{b\}\times]0,i(b)[$$

we conclude that the topology of $\|G\|$ is *not* T_1 (i.e. it contains a finite subset that is not closed, for example every single traversal at a singular vertex).

Lemma 2.14. The image of a tile compatible set (in the sense of Definition 2.11) under $\beta_1 \times \cdots \times \beta_n$ is tile compatible (in the sense of Definition 2.8).

Proof. Suppose that $Y = (\beta_1 \times \cdots \times \beta_n)^{-1}X$ for some tile compatible set X and that $q \in Y \cap (\beta_1 \times \cdots \times \beta_n)^{-1}\tau$ for some tile τ. Let $p \in \tau$. We have $(\beta_1 \times \cdots \times \beta_n)(q) \in X$ therefore $\tau \subseteq X$ because X is tile compatible; hence $p \in X$. For $i \in \{1,\dots,n\}$ if q_i is a traversal, then put $q'_i = q_i$; otherwise $q'_i = p_i$. Then $(\beta_1 \times \cdots \times \beta_n)(q') = p$.

Examples 2.15. We denote by $G_{\mathbb{Z}}$ the graph whose arrows (all of length 1) are $(n, n{+}1)$ with $n \in \mathbb{Z}$, the source and the target being n and $n{+}1$ respectively.

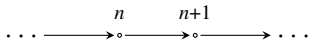

The map

$$\begin{aligned} |G_{\mathbb{Z}}| &\to \mathbb{R} \\ n &\mapsto n \\ ((n,n+1),t) &\mapsto n+t \end{aligned}$$

is a bijection, as well as the desingularizator of $G_{\mathbb{Z}}$ since every vertex has a unique traversal. So we have a canonical bijection from $\|G_{\mathbb{Z}}\|$ to $\mathbb{R}$.

Examples 2.16. Assuming that the length of every arrow of G is 1 we have:

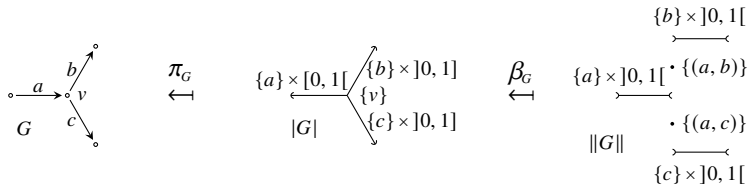

We will see that every set $\|G\|$ carries a standard atlas (Definition 3.14). Connectedness of this atlas, as well as non-Hausdorffness of its underlying topology (Definition 3.9), derive form the fact that any neighborhood of (a,b) meets any neighborhood of (a',b') precisely when $a = a'$ or $b = b'$.

We have already observed that the endpoints of G are 'forgotten' by $\|G\|$. One may think of this as an issue since the *starting point* of the automaton associated with a process is an endpoint. We circumvent the difficulty taking advantage of the following control flow graph feature: the starting

point has a unique outgoing arrow a. So we can harmlessly declare any point on $\{a\}\times]0,\ell(a)[$ to be the starting point instead of the source of a. Another way to solve the problem, and maybe a more canonical one, consists of adding a fresh arrow ending at the starting point we wish to save. The same remarks hold (and the same tricks work) for the *final points* of the graph.

Definition 2.17. The *support* for the *smooth model* of P is the set

$$\|P\| \quad = \quad (\beta_1\times\cdots\times\beta_n)^{-1}\,|P|$$

with $|P|$ as in Definition 2.9. Following Definition 2.11, the set $\|P\|$ is tile compatible.

3. Differential calculus on atlases

The notion of an atlas allows to apply Differential Calculus beyond the class of maps between open subsets of (finite dimensional) normed spaces[(2)]. We recall some basic definitions without using local coordinates[(3)], and atlases are defined on sets without presupposing any topology (the latter can indeed be recovered from the charts).

Manifolds are equivalence classes of atlases. Such a class contains a maximal element (with respect to inclusion) to which the manifold is usually identified. However, each manifold met in this article comes with canonical representative that is much more tractable than the maximal one. Indeed, the representation of any tangent vector to a *standard atlas* (Definition 3.14) does not depend on the *standard chart* (Definition 3.12) in which it is represented. For this reason, we will only deal with atlases.

Last but not least, we are not concerned with global properties of manifolds, so we can harmlessly let them be non-Hausdorff.[(4)] This seemingly anodyne weakening is of crucial importance here for it allows *branchings*, which is precisely what we need to build manifolds from graphs. The relation between non-Hausdorff manifolds and branching, non-determinism, and bifurcation, is well-known (Müller (2013)). Such manifolds naturally appear in the study of dynamical systems (Goel (1987)), foliations of the plane (Haefliger & Reeb (1957)), (Gauld, 2014, Chap.9) and general relativity (Hájíček (1971); Luc & Placek (2020)). The novelty is that we apply it to formalize the intuition that a graph is a kind of 'discrete vector field'.

3.1 Differential calculus

We assume that Differential Calculus for smooth maps $f:U\to\mathbb{R}^m$ with U open subset of $\mathbb{R}^n$ is known, see (Lang, 2002, Chap. I, §2-4) or (Nachbin, 1981, §11,13,14,16). The derivative of f at $x\in U$ is a linear map $Df_x:\mathbb{R}^n\to\mathbb{R}^m$ which we identify with an $m\times n$ matrix with entries in $\mathbb{R}$, i.e. an element of $\mathrm{Mat}_n^m(\mathbb{R})$. All one really needs here is the derivative of identities and the *chain rule*

$$D(id_U)_x \quad = \quad id_{\mathbb{R}^n} \qquad \text{and} \qquad D(g\circ f)_x \quad = \quad Dg_{fx}\circ Df_x \tag{1}$$

for every element x of an open subset U of $\mathbb{R}^n$. If the map f has a smooth inverse f^{-1} with $fx=y$, then the above relations imply that

$$D(f^{-1})_y \quad = \quad \big(Df_x\big)^{-1}\,. \tag{2}$$

(2) « Differentiability is a local phenomenon, so to talk about it, we need only a space which is locally like euclidean space» / «a manifold is locally like the arena of calculus» (Gauld, 1982, p.28, 53). See also (do Carmo, 1992, p.1).

(3) This approach is advocated in (Lang, 1999, Foreword) and (Nachbin, 1981, chap.13), in accordance with the principle that local coordinates are only meant to be a tool for computations.

(4) « The concept of Hausdorffness is irrelevant for much of local differential geometry » (Hicks, 1965, p.3). A similar point of view is adopted in (Lang, 1999, p.23).

3.2 Atlases

A *chart* of dimension $n \in \mathbb{N}$, or *n-chart*, is a *bijection* ϕ whose codomain is an *open* subset of $\mathbb{R}^n$. The n-dimensional chart $\tilde{\phi}$ is a *subchart* of ϕ, which we denote by $\tilde{\phi} \subseteq \phi$, when $\tilde{\phi}$ and ϕ agree on $dom\,\tilde{\phi} \subseteq dom\,\phi$ (note that we necessarily have $cod\,\tilde{\phi} \subseteq cod\,\phi$ since $\tilde{\phi}$ is onto). The subcharts of ϕ are canonically identified with the subsets of $dom\,\phi$ whose images under ϕ are open, and therefore with the open subsets of $cod\,\phi$. The initial topology on $dom\,\phi$ is the only one that makes ϕ a homeomorphism. A subset $U \subseteq dom\,\phi$ is open in the initial topology if, and only if, its image under ϕ is open. We denote by ϕ_U the corresponding subchart of ϕ.

Definition 3.1. Given the charts ϕ and ψ, a map f is said to be *smooth around* $p \in dom\,f$ from ϕ to ψ when there are open subsets U and V of $dom\,\phi$ and $dom\,\psi$ such that $p \in U \subseteq dom\,f$, $f(U) \subseteq V$, and the composite $\psi_V \circ f \circ \phi_U^{-1}$ is smooth as a map from ϕU to ψV. The sets U and V are called *witnesses of smoothness*. The map $\psi_V \circ f \circ \phi_U^{-1}$ is referred to as the *representation* of f in the charts ϕ_U and ψ_V.

Charts are meant to allow differential calculus for mappings whose (co)domains may not be subsets of $\mathbb{R}^n$. Indeed, with the denotation from Definition 3.1, it is natural to think of the linear map $D(\psi \circ f \circ \phi^{-1})_{\phi p}$ as the representation of 'the' derivative of f at p (Definition 3.22). However, such a representation essentially depends on the charts ϕ and ψ. The concept of an atlas arises from the need to regulate this dependency:

Definition 3.2. The n-charts ϕ and ψ are said to be *compatible*, which denote by $\phi \sim \psi$, when:

i) the intersection $U = dom\,\phi \cap dom\,\psi$ is open in both $dom\,\phi$ and $dom\,\psi$, and
ii) the composites $\phi_U \circ \psi_U^{-1}$ and $\psi_U \circ \phi_U^{-1}$, which we call *transition maps*, are smooth.

Following the common usage, we 'forget' the subscript U and just write $\phi \circ \psi^{-1}$ and $\psi \circ \phi^{-1}$ instead of $\phi_U \circ \psi_U^{-1}$ and $\psi_U \circ \phi_U^{-1}$.

Remark 3.3. With the denotation from Definition 3.2, for every $X \subseteq U$, if $\phi_U X$ is open in $\mathbb{R}^n$, then so is $\psi_U X$ because we have $\psi_U X = (\psi_U \circ \phi_U^{-1})(\phi_U X)$ and $\psi_U \circ \phi_U^{-1}$ is a diffeomorphism so the *inverse mapping theorem* applies (Lang, 2002, Thm.5.2, p.13). Consequently, the charts ϕ_U and ψ_U induce the same topology on U, so it suffices to require that U is open in $dom\,\phi$ (resp. $dom\,\psi$) in the first point of Definition 3.2.

One switches from one representation of the derivative of f to another by means of the following result:

Lemma 3.4. Suppose that $\phi_0 \sim \phi_1$ and $\psi_0 \sim \psi_1$. Let f be a map such that $dom\,\phi_0 \cap dom\,\phi_1 \subseteq dom\,f$ and $f(dom\,\phi_0 \cap dom\,\phi_1) \subseteq dom\,\psi_0 \cap dom\,\psi_1$. Given any $p \in dom\,\phi_0 \cap dom\,\phi_1$, the map f is smooth nearby p with respect to ϕ_0, ψ_0 if, and only if, so it is with respect to ϕ_1, ψ_1. In particular we have

$$\psi_1 \circ g \circ \phi_1^{-1} = (\psi_1 \circ \psi_0^{-1}) \circ (\psi_0 \circ g \circ \phi_0^{-1}) \circ (\phi_0 \circ \phi_1^{-1})$$

being understood that the ϕ_i's and the ψ_i's in the above equation are subcharts of ϕ_i and ψ_i whose domains respectively contain p and $f(p)$.

Definition 3.5. An *atlas* is a collection $\mathcal{A}$ of pairwise compatible charts, its *support* is the set

$$|\mathcal{A}| \quad = \quad \bigcup \{ dom\,\phi \mid \phi \in \mathcal{A} \} \quad .$$

Following Definition 3.2, all the charts of $\mathcal{A}$ have the same dimension, which is the *dimension* of $\mathcal{A}$.

Definition 3.6. Given the atlases $\mathcal{A}$ and $\mathcal{B}$, a map $f : |\mathcal{A}| \to |\mathcal{B}|$ is said to be *smooth nearby* $p \in |\mathcal{A}|$ when there exist two charts $\phi \in \mathcal{A}$ and $\psi \in \mathcal{B}$ such that f is smooth nearby p from ϕ to ψ (Definition 3.1). It is said to be *smooth* when it is so nearby every point of $\mathcal{A}$. Smooth maps are the morphisms of the category of atlases, which we denote by **Atl**. The cartesian product of $\mathcal{A}$ and $\mathcal{B}$ in **Atl** is the atlas

$$\mathcal{A} \times \mathcal{B} \quad = \quad \left\{ \phi \times \psi \;\middle|\; \phi \in \mathcal{A} \,;\, \psi \in \mathcal{B} \right\} \quad .$$

The *transition map* between the charts $\phi \times \psi$ and $\phi' \times \psi'$ of $\mathcal{A} \times \mathcal{B}$ is the product $\tau \times \tau'$ of the transition map τ between ϕ and ϕ' and the transition map τ' between ψ and ψ'.

As an immediate consequence of Lemma 3.4, smoothness nearby p does not depend on the choice of ϕ and ψ (with the denotation from Definition 3.6).

Examples 3.7. Unless otherwise stated, any open subset U of $\mathbb{R}^n$ is equipped with the atlas $\{id_U\}$. For any chart ϕ, the one element set $\{\phi\}$ is an atlas, and ϕ induces a smooth map from $\{\phi\}$ to $\{id_{cod\,\phi}\}$.

Definition 3.8. A *curve* is a smooth map defined on open interval of $\mathbb{R}$.

Two atlases with the same support are said to be *equivalent* if their union is still an atlas. A *manifold* is an equivalence class of atlases.

Definition 3.9. The support of an atlas $\mathcal{A}$ is endowed with the initial topology, i.e. the least one making all the elements of $\mathcal{A}$ continuous.

According to the above definition, any chart of $\mathcal{A}$ is a homeomorphism. As in (Hicks, 1965, p.3) and (Lang, 1999, Chap.II), we *do not require* the topology of an atlas to be Hausdorff; yet we have:

Lemma 3.10. (Gauld, 2014, Proposition 9.2, p.153). The topology of any atlas is T_1.

Definition 3.11. A path on the atlas $\mathcal{A}$ is said to be *smooth* when it is the restriction of a curve on $\mathcal{A}$ (Definition 3.8). A *piecewise smooth* path on $\mathcal{A}$ is a concatenation of smooth paths on $\mathcal{A}$.

The standard atlas of a tuple of graphs. Given a graph G with $\ell : G^{(1)} \to [R, \infty[$, the set $\|G\|$ (Definition 2.11) admits a standard atlas that we now describe. Given $x \in G^{(1)}$ we define

$$I_x \;=\;]0, \tfrac{\ell(x)}{2}[\,, \qquad J_x \;=\;]\tfrac{\ell(x)}{2}, \ell(x)[\,, \tag{3}$$

and keep in mind the bijection $t \mapsto t - \ell(x)$ from J_x to $-I_x$.

Definition 3.12. The *standard charts* of G are the following bijections

$$\begin{array}{rcl} \phi_a & : & \{a\} \times]0, \ell(a)[\;\to\;]0, \ell(a)[\,, \quad \text{and} \\ \phi_{ab} & : & \{a\} \times J_a \;\cup\; \{(a,b)\} \;\cup\; \{b\} \times I_b \;\to\; -I_a \cup \{0\} \cup I_b \\ & \text{with} & (a,t) \mapsto t - \ell(a)\,, \quad (a,b) \mapsto 0\,, \quad (b,t) \mapsto t \end{array}$$

for all arrows a and all traversals (a, b) of G. The *standard charts* of $(G_1, \ldots, G_n)$ are the products $\phi_1 \times \cdots \times \phi_n$ with ϕ_i standard chart of G_i for $i \in \{1, \ldots, n\}$.

Lemma 3.13. The collection of standard charts of $(G_1, \ldots, G_n)$ is an atlas on the set $\|G_1\| \times \cdots \times \|G_n\|$, and the transition maps between its charts are vector translations.

Proof. The domains of the standard charts cover $\|G\|$. Given two arrows a and a' with $a \neq a'$ the domains of ϕ_a and $\phi_{a'}$ are disjoint. By applying the distributivity of $\cap$ over $\cup$, the intersection of $dom\ \phi_{ab}$ and $dom\ \phi_{a'b'}$ is the disjoint union

$$(\{a\} \cap \{a'\}) \times (J_a \cap J_{a'}) \ \cup \ \big(\{(a,b)\} \cap \{(a',b')\}\big) \ \cup \ (\{b\} \cap \{b'\}) \times (I_b \cap I_{b'})\,.$$

In all the cases, the transition map $\phi_{a'b'} \circ \phi_{ab}^{-1}$ is an identity; but depending on the case, the domain of this identity is

$$\begin{array}{ll} cod\ (\phi_{ab}) & \text{if } (a,b) = (a',b')\,, \\ J_a & \text{if } a = a' \text{ and } b \neq b'\,, \\ I_b & \text{if } a \neq a' \text{ and } b = b'\,, \\ \emptyset & \text{otherwise.} \end{array}$$

Also, if $c \notin \{a,b\}$ then the domains of the charts ϕ_{ab} and ϕ_c are disjoint. So the only nontrivial case is when $c \in \{a,b\}$. We take $\{a\} \times J_a$ as witness of compatibility of ϕ_{ab} and ϕ_a since we have

$$\phi_{ab} \circ \phi_a^{-1} \quad : \quad t \in J_a \quad \mapsto \quad t - \ell(a) \in -I_a\,,$$

and of course J_b as witness of compatibility of ϕ_{ab} and ϕ_b since $\phi_{ab} \circ \phi_b^{-1}$ is just the identity map on J_b. We readily deduce the higher dimensional case from the description of the transition maps between the charts of a product of atlases (Definition 3.6).

Definition 3.14. The *standard atlas* of the tuple of graphs $(G_1, \ldots, G_n)$ is the collection of standard charts of $(G_1, \ldots, G_n)$ (Definition 3.12). Following Definition 3.6, it is the product atlas $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ of the standard atlases $\mathcal{A}_{G_i}$ of G_i.

Examples 3.15. The canonical bijection $\|G_{\mathbb{Z}}\| \cong \mathbb{R}$ (Example 2.15) induces a smooth diffeomorphism between the standard atlas of $G_{\mathbb{Z}}$ and $\{id_{\mathbb{R}}\}$ (Example 3.7).

As an immediate consequence of Lemma 3.10, the topology of $\|G\|$ (Remark 2.13) differs from that of the atlas $\mathcal{X}_G$ (Definition 3.9).

3.3 Tangent bundles and vector fields

Tangent vectors. Let $\mathcal{A}$ be an atlas of dimension n. Given $\phi, \psi \in \mathcal{A}$, $p \in dom\ (\phi)$, $q \in dom\ (\psi)$, and $a, b \in \mathbb{R}^n$, we write $(p, \phi, a) \sim (q, \psi, b)$ when $p = q$ and $D(\psi \circ \phi^{-1})_{\phi p} a = b$. This makes sense because the elements of $\mathcal{A}$ are pairwise compatible. The relation $\sim$ is reflexive, transitive, and symmetric as a consequence of (1) and (2) from §3.1.

Definition 3.16. (Hirsch, 1976, p.16-7). The *tangent vectors* of $\mathcal{A}$ are the elements of the quotient

$$T\mathcal{A} \quad = \quad \big\{(p, \phi, a) \;\big|\; \phi \in \mathcal{A},\ p \in dom\ \phi, \text{ and } a \in \mathbb{R}^n\big\} \,/\, \sim$$

All the tuples contained in an equivalence class v have the same first component p, we say that v is a *tangent vector at* p, or that p is the *attachment point* of v. We denote by $T_p\mathcal{A}$ the set of all such vectors.

Examples 3.17. The tangent vectors of a single element atlas $\{\phi\}$ (Example 3.7) can be identified with the elements of $dom\ \phi \times \mathbb{R}^n$, with n the dimension of the chart ϕ.

For every $v \in T_p\mathcal{A}$ there is a unique $[\![v]\!]^\phi_p \in \mathbb{R}^n$ such that v is the equivalence class of the tuple $(p, \phi, [\![v]\!]^\phi_p)$ (this holds because every mapping $D(\psi \circ \phi^{-1})_{\phi p}$ with $\phi, \psi \in \mathcal{A}$ and $p \in dom\,\phi \cap dom\,\psi$, is a bijection). By definition of the relation $\sim$ we have

$$[\![v]\!]^\psi_p \quad = \quad D(\psi\circ\phi^{-1})_{\phi p} \circ [\![v]\!]^\phi_p \tag{4}$$

for every $\psi \in \mathcal{A}$ with $p \in dom\,\phi$.

Examples 3.18. For any standard charts ϕ and ψ of $(G_1, \ldots, G_n)$, the derivative of $\psi \circ \phi^{-1}$ at any point of $cod\,\phi \cap cod\,\psi$ is the identity map $id_{\mathbb{R}} : \mathbb{R} \to \mathbb{R}$ (Lemma 3.13). It follows from (4) that for all $v \in T_p(\prod_{i=1}^n \mathcal{A}_{G_i})$ the representation $[\![v]\!]^\phi_p$ *does not* depend on the standard chart ϕ. Hence we may as well decide that

$$T(\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}) \quad = \quad \|G_1\| \times \cdots \times \|G_n\| \times \mathbb{R}^n \quad . \tag{5}$$

Tangent spaces. We transfer the (topological) real vector space structure of $\mathbb{R}^n$ to $T_p\mathcal{A}$ in the obvious way: the linear combination $u + \lambda v$ in $T_p\mathcal{A}$ is characterized by the equality

$$[\![u + \lambda v]\!]^\phi_p \quad = \quad [\![u]\!]^\phi_p + \lambda[\![v]\!]^\phi_p \quad .$$

It does not depend on the chart ϕ because we have the relation (4) and the map $D(\psi\circ\phi^{-1})_{\phi p}$ is linear. Given $n \in \mathbb{N}$ there is a unique, up to isomorphism, structure of n-dimensional Hausdorff *topological vector spaces* (*tvs* for short) (Schaefer & Wolff, 1999, Chap.I, 3.2, 3.4). Consequently, the only tvs structure on $T_p\mathcal{A}$ that makes the bijection $[\![\_]\!]^\phi_p : T_p\mathcal{A} \to \mathbb{R}^n$ an isomorphism does not depend on ϕ.

Definition 3.19. The resulting topological vector space $T_p\mathcal{A}$ is called the *tangent space* of the atlas $\mathcal{A}$ at point p.

Examples 3.20. The mappings $[\![\_]\!]^\phi_p$ are tvs isomorphisms.

Examples 3.21. Following the description of $T\mathcal{A}_G$ given in Example 3.18 (5), we can suppose that $T_p(\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}) = \{p\} \times \mathbb{R}^n$ with the obvious topological vector space structure.

Derivative of a smooth map. Given a smooth map $f : \mathcal{A} \to \mathcal{B}$ (Definition 3.6) and the charts $\phi \in \mathcal{A}$ and $\psi \in \mathcal{B}$ with $p \in dom\,\phi$ and $f(p) \in dom\,\psi$, neither the smoothness of f nearby p (Definition 3.1) nor the validity of the following equality

$$[\![Df_p v]\!]^\psi_{fp} \quad = \quad D(\psi\circ f\circ\phi^{-1})_{\phi p}[\![v]\!]^\phi_p \tag{6}$$

depends on ϕ and ψ : it suffices to apply i) the chain rule (§3.1(1)) to the formula from Lemma 3.4, and ii) the relation (4). Hence the following definition sound:

Definition 3.22. The *derivative* of f at p is the only map $Df_p : T_p\mathcal{A} \to T_{fp}\mathcal{B}$ satisfying the equality (6) for every $v \in T_p\mathcal{A}$. It is linear because so are the mappings $[\![\_]\!]^\phi_p$, $[\![\_]\!]^\psi_{fp}$, and $D(\psi\circ f\circ\phi^{-1})_{\phi p}$; the latter is called the *representation* of Df_p in the charts ϕ and ψ, it is denoted by $[\![Df_p]\!]^\psi_\phi$.

We switch from a representation of Df_p to another applying the chain rule (§3.1(1)) to the formula from Lemma 3.4.

Examples 3.23. Let $f : \mathcal{A} \to \mathcal{B}$ be a smooth map between standard atlases (Definition 3.14) of dimensions n and m. The transition maps between standard charts are vector translations (Lemma 3.13) so the representation of Df_p in the charts $\phi \in \mathcal{A}$ and $\psi \in \mathcal{B}$ does not depend on these latter.

As a linear map from $\mathbb{R}^n$ to $\mathbb{R}^m$, it can be identified with an $n \times m$ matrix with entries in $\mathbb{R}$ which we just denote by $f'p$.

Definition 3.24. The matrix $f'p$ from Example 3.23 is called the *standard representation* of Df_p.

Let $\mathbf{Grph}_\ell$ be the category whose objects are graphs equipped with a length map $\ell : G^{(1)} \to [R, \infty[$ and whose morphisms are the graph morphisms $g : G \to G'$ such that the ratio $\frac{\ell(g_1(a))}{\ell(a)}$ does not depend on a; we denote it by r_g. Given two arrows a and b of G, we write $g_1 a = a'$ and $g_1 b = b'$. The map $f : \|G\| \to \|G'\|$ is defined by $f(a,t) = (a', r_g t)$ for every $t \in]0, \ell(a)[$, and by $f(a,b) = (a', b')$ when (a,b) is a traversal. Such a map f is called a *dilation*. Denoting by $\psi_{a'}$ and $\psi_{a'b'}$ the charts of the standard atlas of $\mathcal{A}_{G'}$ associated with a' and (a', b'), we have

$$\begin{array}{rcl} \psi_{a'} \circ f \circ \phi_a^{-1} & = & t \in]0, \ell(a)[\quad \mapsto \quad r_g t \in]0, \ell(a')[\\ \psi_{a'b'} \circ f \circ \phi_{ab}^{-1} & = & t \in]\frac{-\ell(a)}{2}, \frac{\ell(b)}{2}[\quad \mapsto \quad r_g t \in]\frac{-\ell(a')}{2}, \frac{\ell(b')}{2}[\end{array}$$

so the map f is smooth, and with the denotation from Definition 3.24 we have $f'p = r_g$ for all $p \in \|G\|$. In particular the construction $G \mapsto \mathcal{A}_G$ (Definition 3.14) extends to a functor

$$\mathcal{A}_* : \mathbf{Grph}_\ell \to \mathbf{Atl} .$$

Tangent bundle We denote by $GL_n(\mathbb{R})$ the group of $n \times n$ invertible matrices with real entries; its underlying set is open in $\mathrm{Mat}_n(\mathbb{R})$ equipped with the standard topology.

The set $T\mathcal{A}$ (Definition 3.16) is the support of an atlas: for every $\phi \in \mathcal{A}$ one defines the chart

$$T\phi \quad : \quad \bigsqcup_{p \in dom\ \phi} T_p\mathcal{A} \quad \to \quad cod\ \phi \times \mathbb{R}^n \tag{7}$$

by $T\phi(v) = \left(\phi p\, ,\ [\![v]\!]_p^\phi \right)$ for every $v \in T_p\mathcal{A}$. Given ϕ, $\psi \in \mathcal{A}$, the bijection $T\psi \circ (T\phi)^{-1} : cod\ \phi \times \mathbb{R}^n \to cod\ \psi \times \mathbb{R}^n$ satisfies the equality

$$T\psi \circ (T\phi)^{-1}(x,t) \quad = \quad \Big((\psi \circ \phi^{-1})x\, ,\ D(\psi \circ \phi^{-1})_x t \Big)$$

so it is smooth because so are the maps

$$\psi \circ \phi^{-1} : cod\ \phi \to cod\ \psi \qquad \text{and} \qquad D(\psi \circ \phi^{-1}) : cod\ \phi \to GL_n(\mathbb{R}) .$$

This atlas is called the *tangent bundle* of $\mathcal{A}$; it is also denoted by $T\mathcal{A}$. The construction extends to a functor $T : \mathbf{Atl} \to \mathbf{Atl}$: if the map $f : \mathcal{A} \to \mathcal{B}$ is smooth, then so is the map $Tf : T\mathcal{A} \to T\mathcal{B}$ defined by $Tf(v) = Df_p v$ for every $v \in T_p\mathcal{A}$. Indeed, assuming that $T\phi(v) = (x,t)$ we have

$$\left(T\psi \circ Tf \circ (T\phi)^{-1} \right)(x,t) \quad = \quad \left(\psi \circ f \circ \phi^{-1} x\, ,\ D(\psi \circ f \circ \phi^{-1})_x \cdot t \right) ,$$

and $D(\psi \circ f \circ \phi^{-1})$ is smooth as a map from $cod\ \phi$ to $\mathrm{Mat}_n^m(\mathbb{R})$. The tangent bundle functor preserves binary products in the sense that the bijection

$$\big((p,q), \phi \times \psi, (a,b)\big) \quad \leftrightarrow \quad \big((p, \phi, a), (q, \psi, b)\big)$$

induces a smooth diffeomorphism between $T(\mathcal{A} \times \mathcal{B})$ and $T\mathcal{A} \times T\mathcal{B}$.

Examples 3.25. The tangent bundle of a single element atlas $\{\phi\}$ is the single element atlas $\{T\phi\}$. After Example 3.17 we can suppose that $T\phi : (p,v) \in dom\ (\phi) \times \mathbb{R}^n \mapsto (\phi(p), v) \in cod\ (\phi) \times \mathbb{R}^n$.

Definition 3.26. Following Example 3.18 (5), for every standard chart ϕ of $(G_1, \ldots, G_n)$ the chart $T\phi$ from (7) is the product map $\phi \times id_{\mathbb{R}^n}$:

$$\begin{array}{rcl} T\phi : dom\ \phi \times \mathbb{R}^n & \to & cod\ \phi \times \mathbb{R}^n \\ (p,t) & \mapsto & (\phi p, t) \end{array} .$$

Assuming that f is as in Example 3.23 we have

$$Df_p(p,t) \quad = \quad Tf(p,t) \quad = \quad \big(\, fp\,,\ f'p\cdot t\,\big)$$

with $f'p\cdot t$ being the product of the standard matrix of Df_p (Definition 3.24) and the vector $t\in\mathbb{R}^n$. Bijections of the form $T\phi$ are called the *standard tangent charts* of $(G_1,\dots,G_n)$, they form the *standard tangent atlas* of $(G_1,\dots,G_n)$. The *standard representation* of the tangent bundle of $\mathcal{A}_{G_1}\times\cdots\times\mathcal{A}_{G_n}$ is the set $\|G_1\|\times\cdots\times\|G_n\|\times\mathbb{R}^n$ equipped the standard tangent atlas.

Vector fields and parallelizations. The projection $\pi : T\mathcal{A}\to\mathcal{A}$ sends every tangent vector to its attachment point. It is smooth because its local representation in the charts $T\phi$ and ϕ, see (7), is the projection $cod\,\phi\times\mathbb{R}^n\to cod\,\phi$. A *vector field* over $\mathcal{A}$ is a smooth map $f:\mathcal{A}\to T\mathcal{A}$ such that $\pi\circ f = id$. Such an f is called a *smooth section* of π; it associates a vector of $T_p\mathcal{A}$ to every point p of $\mathcal{A}$. A *parallelization* of $\mathcal{A}$ is a tuple $(f_1,\dots,f_n)$ of vector fields over $\mathcal{A}$ such that $(f_1(p),\dots,f_n(p))$ is a vector basis of $T_p\mathcal{A}$ for every point $p\in\mathcal{A}$.

Remark 3.27. For every parallelization $f_\star = (f_1,\dots,f_n)$ of $\mathcal{A}$ the following map is a smooth diffeomorphism

$$\big(\, p\,,\ (a_1,\dots,a_n)\,\big) \ \in\ \mathcal{A}\times\mathbb{R}^n \qquad \mapsto \qquad \sum_{i=1}^{n} a_i f_i(p) \ \in\ T\mathcal{A}$$

Conversely, from any smooth diffeomorphism $\Pi:\mathcal{A}\times\mathbb{R}^n\to T\mathcal{A}$ one recovers a parallelization of $\mathcal{A}$ putting $f_i = \Pi\circ\big(p\mapsto(p,e_i)\big)$ with $p\in\mathcal{A}$ and e_i denoting the i^{th} canonical basis of $\mathbb{R}^n$.

Examples 3.28. The standard parallelization of a single element atlas $\{\phi\}$ of dimension n is the tuple $(f_1,\dots,f_n)$ with $f_i(p) = (p,e_i)$ with e_i denoting the i^{th} vector of the canonical base of $\mathbb{R}^n$ (Examples 3.17 and 3.25).

Definition 3.29. Following Definition 3.26, the *standard parallelization* of $(G_1,\dots,G_n)$ is the tuple of vector fields $(f_1,\dots,f_n)$ with:

$$\begin{array}{rccc} f_i : & \|G_1\|\times\cdots\times\|G_n\| & \to & (\|G_1\|\times\cdots\times\|G_n\|)\times\mathbb{R}^n \\ & p & \mapsto & p\quad ,\ e_i \end{array}$$

for $i\in\{1,\dots,n\}$ and e_i the i^{th} vector of the canonical basis of $\mathbb{R}^n$. The *standard basis* of the tangent space at p is $(f_1(p),\dots,f_n(p))$; its *positive cone* is the set of vectors whose coordinates in the standard basis are non-negative.

Remark 3.30. If $(f_1,\dots,f_n)$ and $(g_1,\dots,g_m)$ are parallelizations of $\mathcal{A}$ and $\mathcal{B}$, then $(f_1,\dots,f_n,g_1,\dots,g_m)$ is a parallelization of $\mathcal{A}\times\mathcal{B}$. Strictly speaking, even if the first two tuples are standard parallelizations (Definition 3.29), the third one is not. Yet, it is related to a standard parallelization by the obvious smooth diffeomorphism

$$(\|G_1\|\times\mathbb{R})\times\cdots\times(\|G_n\|\times\mathbb{R}) \quad\cong\quad (\|G_1\|\times\cdots\times\|G_n\|)\times\mathbb{R}^n \quad .$$

Lemma 3.31. Assume that f is as in Example 3.23. The standard matrix representation $f'p$ (Definition 3.24) is the matrix of the linear map Df_p in the standard bases (Definition 3.29) of the tangent vector spaces at p and fp.

Proof. The matrix $f'p$ represents $D(\psi\circ f\circ\phi^{-1})_{\phi p}$ in the canonical bases of $\mathbb{R}^n$ and $\mathbb{R}^m$ regardless of the chosen standard charts ϕ and ψ provided their domains contain p and fp (Definition 3.24). We have the atlas morphisms $\phi:\{\phi\}\to\{id_{cod\,\phi}\}$, $\psi:\{\psi\}\to cod\,\{id_{cod\,\psi}\}$ (Example 3.7), and

$f : \{\phi\} \to \{\psi\}$, so we can write

$$D(\psi \circ f \circ \phi^{-1})_{\phi p} \quad = \quad D\psi_{fp} \circ Df_p \circ D\phi^{-1}_{\phi p}$$

Assume that $\mathbb{R}^n$ and $\mathbb{R}^m$ are equipped with their canonical bases, and that the tangent spaces at p and fp are equipped with their standard bases. Then the matrices of $D\psi_{fp}$ and $D\phi^{-1}_{\phi p}$ are identities, while the matrix of Df_p is $f'p$.

Remark 3.32. If γ is a curve on the standard atlas $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$, then for every standard chart ϕ whose domain contains $\gamma(t_0)$, we have the curve $\phi \circ \gamma$ in $\mathbb{R}^n$. By Lemma 3.31, the coordinates of the 'speed' vector $(\phi \circ \gamma)'(t_0)$ in the canonical basis of $\mathbb{R}^n$ are the coordinates of the tangent vector $\gamma'(t_0)$ in the standard basis at $\gamma(t_0)$.

4. Metrics

We recall basic notions of metric space theory (§4.1), and thoroughly describe metric graphs: the metric on $|G|$ is lifted to a pseudometric on $\|G\|$ along the desingularizator β_G (§4.2). The pseudometric on $\|G_1\| \times \cdots \times \|G_n\|$ and the standard parallelization on $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ (together with a norm on $\mathbb{R}^n$) induce two natural ways of defining the length of a piecewise smooth path: we prove that they coincide and determine the (pseudo)metrics on $|G_1| \times \cdots \times |G_n|$ and $\|G_1\| \times \cdots \times \|G_n\|$ (§4.3). The abovementioned (pseudo)metrics actually depend on the choice of a norm on $\mathbb{R}^n$, according to execution time consideration, we explain that the ∞-norm (Definition 4.42) should be preferred (§6). We start with a small reminder on paths.

A continuous map $\gamma : [a, b] \to X$ with $a \leqslant b$ is called a *path* on X. The *opposite* of γ is the path $\bar{\gamma} : [a, b] \to X$ with $\bar{\gamma}(t) = \gamma(a + b - t)$. The *endpoints* of γ are $\gamma(a)$ and $\gamma(b)$, more precisely its *starting* and *finishing* points. A *subdivision* of $[a, b]$ is a finite increasing sequence of elements of $[a, b]$ whose first and last elements are a and b. A one-to-one path is called an *arc*. A *pseudo-arc* is a path γ such that $\gamma^{-1}\{p\}$ is connected for every $p \in dom\,(\gamma)$ (the empty set is connected). A *weak reparametrization* is a nondecreasing path onto an interval (equivalently, an nondecreasing surjection between two intervals). A *reparametrization* is an increasing path onto an interval (equivalently, an increasing surjection between two intervals).

Remark 4.1. Two arcs $\alpha : [a, b] \to X$ and $\beta : [c, d] \to X$ with the same image and the same starting point are equal up to reparametrization. Indeed, denote by $\tilde{\alpha}$ and $\tilde{\beta}$ the homeomorphisms induced by α and β on their common image. Let $\theta = \tilde{\beta}^{-1} \circ \tilde{\alpha}$, we have $\alpha = \beta \circ \theta$ and the homeomorphism θ is increasing because $\theta(a) = c$.

Given a path $\delta : [b, c] \to X$ with $\gamma(b) = \delta(b)$, the *concatenation* $\gamma \cdot \delta$ is the path defined on $[a, c]$ such that $\gamma \cdot \delta(t) = \gamma(t)$ if $t \leqslant b$, and $\gamma \cdot \delta(t) = \delta(t)$ if $b \leqslant t$. If $dom\,(\delta) = [b', c]$ then we still write $\gamma \cdot \delta$ to mean $\gamma \cdot (\delta \circ \tau)$ with τ denoting the translation $t \mapsto t + b' - b$ from $[b, c - b' + b]$ to $[b', c]$.

For every path γ on a *Hausdorff* space X, there exists an *arc* α on X whose endpoints are those of γ. Since any subspace of a Hausdorff space is Hausdorff, we can suppose that $img\,(\alpha) \subseteq img\,(\gamma)$. The following result (Douady & Douady, 2020, p.12) states that the arc α can be chosen so that it satisfies extra properties that will be used later:

Lemma 4.2. For every path γ from p to q on a *Hausdorff* space X, there exists an increasing lower semicontinuous map $f : [0, r] \to dom\,(\gamma)$ such that $\gamma \circ f$ is an arc from p to q, and f is continuous if, and only if, γ is an arc. Moreover we can suppose that $f(0) = \min(dom\,(\gamma))$.

Proof. Assume that $dom\,(\gamma) = [x, y]$ and suppose that $\gamma^{-1}\{p\}$ contains more than one element. Let s and s' be the least and the greatest elements of $\gamma^{-1}\{p\}$. The subpath $\gamma|_{[s,s']}$ is a nontrivial loop that

we can shunt by considering $\gamma|_{[x,s]} \cdot \gamma|_{[s',y]}$. In order to obtain an arc, we need to shunt all the loops; the problem is that there might be infinitely many of them, and they might be nested (keep in mind space filling curves). Following (Douady & Douady, 2020, p.12) we consider the set Ω of all the open subsets of $[0,1]$ neither containing x nor y, and whose connected components $]s,s'[$ satisfies $\gamma(s)=\gamma(s')$. Every such connected component corresponds to a loop we wish to shunt. The set Ω ordered by inclusion is inductive (the Hausdorff property of X intervenes here) so it has a maximal element W (Zorn's Lemma). The only possible isolated points of $W^c=([x,y]\setminus W)$ are x and y. For every connected component C of $W^c\setminus\{x,y\}$ the restriction $\gamma|_C$ is an arc. The idea is to obtain an arc α from p to q as the (infinite) concatenation of the arcs $\gamma|_C$ for C connected component of $W^c\setminus\{x,y\}$. In order to formalize this construction, denote by λ the Lebesgue measure on $\mathbb{R}$ (we just need the measure of compact subsets of $\mathbb{R}$ actually). Define the map $h:[x,y]\to[0,\lambda(W^c)]$ by $h(t)=\lambda(W^c\cap[x,x+t])$ and the map $f:[0,\lambda(W^c)]\to[x,y]$ by $f(\tau)=\min(h^{-1}\{\tau\})$. The map f is a section of h (i.e. $h\circ f=id$), it is increasing, lower semicontinuous (but not continuous in general) and $\alpha=\gamma\circ f$ is the expected arc. Observe that W is the interior of $]x,y[\setminus img\,(f)$. Note that if γ is an arc, then $\emptyset$ is the only element of Ω, so both h and f are homeomorphisms.

In the preceding proof, we have $]x,y[\in\Omega$ if, and only if, $\gamma(x)=\gamma(y)$, in which case f is constant equal to 0. The Hausdorff property cannot be omitted: suppose that G is the graph from Example 2.16 and that $\|G\|$ is equipped with the topology of the standard atlas of G (Definition 3.14); any point (b,t) is related to any point (c,t') by a path, but there is no arc joining them.

Remark 4.3. A path on a Hausdorff space preserves closures. An arc on a Hausdorff space induces a homeomorphism on its image (Gauld, 1982, Theorem 6, p.32).

4.1 Metrics and pseudometrics

A *pseudometric* on a set M is a map $d:M\times M\to\mathbb{R}_+\cup\{\infty\}$ such that $d(x,x)=0$, $d(x,y)=d(y,x)$, and $d(x,z)\leqslant d(x,y)+d(y,z)$ for all x, y, and $z\in M$. It is a *metric* when $d(x,y)=0$ implies that $x=y$. A *(pseudo)metric space* is a set equipped with a (pseudo)metric; it comes with a topology whose open sets are the unions of *open balls*, namely the sets of the form

$$\left\{\, p\in M \;\middle|\; d(p,c)<r \,\right\}$$

with $r>0$ and $c\in M$. The *open r-neighborhood* of $X\subseteq M$ is the union of the open balls of radius r centered at some point of X.

Remark 4.4. Any path γ on a pseudometric space (M,d) is *uniformly continuous* in the sense that for every $\varepsilon>0$ there exists $\delta>0$ such that $|t-t'|<\delta$ implies $d(\gamma(t),\gamma(t'))<\varepsilon$ (Rudin, 1976, 4.19, p.91).

Definition 4.5. The *length* of a path γ on the (pseudo)metric space (M,d), which we denote by $L(\gamma)$, is defined as the least upper bound of the values

$$\sum_{k=1}^{n} d\big(\gamma(t_{k-1}),\gamma(t_k)\big)$$

with $t_0\leqslant\cdots\leqslant t_n$ running through all the possible subdivisions of $[a,b]$. We say that γ is *rectifiable* when $L(\gamma)$ is finite.

Remark 4.6. The length of a path is invariant under weak reparametrization: for every path $\gamma:[a,b]\to(M,d)$ and every weak reparametrization $\theta:[x,y]\to[a,b]$, both γ and $\gamma\circ\theta$ have the same length (Papadopoulos, 2013, 1.1.8). For this reason, when γ is rectifiable, it is often convenient to suppose that γ is *arclength parametrized*, i.e. for all $a\leqslant t\leqslant t'\leqslant b$, the length of the subpath $\gamma_{|[t,t']}$ is $t'-t$. Indeed, for any path γ there is a unique reparametrization $\theta:[0,L(\gamma)]\to dom\,(\gamma)$

such that $\gamma \circ \theta$ is arclength parametrized, see (Papadopoulos, 2013, §1.2) or (Bridson & Haefliger, 1999, Proposition 1.20(5), p.13).

Lemma 4.7. Every arclength parametrized pseudo-arc α is an arc.

Proof. Suppose that $t \leqslant t'$ and $\alpha(t) = \alpha(t')$. We have $t' - t = L(\alpha|_{[t,t']})$ because α is arclength parametrized, and the pseudo-arc α is constant on $[t, t']$. Therefore $t = t'$.

Definition 4.8. Let (M, d) be metric space. The path $\gamma : [a, b] \to M$ is a *geodesic* when $d(\gamma(t), \gamma(t')) = |t - t'|$ for all t, $t' \in [a, b]$. It is a *local geodesic* when for all $t \in [a, b]$ there exists $\varepsilon > 0$ such that the restriction of γ to $[a, b] \cap [t - \varepsilon, t + \varepsilon]$ is a geodesic. A subset $X \subseteq M$ is said to be *geodesically stable*[5] when for all p, $q \in X$ there exists a geodesic from p to q and they are all contained in X.

Definition 4.9. The *length (pseudo)metric* d_L associated with d is given by

$$d_L(p, q) \quad = \quad \inf \left\{ L(\gamma) \;\middle|\; \gamma \text{ path on } M \text{ from } p \text{ to } q \right\}.$$

We have $d_L(p, q) = \infty$ when there is no path from p to q.

The inequality $d \leqslant d_L$ is always satisfied: the (pseudo)metric d is said to be *intrinsic* when $d = d_L$. Note that if d is a metric, then so is d_L because $d_L(p, q) = 0$ implies $d(p, q) = 0$, and therefore $p = q$. The converse is also true: assume that d_L is a metric and $d(p, q) = 0$. Any map $\gamma : [a, b] \to M$ whose image is contained in $\{p, q\}$ is continuous with respect to the topology induced on M by d. Moreover we have $L(\gamma) = 0$, therefore $d_L(p, q) = 0$ and $p = q$.

Given the metric spaces (M_i, d_i) for $i \in \{1, \ldots, n\}$, there is no preferred metric d on the product $M_1 \times \cdots \times M_n$ among those making the projection maps continuous. The 1-Lipschitz maps are often taken as metric space morphisms, so we require that $d_i(p_i, q_i) \leqslant d(p, q)$ for all points $p = (p_1, \ldots, p_n)$ and $q = (q_1, \ldots, q_n)$, and all $i \in \{1, \ldots, n\}$. In the context of this article, it also desirable that each slice

$$\sigma_i \;:\; M_i \quad \hookrightarrow \quad \{p_1\} \times \cdots \times \{p_{i-1}\} \times M_i \times \{p_{i+1}\} \times \cdots \times \{p_n\}$$

is an isometry, i.e. $d_i(x, y) = d(\sigma_i(x), \sigma_i(y))$ for all $i \in \{1, \ldots, n\}$ and x, $y \in M_i$. The most well-known example of product metric is the euclidean one

$$\sqrt{\sum_{i=1}^{n} {d_i}^2} \quad ,$$

which is a special instance of α-product metric:

Definition 4.10. For $\alpha \in [1, \infty]$ the *α-product metric* is given by $({d_1}^\alpha + \cdots + {d_n}^\alpha)^{\frac{1}{\alpha}}$ with the usual convention that the preceding formula is replaced by $\max \left\{ d_i \;\middle|\; i \in \{1, \ldots, n\} \right\}$ when $\alpha = \infty$. The α-product metric is also given by

$$(p, q) \quad \mapsto \quad |d_1(p_1, q_1), \ldots, d_n(p_n, q_n)|_\alpha$$

with $|\_|_\alpha$ denoting the *α-norm* on $\mathbb{R}^n$, i.e. $|t|_\alpha = (|t_1|^\alpha + \cdots + |t_n|^\alpha)^{\frac{1}{\alpha}}$ when $\alpha \neq \infty$ and $|t|_\infty = \max \left\{ |t_i| \;\middle|\; i \in \{1, \ldots, n\} \right\}$ for every $t \in \mathbb{R}^n$.

[5] Such a set is said to be *convex* in (Bridson & Haefliger, 1999, p.4); we have reserved this term for another usage (§5.1).

Proposition 4.11. Any (finite) α-product of intrinsic (pseudo)metrics is an intrinsic (pseudo)metric.

Proof. See the proof of (Bridson & Haefliger, 1999, 5.3(1), p.56) for the case $\alpha = 2$. One readily checks that it remains valid for any $\alpha \in [1, \infty[$ instead of 2; see also (Bridson & Haefliger, 1999, 5.5(1), p.58) for $\alpha < \infty$; the case $\alpha = \infty$ just requires to drop the exponent α and replace some occurrences of '+' by 'max'.

4.2 Metric graphs

Given the map $\ell : G^{(1)} \to \mathbb{R}_+$ assigning a length to every arrow of a graph G, we wish to define a pseudometric space in which the distance between the points p and q is the greatest lower bound of the lengths of the paths joining p and q. Such a space is called a *metric graph*, it can be defined as the quotient of a disjoint union of intervals (Bridson & Haefliger, 1999, 5.21(3)). If ℓ has a *positive* lower bound R, the resulting pseudometric is an intrinsic metric (Bridson & Haefliger, 1999, 5.20,28,29). Every metric graph is actually a 1-dimensional *metric simplicial complexe* (Bridson & Haefliger, 1999, 7.40(1)). Under the extra assumption that the set $\{\ell(a) \mid a \in G^{(1)}\}$ is finite, a metric graph inherits strong properties (Bridson & Haefliger, 1999, 7,5,9,10,19,21,24,50,62) from metric simplicial complexes built on finitely many *shapes* (Bridson & Haefliger, 1999, 7.2). Nevertheless, dealing with the latter requires a lot of material that is irrelevant in dimension 1. This is why we have chosen a more elementary approach. The pedestrian construction given here is standard (Bridson & Haefliger, 1999, §1.9, p.6), requires no prerequisites, and allows direct proofs without the finiteness hypothesis about the image of the map ℓ (at least for the results we are concerned with). We have slightly adapted the presentation to our needs.

Each arrow a comes with the canonical map (see Definition 2.3 for $|G|$)

$$\chi_a : t \in [0, \ell(a)] \mapsto \begin{cases} \operatorname{src} a & \text{if } t = 0 \\ \operatorname{tgt} a & \text{if } t = \ell(a) \in |G| \\ (a, t) & \text{otherwise} \end{cases} .$$

Remark 4.12. The *finest* topology making all the maps χ_a continuous is the topology of $|G|$ (Definition 2.7). It suffices to note that if V is an open set of the latter topology containing the vertex v, and if a is an arrow of G such that $\operatorname{src}(a) = v$ (resp. $\operatorname{tgt}(a) = v$), then $\chi_a^{-1}(V)$ contains $[0, \varepsilon[$ (resp. $]\ell(a) - \varepsilon, \ell(a)]$).

Lemma 4.13. If $I =]0, \ell(a)[$ or if I is a proper closed subinterval of $[0, \ell(a)]$, then the map χ_a induces a homeomorphism from I to $\chi_a(I)$. If $\operatorname{src}(a) \neq \operatorname{tgt}(a)$, then χ_a induces a homeomorphism on its image.

Proof. Denote by $\tilde{\chi}_a :]0, \ell(a)[\to \chi_a(]0, \ell(a)[)$ the bijection induced by χ_a. The direct image map $\tilde{\chi}_a^*$ induces a bijection between the open subintervals of $]0, \ell(a)[$ and the segments of G (Definition 2.5) contained in $\chi_a(]0, \ell(a)[)$. The former (resp. the latter) collection is a base of the topology of $]0, \ell(a)[$ (resp. $\chi_a(]0, \ell(a)[)$). In all the other cases, the map induced by χ on I is an arc on the space $|G|$ which is Hausdorff (Remarks 2.6 and 4.12). We conclude by Remark 4.3.

Remark 4.14. Given an arrow a and a traversal (a, b) of a graph G, the standard charts ϕ_a and ϕ_{ab} (Definition 3.12) are related to the mappings χ_a and χ_b : for all $t \in cod\, \phi_a$ we have $\beta_G \circ \phi_a^{-1}(t) =$

$\chi_a(t) = (a, t)$, and for all $t \in cod\ \phi_{ab}$ we have

$$\beta_G \circ \phi_{ab}^{-1} \quad = \quad \begin{cases} \chi_a(\ell(a) + t) = (a, \ell(a) + t) & \text{if } t < 0 \\ \chi_a(\ell(a)) = \text{tgt } a = \text{src } b = \chi_b(0) & \text{if } t = 0 \\ \chi_b(t) = (b, t) & \text{if } t > 0 \end{cases}$$

with β_G denoting the desingularizator of G (Definition 2.11).

We define the standard metric and pseudometric structures on $|G|$ and $\|G\|$:

Lemma 4.15. A path on $|G|$ visits finitely many vertices.

Proof. The union V of stars of radius $\frac{R}{3}$ centered at some vertex of G (Definition 2.4) does not contain any middle point. The set $U = |G| \setminus \{\text{vertex}\}$ is an open neighborhood of $|G| \setminus V$ that does not contain any vertex of G. The family C made of U and the stars of radius $\frac{R}{3}$ centered at some vertex of G *visited* by γ is an open covering of the image of γ. By construction, no *strict* subfamily of C covers the image of γ which is compact, so γ visits finitely many vertices of G.

Definition 4.16. A *step* s is a path of the form $\chi_a \circ \theta$ with θ arc on $[0, \ell(a)]$; it is said to be *affine* if so is θ; *directed* or *antidirected* according to whether θ is increasing or decreasing; *standard* if θ is an inclusion $[x, y] \hookrightarrow [0, \ell(a)]$ or its opposite, in which case the step s is denoted by (a, x, y) or (a, y, x) accordingly. We define the *length* of s, denoted by $\ell(s)$, as the length of the interval *img* θ. The *support* of s is the arrow a. The steps $\chi_a \circ \theta$ and $\chi_b \circ \theta'$ are said to be *overlapping* when $a = b$ and the interior of *img* $\theta \cap$ *img* θ' is nonempty. In higher dimensions, a step of $|G_1| \times \cdots \times |G_n|$ is a product $s_1 \times \cdots \times s_n$ of steps s_i of $|G_i|$.

Remark 4.17. If the step $s = \chi_a \circ \theta$ is written as the concatenation $s_1 \cdots s_n$ then we have the subdivision $t_1 < \cdots < t_n$ of *dom* (s) such that $s_i = \chi_a \circ \theta_i$ with θ_i the restriction of θ to $[t_{i-1}, t_i]$. Hence $\theta(t_{i-1}) \leqslant \theta(t_i)$ for every $i \in \{1, \ldots, n\}$ or $\theta(t_{i-1}) \geqslant \theta(t_i)$ for every $i \in \{1, \ldots, n\}$. In both cases we have $|\theta(t_n) - \theta(t_0)| = \sum_{i=1}^{n} |\theta(t_i) - \theta(t_{i-1})|$ that is to say $\ell(s) = \ell(s_1) + \cdots + \ell(s_n)$.

The next result justifies the notion of a step:

Lemma 4.18. If the concatenations of steps $s_1 \cdots s_n$ and $\tilde{s}_1 \cdots \tilde{s}_m$ are equal, then $\ell(s_1) + \cdots + \ell(s_n) = \ell(\tilde{s}_1) + \cdots + \ell(\tilde{s}_m)$.

Proof. Let $t_0 < \cdots < t_n$ and $\tilde{t}_0 < \cdots < \tilde{t}_m$ with *dom* $(s_i) = [t_{i-1}, t_i]$ and *dom* $(\tilde{s}_i) = [\tilde{t}_{i-1}, \tilde{t}_i]$. Following our notion of concatenation we have $[t_0, t_n] = [\tilde{t}_0, \tilde{t}_m]$. Suppose that $t_0 < \cdots < t_n$ is finer than $\tilde{t}_0 < \cdots < \tilde{t}_m$ (otherwise consider a third subdivision that is finer than both of them) and conclude by Remark 4.17.

Definition 4.19. A path γ on $|G|$ is said to be *admissible* (resp. *piecewise affine*) when it can be written as a (finite) concatenation of (resp. *affine*) steps $s_1, \ldots, s_n$; the *length* of γ is $\ell(\gamma) = \ell(s_1) + \cdots + \ell(s_n)$ (Lemma 4.18). A path on $|G_1| \times \cdots \times |G_n|$ is said to be *admissible* (resp. *piecewise affine*) when so are all its components. A path on $\|G_1\| \times \cdots \times \|G_n\|$ is said to be *piecewise affine* when so is its image under the product of desingularizators $\beta_1 \times \cdots \times \beta_n$ (Definition 2.11).

Remark 4.20. The class of admissible paths is stable under reparametrization; the image and the length of an admissible path are left unchanged under reparametrization.

Remark 4.21. Every step is admissible and for every arrow a, we have $\ell(\chi_a) = \ell(a)$.

Lemma 4.22. Every arc on $|G|$ is admissible.

Proof. The set of vertices visited by α is finite (Lemma 4.15) and α is one-to-one, so the set $\gamma^{-1}(\{\text{vertices}\})$ is finite: let $t_0 < \cdots < t_k$ be its enumeration. Given $i \in \{1, \ldots, k\}$, the map $\theta : \xi \in [t_{i-1}, t_i] \to \chi_a^{-1}(\alpha(\xi)) \in [0, \ell(a)]$ is an arc (Remark 4.3 and Lemma 4.13) such that the restriction of α to $[t_{i-1}, t_i]$ is $\chi_a \circ \theta$.

Lemma 4.23. If the image of the arc α is contained in the image of the admissible path γ, then $\ell(\alpha) \leqslant \ell(\gamma)$. Equality holds if, and only if, γ is an arc and $img\,(\alpha) = img\,(\gamma)$.

Proof. We write γ as a concatenation of steps $\gamma_1 \cdots \gamma_n$ (Definition 4.19). If α is reduced to a single step, then we can write it as $\alpha_1 \cdots \alpha_m$ so that for every $j \in \{1, \ldots, m\}$ and every $i \in \{1, \ldots, n\}$, if α_j and γ_i are overlapping (Definition 4.16), then $img\,(\alpha_j) \subseteq img\,(\gamma_i)$; in particular there is some $i \in \{1, \ldots, n\}$ such that $img\,(\alpha_j) \subseteq img\,(\gamma_i)$. As a consequence of Lemma 4.22, the same holds for any arc α. Hence we have the map $f : \{1, \ldots, m\} \to \{1, \ldots, n\}$ defined by

$$f(j) \quad = \quad \min\left\{ i \in \{1, \ldots, n\} \;\middle|\; \mathrm{tr}(\alpha_j) \subseteq \mathrm{tr}(\gamma_i) \right\}$$

and we observe that

$$\ell(\alpha) \quad = \quad \sum_{i=1}^{n} \left(\sum_{j \in f^{-1}\{i\}} \ell(\alpha_j) \right). \tag{8}$$

For any $i \in \{1, \ldots, n\}$ and any distinct $j, j' \in f^{-1}\{i\}$ the steps α_j and $\alpha_{j'}$ do not overlap so we have

$$\sum_{j \in f^{-1}\{i\}} \ell(\alpha_j) \quad \leqslant \quad \ell(\gamma_i)\,. \tag{9}$$

Since $\ell(\gamma) = \ell(\gamma_1) + \cdots + \ell(\gamma_n)$ (Definition 4.19), we deduce from (8) and (9) that $\ell(\alpha) \leqslant \ell(\gamma)$. Under the assumption that γ is an arc such that $img\,(\alpha) = img\,(\gamma)$, the roles of α and γ can be swapped, so we have $\ell(\alpha) = \ell(\gamma)$.

If the image of γ strictly contains that of α, then some step γ_i is not covered by α and consequently, we have some index i for which the inequality (9) is strict.

Suppose that γ is not an arc and $img\,(\alpha) = img\,(\gamma)$. We have a point $p \in img\,(\gamma)$ such that $\gamma^{-1}\{p\}$ is not a singleton, and (a unique) $\tau \in dom\,(\alpha)$ such that $\alpha(\tau) = p$. Let $j \in \{1, \ldots, m\}$ be such that $\tau \in dom\,(\alpha_j)$ (there is at most two such indices, and they are consecutive). In particular $img\,(\alpha_j) \subseteq img\,(\gamma_k)$ with $k = f(j)$, and $t \in dom\,(\gamma_k)$ with $t = \min(\gamma^{-1}\{p\})$. Let $t' = \max(\gamma^{-1}\{p\})$, we have $t < t'$ and $\gamma(t') = \gamma(t)$. Denote by a the support of the step γ_k. If the image of χ_a is covered by γ, then it is also covered by the arc α, so $\mathrm{src}(a) \neq \mathrm{tgt}(a)$. In any case the step γ_k is an arc (Lemma 4.13), so $t' \notin dom\,(\gamma_k)$. Let $k' \in \{1, \ldots, n\}$ such that $t' \in dom\,(\gamma_{k'})$; one of the steps $\gamma_{k'-1}$, $\gamma_{k'}$, and $\gamma_{k'+1}$ is overlapping α_j. Up to renaming, suppose that it is $\gamma_{k'}$. We have $img\,(\alpha_j) \subseteq img\,(\gamma_{k'})$ and $k < k'$. It follows that the inequality (9) is strict for $i = k'$. In both cases ($img\,(\alpha) \neq img\,(\gamma)$ or γ not one-to-one), we have $\ell(\alpha) < \ell(\gamma)$.

Definition 4.24. The *standard pseudometric* (one readily checks that it is a pseudometric) on $|G|$ is the map $d_{|G|} : |G| \times |G| \to \mathbb{R}_+ \cup \{\infty\}$ defined by

$$d_{|G|}(p, q) \quad = \quad \inf\left\{ \ell(\gamma) \;\middle|\; \gamma : \text{admissible path from } p \text{ to } q \right\}$$

with the convention that $\inf \emptyset = \infty$; p and q are said to be *neighbors* when $d_{|G|}(p, q) < \frac{R}{2}$.

Remark 4.25. Let γ be an admissible path from p to q. The topology of $|G|$ (Definition 2.7) is Hausdorff (Remark 2.6) so we have an arc α from p to q such that $img\,(\alpha) \subseteq img\,(\gamma)$ (Lemma 4.2).

By Lemmas 4.22 and 4.23 we have $\ell(\alpha) \leqslant \ell(\gamma)$. Hence we can write 'arc' instead of 'admissible path' in Definition 4.24.

Remark 4.26. Let γ be an admissible path from (a, t) to some vertex. By a connectedness argument, we have $\chi_a[0, t] \subseteq img\,(\gamma)$ or $\chi_a[t, \ell(a)] \subseteq img\,(\gamma)$. In particular the restrictions of χ_a to $[0, t]$ and $[t, \ell(a)]$ are arcs (Lemma 4.13). It follows from Lemma 4.23 that $\ell(\gamma) \geqslant \min\{t, \ell(a) - t\}$. If γ visits two *distinct* vertices, then by a connectedness argument we have $img\,(\chi_a) \subseteq img\,(\gamma)$ for some arrow a, so $\ell(\gamma) \geqslant \ell(a) \geqslant R$.

Let $p = \chi_a(s)$ and $q = \chi_a(t)$ with $\{s, t\} \not\subseteq \{0, \ell(a)\}$, which amounts to say that at least one of the points p and q is not a vertex.

i) If $\{s, t\} \subseteq]0, \ell(a)[$ or $\operatorname{src} a \neq \operatorname{tgt} a$, then (a, s, t) is the unique standard step from p to q (Definition 4.16),

ii) If $\{s, t\} \not\subseteq]0, \ell(a)[$ and $\operatorname{src} a = \operatorname{tgt} a$, then we have exactly two standard steps σ and σ' from p to q, namely:

- $(a, s, 0)$ and $(a, s, \ell(a))$ if $t \in \{0, \ell(a)\}$, or
- $(a, 0, t)$ and $(a, \ell(a), t)$ if $s \in \{0, \ell(a)\}$.

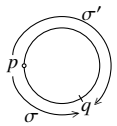

In both cases the sum of their length is $\ell(a)$, and only one of them is directed. From the extra assumption that $\frac{\ell(a)}{2} \notin \{s, t\}$ (i.e. neither p nor q stands on the midpoint of $\{a\} \times]0, a[$) we deduce that one of the steps σ and σ' is strictly shorter than the other.

Definition 4.27. Let $p, q \in img\,(\chi_a)$, one of them not being vertex. The *dash* from p to q, denoted by $p \to q$, is the shortest standard step from p to q (Definition 4.16) with the convention that if there are two such steps, then $p \to q$ is the directed one. When it is not the only step from p to q, the other one is called the *complement* of $p \to q$. In higher dimension, for p and $q \in |G_1| \times \cdots \times |G_n|$, we write $p \to q$ to denote the product $(p_1 \to q_1) \times \cdots \times (p_n \to q_n)$ when every component is well-defined.

Remark 4.28. The step $\chi_a \circ \theta$ in an arc if, and only if, $\operatorname{src}(a) \neq \operatorname{tgt}(a)$ or $img\,\theta \neq dom\,(\chi_a)$. It follows that every dash is an arc.

Any two neighbors (Definition 4.24) are related by a unique geodesic which is made of at most two dashes (Definition 4.27):

Lemma 4.29. Suppose that $p = \chi_a(s)$ and $q = \chi_b(t)$ with $d_{|G|}(p, q) < \frac{R}{2}$ and $p \neq q$.

If $a = b$ and ($\operatorname{src} a \neq \operatorname{tgt} a$ or $|t - s| < \frac{1}{2}\ell(a)$), then $\ell(p \to q) = d_{|G|}(p, q) = |t - s|$.

If $a = b$ and $\operatorname{src} a = \operatorname{tgt} a$ and $|t - s| \geqslant \frac{1}{2}\ell(a)$, or if $a \neq b$, then there is a unique vertex v (which is an endpoint shared by a and b) such that $\ell(p \to v \to q) = d_{|G|}(p, q)$, i.e. $\ell(a) - |t - s|$ in the case where $a = b$.

Moreover, the only admissible path of length $d_{|G|}(p, q)$ from p to q is, up to reparametrization, the arc $p \to q$ or the arc $p \to v \to q$ accordingly.

Proof. The arrows a and b share at least one endpoint, and one (at least) of the points p and q is not a vertex (Remark 4.26). We have an admissible path γ from p to q with $\ell(\gamma) < R/2$ because $d_{|G|}(p, q) < R/2$ (Definition 4.24). There are three cases to examine:

Case 1. Assume that $a \neq b$.

By Remark 4.26 and because $\ell(\gamma) < R/2$ we have $img\,(p \to v) \subseteq img\,(\gamma)$ with $v \in \{\operatorname{src} a, \operatorname{tgt} a\}$ (Definition 4.27). Similarly we have $img\,(v' \to q) \subseteq img\,(\gamma)$ with $v' \in \{\operatorname{src} b, \operatorname{tgt} b\}$. We deduce that $v = v'$ from Remark 4.26. The admissible path $\alpha = (p \to v \to q)$ is an arc because so are the

dashes $p \to v$ and $v \to q$ (Remark 4.28) and $img\,(p \to v) \cap img\,(v \to q) = \{v\}$. We apply Lemma 4.23 to conclude that $\ell(\alpha) \leqslant \ell(\gamma)$ with equality if, and only if, γ is an arc having the same image as α, in which case γ and α are equal up to reparametrization (Remark 4.1). By Definition 4.24 we have $\ell(\alpha) = d_{|G|}(p, q)$.

The two remaining cases share the assumption that $a = b$. Up to exchanging the roles of s and t we can suppose that $s \leqslant t$. By Remark 4.26, the image of γ either contains the set $B = \{a\} \times [s, t]$ or the set $C = \{a\} \times ([0, s] \cup [t, \ell(a)])$.

Case 2. Assume that $\mathrm{src}\, a = \mathrm{tgt}\, a$, which we denote by v, and $t - s \geqslant \frac{1}{2}\ell(a)$.
The standard steps $(a, s, 0)$ and $(a, \ell(a), t)$ are $p \to v$ and $v \to q$ (Definition 4.27). The path γ does not cover B, otherwise we would have $\ell(\gamma) \geqslant R/2$ (Lemma 4.23). Therefore $img\,(\gamma)$ contains C which is the image of $p \to v \to q$. We have $img\,(p \to v) \cap img\,(v \to q) = \{v\}$, from which we deduce, as in the first case, that $\ell(p \to v \to q) \leqslant \ell(\gamma)$ with equality if, and only if, γ and $p \to v \to q$ are equal up to reparametrization.

Case 3. Assume that $\mathrm{src}\, a \neq \mathrm{tgt}\, a$ or $t - s < \frac{1}{2}\ell(a)$.
The dash $p \to q$ is well-defined (Definition 4.27), and its length is $t - s$. If $C \subseteq img\,(\gamma)$ then one of the following situations occurs:

- γ visits $\mathrm{src}\, a$ and $\mathrm{tgt}\, a$ (which are distinct), so $\ell(\gamma) > R$ (Remark 4.26), or
- $\ell(\gamma) \geqslant \ell(a) - (t - s) > \frac{1}{2}\ell(a)$ (Lemma 4.23).

In both situations the inequality $\ell(\gamma) < R/2$ is not satisfied, therefore $img\,(\gamma)$ contains B which is the trace of the arc $p \to q$ (Remark 4.28). With the same arguments as in the two first cases, we deduce that $\ell(p \to q) \leqslant \ell(\gamma)$ with equality if, and only if, γ and $p \to q$ are equal up to reparametrization.

Lemma 4.30. For every admissible path γ on $|G|$ we have $L(\gamma) = \ell(\gamma)$, i.e. the length in the sense of Definition 4.5 is equal to the length in the sense of Definition 4.19.

Proof. Given a step s, if the subdivision $t_0 < \cdots < t_n$ of $dom\,(s)$ is fine enough, we have $\ell(s_i) < \frac{R}{2}$ with s_i denoting the restriction of s to $[t_{i-1}, t_i]$. It follows that $d_{|G|}(s(t_i), s(t_{i-1})) < \frac{R}{2}$ (Definition 4.24). From Lemma 4.29 we deduce that $\ell(s_i) = d_{|G|}(s(t_i), s(t_{i-1}))$, and therefore

$$\sum_{i=1}^{n} d_{|G|}(s(t_i), s(t_{i-1})) \quad = \quad \sum_{i=1}^{n} \ell(s_i) \quad = \quad \ell(s)\,,$$

the second equality being given by Remark 4.17. Hence $L(s) = \ell(s)$. Given the concatenation of steps $\gamma_1 \cdots \gamma_n$ we have

$$L(\gamma_1 \cdots \gamma_n) \quad = \quad \sum_{i=1}^{n} L(\gamma_i) \quad = \quad \sum_{i=1}^{n} \ell(\gamma_i) \quad = \quad \ell(\gamma_1 \cdots \gamma_n)$$

the last equality being given by Definition 4.19.

Definition 4.31. Given two neighbors p and q (Definition 4.24), the *shortest arc* from p to q is the one provided by Lemma 4.29.

Corollary 4.32. The map $d_{|G|}$ (Definition 4.24) is a metric.

Proof. Suppose that $d_{|G|}(p, q) = 0$ so p and q are neighbors and the shortest admissible path from p to q is of null length. Following Definition 4.19 we deduce that $p = q$.

Remark 4.33. If we allow the lengths of arrows to be taken in $\mathbb{R}_+$ (or even $]0,\infty[$) instead of $[R,\infty[$ there might be arbitrarily short paths joining two distinct points, resulting in a pseudometric space instead of a metric one.

The length metric is approximated by arcs:

Lemma 4.34. For every path γ from p to q on the metric graph $|G|$, there exists an arc α from p to q such that $img\,(\alpha) \subseteq img\,(\gamma)$, $L(\alpha) \leqslant L(\gamma)$, and γ is a pseudo-arc if, and only if $L(\alpha) = L(\gamma)$.

Proof. Let γ be a path on $|G|$ from p to q. If γ is not rectifiable, then consider any arc α from p to q such that $img\,(\alpha) \subseteq img\,(\gamma)$. By Lemma 4.22 we know that α is admissible, hence it is rectifiable (Lemma 4.30).

If γ is rectifiable, then we can suppose that it is arclength parametrized (Remark 4.6). Let f be as in Lemma 4.2, so $\alpha = f \circ \gamma$ is an arc from p tp q. We have $img\,(\alpha) \subseteq img\,(\gamma)$ and given a subdivision $\tau_0 < \cdots < \tau_n$ of $dom\,(\alpha) = dom\,(f)$, the sequence $f(\tau_0) < \cdots < f(\tau_n)$ of elements of $dom\,(\gamma)$ satisfies $\gamma(f(\tau_i)) = \alpha(\tau_i)$. It follows that $L(\alpha) \leqslant L(\gamma)$. If γ is an arc, then f is a reparametrisation, and therefore $L(\alpha) = L(\gamma)$; otherwise f is not continuous, which means here that we have $\tau \in dom\,(f)$ such that

$$r \quad = \quad \inf\left\{ f(\tau+\varepsilon) - f(\tau) \;\middle|\; \varepsilon > 0;\ t+\varepsilon \in dom\,(\gamma) \right\} \quad > \quad 0$$

Let $t = f(\tau)$ and $t' = t + r$, the subpath $\gamma|_{[t,t']}$ is a loop (it has to be so because $\gamma \circ f$ is continuous). We have $L(\gamma|_{[t,t']}) = r > 0$ because γ is arclength; it follows that for any $\varepsilon > 0$ we have a subdivision $t_0 < \cdots < t_m$ of $[t,t']$ such that

$$r - \varepsilon \quad \leqslant \quad \sum_{i=1}^{m} d_{|G|}(\gamma(t_i), \gamma(t_{i-1}))\ .$$

There is $i \in \{0,\dots,n\}$ such that $f(\tau_i) \leqslant t_0 < \cdots < t_m < f(\tau_{i+1})$ with the convention that $f(\tau_{n+1})$ stands for M. By inserting $t_0 < \cdots < t_m$ between $f(\tau_i)$ and $f(\tau_{i+1})$ (do not duplicate t_0 if $t_0 = f(\tau_i)$) we obtain a subdivision $\sigma_0 < \cdots < \sigma_N$ of $dom\,(\gamma)$ such that

$$\sum_{i=1}^{n} d_{|G|}(\alpha(\tau_i), \alpha(\tau_{i-1})) + r - \varepsilon \quad \leqslant \quad \sum_{i=1}^{N} d_{|G|}(\gamma(\sigma_i), \gamma(\sigma_{i-1}))\ .$$

It follows that $L(\alpha) + r \leqslant L(\gamma)$ and therefore $L(\alpha) < L(\gamma)$.

Lemma 4.35. The metric $d_{|G|}$ is intrinsic.

Proof. We have

$$\begin{aligned}
d_{|G|}(p,q) &= \inf\left\{ \ell(\gamma) \;\middle|\; \gamma : \text{admissible path from } p \text{ to } q \right\} && \text{Definition 4.24}\\
&= \inf\left\{ \ell(\gamma) \;\middle|\; \gamma : \text{arc from } p \text{ to } q \right\} && \text{Remark 4.25}\\
&= \inf\left\{ L(\gamma) \;\middle|\; \gamma : \text{arc from } p \text{ to } q \right\} && \text{Lemmas 4.22}\\
& && \text{and 4.30}\\
&= \inf\left\{ L(\gamma) \;\middle|\; \gamma : \text{path } p \text{ to } q \right\} && \text{Lemma 4.34}
\end{aligned}$$

An open ball with 'small' radius is geodesically stable (Definition 4.8):

Corollary 4.36. Every open ball B of radius $r \leqslant \frac{R}{4}$ is geodesically stable. If the center of B is an isolated vertex then B is a singleton; otherwise the center is $\chi_c(\rho)$ for some arrow c and

$\rho \in [0, \ell(c)]$, and

$$B = \begin{cases} \{c\} \times]\rho - r, \rho + r[& \text{if } r \leqslant \rho \leqslant \ell(c) - r, \\ S(\operatorname{src} c, r - \rho) \cup \{c\} \times]0, \rho + r[& \text{if } r > \rho, \text{ and} \\ \{c\} \times]\rho - r, \ell(c)[\cup S(\operatorname{tgt} c, \rho + r - \ell(c)) & \text{if } \rho + r > \ell(c). \end{cases} \tag{10}$$

with $S(v, x)$ denoting the star centered at the vertex v with radius x (Definition 2.4).

Proof. The case where the center of B is an isolated vertex is obvious, so we assume it is not. First we prove that the right hand part of each equality appearing in (10) is a geodesically stable set. Let p and q be two distinct points of such a set, which we denote by X in the sequel. If X is described by the 2^{nd} (resp. 3^{rd}) case of (10), then it contains a unique vertex v, which is the center of the star appearing in this description, i.e. $v = \operatorname{src} c$ (resp. $v = \operatorname{tgt} c$). In particular, the connected components of $X \setminus \{v\}$ are of the form $\{x\} \times I$ with I open interval of $]0, \ell(x)[$ of length at most $2r$ which is either initial (in which case $\operatorname{src} x = v$) or final (in which case $\operatorname{tgt} x = v$). Since p and q are neighbors ($d_{|G|}(p, q) < R/2$), the shortest arc α from $p = \chi_a(s)$ to $q = \chi_b(t)$ (Definition 4.31) is the only admissible path of length $d_{|G|}(p, q)$ up to reparametrization (Lemma 4.29). We have one of the following situations:

- There is a connected component $\{x\} \times I$ of $X \setminus \{v\}$ such that $p, q \in \{v\} \cup \{x\} \times I$ (therefore $x \in \{a, b\}$). Since the length of I is at most $2r$, the arc α is reduced to the dash $p \to q$ (Definition 4.27) and its image is included in $\{v\} \cup \{x\} \times I$.
- The points p and q respectively belong to the connected components $\{a\} \times I$ and $\{b\} \times J$ of $X \setminus \{v\}$. We observe that $\operatorname{length}(I) + \operatorname{length}(J) \leqslant 2r$. If $a = b$, then $I \cap J = \emptyset$ and
$$|t - s| \geqslant \ell(a) - (\ell(I) + \ell(J)) \geqslant \ell(a) - 2r \geqslant \ell(a) - \tfrac{R}{2} \geqslant \tfrac{1}{2}\ell(a)\,.$$
Hence no matter that $a \neq b$ or $a = b$, the arc α is $p \to v \to q$ and its image is contained in $(\{a\} \times I) \cup \{v\} \cup (\{b\} \times J)$.

By Lemma 4.34 any geodesic γ from p to q is an arc up to reparametrization. Therefore α and γ are equal up to weak reparametrization, and we have in particular $img(\gamma) = img(\alpha)$. Hence X is geodesically stable.

The open ball B obviously contains X, it remains to check the converse inclusion. Any point p of B is a neighbor of the center q of B. Let α be the shortest arc from p to q (Definition 4.31), we have $L(\alpha) < r$. According to the description of α given in Lemma 4.29 we have $p = \chi_a(t)$ and $q = \chi_c(\rho)$ with $t \in [0, \ell(a)]$ and $\rho \in [0, \ell(c)]$. We have three cases to deal with according to the inequalities relating r and ρ.

Suppose that $r \leqslant \rho \leqslant \ell(c) - r$. The length of any path whose image covers $\{c\} \times]0, r]$ or $\{c\} \times [\ell(c) - r, \ell(c)[$ is at least r. Hence the image of α is included in $\{c\} \times]\rho - r, \rho + r[$.

Now assume that $r > \rho$. If p and q are on the same arrow of $|G|$ (i.e. $a = c$) then α is reduced to the dash $p \to q$ whose image is $\chi_c[t \wedge \rho, t \vee \rho]$. This latter is included in $\{\operatorname{src} c\} \cup (\{c\} \times]0, \rho + r[)$. If p and q are not on the same arrow of $|G|$ (i.e. $a \neq c$), then α is $p \to v \to q$ with v the only vertex of B. The length of $v \to q$ is ρ so the length of $p \to v$ is $d_{|G|}(p, q) - \rho < r - \rho$. Moreover $v = \operatorname{src} c$ because $r < R \leqslant \ell(c)$. Hence $p \in S(v, r - \rho)$. The last case is dealt with the same way.

A similar result holds for closed balls with the constraint $r < \frac{R}{4}$ and the boundaries in the description (10) adapted accordingly. We lift the distance $d_{|G|}$ along the desingularizator β_G:

Definition 4.37. The *standard pseudometric* on $\|G\|$ is $d_{|G|} \circ (\beta_G \times \beta_G)$ – see Definition 2.11, we denote it by $d_{\|G\|}$.

Remark 4.38. A map f from a topological space X to $\|G\|$ (equipped with the pseudo-metric topology) is continuous if, and only if, so is the composite $\beta_G \circ f$. Suppose that $\beta_G \circ f$ is continuous at $x_0 \in X$, and let $\varepsilon > 0$. We have a neighborhood V of x_0 such that $\beta_G(f(V)) \subseteq B(\beta_G(f(x_0)), \varepsilon)$. Given $x \in V$ we have $d_{\|G\|}(f(x_0), f(x)) = d_{|G|}(\beta_G(f(x_0)), \beta_G(f(x)))$ (Definition 4.37) therefore $f(V) \subseteq B(f(x_0), \varepsilon)$, and f is continuous. The converse is immediate. In particular the pseudometric $d_{\|G\|}$ is *intrinsic* because so is $d_{|G|}$ (Lemma 4.35).

Remark 4.39. The topology induced by the standard atlas $\mathcal{A}_G$ (Definition 3.14), which we call the *atlas topology*, is strictly finer than the topology induced by the standard pseudometric on $\|G\|$, which we call the *metric topology*, since the distance between two traversals at the same point is null.

For the rest of §4 we fix $\alpha \in [1, \infty]$ and a tuple of graphs $(G_1, \ldots, G_n)$.

Definition 4.40. The *standard α-distance* on $(G_1, \ldots, G_n)$ both refers to the α-product $d^{(\alpha)}$ of the standard metrics $d_{|G_i|}$ (Definition 4.16), and to the α-product $\tilde{d}^{(\alpha)}$ of the standard pseudometrics $d_{\|G_i\|}$ (Definition 4.37); see Definition 4.10 for α-products.

Remark 4.41. It readily derives from Definition 4.40 that the product of desingularizators $\beta_{G_1} \times \cdots \times \beta_{G_n}$ preserves pseudometrics.

Following Definition 4.5 we have:

Definition 4.42. The *α-length* of a path γ on $|G_1| \times \cdots \times |G_n|$ (resp. $\|G_1\| \times \cdots \times \|G_n\|$) which we denote by $L^\alpha(\gamma)$, is the least upper bound of the sums $\sum_{i=1}^{k} d^{(\alpha)}\big(\gamma(t_{i-1}), \gamma(t_i)\big)$ for $t_0 < \cdots < t_k$ subdivisions of $dom\, \gamma$.

Remark 4.43. Every path γ on $\|G_1\| \times \cdots \times \|G_n\|$ is continuous with respect to the pseudometric from Definition 4.37, which is equivalent to have $\beta_{G_i} \circ \gamma_i$ continuous on the metric graph $|G_i|$ for every $i \in \{1, \ldots, n\}$ (Remark 4.38); in that case we have $L^\alpha(\gamma) = L^\alpha((\beta_{G_1} \times \cdots \times \beta_{G_n}) \circ \gamma)$.

Lemma 4.44. The pseudometric $d^{(\alpha)}$ is intrinsic.

Proof. By Proposition 4.11, Lemma 4.35, and Remark 4.38.

The next result states that each standard chart induces a 'local isometry', the codomain being understood as a subspace of $\mathbb{R}^n$ equipped with the α-norm:

Proposition 4.45. For every standard chart ϕ of $G_1 \times \cdots \times G_n$ (Definition 3.12) and all $p, q \in dom\, \phi$, if $|\phi(p) - \phi(q)|_\alpha < \frac{R}{2}$, then $\tilde{d}^{(\alpha)}(p, q) = |\phi(p) - \phi(q)|_\alpha$.

Proof. By Definition 4.37 we have $d_{\|G\|}(p, q) = d_{|G|}((\beta_G \circ \phi^{-1})(\phi p), (\beta_G \circ \phi^{-1})(\phi q))$. Applying Remark 4.14, $(\beta_G \circ \phi^{-1})(\phi p)$ is equal to $\chi_a(\phi p)$ if $\phi = \phi_a$ for some arrow a; to $\chi_a(\ell(a) - \phi p)$ or $\chi_b(\phi p)$ depending on the sign of ϕp if $\phi = \phi_{ab}$ for some traversal (a, b). The same holds for q. We end up with 5 cases to examine, in each of which Lemma 4.29 applies and gives $d_{\|G\|}(p, q) = |\phi(p) - \phi(q)|$. In higher dimension, we follow Definition 4.40 and apply the one-dimension case to each component: $\tilde{d}^{(\alpha)}(p, q) = |\ldots, d_{G_i}(\beta_{G_i}(p_i), \beta_{G_i}(q_i)), \ldots|_\alpha = |\ldots, \phi_i(p_i) - \phi_i(q_i), \ldots|_\alpha = |\phi(p) - \phi(q)|_\alpha$.

4.3 Length of piecewise smooth paths

For every $t \in]a,b[$ the derivative of a smooth path $\gamma : [a,b] \to \mathcal{A}$ at t is a linear map $D\gamma_t : T_t]a,b[\to T_{\gamma(t)}\mathcal{A}$. Hence $D\gamma_t$ can be identified with the tangent vector $D\gamma_t(1)$, which we denote by $\gamma'(t)$. We obtain a continuous map $\gamma' : [a,b] \to T\mathcal{A}$ setting $\gamma'(a) = c'(a)$ and $\gamma'(b) = c'(b)$ with c being any curve extending γ.

Assuming that every tangent space of the atlas $\mathcal{A}$ is equipped with a norm, one defines the length of a smooth path γ on $\mathcal{A}$ as the sum of the norms of the speed vectors along γ: this is the infinitesimal version of Definition 4.5. *Riemannian metrics* (do Carmo, 1992, 2.1) are such structures, nevertheless, they do not fit our needs (§7, *Finsler geometry*).

Definition 4.46. Let $(f_1, \ldots, f_n)$ be a parallelization of the atlas $\mathcal{A}$, and let $|\_|_\star$ be a *norm* on $\mathbb{R}^n$; the mapping

$$v \in T_p\mathcal{A} \quad \mapsto \quad |(a_1, \ldots, a_n)|_\star \in \mathbb{R}_+ \tag{11}$$

with $v = a_1 f_1(p) + \cdots + a_n f_n(p)$ is a norm on $T_p\mathcal{A}$ which we still denote by $|\_|_\star$. The *length* of a smooth path $\gamma : [x,y] \to \mathcal{A}$ is defined as

$$\mathcal{L}(\gamma) \quad = \quad \int_x^y |\gamma'(t)|_\star \, dt \ .$$

The length of a piecewise smooth path $\gamma_1 \cdots \gamma_k$ is $\mathcal{L}(\gamma_1) + \cdots + \mathcal{L}(\gamma_k)$. The *distance* between two points p and q of $\mathcal{A}$ is then defined as

$$d_{\mathcal{A}} \quad = \quad \inf \left\{ \mathcal{L}(\gamma) \;\middle|\; \gamma : \text{(piecewise) smooth path from } p \text{ to } q \right\} .$$

One can indifferently consider smooth or piecewise smooth paths in the above definition. Indeed, for every piecewise smooth path γ one has a finite set $F \subseteq dom\,\gamma$ such that $\gamma'(t)$ exist for all $t \notin F$. If $\theta : [a,b] \to [a,b]$ is a non-decreasing surjective smooth path whose derivatives $D^{(k)}\theta$ vanish at each point of F and at every order $k \geqslant 1$, then the composite $\gamma \circ \theta$ is a smooth path. By a mere change of variable we have $L(\gamma) = L(\gamma \circ \theta)$. The map θ is obtained, for example, as

$$t \quad \mapsto \quad a + M \int_a^t \left(1 - \sum_{x \in F} \beta_x \right)$$

with $\{\beta_x : \mathbb{R} \to [0,1] \mid x \in F\}$ a family of *bump functions* such that $\beta_x(x) = 1$, $D^{(k)}\beta_x(x) = 0$ at every order $k \geqslant 1$, and $\beta_x \beta_y = 0$ when $x \neq y$. The constant M has to be chosen in a way that $\theta(b) = b$, in other words

$$M \quad = \quad \frac{b-2a}{S} \qquad \text{with} \qquad S \quad = \quad \int_a^b \left(1 - \sum_{x \in F} \beta_x \right) \quad .$$

Definition 4.47. The *smoothed α-length* of a piecewise smooth path γ on the atlas $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$, which we denote by $\mathcal{L}^\alpha\gamma$, is given by Definition 4.46 with the standard parallelization $(f_1, \ldots, f_n)$ of $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ (Definition 3.29) and the α-norm $|\_|_\star$ on $\mathbb{R}^n$. We just write $\mathcal{L}$ when $\alpha = 1$. We denote by $d_{\mathcal{A}}^{(\alpha)}$ the metric induced by $\mathcal{L}^\alpha$.

Remark 4.48. We have the standard inequality $\mathcal{L}^\alpha(\gamma) \leqslant \mathcal{L}(\gamma)$ as an immediate consequence of the *Jensen inequality* (Brokate & Kersting, 2015, 5.7).

The distance mapping $d : \mathcal{A} \times \mathcal{A} \to \mathbb{R}_+ \cup \{\infty\}$ (Definition 4.46) is a pseudometric, one may indeed have $d(p,q) = 0$ and $p \neq q$:

Examples 4.49. Let $\mathcal{L}$ be the smoothed length on $\mathcal{A}_G$ with (a,b) and (c,d) be two traversals at v with a, b, c, and d pairwise distinct. There is a piecewise smooth path on $\mathcal{A}_G$ from (a,b) to (c,d)

of arbitrarily small length. It suffices to concatenate the four smooth paths of length ε represented by the arrows on the diagram

$$(a,b) \;\to\; (a,1-\varepsilon) \;\to\; (a,d) \;\to\; (d,\varepsilon) \;\to\; (c,d)$$

and illustrated here below:

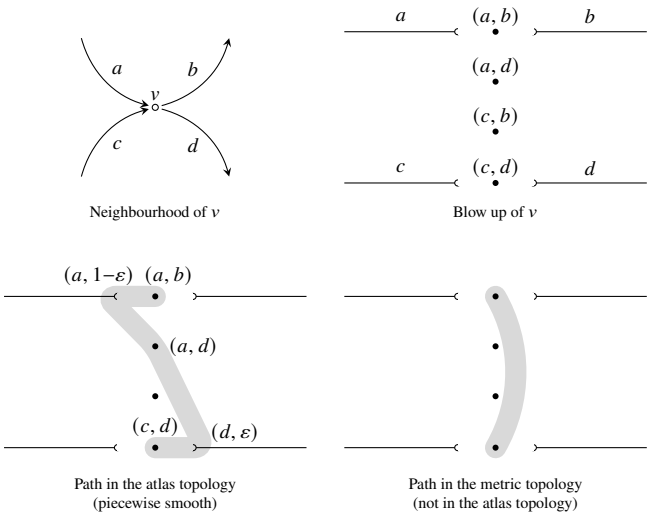

In particular, the greatest lower bound defining the distance between (a,b) and (c,d) – following Definition 4.46 – is not reached.

Lemma 4.50. For every smooth path γ on $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ we have $\mathcal{L}^\alpha \gamma = L^\alpha \gamma$.

Proof. First observe that γ is continuous with respect to the topology of $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ which is finer than the topology of the pseudometric $d^{(\alpha)}$ (Remark 4.39). Assume that $img\,(\gamma) \subseteq dom\,(\phi)$ for some standard chart ϕ, by Remark 3.32 we have

$$\int_x^y |\gamma'(t)|_\alpha \, dt \quad = \quad \int_x^y |(\phi \circ \gamma)'(t)|_\alpha \, dt \quad . \tag{12}$$

For a sufficiently fine subdivision $t_0 < \cdots < t_k$ of $dom\,(\gamma)$ we have (Proposition 4.45)

$$\sum_{i=1}^{k} d^{(\alpha)}(\gamma(t_i), \gamma(t_{i-1})) \quad = \quad \sum_{i=1}^{k} |\phi \circ \gamma(t_i) - \phi \circ \gamma(t_{i-1})|_\alpha \quad . \tag{13}$$

By a standard result about smooth paths on $\mathbb{R}^n$, the sum $\int_x^y |(\phi \circ \gamma)'(t)|_\alpha dt$ (right member of (12)) is the least upper bound of the sums $\sum_{i=1}^k |\phi \circ \gamma(t_i) - \phi \circ \gamma(t_{i-1})|_\alpha$ (right member of (13)) for $t_0 < \cdots < t_k$ running through the set of subdivisions of $dom\,(\gamma)$. The abovementioned result is given in (Papadopoulos, 2013, Proposition 1.3.1) for $\alpha = 2$, though the proof is obviously valid for any $\alpha \in [1, \infty]$.

We no longer assume that $dom\,(\gamma)$ is contained in the domain of a standard chart. Nevertheless, since $dom\,(\gamma)$ is compact and the domains of the standard charts form an open covering of $\|G_1\| \times \cdots \times \|G_n\|$, we have a subdivision $t_0 < \cdots < t_k$ of $dom\,(\gamma)$ such that for all $i \in \{1, \ldots, k\}$ the restriction of γ to $[t_i, t_{i-1}]$ is contained in the domain of some standard chart.

So far we have introduced two metrics on the set $\|G\|$, both of equal interest. They are actually equal:

Lemma 4.51. For every piecewise affine map γ on $|G_1| \times \cdots \times |G_n|$ and every $\varepsilon > 0$, there exists a piecewise smooth path δ on $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ such that $\mathcal{L}^\alpha(\delta) \leqslant L^\alpha(\gamma) + \varepsilon$. If γ is directed, then we can suppose that $\mathcal{L}^\alpha(\delta) = L^\alpha(\gamma)$.

Proof. Fix $t_0 \in dom\,(\gamma)$ and choose a compact interval I, neighborhood of t_0, such that $\gamma(I \setminus \{t_0\})$ does not contain any vertex.

If $\gamma(t_0)$ is not a vertex, then for every $t \in I$ both $|G|$ and $\|G\|$ contains $\gamma(t)$, which it is the only element of $\|G\|$ satisfying $\beta_G(\gamma(t)) = \gamma(t)$. Hence the mapping $\delta : t \in I \mapsto \gamma(t) \in \|G\|$ is smooth (because it is affine) and satisfies $\beta_G \circ \delta = \gamma|_I$.

Now suppose that $\gamma(t_0)$ is a vertex. It occurs at the junction between two consecutive steps s and s' of the piecewise affine path γ. Without loss of generality, we assume that all the steps appearing in γ are nondegenerate. Hence the supports a and a' of s and s' are uniquely defined, and each of the steps s and s' is either directed or antidirected (Definition 4.16).

If both s and s' are directed (resp. antidirected), then $\operatorname{tgt} a = \operatorname{src} a'$ (resp. $\operatorname{src} a = \operatorname{tgt} a'$). The mapping $\delta : I \to \|G\|$ sending $t \neq t_0$ to the unique lifting of $\gamma(t)$, and t_0 to the traversal (a, a') (resp. (a', a)) is piecewise smooth. As before we have $\beta_G \circ \delta = \gamma|_I$.

If s is directed while s' is antidirected (the 'dual' case is dealt with the same way), then we have $\operatorname{tgt} a = \operatorname{tgt} s = \operatorname{src} s' = \operatorname{tgt} a'$. If $a = a'$ then we can shorten both s and s' so that they no longer visit $\operatorname{tgt} a$, this shortens the path γ without changing its source nor its target. If $a \neq a'$ then the vertex $\operatorname{tgt} a$ has at least two in ingoing arrows, and by the assumption made in Definition 2.1, it also has an outgoing arrow b. On $I_- = I \cap]-\infty, t_0]$ we define δ_1 the path sending t to $\gamma(t)$ for $t \neq t_0$; and t_0 to (a, b). We have $\beta_G \circ \delta_1 = \gamma|_{I_-}$. On $I_+ = I \cap [t_0, \infty[$ we define the path δ_3 sending t to $\gamma(t)$ for $t \neq t_0$; and t_0 to (a', b). We have $\beta_G \circ \delta_3 = \gamma|_{I_+}$. The paths δ_1 and δ_3 are smooth and we have $(\beta_G \circ \delta_1) \cdot (\beta_G \circ \delta_3) = \gamma|_I$.

As in Example 4.49, we have an *arbitrarily short* piecewise smooth path δ_2 joining the traversals (b, a) and (b, a') (apart from these traversals, δ_2 covers $\{b\} \times]0, \varepsilon'/2[$ with ε' being the length of δ_2). A path arising in that context is called a *patch*.

We have thus three situations: i) $\gamma(t_0)$ is not a vertex, ii) $\gamma(t_0)$ is a vertex and both a and a' goes in the same direction, and iii) $\gamma(t_0)$ is a vertex and both a and a' goes in opposite directions. *Note that if γ is directed, then the case iii) does not occur.* In the cases i) and ii) we define $\delta_1 = \delta|_{I_-}$ and $\delta_3 = \delta|_{I_+}$, so the concatenation $\delta_1\delta_3 = \delta$ makes sense. In the case iii) we need a patch δ_2 to form the piecewise smooth concatenation $\delta_1\delta_2(\delta_3 \circ (t \in J \mapsto t - \varepsilon' \in I_+))$ with $J = I_+ + \varepsilon'$. Hence we have piecewise smooth paths δ_1, δ_2, and δ_3 such that

(1) $(\beta_G \circ \delta_1) = \gamma|_{I_-}$ and $(\beta_G \circ \delta_3) = \gamma|_{I_+}$,
(2) in the cases i) and ii), $dom\,(\delta_2) = \{t_0\}$ and $\delta = \delta_1\delta_3$ is piecewise smooth, and
(3) in the case iii), $\mathcal{L}^\alpha(\delta_2)$ is arbitrarily small and $\delta = \delta_1\delta_2(\delta_3 \circ \theta))$ is piecewise smooth with $J = I_+ + \mathcal{L}^\alpha(\delta_2)$ and $\theta : t \in J \mapsto t - \mathcal{L}^\alpha(\delta_2) \in I_+$.

In higher dimension, we have $\gamma = \gamma_1 \times \cdots \times \gamma_n$. We can choose I so that for every $k \in \{1, \ldots, n\}$ the set $\gamma_k(I \setminus \{t_0\})$ does not contain any vertex. By applying the above reasoning for every dimension $k \in \{1, \ldots, n\}$ we obtain the piecewise smooth paths $\delta_1^{(k)}$, $\delta_2^{(k)}$, $\delta_3^{(k)}$ satisfying 1), 2), and 3) with respect to γ_k. We form the products

$$\delta_1 = \delta_1^{(1)} \times \cdots \times \delta_1^{(n)}, \quad \delta_2 = \delta_2^{(1)} \times \cdots \times \delta_2^{(n)}, \quad \text{and} \quad \delta_3 = \delta_3^{(1)} \times \cdots \times \delta_3^{(n)}.$$

We have $(\beta_{G_1} \times \cdots \times \beta_{G_n}) \circ \delta_1 = \gamma|_{I_-}$ and $(\beta_{G_1} \times \cdots \times \beta_{G_n}) \circ \delta_3 = \gamma|_{I_+}$ because the equalities hold componentwise. We have $L^\alpha((\beta_{G_1} \times \cdots \times \beta_{G_n}) \circ \delta_j) = L^\alpha(\delta_j) = \mathcal{L}^\alpha(\delta_j)$ for $j \in \{1, 2, 3\}$ (Remark 4.43 and Lemma 4.50). Therefore $L^\alpha(\gamma|_I) = \mathcal{L}^\alpha(\delta_1) + \mathcal{L}^\alpha(\delta_3)$. If γ is directed then so are all the paths γ_k, consequently the case iii) does not occur, and we have $L^\alpha(\gamma|_I) = \mathcal{L}^\alpha(\delta_1\delta_3)$. Otherwise the case iii) may occur so we cannot ignore the patch δ_2. Yet, in each coordinate, $\mathcal{L}^\alpha(\delta_2^{(k)})$ can be made arbitrarily small. We choose a domain of definition common to all the patches $\delta_2^{(k)}$ (some of them are

'useless' loops if we are not in the case iii) for γ_k at t_0). In particular we have the shift $\theta = \theta_1 \times \cdots \times \theta_n$, and $\mathcal{L}^\alpha(\delta_1\delta_2(\delta_3 \circ \theta)) = L^\alpha(\gamma|_I) + \mathcal{L}^\alpha(\delta_2)$ (keeping in mind that $\mathcal{L}^\alpha(\delta_3 \circ \theta) = \mathcal{L}^\alpha(\delta_3)$, see Remark 4.6). We have $\mathcal{L}^\alpha(\delta_2) \leqslant \mathcal{L}(\delta_2)$ (Remark 4.48) so making each component $\delta_2^{(k)}$ arbitrarily small guarantees that so is $\mathcal{L}^\alpha(\delta_2)$.

Theorem. *Let p and q be the images under $\beta_{G_1} \times \cdots \times \beta_{G_n}$ of p' and q', we have*

$$d^{(\alpha)}(p,q) \quad = \quad \inf\left\{ \mathcal{L}^\alpha(\delta) \;\middle|\; \delta \text{ piecewise smooth path from } p' \text{ to } q' \right\}.$$

Proof. Since $d^{(\alpha)}$ is intrinsic (Lemma 4.44) we have

$$d^{(\alpha)}(p,q) \quad = \quad \inf\left\{ L^\alpha(\gamma) \;\middle|\; \gamma \text{ path from } p \text{ to } q \right\}.$$

We have $L^\alpha((\beta_{G_1} \times \cdots \times \beta_{G_n}) \circ \delta) = L^\alpha(\delta) = \mathcal{L}^\alpha(\delta)$ for every piecewise smooth path δ from p' to q' (Remark 4.43 and Lemma 4.50). Therefore $d^{(\alpha)}(p,q)$ is less than the infimum of the values $\mathcal{L}^\alpha(\delta)$; the converse inequality holds by Lemma 4.51.

5. Local orders

Local orders are similar to 'finite causal orientations' (Segal, 1976, p.23) which would be called 'local preorders' in this article. The importance of local antisymmetry was coined in Lawson (1989) – see also §5.4. The usage of local orders in concurrency theory was initiated in Fajstrup et. al. (2006).

Local orders are based on partially ordered spaces (pospaces) as atlases on charts, see §5.1 and §5.2. The desingularizator β_G induces a morphism from $\mathcal{X}_G$ to $\tilde{\mathcal{X}}_G$, the standard local orders on $|G|$ and $\|G\|$ (Definition 5.6). We lift directed paths along β_G (Theorem 5.3), and prove they admit piecewise affine approximations on tile compatible subspaces of $|G_1| \times \cdots \times |G_1|$ (Theorem 6.2), resp. $\|G_1\| \times \cdots \times \|G_1\|$ (Corollary 6.10). The desingularizator β_G is characterized by a universal property based on local orders (Theorem 5.3). The standard local order $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$ is related to the standard parallelization on $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ (Theorem 5.4).

5.1 Pospaces

A partial order $\leqslant$ on the underlying set of a topological space X is said to be *closed* when so is the subset $\{(a,b) \in X \times X \mid a \leqslant b\}$ of the product space $X \times X$. A *partially ordered space*, or just *pospace*, is a topological space with a *closed* partial order on it. (The closedness condition follows (Hilgert et. al, 1989, p.368), (Gierz et. al., 2003, VI-1.1), and (Goubault-Larrecq, 2013, 9.1.11).) The underlying set of a pospace is referred to as its *support* (which is often denotated as the pospace). The real line $\mathbb{R}$ with its standard topology and order is the prototypical example of a pospace. A *sub-pospace* is a subset of the support with the inherited topology and order.

A subset C of a poset $(X, \leqslant)$ is said to be *order convex* when it contains any $x \in X$ such that $a \leqslant x \leqslant b$ for some $a, b \in C$. A pospace $(X, \leqslant)$ is said to be *locally order convex* when its underlying topology has a basis of order convex open subset, see (Nachbin, 1965, p.26), (Hilgert et. al, 1989, p.371) or (Gierz et. al., 2003, VI-1.5).

A *morphism of pospaces* is an order preserving continuous map. An isomorphism of pospaces is called a *dihomeomorphism*. Pospaces and their morphisms form the category **P**. The product of X and Y in **P** is the product of their topological spaces together with the product order. Let X' be a pospace whose support is contained in the support of X. Then X' is said to be a *subpospace* of X when X' is both a *subspace* and a *subposet* of X, i.e.

- the open subsets of X' are the traces on X' of the open subsets of X, and
- the partial order of X' is the restriction to X' of the partial order of X.

Every subset of the support of X induces a subpospace of X. If this subset is open in X, then the subpospace is said to be *open*. A morphism of pospaces is said to be an *embedding* when it induces a dihomeomorphism on its image.

Definition 5.1. The left action of the n^{th} symmetric group $\mathfrak{S}_n$ on n-tuples is defined by

$$\sigma \cdot (t_1, \ldots, t_n) \quad = \quad (t_{\sigma^{-1}1}, \ldots, t_{\sigma^{-1}n}) \, .$$

By extension, for every n-tules of sets $A_1, \ldots, A_n$, we define the set $\sigma \cdot (A_1 \times \cdots \times A_n)$ as the product $A_{\sigma^{-1}1} \times \ldots \times A_{\sigma^{-1}n}$ and the map $\pi_\sigma : A_1 \times \cdots \times A_n \to \sigma \cdot (A_1 \times \cdots \times A_n)$ by $\pi_\sigma(t) = \sigma \cdot t$. In particular, for every n-tules of mappings $f_1, \ldots, f_n$ we have

$$\pi_\sigma \circ (f_1 \times \cdots \times f_n) \quad = \quad (f_{\sigma^{-1}1} \times \cdots \times f_{\sigma^{-1}n}) \circ \pi_\sigma \tag{14}$$

taking care that π_σ on the right hand side of (14) permutes the domains of the mappings f_k while π_σ on the left hand side permutes their codomains.

The next result is an immediate consequence of the second point of (Schröder, 2003, 10.4.10), which is about (possibly infinite) products of posets.

Lemma 5.2. The group of poset automorphisms of $\mathbb{R}^n$ (with the standard product order) is the product $\mathrm{Aut}(\mathbb{R})^n \times \mathfrak{S}_n$ in the sense that any poset automorphism of $\mathbb{R}^n$ can be written, in a unique way, as $(x_1 \times \cdots \times x_n) \circ \pi_\sigma$ with $x_k \in \mathrm{Aut}(\mathbb{R})$ for $k \in \{1, \ldots, n\}$.

We denote by $\wedge$ and $\vee$ the binary greatest lower bound and least upper bound operators in $\mathbb{R}^n$.

Lemma 5.3. A pospace embedding $\theta : \mathbb{R}^n \to \mathbb{R}^n$ preserves $\vee$ (resp. $\wedge$) if, and only if, its image is a set product of interval of $\mathbb{R}$. Moreover these intervals are open.

Proof. We denote by $\leqslant$ the product order in $\mathbb{R}^n$ without specifying the dimension. Assume that $img\,(\theta)$ is a product of intervals. The operator $\vee$ is computed pointwise in $\mathbb{R}^n$, therefore $img\,(\gamma)$ is stable under $\vee$. Since the inequality $\theta(t \vee t') \geqslant \theta(t) \vee \theta(t')$ is always satisfied and $\theta(t \vee t')$ belongs to $img\,(\gamma)$, we actually have $\theta(t \vee t') = \theta(t) \vee \theta(t')$. The same holds for the $\wedge$ operator.

Assume that θ preserves the $\vee$ operator. Let $\theta_k = \mathrm{proj}_k \circ \theta$ and $I_k = img\,(\theta_k)$ for every $k \in \{1, \ldots, n\}$. Let $x \in I_1 \times \cdots \times I_n$ and $A = \theta^{-1}(\downarrow x)$; we have $\theta(t) \leqslant x$ if, and only if $\theta_k(t) \leqslant x_k$ for every $k \in \{1, \ldots, n\}$. For every $k \in \{1, \ldots, n\}$ we have $t^{(k)} \in \mathbb{R}^n$ such that $\theta_k(t^{(k)}) = x_k$, so

$$\left\{ \begin{array}{ccccc} \theta_k(t^{(1)} \wedge \cdots \wedge t^{(n)}) & \leqslant & \theta_k(t^{(1)}) \wedge \cdots \wedge \theta_k(t^{(n)}) & \leqslant & x_k \\ \theta_k(t^{(1)} \vee \cdots \vee t^{(n)}) & \geqslant & \theta_k(t^{(1)}) \vee \cdots \vee \theta_k(t^{(n)}) & \geqslant & x_k \end{array} \right.$$

It follows that $t^{(1)} \vee \cdots \vee t^{(n)}$ is an upper bound of A which is nonempty because it contains $t^{(1)} \wedge \cdots \wedge t^{(n)}$. Therefore $\sup(A)$ exists and we have $\theta(\sup A) = \tilde{\sup}(\theta(A))$ the least upper bound being taken in the image of θ (this holds because θ is a pospace embedding). Since $\tilde{\sup}(\theta(A))$ is an upper bound of $\theta(A)$ the least upper bound $\sup(\theta(A))$ taken in $\mathbb{R}^n$ exists and $\sup(\theta(A)) \leqslant \tilde{\sup}(\theta(A))$.

Since A is closed in $\mathbb{R}^n$ (because θ is continuous and $\downarrow x$ is closed in $\mathbb{R}^n$) and has an upper bound, each of its elements is below some maximal element of A. If a and b are two such elements, then we have $\theta(a) \vee \theta(b) \leqslant x$, and therefore $\theta(a \vee b) \leqslant x$ because θ preserves $\vee$, hence $a = b$, and $\sup(A) \in A$.

By invariance of domain (Hatcher, 2002, Theorem 2B.3, p.172) the image of θ is open in $\mathbb{R}^n$, therefore we have a set product $J_1 \times \cdots \times J_n$ of open intervals of $\mathbb{R}$, that contains $\tilde{\sup}(\theta(A))$ and that is included in $img\,(\theta)$. If the inequality $\tilde{\sup}(\theta(A)) \leqslant x$ was strict, then we would have $y \in img\,(\theta)$ such that $\tilde{\sup}(\theta(A)) < y < x$, and therefore $\theta^{-1}(y) \in A$ with $\sup(A) < \theta^{-1}(y)$: a contradiction.

5.2 Local orders

A *witness of compatibility* of the pospaces X and X' nearby $p \in X \cap X'$ is a subset W of $X \cap X'$ containing p, on which both X and X' induce the same open subpospace. We say that X and X' are *compatible around* the point p, which we denote by $X \overset{p}{\sim} X'$, when $p \notin X \cap X'$ or such a witness exists. We say that X and X' are *compatible* when they are so around every point, in that case write $X \sim X'$.

A *local order* $\mathcal{X}$ is a collection of pairwise compatible pospaces. The *canonical* local order induced by a pospace X is the one element collection $\{X\}$. The *support* of $\mathcal{X}$ is the union of the supports of its elements, we denote it by $|\mathcal{X}|$.

A *morphism of local orders* from $\mathcal{X}$ to $\mathcal{Y}$ is a mapping $f : |\mathcal{X}| \to |\mathcal{Y}|$ such that for every $p \in |\mathcal{X}|$, every $Y \in \mathcal{Y}$ and every open subset V of Y containing $f(p)$, there exists $X \in \mathcal{X}$ and an open subset U of X containing p such that $f(U) \subseteq V$, and the restriction $f|_U : U \to V$ is order preserving (i.e. for all $a, b \in U$, $a \sqsubseteq_X b \;\Rightarrow\; fa \sqsubseteq_Y fb$). Local orders and their morphisms form the category $\mathbf{L}$.

The cartesian product of $\mathcal{X}$ and $\mathcal{Y}$ in $\mathbf{L}$ is the local order

$$\mathcal{X} \times \mathcal{Y} \quad = \quad \left\{ X \times Y \;\middle|\; X \in \mathcal{X} ;\, Y \in \mathcal{Y} \right\}.$$

The *topology* (on the support) of the local order $\mathcal{X}$ is generated by the subsets U that are open in some pospace of the collection $\mathcal{X}$.

Two local orders $\mathcal{X}$ and $\mathcal{Y}$ are said to be *equivalent*, which we denote by $\mathcal{X} \sim \mathcal{Y}$, when they have the same underlying set S and the identity map id_S induces an isomorphism between them. Actually, we have $\mathcal{X} \sim \mathcal{Y}$ when $\mathcal{X} \cup \mathcal{Y}$ is still a local order: in the end, it amounts to say that any $X \in \mathcal{X}$ is equivalent to any $Y \in \mathcal{Y}$.

A *directed path* on $\mathcal{X}$ is a local order morphism from $\{[a, b]\}$ to $\mathcal{X}$, with $a \leqslant b$.

The standard local orders. The sets $|G|$ and $\|G\|$ respectively come with the standard local orders $\mathcal{X}_G$ and $\tilde{\mathcal{X}}_G$ described in this section.

Definition 5.4. For every vertex v we denote by U_v the canonical star centered at v (Definition 2.4), i.e.

$$U_v \quad = \quad \bigcup_{\mathrm{tgt}\, a = v} \{a\} \times J_a \quad \cup \quad \{v\} \quad \cup \quad \bigcup_{\mathrm{src}\, b = v} \{b\} \times I_b \tag{15}$$

with the intervals $J_a =]0, \frac{\ell(a)}{2}[$ and $I_b =]\frac{\ell(b)}{2}, \ell(b)[$ (as in §3.2 (3)). The set U_v is provided with the *greatest* topology and the *least* partial order $\sqsubseteq_v$ making the maps

$$\gamma_a : t \in -I_a \cup \{0\} \quad \mapsto \quad \begin{cases} (a, t + \ell(a)) & \text{if } t \neq 0 \\ v & \text{if } t = 0 \end{cases} \in U_v$$

$$\delta_b : t \in \{0\} \cup I_b \quad \mapsto \quad \begin{cases} (b, t) & \text{if } t \neq 0 \\ v & \text{if } t = 0 \end{cases} \in U_v$$

continuous and order-preserving for all arrows a such that $\mathrm{tgt}(a) = v$, and all arrows b such that $\mathrm{src}(b) = v$. In particular, the maps γ_a and δ_b are pospace embeddings. Note that for every traversal (a, b) at v the union of the mappings γ_a and δ_b is the pospace embedding

$$\xi_{ab} : t \in -I_a \cup \{0\} \cup I_b \quad \mapsto \quad \begin{cases} (a, t + \ell(a)) & \text{if } t < 0 \\ v & \text{if } t = 0 \\ (b, t) & \text{if } t > 0 \end{cases} \in U_v \quad . \tag{16}$$

The set U_v comes with the map $\tau_v : U_v \to \mathbb{R}$ defined by $\tau_v(v) = 0$, and

$$\tau_v(x,t) \quad = \quad \begin{cases} t - \ell(x) & \text{if } t \in J_x \\ t & \text{if } t \in I_x \end{cases} .$$

Given two points p, q of U_v, we have $p \sqsubseteq_v q$ if, and only if

$$\tau_v p \leqslant 0 \leqslant \tau_v q \qquad \text{or} \qquad \pi_G p = \pi_G q \quad \text{and} \quad \tau_v p \leqslant \tau_v q \,, \tag{17}$$

thus defining a closed the partial order on U_v (closedness of this partial order is proven from an easy – but tedious – case disjunction readily obtained from the negation of the formula (17). The *standard pospace* at v is $(U_v, \sqsubseteq_v)$.

Proposition 5.5. The pospace $(U_v, \sqsubseteq_v)$ is locally order convex.

Proof. By Remark 2.6 and the fact that the stars centered at a vertex and the sets of the form $\{a\} \times]t-\varepsilon, t+\varepsilon[$ are order convex.

Definition 5.6. The *standard local order* on $|G|$, which we denote by $\mathcal{X}_G$, is the collection of pospaces $(U_v, \sqsubseteq_v)$ and $\{a\} \times]0, \ell(a)[$ (equipped with the obvious pospace structure) with v and a ranging in the sets of vertices and arrows of G respectively. The *standard local order* on the set $|G_1| \times \cdots \times |G_n|$ is the product $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$.

Every vertex and every middle point $(a, \ell(a)/2)$ belongs to a single element of the collection $\mathcal{X}_G$. Every point (a,t) with $t \neq \ell(a)/2$ only belongs to $\{a\} \times]0, \ell(a)[$ and U_v with $v = \operatorname{src} \alpha$ or $v = \operatorname{tgt} \alpha$ depending on whether $t < \ell(a)/2$ or $t > \ell(a)/2$. Hence the (nontrivial) witnesses of compatibility are the open sets $\{\alpha\} \times]0, \ell(a)/2[$ and $\{\alpha\} \times]\ell(a)/2, \ell(a)[$ with the obvious topology and order.

Remark 5.7. Assuming that $|G_1| \times \cdots \times |G_n|$ is equipped with the d_∞ metrics, every open ball of radius $r \leqslant \frac{R}{4}$ is contained in some pospace $X \in \mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$.

The standard charts ϕ_a and ϕ_{ab} (Definition 3.12) with a and (a,b) arrows and traversals, are bijections towards open subintervals of $\mathbb{R}$. Their domains, namely $dom\ \phi_a$ and $dom\ \phi_{ab}$, are equipped with the pospace structures so that ϕ_a and ϕ_{ab} become dihomeomorphisms. All these pospaces are locally order convex because so are the open intervals of $\mathbb{R}$.

Definition 5.8. The collection of pospaces $dom\ \phi_a$ and $dom\ \phi_{ab}$ (Definition 3.12) which we denote by $\tilde{\mathcal{X}}_G$, forms the *standard local order* on $\|G\|$ observing that the witnesses of compatibility are of the form $\{a\} \times]0, \ell(a)/2[$ and $\{a\} \times]\ell(a)/2, \ell(a)[$ with a arrow of G. The *standard local order* on the set $\|G_1\| \times \cdots \times \|G_n\|$ is $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$.

Remark 5.9. Given a traversal (a,b) at v, $\varepsilon \in]0, \ell(a)[$, and $\varepsilon' \in]0, \ell(b)[$, the desingularizator β_G (Definition 2.11) induces the canonical pospace isomorphism

$$\begin{array}{c} \{a\} \times]\ell(a) - \varepsilon, \ell(a)[\cup \{(a,b)\} \cup \{b\} \times]0, \varepsilon'[\\ \downarrow \cong \\ \{a\} \times]\ell(a) - \varepsilon, \ell(a)[\cup \{v\} \cup \{b\} \times]0, \varepsilon'[\end{array}$$

hence a pospace embedding from $dom\ \phi_{ab}$ to U_v – see Definition 3.12. In particular β_G induces a morphism of local orders from $\tilde{\mathcal{X}}_G$ to $\mathcal{X}_G$.

Remark 5.10. The domains of the isomorphisms from Remark 5.9 together with sets of the form $\{a\} \times I$ with I open interval of $]0, \ell(a)[$ form a base of the topology of the local order $\tilde{\mathcal{X}}_G$.

Remark 5.11. Denote by **S** the category of topological spaces, and by **H** the full subcategory of Hausdorff spaces. Denote by $H : \mathbf{S} \to \mathbf{H}$ and $h : id_{\mathbf{S}} \to IH$ the left adjoint to inclusion functor $I : \mathbf{H} \hookrightarrow \mathbf{S}$ and its unit. Putting X for the underlying topological space of $\mathcal{X}_G$ we have $H(\beta_G) = id_X$.

5.3 The lifting properties of desingularizators

Lifting directed paths. Every *directed* path on the image of β_G can be lifted along it (Theorem 5.3). Nevertheless, such liftings may not be unique when the starting or the finishing point is a vertex. We begin with a simple yet useful observation:

Remark 5.12. The vertex v is the only element of U_v (Definition 5.4 (15)) such that for every $x \in U_v$ we have $x \sqsubseteq_v v$ or $v \sqsubseteq_v x$. Let C be a *chain* (i.e. a totally ordered set, which is said to be *unbounded* when it has no greatest nor least element) of U_v containing at least two elements. If v is an upper bound of C, then we have a unique arrow a of G such that $C \setminus \{v\} \subseteq \{a\} \times]\ell(a)/2, \ell(a)[$. Moreover $\mathrm{tgt}\, a = v$. Similarly, if v is a lower bound of C, then we have a unique arrow b of G such that $C \setminus \{v\} \subseteq \{b\} \times]0, \ell(a)/2[$. Moreover $\mathrm{src}\, b = v$. If v is neither a lower nor an upper bound of C, then v is both an upper bound of $\downarrow v \cap C$, and a lower bound of $\uparrow v \cap C$. Therefore we have a unique traversal (a, b) of G at v such that C is entirely contained in the image of ξ_{ab} – see Definition 5.4(16). In addition, if C is unbounded and connected (as a subset of U_v which carries a topology), then it is isomorphic to the pospace $\mathbb{R}$.

Given $t \in dom\, \gamma$ for a path γ, we say that γ is constant *before* (resp. *after*) t when $\gamma\{s \in dom\, \gamma \mid s \leqslant t\}$ (resp. $\gamma\{s \in dom\, \gamma \mid s \geqslant t\}$) is reduced to $\{\gamma(t)\}$.

Lemma 5.13. Let γ be a directed path on $\mathcal{X}_G$ and $t \in dom\, \gamma$ such that $\gamma(t)$ is a vertex v. Suppose that $[t', t'']$ is the connected component of $\gamma^{-1}\{v\}$ containing t. If γ is not constant before t (i.e. t' is not the least element of $dom\, \gamma$) then we have a *unique* arrow a such that for every $\varepsilon \in]0, \ell(a)[$ there exists $s \in dom\, \gamma$ such that $s < t'$ and $\gamma[s, t] \setminus \{v\} \subseteq \{a\} \times]\ell(a) - \varepsilon, \ell(a)[$. Moreover $\mathrm{tgt}\, a = v$. Similarly, if γ is not constant after t (i.e. t'' is not the greatest element of $dom\, \gamma$) then we have a *unique* arrow b such that for every $\varepsilon \in]0, \ell(a)[$ there exists $s \in dom\, \gamma$ such that $t'' < s$ and $\gamma[t, s] \setminus \{v\} \subseteq \{b\} \times]0, \varepsilon[$. Moreover $\mathrm{src}\, b = v$.

Proof. We have an open interval I' containing t' such that γ induces a pospace morphism from I' to U_v because γ is a local order morphism. Let $I = \{s \in I' | s \leqslant t\}$ so v is an upper bound of the chain $\gamma(I)$ of U_v. We readily deduce from the definition of t' that $\gamma(I)$ is not reduced to $\{v\}$, so Remark 5.12 applies and we have a unique arrow a such that $\gamma(I) \setminus \{v\} \subseteq \{a\} \times]0, \ell(a)[$, moreover $\mathrm{tgt}\, a = v$. For $\varepsilon \in]0, \ell(a)/2[$ the star $S(v, \varepsilon)$ (Definition 2.4) is an open subset of U_v (Remark 2.6). Since γ is continuous, we have a neighborhood J of t' (with $J \subseteq I'$) such that $\gamma(J)$ is contained in B, therefore in $\{a\} \times]\ell(a) - \varepsilon, \ell(a)[$. Given $s \in I \cap J$ such that $s < t'$ we have $\gamma[s, t'[\subseteq \{a\} \times]\ell(a) - \varepsilon, \ell(a)[$. The case where γ is not constant after t is deduced from the previous one by reversing the local order, or by a direct similar proof.

Definition 5.14. The arrows a and b given by Lemma 5.13 are referred to as the *last* arrow visited *before* t, and the *first* arrow visited *after* t (by γ). If $dom\, \gamma = [t_0, t_1]$ the first arrow visited by γ is the first one it visits after t_0. Similarly, the last arrow visited by γ is the last one it visits before t_1.

Definition 5.15. If $\gamma(t)$ is a vertex, a traversal (a, b) at $\gamma(t)$ is said to be *compatible* with γ at t when the following are satisfieed: i) if the last arrow visited before t exists, then it is a; and ii) if the first arrow visited after t exists, then it is b.

Remark 5.16. If γ is a constant path standing on a vertex v, then any traversal at v is compatible with γ at every $t \in dom\,\gamma$. If γ is not constant and starts on a vertex at t_0, then any traversal (a,b) such that b is the first arrow visited by γ is compatible at t_0. Dually, if γ finishes on a vertex at t_1, then any traversal (a,b) such that a is the last arrow visited by γ is compatible at t_1.

Lemma 5.17. Suppose that $\gamma(t)$ is a vertex v with γ directed path on $\mathcal{X}_G$. For any traversal (a,b) compatible with γ at t, and any $\varepsilon \in]0,R[$, there exists an interval I open in $dom\,\gamma$, containing t, and such that γ induces a pospace morphism

$$\tilde{\gamma} \quad : \quad I \quad \to \quad \{a\} \times]\ell(a) - \varepsilon, \ell(a)[\ \cup \ \{v\} \ \cup \ \{b\} \times]0, \varepsilon[\quad .$$

Moreover, if γ is not constant before (resp. after) t, then $\tilde{\gamma}$ visits $\{a\} \times]\ell(a) - \varepsilon, \ell(a)[$ (resp. $\{b\} \times]0, \varepsilon[$).

Proof. Let t' be the least element of the connected component of $\gamma^{-1}\{v\}$ containing t. Note that $\gamma[t', t] = \{v\}$. If γ is constant before t then t' is the least element of $dom\,\gamma$. In that case we can suppose that t' is also the least element of I. Otherwise a is the last arrow visited by γ before t (Definition 5.14). By Lemma 5.13 we have $s < t'$ such that $\gamma(]s, t']) \setminus \{v\} \subseteq \{a\} \times]\ell(a) - \varepsilon, \ell(a)[$. By definition of a local order morphism (§5.2) the map γ induces a pospace morphism from $]s,t]$ to $\{a\} \times]\ell(a) - \varepsilon, \ell(a)[\cup \{v\}$ for s sufficiently close to t. From the definition of t' we deduce that $\gamma(]s, t'])$ is not reduced to $\{v\}$, therefore it meets $\{a\} \times]\ell(a) - \varepsilon, \ell(a)[$. Whether γ is constant before t or not, we found an interval I_0 whose greatest element is t, that is not reduced to a single element unless t is the least element of $dom\,\gamma$, and such that γ induces a pospace morphism from I_0 to $\{a\} \times]\ell(a) - \varepsilon, \ell(a)[\cup \{v\}$. By similar arguments we obtain an interval I_1 whose least element is t, that is not reduced to a single element unless t is the greatest element of $dom\,\gamma$, and such that γ induces a pospace morphism from I_1 to $\{v\} \cup \{b\} \times]\ell(b) - \varepsilon, \ell(b)[$. The interval $I = I_0 \cup I_1$ matches the requirements.

Definition 5.18. A *lifting* (along $\beta_{G_1} \times \cdots \times \beta_{G_n}$) of a directed path γ on $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$ is a directed path δ on $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$ such that $(\beta_{G_1} \times \cdots \times \beta_{G_n}) \circ \delta = \gamma$.

Proposition 5.19. Let γ be a directed path on $\mathcal{X}_G$. Given a set map δ from $dom\,\gamma$ to $\|G\|$ (Definition 2.11) the following are equivalent:

(1) The map δ is a lifting of γ (Definition 5.18).
(2) The map δ is continuous and satisfies $\beta_G \circ \delta = \gamma$.
(3) For every $t \in dom\,\gamma$, if $\gamma(t)$ is not a vertex, then $\delta(t) = \gamma(t)$, otherwise $\delta(t)$ is a traversal compatible with γ at t (Definition 5.15).

Proof. The first point implies the second one by Definition 5.15. Assume the second point is satisfied. Since $\beta_G \circ \delta = \gamma$, we have $\gamma(t) = \delta(t)$ each time $\gamma(t)$ is not a vertex (Definition 2.11). For the same reason, if $\gamma(t)$ is a vertex v, then $\delta(t)$ is a traversal (a,b) at v. Assume we are in the latter case. Since δ is continuous, and according to the topology of the local order $\tilde{\mathcal{X}}_G$ (Remark 5.10), we have an open interval I containing t and satisfying

$$\delta(I) \quad \subseteq \quad (\{a\} \times]\ell(a) - \varepsilon, \ell(a)[) \ \cup \ \{(a,b)\} \ \cup \ (\{b\} \times]0, \varepsilon[) \quad .$$

If γ is not constant before t, then let a' be the last arrow visited by γ before t (Lemma 5.13). We have $s < t$ such that $\gamma[s,t] \setminus \{v\}$ is nonempty and included in $\{a'\} \times]\ell(a') - \varepsilon, \ell(a')[$. Since $\gamma(s) = \delta(s)$ when $\gamma(s)$ is not a vertex, we have $a = a'$. Similarly, if γ is not constant after t, then b is the first arrow visited by γ after t. Hence (a,b) is compatible with γ at t.

Assume the third point is satisfied. One readily deduces that $\beta_G \circ \delta = \gamma$ from Definitions 2.11 and 5.15. It remains to check that δ is a local order morphism. The mappings γ and δ agree on

$\delta^{-1}(\|G\| \setminus \{\text{vertex of } G\})$ because β_G leaves unchanged any point of $\|G\|$ that is not a vertex – see Definition 2.11. Let t be an element of $dom\,\gamma$.

If $\gamma(t)$ is not a vertex, then $\gamma(t)$ belongs to $\{a\} \times]0, \ell(a)[$ for a unique arrow a. Since γ is a local order morphism, we have an open interval I containing t such that the restriction of γ to I induces a pospace morphism from I to $\{a\} \times]0, \ell(a)[$. Hence the restriction of δ to I coincides with the restriction of γ to I, therefore it is a pospace morphism (the standard pospace $\{a\} \times]0, \ell(a)[$ belongs to both $\mathcal{X}_G$ and $\tilde{\mathcal{X}}_G$).

If $\gamma(t)$ is a vertex v, then $\delta(t)$ is a traversal (a, b) compatible with γ at t (we have supposed that the third point is satisfied). Let $\varepsilon > 0$ be small enough so that we can apply Lemma 5.17 (with the traversal (a, b)) to obtain the pospace morphism

$$\tilde{\gamma} \quad : \quad I \quad \to \quad \{a\} \times]\ell(a) - \varepsilon, \ell(a)[\ \cup \ \{v\} \ \cup \ \{b\} \times]0, \varepsilon[$$

induced by γ with I open interval of $dom\,\gamma$ containing t. Let $\tilde{\beta}_G$ be the pospace isomorphism induced by β_G (Remark 5.9 with the traversal (a, b) at v). Then $\tilde{\beta}_G^{-1} \circ \tilde{\gamma}$ is a local order morphism defined on I satifying $\beta_G \circ \tilde{\beta}_G^{-1} \circ \tilde{\gamma} = \tilde{\gamma}$. The only vertex that $\tilde{\gamma}$ reaches is v and we know that δ and $\tilde{\beta}_G^{-1} \circ \tilde{\gamma}$ agree on $\{s \in I \mid \gamma(s) \neq v\}$. If $\tilde{\gamma}(s) = v$ then $\gamma(s) = \tilde{\gamma}(s) = v$ and $\tilde{\beta}_G^{-1}(v) = (a, b) = \delta(s)$. Hence δ and $\tilde{\beta}_G^{-1} \circ \tilde{\gamma}$ agree on I.

Theorem. *Any directed path on the image of $\beta_{G_1} \times \cdots \times \beta_{G_n}$ (seen as a sub-local order of $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$) admits a lifting.*

Proof. Let γ be a directed path on $img\,(\beta_G)$ and $t \in dom\,\gamma$. If $\gamma(t)$ is not a vertex put $\delta(t) = \gamma(t)$. Otherwise, there is at least one traversal (a, b) at $\gamma(t)$ (Definition 2.11). If γ is not constant before t then we can suppose that a is the last arrow visited by γ before t (Definition 5.14). Similarly, if γ is not constant after t then we can suppose that b is the first arrow visited by γ after t. Put $\delta(t) = (a, b)$. The map δ is a lifting of γ by the third point of Proposition 5.19. The lifting of $\gamma_1 \times \cdots \times \gamma_n$ is obtained by applying the above reasoning to γ_i for every $i \in \{1, \ldots, n\}$.

Universal lifting property. Let $x : \mathbb{R} \to X \in \mathcal{X}_G$ be a *continuous* map whose image does not contain any vertex of G (the underlying set of $\mathcal{X}_G$ is $|G|$, see Definition 2.3). The map x thus takes its values in the space $|G| \setminus G^{(0)}$, whose connected components are the segments $\{a\} \times]0, \ell(a)[$. Consequently, there is a unique arrow a such that $img\,(x) \subseteq dom\,(\phi_a)$, see Definition 3.12. Also the map β_G induces the identity on $|G| \setminus G^{(0)}$, therefore we have a unique continuous map $\tilde{x} : \mathbb{R} \to \{a\} \times]0, \ell(a)[$ such that $x(t) = \beta_G(\tilde{x}(t))$ for all $t \in \mathbb{R}$. If x is a pospace embedding, then so is $\tilde{x}$, and its image is an open subset of $\{a\} \times]0, \ell(a)[$. Extending the codomain of $\tilde{x}$ to the whole domain of β_G, we obtain an embedding y such that $x = \beta_G \circ y$ and $img\,(y)$ is open. Note that $\xi_{ab} \circ \phi_{ab}$ is the inclusion $dom\,(\phi_{ab}) \hookrightarrow \beta_G^{-1} U_v$ – see Definitions 3.12 and 5.4 (16). The next result deals with the case where the image of x contains a vertex:

Proposition 5.20. For every pospace embedding $x : \mathbb{R} \to U_v \in \mathcal{X}_G$ with $v \in img\,(x)$, there is a unique traversal (a, b) at v such that $img\,(x) \subseteq img\,(\xi_{ab})$. The map

$$y \ : \ t \in \mathbb{R} \quad \mapsto \quad \begin{cases} (a, b) & \text{if } x(t) = v \\ x(t) & \text{if } x(t) \neq v \end{cases} \quad \in \quad dom\,(\phi_{ab}) \tag{18}$$

is the only *set map* y from $\mathbb{R}$ to $dom\,(\phi_{ab})$ such that

$$x \quad = \quad \beta_G \circ \xi_{ab} \circ \phi_{ab} \circ y\,. \tag{19}$$

Moreover, the mapping $\tilde{x} = \xi_{ab} \circ \phi_{ab} \circ y$ is the only *continuous map* from $\mathbb{R}$ to $\beta_G^{-1} U_v$ satisfying $x = \beta_G \circ \tilde{x}$. Furthermore $\tilde{x}$ is a pospace embedding whose image is open.

Proof. The map x induces a pospace isomorphism on its image therefore $img\,(x)$ is an unbounded connected chain of U_v. So we have a unique traversal (a,b) at v such that $img\,(x) \subseteq img\,(\xi_{ab})$ (Remark 5.12). The map defined at line (18) readily satisfies equality (19). It is the only one because $\beta_G p = p$ for all $p \in dom\,(\phi_{ab}) \setminus \{(a,b)\}$ and $\beta_G(a,b) = v$. To check the uniqueness of $\tilde{x}$ it suffices to observe that $\beta_G^{-1}U_v$ is the union of the sets $dom\,(\phi_{a'b'})$ for all traversal (a',b') at v, and consequently that we have $\beta_G p = p$ for all $p \in \beta_G^{-1}U_v$ that is not a traversal. However, there may be more than one traversal in $\beta_G^{-1}U_v$ and the image of any of them under β_G is v (it is actually the case when v is singular – see Definition 2.10). Denote by t the only element of $\mathbb{R}$ such that $x(t) = v$. We already know that the image of $\mathbb{R} \setminus \{t\}$ under x is contained in $dom\,(\phi_{ab})$, which is a neighborhood of (a,b) not containing any other traversal. Since $\tilde{x}$ is continuous and sends t to a traversal, this latter must be (a,b). Denoting by $\tilde{\xi}_{ab}$ and x' the corestrictions of the pospace embeddings ξ_{ab} and x to their images, we have

$$\tilde{x} \quad = \quad \left(cod\,(\phi_{ab}) \hookrightarrow \beta_G^{-1}U_v\right) \circ \phi_{ab}^{-1} \circ \tilde{\xi}_{ab}^{-1} \circ x'$$

with $cod\,(\phi_{ab}) \hookrightarrow \beta_G^{-1}U_v$ pospace embedding, while ϕ_{ab}^{-1}, $\tilde{\xi}_{ab}^{-1}$ and x' are dihomeomorphisms; therefore $\tilde{x}$ is a pospace embedding. Hence $img\,(\tilde{x})$ is an unbounded connected chain of $dom\,(\phi_{ab})$, so it is open in the open subset $dom\,(\phi_{ab})$ of $\beta_G^{-1}U_v$.

Lemma 5.21. Let $\theta : \mathbb{R}^n \to X_1 \times \cdots \times X_n$ be a pospace embedding with $X_k \in \mathcal{X}_{G_k}$ for every $k \in \{1,\ldots,n\}$. If θ preserves $\vee$ (resp. $\wedge$), then $img\,(\theta) = C_1 \times \cdots \times C_n$ with C_k unbounded connected chain of X_k.

Proof. Assume that $X_k = U_v$ (Definition 5.4 (15)) for some k and that v is a vertex of G_k. Any element $x \neq v$ is either smaller or greater than v (Remark 5.12). Moreover, given two elements of U_v that are strictly smaller (resp. greater) than v the following are equivalent:

- they are comparable (i.e. one of them is smaller than the other),
- they have a lower bound (resp. an upper bound), and
- they belong to $\{a\} \times \left]\frac{1}{2}\ell(a), \ell(a)\right[$ for some arrow a of G such that $\operatorname{tgt} a = v$ (resp. $\{b\} \times \left]0, \frac{1}{2}\ell(b)\right[$ for some arrow b of G such that $\operatorname{src} b = v$).

Given t and t' in $\mathbb{R}^n$ we have $\theta(t \wedge t') \leqslant \theta(t), \theta(t') \leqslant \theta(t \vee t')$ because θ is order preserving. It follows that either

- $\theta_k(t)$ and $\theta_k(t')$ belong to $\{a\} \times]0, \ell(a)[$ for some arrow a of G, or
- we have $\theta_k(t) \leqslant v \leqslant \theta_k(t')$ or $\theta_k(t') \leqslant v \leqslant \theta_k(t)$.

In both cases we deduce that $img\,(\theta) \subseteq J_1 \times \cdots \times J_n$ with J_k either equal to $\chi_{a_k}\,]0, \ell(a_k)[$ or $\chi_{a_k}\,]0, \ell(a_k)[\cup \{(a_k, b_k)\} \cup \chi_{b_k}\,]0, \ell(b_k)[$ for a unique traversal (a_k, b_k).

We have a dihomeomorphism $\psi : J_1 \times \cdots \times J_n \cong \mathbb{R}^n$ obtained as products of dihomeomorphisms $\psi_k : J_k \cong \mathbb{R}$. The corestriction $\tilde{\theta}$ of θ to $img\,(\theta)$ is a dihomeomorphism. In particular $img\,(\theta)$ is a lattice. The $\vee$-preservation hypothesis about θ means that for all $t, t' \in dom\,(\theta)$, $\theta(t) \vee \theta(t')$ exists in $X_1 \times \cdots \times X_n$ and is equal to $\theta(t \vee t')$. It is clear that the least upper bound of $\theta(t)$ and $\theta(t')$ exists in $X_1 \times \cdots \times X_n$ if, and only if, it exists in $J_1 \times \cdots \times J_n$, and that both coincide. In other words $img\,(\theta)$ is stable under $\vee$, the least upper bound operator of $J_1 \times \cdots \times J_n$, which means that the embedding $img\,(\theta) \hookrightarrow J_1 \times \cdots \times J_n$ preserves $\vee$. Hence the composite

$$\mathbb{R}^n \xrightarrow{\tilde{\theta}} img\,(\theta) \hookrightarrow J_1 \times \cdots \times J_n \xrightarrow{\psi} \mathbb{R}^n$$

is a pospace embedding preserving $\vee$, so its image is a product of open intervals $I_1 \times \cdots \times I_n$ (Lemma 5.3) and for every $k \in \{1, \ldots, n\}$, the chain C_k is $\psi_k^{-1}(I_k)$.

Proposition 5.22. If θ is a pospace embedding that preserves $\vee$ (resp. $\wedge$) as in Lemma 5.21, then there is a unique permutation $\sigma \in \mathfrak{S}_n$ and a unique tuple of pospace embedding $(x_1, \ldots, x_n)$ with $x_k : \mathbb{R} \to X_k$ for $k \in \{1, \ldots, n\}$ such that $\theta = (x_1 \times \cdots \times x_n) \circ \pi_\sigma$ – see Definition 5.1.

Proof. From Lemma 5.21 we know that the corestriction of θ to its image is a pospace isomorphism of the form $\tilde{\theta} : I_1 \times \cdots \times I_n \to C_1 \times \cdots \times C_n$ with I_k open interval of $\mathbb{R}$ and C_k connected unbounded chain of X_k. Then for every $k \in \{1, \ldots, n\}$ we have pospace isomorphisms $\phi_k : \mathbb{R} \cong I_k$ and $\psi_k : C_k \cong \mathbb{R}$. Denoting by ϕ and ψ the pospace isomorphisms $\phi_1 \times \cdots \times \phi_n$ and $\psi_1 \times \cdots \times \psi_n$ the composite $\psi \circ \tilde{\theta} \circ \phi$ is a pospace automorphism of $\mathbb{R}^n$. From Lemma 5.2 we deduce that

$$\psi \circ \tilde{\theta} \circ \phi \quad = \quad (\chi_1 \times \cdots \times \chi_n) \circ \pi_\sigma$$

for a (unique) family $\chi_1, \ldots, \chi_n$ of automorphisms of the poset $\mathbb{R}$ and a (unique) $\sigma \in \mathfrak{S}_n$. It follows from Definition 5.1 (14) that

$$\tilde{\theta} \quad = \quad \Big(\big(\underbrace{\psi_1^{-1} \circ \chi_1 \circ (\phi_{\sigma^{-1}1})^{-1}}_{y_1} \big) \times \cdots \times \big(\underbrace{\psi_n^{-1} \circ \chi_n \circ (\phi_{\sigma^{-1}n})^{-1}}_{y_n} \big) \Big) \circ \pi_\sigma \ .$$

The n-tuple of mappings $(y_1, \ldots, y_n)$ is uniquely defined because π_σ is (in particular) a bijection. We conclude setting $x_k = (C_k \hookrightarrow X_k) \circ y_k$ for $k \in \{1, \ldots, n\}$. Assume that we have permutations $\sigma, \tau \in \mathfrak{S}_n$ and tuples $(x_1, \ldots, x_n)$, $(y_1, \ldots, y_n)$ as in the statement that satisfy $(x_1 \times \cdots \times x_n) \circ \pi_\sigma = \theta = (y_1 \times \cdots \times y_n) \circ \pi_\tau$. In particular we have $(y_1 \times \cdots \times y_n) = (x_1 \times \cdots \times x_n) \circ \pi_{\sigma \circ \tau^{-1}}$ and therefore $\sigma = \tau$.

Each desingularizator β_G is characterized by a universal property (Theorem 5.3) whose statement requires some terminology we now introduce:

Definition 5.23. A *local $\vee$-embedding* of dimension $n \in \mathbb{N}$ is a local order morphism $f : \mathcal{E} \to \mathcal{Y}$ such that for every point p of $\mathcal{E}$, and every $Y \in \mathcal{Y}$ containing $f(p)$, there exists $U \in \mathcal{E}$ and E open subset of U containing p, dihomeomorphic with $\mathbb{R}^n$, such that $f(E) \subseteq Y$ and the restriction $f_E : E \to Y$ is a pospace embedding that preserves $\vee$.

Theorem. *For every $\vee$-embedding $f : \mathcal{E} \to \mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$ of dimension n, there is a unique* continuous *map $g : \mathcal{E} \to \tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$ such that $f = \bar{\beta} \circ g$ with $\bar{\beta} = \beta_{G_1} \times \cdots \times \beta_{G_n}$; moreover g is a local $\vee$-embedding of dimension n.*

Proof. Let p be a point of $\mathcal{E}$ and $X_1 \times \cdots \times X_n$ containing $f(p)$ with $X_k \in \mathcal{X}_{G_k}$ for every $k \in \{1, \ldots, n\}$. According to the description of the local order $\mathcal{X}_{G_k}$ (Definition 5.6) we can suppose that for every $k \in \{1, \ldots, n\}$ we have

$$X_k \quad = \quad \begin{cases} \{a\} \times]0, \ell(a)[& \text{if } \mathrm{proj}_k(f_E(p)) = (a, s) \text{ with} \\ & a \text{ arrow of } G_k \text{ and } s \in]0, \ell(a)[\ . \\ U_v & \text{if } \mathrm{proj}_k(f_E(p)) = v \text{ vertex of } G_k \ . \end{cases} \tag{20}$$

Since f is a local $\vee$-embedding of dimension n, we have $U \in \mathcal{E}$ together with an open subpospace E of U containing p, dihomeomorphic to $\mathbb{R}^n$, such that the restriction $f_E : E \to X_1 \times \cdots \times X_n$ induces a pospace embedding that preserves $\vee$. Given a dihomeomorphism $\varphi_E : E \cong \mathbb{R}^n$ the embedding

$f_E \circ \varphi_E^{-1}$ preserves $\vee$ so Proposition 5.22 applies to it and we obtain

$$f_E \quad = \quad (x_1 \times \cdots \times x_n) \circ \pi_\sigma \circ \varphi_E$$

with $\sigma \in \mathfrak{S}_n$ and $x_k : \mathbb{R} \to X_k$ pospace embedding for every $k \in \{1, \ldots, n\}$. One can actually suppose that σ is the identity even if it means replacing φ_E by $\pi_\sigma \circ \varphi_E$, in which case we have

$$f_E \quad = \quad (x_1 \times \cdots \times x_n) \circ \varphi_E \,. \tag{21}$$

Let k be an element of $\{1, \ldots, n\}$. If $X_k = \{a\} \times]0, \ell(a)[$ – see (20), then let $\tilde{x}_k$ be the identity map on $\{a\} \times]0, \ell(a)[$. If $X_k = U_v$ then apply Proposition 5.20 to the pospace embedding x_k to obtain the unique continuous map $\tilde{x}_k : \mathbb{R} \to \beta_{G_k}^{-1} U_v$ such that

$$x_k \quad = \quad \beta_{G_k} \circ \tilde{x}_k \,. \tag{22}$$

We denote the map $(\tilde{x}_1 \times \cdots \times \tilde{x}_n) \circ \varphi_E$ by g_p; it induces a pospace embedding that preserves $\vee$ because each $\tilde{x}_k$ is an (open) pospace embedding (Proposition 5.20) defined on a totally ordered set (so $\tilde{x}_k$ preserves $\vee$). As a consequence of (21) and the equalities (22) for $k \in \{1, \ldots, n\}$, we have

$$f_E \quad = \quad \bar{\beta} \circ (\tilde{x}_1 \times \cdots \times \tilde{x}_n) \circ \varphi_E \quad = \quad \bar{\beta} \circ g_p \,. \tag{23}$$

Given $k \in \{1, \ldots, n\}$ we apply the k^{th} projection to the relations (21) and (23) to obtain

$$\mathrm{proj}_k \circ f_E \quad = \quad x_k \circ \mathrm{proj}_k \circ \varphi_E \quad = \quad \beta_{G_k} \circ \tilde{x}_k \circ \mathrm{proj}_k \circ \varphi_E \,. \tag{24}$$

Since $\mathrm{proj}_k \circ \varphi_E$ is onto, we have $img\,(x_k) = img\,(\mathrm{proj}_k \circ f_E)$. In the light of relation (24), for $e \in E$, we deduce from Proposition 5.20 (18) that $\tilde{x}_k(\mathrm{proj}_k(\varphi(e))) = \mathrm{proj}_k(f_E(e))$ if $\mathrm{proj}_k(f_E(e)) \neq v$, and $\tilde{x}_k(\mathrm{proj}_k(\varphi(e))) = (a, b)$ if $\mathrm{proj}_k(f_E(e)) = v$, with (a, b) being the unique traversal at v in the graph G_k such that $img\,(\mathrm{proj}_k \circ f_E) \subseteq img\,(\xi_{ab})$ – see Definition 5.4 (16) for ξ_{ab}. What really matters here is that $\tilde{x}_k(\mathrm{proj}_k(\varphi(e)))$ only depends on f. Its image is open because each map $\tilde{x}_k$ is open (Proposition 5.20) and φ_E is (in particular) onto. Thus we have a family of $\vee$-preserving open pospace embeddings

$$\left\{ g_p : E_p \to \beta_{G_1}^{-1} X_1^{(p)} \times \cdots \times \beta_{G_n}^{-1} X_n^{(p)} \;\middle|\; p \in \mathcal{E} \right\}$$

such that:

- E_p is an open neighborhood of p that is dihomeomorphic to $\mathbb{R}^n$,
- $X_k^{(p)} \in \mathcal{X}_{G_k}$ for every $k \in \{1, \ldots, n\}$, and
- for all $p, q \in \mathcal{E}$, the mappings g_p and g_q coincide on $E_p \cap E_q$.

The mapping g defined by $g(p) = g_p(p)$ for every $p \in \mathcal{E}$ fulfills the requirements of the statement. It remains to see that it is the only one. To this aim, assume that the continuous map $h : \mathcal{E} \to \tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$ satisfies $\bar{\beta} \circ h = \bar{\beta} \circ g$. Given $p \in \mathcal{E}$ we have $\bar{\beta} \circ (\tilde{x}_1 \times \cdots \times \tilde{x}_n) \circ \varphi_E = \bar{\beta} \circ h_{E_p}$ with h_{E_p} the restriction of h to E_p. We deduce that we have $\beta_{G_k} \circ \tilde{x}_k \circ \mathrm{proj}_k \circ \varphi_E = \beta_{G_k} \circ \mathrm{proj}_k \circ h_{E_p}$ for every $k \in \{1, \ldots, n\}$. Since h is continuous, so is $\mathrm{proj}_k \circ h_{E_p}$; we deduce from the uniqueness property in Proposition 5.20 that $\tilde{x}_k \circ \mathrm{proj}_k \circ \varphi_E = \mathrm{proj}_k \circ h_{E_p}$ for every $k \in \{1, \ldots, n\}$. In other words we have $g_p = h_{E_p}$. Since the latter equality holds for every $p \in \mathcal{E}$ we have $h = g$.

5.4 Cone fields and local orders

For any manifold M, there is an equivalence (in the categorical sense) between the local orders on (the support of) M whose elements are locally order convex, and the upper semicontinuous conal fields on M admitting sections at every point (Lawson, 1989, 2.7). Strictly speaking, the above statement involves a manifold instead of an atlas, and Lawson (1989) does not make the topological properties required on M explicit. However, this statement is local by nature, and its proof does

not seem to suppose anything about the underlying topology of M; assuming that it holds for the atlas $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$, the conal field associated to the standard parallelization and the standard local order on $\tilde{X}_{G_1} \times \cdots \times \tilde{X}_{G_n}$ are related by this equivalence (Proposition 5.26 and Theorem 5.4). Except for the latter results, this section is entirely based on (Lawson, 1989, §1 and §2).

A *wedge* is a topologically closed subset W of a finite dimensional real vector space such that $\mathbb{R}_+ W \subseteq W$ and $W + W \subseteq W$. A *cone* is a wedge C such that $C \cap -C = \{0\}$. A *cone field* on an atlas $\mathcal{A}$ is a map assigning a cone $C(p)$ of the tangent space $T_p\mathcal{A}$ to every point p of $\mathcal{A}$. A *conal atlas* is an atlas together with a cone field. If C and D are cone fields on the atlases $\mathcal{A}$ and $\mathcal{B}$ then the mapping $C \times D$ is a cone field on $\mathcal{A} \times \mathcal{B}$.

Examples 5.24. Any parallelization $(f_1, \ldots, f_n)$ canonically induces a cone field:

$$C(p) \quad = \quad \left\{ \lambda_1 f_1(p) + \cdots + \lambda_n f_n(p) \;\middle|\; \lambda_i \geqslant 0 \text{ for all } i \in \{1, \ldots, n\} \right\}.$$

Such cone fields are said to be *cartesian*. If C and D are the cartesian cones induced by the parallelizations $(f_1, \ldots, f_n)$ and $(g_1, \ldots, g_m)$, then $C \times D$ is the cartesian cone induced by the parallelization $(f_1, \ldots, f_n, g_1, \ldots, g_m)$.

Examples 5.25. The *standard cone field* C_G on $\mathcal{A}_G$ is the cartesian cone field induced by the vector field of Definition 3.29, that is to say:

$$C_G \quad : \quad p \in \|G\| \quad \mapsto \quad \left\{ (p,t) \;\middle|\; t \in \mathbb{R}_+ \right\} \quad \text{cone of } \{p\} \times \mathbb{R}\,.$$

Conal preorder. For every curve c extending a smooth path γ (Definition 3.11) and every $t \in dom\,(\gamma)$ we have the *left hand* and the *right hand* derivatives of γ at t, namely

$$\lim_{x \to t^-} D\gamma_x = Dc_t \qquad \text{and} \qquad \lim_{x \to t^+} (D\gamma)_x = Dc_t\,;$$

which we denote by $D\gamma_{t^-}$ and $D\gamma_{t^+}$. A piecewise smooth path admits derivatives on both sides at every t, although there is finitely many t's at which they may not coincide.

A piecewise smooth path γ is said to be *conal* on $(\mathcal{A}, C)$ if $D\gamma_{t^+} \in C(\gamma(t))$ for all $t \in dom\,(\gamma)$. Given the points p and q of an open subset U of $\mathcal{A}$, we write $p \preccurlyeq_U q$ when there exists a conal curve on U from p to q. The relation $\preccurlyeq_U$ thus defined is the *conal preorder* on U.

Infinitesimal preorder. A *partial curve* on an atlas $\mathcal{A}$ is a map $\gamma : D \to |\mathcal{A}|$ with $0 \in D \subseteq \mathbb{R}_+$ and 0 cluster point of $D \setminus \{0\}$. Given $X \subseteq |\mathcal{A}|$, we say that $v \in T_p\mathcal{A}$ is a *subtangent vector* of X at p when there exists a partial curve $\gamma : D \to |\mathcal{A}|$ with $\gamma(0) = p$ and $\gamma(D \setminus \{0\}) \subseteq X$, and a chart $\phi \in \mathcal{A}$, such that

$$[\![v]\!]^{\phi}_p \quad = \quad \lim_{0 \leftarrow t \in D} \tfrac{1}{t}\big(\phi(\gamma(t)) - \phi(p)\big)\,.$$

Given $\psi \in \mathcal{A}$ with $p \in dom\,\psi$, it suffices to write $\psi \circ \gamma = (\psi \circ \phi^{-1}) \circ (\phi \circ \gamma)$ and apply the chain rule (i.e. a form of the chain rule adapted to partial curves) to check that the above equality still holds with ψ instead of ϕ.

Assume that the local order $\mathcal{X}$ and the atlas $\mathcal{A}$ are based on the same topological space. The *infinitesimal preorder* of $\mathcal{X}$ is the map assigning to each point $p \in \mathcal{A}$ the smallest wedge W_p of $T_p\mathcal{A}$ containing all the subtangent vectors of the set $\{p \preccurlyeq_U u \mid u \in U\}$ at p for some pospace $(U, \leqslant) \in \mathcal{X}$ such that $p \in U$. Because of the coherence condition satisfied by a local order, the wedge W_p does not depend on $(U, \leqslant)$.

Lawson correspondence. We say that the cone C' *surrounds* the cone C when $C \setminus \{0\}$ is contained in the interior of C' (see (Hilgert et. al, 1989, IV.6.4, p.342) for a more conceptual definition of surrounding cones). Given a chart $\phi \in \mathcal{A}$ and $p, q \in dom\,\phi$, the composite $([\![\_]\!]^{\phi}_q)^{-1} \circ [\![\_]\!]^{\phi}_p : T_p\mathcal{A} \to$

$T_q\mathcal{A}$ is a tvs isomorphism (Example 3.20) which we denote by $[\![\_]\!]^{\phi}_{pq}$; its inverse is $[\![\_]\!]^{\phi}_{qp}$, and $[\![\_]\!]^{\phi}_{pp}$ is the identity map.

A cone field C on $\mathcal{A}$ is said to be *upper semicontinuous* at point p when for every chart $\phi \in \mathcal{A}$ around p, and every cone C' that surrounds $C(p)$, there exists a neighborhood U of p such that C' surrounds $[\![C(q)]\!]^{\phi}_{qp}$ for every $q \in U$. We say that C *admits sections at* p when for every $v \in C(p)$ there exists a vector field f on $\mathcal{A}$ such that $f(p) = v$.

An open cover $\mathcal{U}$ of the atlas $\mathcal{A}$ is said to be *admissible* for the cone field C on $\mathcal{A}$ when for every $U \in \mathcal{U}$ the conal preorder $\preccurlyeq_U$ is antisymmetric, the pospace $(U, \preccurlyeq_U)$ is locally order convex (§5.1), and the collection $\{(U, \preccurlyeq_U) \mid U \in \mathcal{U}\}$ is a local order. Given an atlas $\mathcal{A}$, Lawson's result (Lawson, 1989, 2.7) states that:

(1) For every upper semicontinuous cone field C on $\mathcal{A}$ admitting sections at every point, there exists an open cover $\mathcal{U}$ admissible for C. Moreover, for every open cover $\mathcal{U}$ admissible for C, the infinitesimal preorder of $\{(U, \preccurlyeq_U) \mid U \in \mathcal{U}\}$ is C.
(2) For every local order $\mathcal{X}$ made of locally convex pospaces, the infinitesimal preorder of $\mathcal{X}$ is an upper semicontinous cone field C on $\mathcal{A}$ admitting sections at every point. Moreover, the local order $\mathcal{X}$ is equivalent (in the sense of §5.2) to $\{(U, \preccurlyeq_U) \mid U \in \mathcal{U}\}$ for every open cover $\mathcal{U}$ admissible for C.

Proposition 5.26. The standard local order $\tilde{\mathcal{X}}_G$ on $\|G\|$ (Definition 5.8) and the standard cone field C_G on $\mathcal{A}_G$ (Example 5.25) are deduced from each other through Lawson's correspondence.

Proof. The fact that C_G is upper semicontinuous and admits sections at every point is immediate. Denote by $\sqsubseteq_{ab}$ the partial order on $dom\,(\phi_{ab})$ that comes from $\phi_{ab} \in \tilde{\mathcal{X}}_G$. Denote by $\preccurlyeq_{ab}$ the conal preorder on $dom\,(\phi_{ab})$ induced by C_G. Let $p, q \in dom\,\phi_{ab}$. The inequality $p \sqsubseteq_{ab} q$ amounts to $\phi_{ab}p \leqslant \phi_{ab}q$ in $\mathbb{R}$. The restriction of ϕ_{ab}^{-1} to the segment $[\phi_{ab}p, \phi_{ab}q]$ is a conal curve on $dom\,\phi_{ab}$, which implies that $p \preccurlyeq_{ab} q$. The other way round, if $p \preccurlyeq_{ab} q$, then we have a piecewise smooth conal curve γ on $dom\,\phi_{ab}$ from p to q. The fact that γ is conal means that for every $t \in dom\,(\gamma)$ we have $D\gamma_t(1) = rC_G(p)$ with $r > 0$, and therefore the map $\phi_{ab}^{-1} \circ \gamma$ is a piecewise derivative map between intervals of $\mathbb{R}$ whose derivative is non-negative; it is thus non-decreasing. The map ϕ_{ab} (Definition 3.12) is a dihomeomorphism (Definition 5.8 and its preamble), hence γ is a pospace morphism; we deduce that $p \sqsubseteq_{ab} q$. Given $\phi_a \in \tilde{\mathcal{X}}_G$ the same reasoning holds for the partial order $\sqsubseteq_a$ on $dom\,(\phi_a)$, and $\preccurlyeq_a$ the conal preorder on $dom\,(\phi_a)$ induced by C_G. So we have proven that $\sqsubseteq_{ab}$ (resp. $\sqsubseteq_a$) is the conal preorder on $dom\,(\phi_{ab})$ (resp. $dom\,(\phi_a)$) induced by C_G. It follows that the collection

$$\mathcal{U} \quad = \quad \big\{dom\,(\phi_a),\ dom\,(\phi_{ab}) \;\big|\; a \text{ and } (a,b) \text{ arrow and traversal of } G\big\}$$

is admissible for C_G, and that the local order $\{(U, \preccurlyeq_U) \mid U \in \mathcal{U}\}$ is $\tilde{\mathcal{X}}_G$. From the first point of Lawson's result, we deduce that C_G is the infinitesimal preorder of $\tilde{\mathcal{X}}_G$.

The following lemmas derive from the fact that given atlases $\mathcal{A}$ and $\mathcal{B}$, partial curves (resp. piecewise smooth paths) on $\mathcal{A} \times \mathcal{B}$ are of the form $\alpha \times \beta$ with α and β partial curves (resp. piecewise smooth paths) on $\mathcal{A}$ and $\mathcal{B}$.

Lemma 5.27. If C and D are the infinitesimal preorders of the local orders $\mathcal{X}$ and $\mathcal{Y}$ (whose elements are supposed to be locally convex), then $C \times D$ is the infinitesimal preorder of the local order $\mathcal{X} \times \mathcal{Y}$.

Lemma 5.28. If $\mathcal{X}$ and $\mathcal{Y}$ are local orders associated to the upper semicontinuous cone fields C and D, then the local order $\mathcal{X} \times \mathcal{Y}$ is associated to the cone field $C \times D$ (which is upper semicontinuous).

Let $G_1, \dots, G_n$ be a tuple of graphs. For each $k \in \{1, \dots, n\}$ we have the standard local order $\tilde{\mathcal{X}}_{G_k}$ (Definition 5.8) and the standard cone field C_{G_k} (Example 5.25) both admitting $\|G_k\|$ as underlying set. As an immediate consequence of Proposition 5.26 and Lemmas 5.27 and 5.28 we have:

Theorem. *The cone field $C_{G_1} \times \cdots \times C_{G_n}$ and the local order $\tilde{\mathcal{X}}_{G_1} \times \cdots \times \tilde{\mathcal{X}}_{G_n}$ are deduced from each other through Lawson's correspondence.*

6. Execution time

The smooth model $\|P\|$ of a program P (Definition 2.17) is equipped with a pseudometric inherited from that of $\|G_1\| \times \cdots \times \|G_n\|$; while the pseudometric of each $\|G_i\|$ is derived from $\ell_i : \{\text{arrows of } G_i\} \to [R, \infty[$ (Definitions 4.19, 4.24, and 4.37) that of $\|G_1\| \times \cdots \times \|G_n\|$ depends on $\alpha \in [1, \infty]$ (Definitions 4.40 and 4.47, and Theorem 4.3). In order that the length of any directed path on $\|P\|$ coincides with the execution time of its sequence of multi-instructions, we should take $\alpha = \infty$. We explain why in §6.1 and prove in §6.2 that directed paths admit piecewise affine approximations (Theorem 6.2 and Corollary 6.10).

6.1 The parallel execution time principle

We are concerned with concurrency theory so we wish to fix the parameter α according to the *parallel execution time principle*, which states that for every multi-instruction μ, if T_i is the time required by the i^{th} process to execute the instruction $\mu(i)$ (for $i \in dom\, \mu$), then the time T required by P to execute μ is

$$T \quad = \quad \max \left\{ T_i \;\middle|\; i \in dom\, \mu \right\}. \tag{25}$$

Let us examine the consequences of this principle on the discrete model first. Following (Haucourt, 2018, 3.4) a directed path on $G_1 \times \cdots \times G_n$ is a sequence of points $p_0, \dots, p_N$ satisfying for $i \in \{1, \dots, N\}$

$$\begin{cases} \left(\mathrm{pr}_k(p_i) \neq \mathrm{pr}_k(p_{i-1}) \;\Rightarrow\; \mathrm{src}(\mathrm{pr}_k(p_i)) = \mathrm{pr}_k(p_{i-1}) \text{ for all } k \in \{1, \dots, n\} \right), \text{ or} \\ \left(\mathrm{pr}_k(p_i) \neq \mathrm{pr}_k(p_{i-1}) \;\Rightarrow\; \mathrm{pr}_k(p_i) = \mathrm{tgt}(\mathrm{pr}_k(p_{i-1})) \text{ for all } k \in \{1, \dots, n\} \right) \end{cases}$$

with $\mathrm{pr}_k : G_1 \times \cdots \times G_n \to G_k$ denoting the k^{th} projection (this definition is motivated by topological arguments (Haucourt, 2018, Lemma 6.1)). Depending on the case, the indices i such that $\mathrm{pr}_k(p_i) \neq \mathrm{pr}_k(p_{i-1})$ are those of the processes that start or finish the execution of an instruction. The 'geometric length' of the path $p_0, \dots, p_N$ is N, which means that each step $p_{i-1} \to p_i$ requires 'one unit of time'. Indeed the arrows that have appeared (in the first case), or disappeared (in the second one), represent actions that have been executed *simultaneously*.

Every discrete path $p_0, \dots, p_N$ induces (in the sense of (Haucourt, 2018, Definition 6.2)) a piecewise affine path on the continuous model $|P|$ (Definition 2.9); this path is lifted to a piecewise affine path γ on $\|P\|$ according to Theorem 5.3. The infinitesimal step of γ at time t is the tangent vector $\gamma'(t)$ (it is required to belong to the positive cone of the tangent space $T_{\gamma(t)}(\mathcal{A}_1 \times \cdots \times \mathcal{A}_n)$ – Definition 3.29). In order to compare with the discrete case, we assume that the scheduler is only allowed to decide whether a given process of P is active or not, its speed being fixed at 1. In other words we assume that $\gamma'(t) \in \{0, 1\}^n$. Following these observations, the norm $|v|_\infty = \max\{|v_1|, \dots, |v_n|\}$ for $v \in \mathbb{R}^n$ (i.e. $\alpha = \infty$) is the one for which the length of any directed path corresponds to the parallel execution time of its sequence of multi-instructions.

Analogously, if we had been to measure the energy consumed along $p_0, \dots, p_N$, each step would have contribute with the number of appearing/disappearing arrows. The *parallel execution*

energy principle indeed states that given a multi-instruction μ, if E_i is the energy consumed by the i^{th} process to execute the instruction $\mu(i)$ (for $i \in dom\ \mu$) then the energy E consumed by P to execute μ is

$$E \quad = \sum_{i \in dom\, \mu} E_i \quad .$$

Compare to the parallel execution time principle (25). Suppose that the energy consumed by the i^{th} process at time t is the i^{th} component of $\gamma'(t)$. Accordingly, the energy consumed along γ is given by its length relative to the distance induced by the norm $|v|_1 = |v_1| + \cdots + |v_n|$ for $v \in \mathbb{R}^n$ (i.e. $\alpha = 1$). We observe that

$$\int_0^T |\gamma'(t)|_1 dt \quad = \quad \int_0^T \Big(\sum_{i=1}^n |\gamma_i'(t)| dt \Big) \quad = \quad \sum_{i=1}^n \Big(\int_0^T |\gamma_i'(t)| dt \Big) .$$

The energy consumed depends less on the scheduling than the execution time does.

6.2 Piecewise affine approximations

As long as we are only interested in minimizing the length of directed paths between two points of a tile compatible set (Definitions 2.8 and 2.11) we can restrict our attention to piecewise affine paths (Definition 4.19): this claim is formalized in Theorem 6.2 and Corollary 6.10.

The sets $|G_1| \times \cdots \times |G_n|$ and $\|G_1\| \times \cdots \times \|G_n\|$ are equipped with their standard local orders (Definitions 5.6 and 5.8) from which their topologies are inherited (§5.2). The pseudometrics d_i on $|G_i|$ and $\|G_i\|$ are given by Definitions 4.24 and 4.37. Following the previous execution time considerations, the chosen product metric is

$$d_\infty\big((p_1, \ldots, p_n), (q_1, \ldots, q_n)\big) \quad = \quad \max \big\{ d_i(p_i, q_i) \,\big|\, i \in \{1, \ldots, n\} \big\} .$$

Let $\xi : I \to X$ be a path with X subset of $|G_1| \times \cdots \times |G_n|$ or $\|G_1\| \times \cdots \times \|G_n\|$.

Definition 6.1. An *ε-approximation* of ξ on X is a path $\alpha : I \to X$ such that:

- α and ξ have the same starting and the same finishing points,
- α and ξ are ε-close, i.e. $d_\infty(\alpha(t), \xi(t)) < \varepsilon$ for every $t \in I$,
- α is shorter than ξ, and
- if ξ is directed, then so is α.

Of course any path is its own ε-approximation whatever ε and X are. From now on we assume that X is tile-compatible; our goal is to prove that any directed path ξ on X has a piecewise affine ε-approximation for arbitrarily small $\varepsilon > 0$. We first deal with the case where $X \subseteq |G_1| \times \cdots \times |G_n|$, which requires more technicalities.

Definition 6.2. The closure of a tile τ is denoted by $\bar{\tau}$. The *starred neighborhood* $\mathcal{N}_\tau$ of τ is the union of all the tiles τ' such that $\tau \subseteq \bar{\tau}'$.

Lemma 6.3. The closure of the tile over $p \in G_1 \times \cdots \times G_n$ (Definition 2.8) is the union of the tiles over the points q such that $q_i = p_i$ or $q_i \in \{\mathrm{src}(p_i), \mathrm{tgt}(p_i)\}$ for every $i \in \{1, \ldots, n\}$.

Proof. It suffices to remind that the tile τ is the product $\tau_1 \times \cdots \times \tau_n$ (Definition 2.8) so $\bar{\tau} = \bar{\tau}_1 \times \cdots \times \bar{\tau}_n$. The closure of $\{a\} \times]0, \ell(a)[$, for any arrow a, is the set $\{a\} \times]0, \ell(a)[\cup \{\mathrm{src}(a), \mathrm{tgt}(a)\}$ (Remark 2.6).

Lemma 6.4. Let $p \in |G_1| \times \cdots \times |G_n|$ and $(y_1, \ldots, y_n) \in G_1 \times \cdots \times G_n$ with $y_i = \pi_{G_i} p$ (Definition 2.3). If τ is the tile over $(x_1, \ldots, x_n)$ then we have:

- $p \in \mathcal{N}_\tau$ if, and only if, for every $i \in \{1, \ldots, n\}$ we have $y_i = x_i$ or $x_i \in \{\mathrm{src}(y_i), \mathrm{tgt}(y_i)\}$, i.e. $\mathcal{N}_\tau = I_1 \times \cdots \times I_n$ with $I_k = \{x_k\} \times]0, \ell(x_k)[$ if $x_k \in G^{(1)}$, and $I_k = S(x_k)$ the star centered at x_k (Definition 2.4) if $x_k \in G^{(0)}$.
- $p \in \bar{\mathcal{N}}_\tau$ if, and only if, for every $i \in \{1, \ldots, n\}$ we have $y_i = x_i$ or $x_i \in \{\mathrm{src}(y_i), \mathrm{tgt}(y_i)\}$ or $y_i \in \{\mathrm{src}(x_i), \mathrm{tgt}(x_i)\}$ or there exists an arrow a such that $\{\mathrm{src}(a), \mathrm{tgt}(a)\} = \{x_i, y_i\}$.

Proof. The first equivalence is an immediate consequence of Lemma 6.3. It follows that $\bar{\mathcal{N}}_\tau = \bar{I}_1 \times \cdots \times \bar{I}_n$, each closure $\bar{I}_k$ being given by Remark 2.6.

Corollary 6.5. The starred neighborhood of a tile is open (in the metric topology).

Proof. It is indeed a finite product (Lemma 6.4) of open sets (Remark 2.6).

Remark 6.6. The starred neighborhood of a tile τ is actually the least tile compatible open subset containing τ.

Uniqueness of geodesics and stability under geodesics are not preserved under ∞-product metric: for example, the length of the broken line γ (on the following picture) is equal to the length of the vertical segment from p to q. In particular γ is a geodesic joining two points of the ball B which gets out of B.

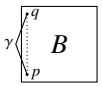

Let $p, q \in |G_1| \times \cdots \times |G_n|$ such that $d_\infty(p, q) < \frac{R}{2}$. For every $k \in \{1, \ldots, n\}$, the unique geodesic γ_k on $|G_k|$ from p_k to q_k is given by $p_k \to q_k$ or $p_k \to v_k \to q_k$ depending on whether we are in the first or the second case of Lemma 4.29.

Lemma 6.7. Denote by p the centre of a closed ball B of radius $r < \frac{R}{4}$ such that for every $k \in \{1, \ldots, n\}$, if the k^{th} projection p_k is not a vertex, then B_k contains no vertex. Then, for every $q \in B$ and every $k \in \{1, \ldots, n\}$, the shortest path from p_k to q_k is the dash $p_k \to q_k$. The higher dimensional dash $p \to q = (p_1 \to q_1) \times \cdots \times (p_n \to q_n)$ (with affine parametrization) is of length $d(p, q)$ and its image is contained in the union of the tiles containing p and q respectively. Moreover, if $p \leqslant_U q$ (resp. $q \leqslant_U p$) for some pospace U of the standard local order $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$ (Definition 5.6), then $p \to q$ (resp. $q \to p$) induces (as a map) a pospace morphism with values in U, and therefore a directed path on $\mathcal{X}_{G_1} \times \cdots \times \mathcal{X}_{G_n}$.

Proof. Let $q \in B$ and $k \in \{1, \ldots, n\}$. The (image of the) shortest arc α from p_k to q_k (Definition 4.31) is contained in B (the distance from any point along it to p is less that $d_{|G|_k}(p_k, q_k) < r$). The arc α is either of the form $p_k \to q_k$ or $p_k \to v_k \to q_k$ for a unique vertex v_k (Lemma 4.29); however, the second form is excluded by hypothesis on B (if p_k is not a vertex then B_k contains no vertex). We have $img\,(\gamma) \setminus \{p\} \subseteq \tau$ with τ the tile containing q (Corollary 4.36). The last point is an immediate consequence of Definitions 4.27 and 5.6.

Definition 6.8. An *ε-discretization* ($\varepsilon > 0$) of the path $\gamma : [x, y] \to |G_1| \times \cdots \times |G_n|$ is a sequence $x = t_0 < \cdots < t_k \leqslant y$ with open balls $B_0, \ldots, B_k$ such that i) each B_i is centered at $\gamma(t_i)$, with radius at most ε, and satisfies the requirements of Lemma 6.7, ii) for every $i \in \{1, \ldots, k\}$, $]t_{i-1}, t_i[\cap \gamma^{-1}(B_{i-1} \cap B_i) \neq \emptyset$, and iii) $img\,(\gamma) \subseteq B_0 \cup \cdots \cup B_k$.

Lemma 6.9. For every $\varepsilon > 0$, every path on $|G_1| \times \cdots \times |G_n|$ admits an ε-discretization.

Proof. Suppose that $dom\,(\gamma) = [x, y]$ and denote by t_∞ the least upper bound of the set of elements $t \in dom\,(\gamma)$ such that the restriction $\gamma|_{[x,t]}$ has an *ε-discretization*. Choose an open ball B_∞ centered at $\gamma(t_\infty)$ with radius less than ε and satisfying the requirements of Lemma 6.7. Let $t < t_\infty$ such that $\gamma[t, t_\infty] \subseteq B_\infty$. Let $t_0 < \cdots < t_k \leqslant t$ with open balls $B_0, \ldots, B_k$ be an *ε-discretization* of the restriction $\gamma|_{[x,t]}$. Even if it means taking t a bit bigger, we can suppose that $t_k < t < t_\infty$ with $\gamma(t) \in B_k \cap B_\infty$. Then $t_0 < \cdots < t_k < t_\infty \leqslant t_\infty$ with open balls $B_0, \ldots, B_k, B_\infty$ is an *ε-discretization* of the restriction $\gamma|_{[x,t_\infty]}$. If we had $t_\infty < y$ then we would have $t' \in]t_\infty, y[$ such that $\gamma[t_\infty, t'] \subseteq B_\infty$ so the precedent ε-discretization of $\gamma|_{[x,t_\infty]}$ would also be an ε-discretization of $\gamma|_{[x,t']}$, which would contradict the definition of t_∞.

We deduce the approximation result:

Theorem. *Every path on X admits a piecewise affine ε-approximation.*

Proof. Let $\varepsilon > 0$ and suppose that $x = t_0 < \cdots < t_k \leqslant y$ with open balls $B_0, \ldots, B_k$ is an ε-discretization of γ defined on $[x, y]$. For every $i \in \{1, \ldots, k\}$ choose $t'_i \in]t_{i-1}, t_i[$ such that $\gamma(t'_i) \in B_{i-1} \cap B_i$ and define $t'_{k+1} = y$. We have the sequence

$$\gamma(t_0)\,,\ \gamma(t'_1)\,,\ \gamma(t_1)\,,\ \ldots\ ,\ \gamma(t'_k)\,,\ \gamma(t_k)\,,\ \gamma(y)$$

and for every $i \in \{0, \ldots, k\}$ we apply Lemma 6.7 from $\gamma(t'_i)$ to $\gamma(t_i)$ and from $\gamma(t_i)$ to $\gamma(t'_{i+1})$ in the ball B_i (omitting $\gamma(t'_0)$ which is not defined). The concatenation of dashes

$$\gamma(t_0) \to \gamma(t'_1) \to \gamma(t_1) \to \cdots \to \gamma(t'_k) \to \gamma(t_k) \to \gamma(y)$$

is the expected ε-approximation.

Now suppose that Y is a tile compatible subset of $\|G_1\| \times \cdots \times \|G_n\|$.

Corollary 6.10. Every directed path on Y admits a piecewise affine ε-approximation on Y.

Proof. Let ξ be a directed path on Y. According to Lemma 2.14 the set $X = (\beta_1 \times \cdots \times \beta_n)(Y)$ is tile compatible. The composite $(\beta_1 \times \cdots \times \beta_n) \circ \xi$ is a directed path on X which admits a (directed) piecewise affine ε-approximation α for any $\varepsilon > 0$ (Theorem 6.2). The expected approximation is the lifting of α (Theorem 5.3).

7. Afterword

A glimpse of Finsler geometry. Finsler metrics are motivated by the observation that measuring the length of a curve on a manifold $\mathcal{M}$ only requires a way of measuring the length of tangent vectors, see (Papadopoulos, 2013, p.40), (Chern & Shen, 2005, p.1). In this context, we call *(geometric) action functional* any map F from $T\mathcal{M}$ to $\mathbb{R}_+$ such that $F(x + y) \leqslant F(x) + F(y)$ and $F(\lambda x) = \lambda F(x)$ for all $p \in \mathcal{M}$, $x, y \in T_p\mathcal{M}$, and $\lambda > 0$. As suggested by the physical motivations described in (Bao et. al., 2000, p.1,2), the action functional F can be chosen so that the length of every curve γ with respect to F has a physical meaning. Hence minimizing length amounts to minimizing action. Any map arising from Definition 4.46 (11) is an action functional because $|\_|_\star$ is a norm. When the latter is $|\_|_\infty$ the length of γ is the time required to execute the sequence of multi-instructions that γ represents – see §6.1, and compare to the navigation problem (Chern & Shen, 2005, 1.4).

Following this idea, a *Finsler metric* is a family of Minkowski norms on $T_p\mathcal{M}$ smoothly varying with $p \in \mathcal{M}$ (Bao et. al., 2000, p.5). More formally, it is a map $F : T\mathcal{M} \to \mathbb{R}_+$ which is smooth on the *slit tangent bundle* (i.e. $T\mathcal{M} \setminus \{\text{zero tangent vectors}\}$) and whose restrictions to tangent spaces

are Minkowski norms (Chern & Shen, 2005, 1.2.1). These axioms are motivated by the notion of *curvature*, which plays a central role in the study of both Riemannian and Finsler geometries. Indeed, the latter are meant to generalize the former, for which the importance of curvature is well-known (Bao et. al., 2000, Introduction).

Although our mathematical framework arises from the same motivations, technicalities rule it out of the scope of standard Finsler geometry. First, manifolds of the form $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ are non-Hausdorff (in all relevant cases) so their topologies are not induced by their (pseudo)metrics. This defect is related to the non-determinism specific to our topic, and it seems impossible to get rid of it. The other issue is about the regularity of the action functional introduced in Definition 4.46 (11): it fails to be smooth on the slit tangent bundle when it is built from the norm we are interested in, namely $|\_|_\infty$ (§6.1). The reason is that the norm $|\_|_\infty$ is derivative at $(x_1, \ldots, x_n)$ if, and only if, $x_i \neq x_j$ for all $i \neq j$. About this last point, one may argue that 'ill-behaved' points form a neglectable closed subset of $\mathbb{R}^n$ (with respect to Lebesgue measure), so they may be manageable. Indeed, the problem of finding minimum action curves in degenerate (i.e. non-smooth) Finsler metrics is studied in Heymann (2015) whose preface indicate concrete situations in which they naturally occur. Heymann's theory is designed for action functional defined on $\mathbb{R}^n$, yet we are confident that this limitation can be easily overcome, at least for manifolds of the form $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ (although we will have to renounce uniqueness of geodesics, which is not prohibitive in computer science). Apart from this problem, Heymann theory satisfies many desirable properties (Heymann, 2015, p.6), in particular it applies to any action functional F on $\mathbb{R}^n$ that follows Definition 4.46 (11) (Heymann, 2015, Definition 2.4). In view of applications to geometric models of concurrency, we would like to adapt it to find *directed* minimizing curves.

More about Non-Hausdorff manifolds. Differential geometry (resp. topology) is really about metrisable manifolds, which are precisely those that can be embedded into $\mathbb{R}^n$ (Whitney's embedding theorem (Hirsch, 1976, Theorem 2.14, p.55)). Indeed, the standard tools of these two branches of mathematics are of little help as it is about global properties of non-Hausdorff manifolds. Some easy facts illustrate this claim.

A Hausdorff manifold is metrisable if, and only if, it satisfies any of the 119 topological property listed in (Gauld, 2014, §2.2, p.27) (second countability and paracompactness are those usually given in textbooks). The study of non-metrisable Hausdorff manifolds can lead to subtle set theoretic considerations (Gauld, 2014, Chapter 6). Non-Hausdorff manifolds fail to satisfy the most elementary results of Differential Geometry. For example they cannot be embedded in $\mathbb{R}^n$ (because any subspace of a Hausdorff space is Hausdorff). Also, while the circle and the real line are the only 1-dimensional metrisable manifolds (Hirsch, 1976, Exercise 6, p.20), «a reasonable classification of non-Hausdorff manifolds seems infeasible even in dimension 1» (Gauld, 2014, p.153).

Dealing with non-Hausdorff manifolds also has an impact on cobordism. A *cobordism* between the compact n-manifolds M_0 and M_1 is a $(n+1)$-manifold W whose boundary ∂W is isomorphic to the disconnected union $M_0 \cup M_1$ (Hirsch, 1976, §7.1,p.169). This definition tacitly assumes that W is metrisable. In this case, two 0-manifolds (i.e. finite sets) are related by a cobordism if, and only if, their cardinals have the same parity. Indeed, a set with $2k$ elements (seen as a 0-manifold) is the boundary of the disconnected union of k copies of $[0, 1]$. The only compact 1-manifold is the circle $\mathbb{S}^1$, and the well-known 'trouser manifold' (Gauld, 1982, Figure 55, p.166) is a cobordism between $\mathbb{S}^1$ and two copies of $\mathbb{S}^1$; therefore every compact 1-manifold is equivalent to the circle (up to cobordism).

What if we allow the manifold W (from the definition of cobordism) to be non-Hausdorff ? For $k \in \mathbb{N}$ with $k \geqslant 2$ let G_k be the graph with vertices $\{0, \ldots, k\}$ and arrows $0 \to 1$ and $1 \to i$ for every $2 \leqslant i \leqslant k$ (see Example 2.16 for $k = 3$). Provided we adapt our presentation to encompass manifolds with boundary, the manifold associated to the standard atlas $\mathcal{A}_{G_k}$ (Definition 3.14) is

non-Hausdorff, connected, and its boundary has k elements. Moreover, the product $\mathcal{A}_{G_k} \times \mathbb{S}^1$ is a non-Hausdorff manifold whose boundary is the disjoint union of k copies of the circle.

A standard way to build new metrisable manifolds from existing ones is to glue them along (certain connected components of) their boundaries (Wall, 2016, 2.7, p.63). Dealing with non-Hausdorff manifolds allows a more drastic construction: one can glue two n-manifolds by making them share an n-submanifold. For example, consider two copies of the segment $]0, 2[$ and identify their initial segments $]0, 1[$; the resulting manifold is isomorphic to the one described in Example 2.16, see also (Baillif & Gabard, 2008, Figure 1, p.1106).

Non-Hausdorff manifolds are not much studied in mathematics, nevertheless they could play some role in computer science, where graphs are pervasive (because, for example, automata are based on graphs). Moreover, graphs can be seen as 1-dimensional precubical sets; Higher Dimensional Automata (HDA) are built on precubical sets as automata are built on graphs (Haucourt, 2016, 2.4.9, p.45). Due to their expressiveness, HDA are important in concurrency theory van Glabbeek (2006*a*,b), so it would be interesting to determine which precubical sets can be blown up. The present paper already contains a partial answer: the blowup of a tensor product of graphs $G_1 \otimes \cdots \otimes G_n$ (Haucourt, 2016, 2.4.12) is the atlas product $\mathcal{A}_{G_1} \times \cdots \times \mathcal{A}_{G_n}$ (Definition 3.14).

Competing interests: The author declare none.

References

Aho, A. V., Lam, M. S., Sethi, R., and Ullman, J. D. (2007). *Compilers: Principles, Techniques, and Tools (2nd ed)*. Pearson.

Bao, D., Chern, S.-S., and Shen, Z. (2000). *An Introduction to Riemann-Finsler Geometry*. Springer-Verlag.

Baillif, M., and Gabard, A. (2008). Manifolds: Hausdorffness versus homogeneity. *Proceedings of the AMS*, 136(3):1105–1111.

Bridson, M. R., and Haefliger, A. (1999). *Metric Spaces of Non-Positive Curvature*. Springer-Verlag.

Brokate, M. and Kersting, G. (2015). *Measure and Integral*. Birkhäuser.

Coffman, E. G. , Elphick, M., and Shoshani, A. (1971). System deadlocks. *ACM Computing Surveys*, 3(2):67–78.

Carson, S. D. and Reynolds Jr, P. F. (1987). The geometry of semaphore programs. *ACM Transactions on Programming Languages and Systems*, 9(1):25–53.

Chern, S.-S., and Shen, Z. (2005). *Riemann-Finsler Geometry*. World Scientific.

do Carmo, M. P. (1992). *Riemannian Geometry (2nd printing)*. Birkhäuser.

Douady, A. and Douady, R. (2020). *Algebra and Galois Theories*. Springer.

Dijkstra, E. W. (1968). Cooperating sequential processes. In F. Genuys, ed., *Programming Languages*, pp. 43–112. Academic Press.

Fajstrup, L., Goubault, É., and Raussen, M. (2006). Algebraic Topology and Concurrency. *Theoretical Computer Science*, 357(1):241–278.

Gauld, D. B. (1982). *Differential Topology*. Marcel Dekker.

Gauld, D. B. (2014). *Non-metrisable Manifolds*. Springer.

Gierz, G., Hofmann, K. H., Keimel, K., Lawson, J. D., Mislove, M. W., and Scott, D. S. (2003). *Continuous Lattices and Domains*. Cambridge University Press.

Goubault-Larrecq, J. (2013). *Non-Hausdorff Topology and Domain Theory: Selected Topics in Point-Set Topology*. Cambridge University Press.

Goel, S. K. (1987). Nonseparated manifolds and completely unstable flows. *International Journal of Mathematics and Mathematical Sciences*, 10(4):745–756.

Hartshorne, R. (1977). *Algebraic Geometry*. Springer.

Hatcher, A. (2002). *Algebraic Topology*. Cambridge University Press.

Haucourt, E. (2016). Some Invariants of Directed Topology towards a Theoretical Base for a Static Analyzer Dealing with Fine-Grain Concurrency. Habilitation thesis.

Haucourt, E. (2018). The geometry of conservative programs. *Mathematical Structures in Computer Science*, 28(10):1723–1769.

Heymann, M. (2015). *Minimum Action Curves in Degenerate Finsler Metrics*. Springer.

Hilgert, J., Hofmann, K. H., and Lawson, J. D. (1989). *Lie Groups, Convex Cones, and Semigroups*. Oxford University Press.

Hicks, N. J. (1965). *Notes on Differential Geometry*. Van Nostrand Reinhold.

Hirsch, M. W. (1976). *Differential Topology (1997 corr. 6th printing)*. Springer-Verlag.

Haefliger, A. and Reeb, G. (1957). Variétés (non séparées) à une dimension et structures feuilletées du plan. *L'Enseignement Mathématique*, 3(2):107–125.

Hájíček, P. (1971). Causality in non-Hausdorff space-times. *Communications in Mathematical Physics*, 21:74–84.

Kung, H. T. , Papadimitriou, C. H., and Yannakakis, M. Z. (1979). Locking policies: safety and freedom for deadlock. In *Proc. 20th Annual Symposium on Foundations of Computer Science*, pp. 286–297.

Lang, S. (1999). *Fundamentals of Differential Geometry (corr. 2nd printing)*. Springer.

Lang, S. (2002). *Introduction to Differentiable Manifolds (2nd ed)*. Springer.

Lawson, J. D. (1989). Ordered manifolds, invariant cone fields, and semigroups. *Forum Mathematicum*, 1:273–308.

Lipski, W., and Papadimitriou, C. H. (1981). A fast algorithm for testing for safety and detecting deadlocks in locked transaction systems. *Journal of Algorithms*, pp. 211–226.

Luc, J., and Placek, T. (2020). Interpreting non-Hausdorff (Generalized) manifolds in general relativity. *Philosophy of Science*, 87(1):21–42.

Müller, T. (2013). A generalized manifold topology for branching space-times. *Philosophy of Science*, 80(5):1089–1100.

Nachbin, L. (1965). *Topology and Order*. Van Nostrand.

Nachbin, L. (1981). *Introduction to Functional Analysis: Banach Spaces and Differential Calculus*. Marcel Dekker.

Papadopoulos, A. (2013). *Metric Spaces, Convexity, and Nonpositive Curvature (2nd ed)*. European Mathematical Society.

Rudin, W. (1976). *Principles of Mathematical Analysis (3rd ed)*. McGraw-Hill.

Schröder, B. (2003). *Ordered Sets, An Introduction*. Birkhäuser.

Segal, I. E. (1976). *Mathematical Cosmology and Extragalacic Astronomy*. Academic Press.

Schaefer, H. H., and Wolff, M. P. (1999). *Topological Vector Spaces (2nd ed)*. Springer-Verlag.

van Glabbeek, R. J. (2006). On the Expressiveness of Higher Dimensional Automata. *Theoretical Computer Science*, 356(3):265–290.

van Glabbeek, R. J. (2006). Erratum to 'On the Expressiveness of Higher Dimensional Automata'. *Theoretical Computer Science*, 368(1-2):168–194.

Wall, C. T. C. (2016). *Differential Topology*. Cambridge University Press.