Singularities of Translation Manifolds

Sven Caspart

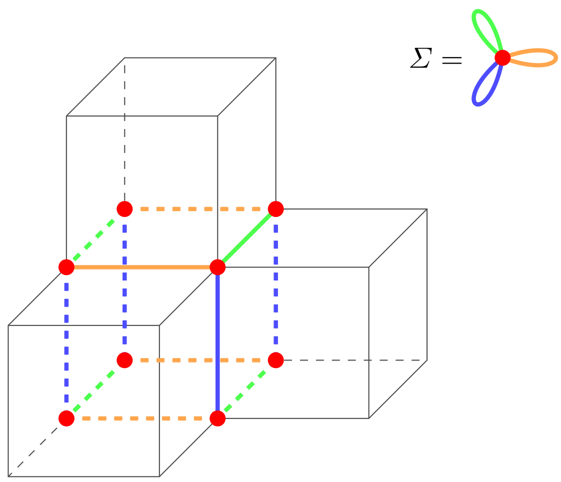

Dissertation

Karlsruhe Institut für Technologie (KIT)
Fakultät für Mathematik

Tag der mündlichen Prüfung: 1. Juli 2020
Referenten: Prof. Dr. Frank Herrlich, Prof. Dr. Gabriela Weitze-Schmithüsen

Contents

0 Introduction

One of the most important, if not the most important, geometric spaces is the Euclidean space $\mathbb{R}^m$. Throughout the history of mathematics its concepts have been generalised in various ways yielding for example spherical and hyperbolic geometry or the more abstract Riemannian geometry. The planar case $\mathbb{R}^2$ was especially fruitful and resulted, among others, in the classification of closed Riemann surfaces.

A geometric object which is very close to resembling $\mathbb{R}^2$ is a translation surface. A rough description is that a translation surface is glued out of polygons where two sides are glued together by a translation. Because it is built out of polygons, it is locally isometric to $\mathbb{R}^2$ and since the gluing is done by translations, i.e. no rotations, we have at each point the directions 'North', 'West', 'South', and 'East' and these are well-defined across the whole translation surface unlike for example on the Möbius strip.

In this work we will describe the generalisation of a translation surface to any dimension m and we will fittingly call it *translation manifold*. The generalisation to higher dimension is a very natural question and the definition given above can easily be adapted to higher dimensions by gluing polytopes instead of polygons. However, the implications are more drastic as most of the theory for translation surfaces relies on two-dimensional concepts like a complex structure, or the classification of surfaces.

Some points of a translation surface are especially interesting, the so called *singularities*. Walking around a singularity does not yield an angle of 2π but of $2\pi k$ for some $k \in \mathbb{N}$ and we can classify the (tame) singularities by this number k. These special points also occur in the higher-dimensional translation manifolds. However, in two dimensions we essentially have only one way (or its inverse) to walk around a singularity. In higher dimensions we have multiple ways to do that, so we can no longer use this approach to understand and classify them. Instead we describe a new method applicable to all dimensions for identifying and comparing singularities across translation manifolds. The new approach uses translation coverings and the developing map.

Of particular interest is whether a singular point is a 'real singularity' or not. If the singularity is on a translation surface and the angle around it is 2π, i.e. $k = 1$, then we can add the point to the translation surface and still have a translation surface. In this case we call the singularity *removable*. For an angle of $2\pi k$ with $k \geq 2$, the singularity cannot be added to the translation surface without disturbing the translation structure, so that is a 'real singularity'.

For singularities of translation manifolds the same question concerning their removability arises. This time the question is more difficult to answer as we do not have the number k to begin with and as singularities are no longer single points but a collection of points like curves or surfaces. In this work we will show that isolated point singularities do not exist on a translation manifold except for a two dimensional manifold, i.e. a translation surface, and expand that result to any collection of singularities which admit a simply connected neighbourhood.

The structure of this text is as follows. In the first chapter, section 1.1 contains a definition of a translation manifold as well as the three different generalisations resulting from the three different but equivalent definitions of a translation surface and we will discuss these divergent generalisations and their merits.

Also part of chapter 1 is the definition and construction of our two main tools: the translation covering and the developing map. This will take place in sections 1.2 and 1.3, respectively.

The discussion of the singularities of a translation manifold happens in chapter 2. There we start by proving theorem 2.1.1 which allows us to test whether a singularity is removable or not using a translation covering (section 2.1). Section 2.2.2 discusses and solves the question when an isolated singularity on a translation manifold manifold is removable.

The next section (section 2.3) generalises the observation of the previous section and introduces *images* and *shadows* of singularities. Shadows are non-singular points but they behave like singularities with respect to the developing map and should as such be treated like singularities. Using this concept, we are able to provide criteria for the removability of an arbitrary singularity of a translation manifold in form of corollary 2.3.19 and theorem 2.3.21.

In section 2.4, we apply the results of section 2.3 to dimension two and three where some of the prerequisites are automatically fulfilled.

In the last chapter, we examine a particular kind of translation manifold: a cubic translation manifold. As its name suggests it is built out of cubes and this cubic structure allows us to find suitable and well-behaved neighbourhoods around any point in the manifold, in particular around its singularities. As a consequence we can prove theorem 3.2.7 which states that singularities of codimension greater than two are always removable (except when part of a larger, codimension two singularity). Furthermore, the description of these neighbourhoods enables us to give a complete classification of all singularities occurring on a cubic translation manifold and how they intersect in theorem 3.2.29.

Acknowledgements:
The author of this work was supported by a stipend of the Carl-Zeiss-Stiftung.

Notation

Throughout this work we are using the following notations:

$a := b$

the symbol a is defined to be the expression b.

$\mathcal{A}$

atlas of a (translation) manifold, see definition 1.1.1.

$\overline{A}$

closure of a set A in a topological space.

$[A]_{\approxeq}$

denotes the equivalence class of faces of cubic translation manifolds which are isometric-isomorphic to the face A; see definition 3.2.18.

$[A]_{\cong}$

denotes the equivalence class of faces of cubic translation manifolds which are translation-isomorphic to the face A; see definition 3.2.18.

$\coprod_i A_i,\ A_1 \amalg A_2$

coproduct of objects A_i; often denotes the disjoint union of sets/spaces.

$\bigvee_i A_i,\ A_1 \vee A_2$

gluing of topological spaces A_i at a single point; called *wedge sum* or *wedge product*. We have $\pi_1(\bigvee_i A_i) = *_i \pi_1(A_i)$ where on the right-hand side we have the free product of the fundamental groups of A_i.

$A \# B$

connected sum of the topological manifolds A and B. It is obtained by cutting out a m-ball of A and B and glueing the resulting $(m-1)$-spheres via a homeomorphism.

$A \approxeq B$

the faces A and B of two cubic translation manifolds are isometric-isomorphic; see definition 3.2.18.

$A \cong B$

the faces A and B of two cubic translation manifolds are translation-isomorphic; see definition 3.2.18.

$]a,b[,\ ]a,b],\ [a,b[,\ [a,b]$

the open, half-open, and closed intervals in $\mathbb{R}$ also often denoted by (a,b), $(a,b]$, $[a,b)$, $[a,b]$.

$\mathrm{Acc}((x_n)_{n\in\mathbb{N}})$

set of all accumulation points of the sequences $(x_n)_{n\in\mathbb{N}}$.

$A \amalg_X B$

gluing of topological spaces A and B along a common subspace X.

$\bar{B}(x,r)$

closed ball of radius r around x in a metric space X, often but not always the same as the closure of the open ball $\overline{B(x,r)}$; $\bar{B}(x,r) := \{y \in X \mid d(x,y) \leq r\}$.

$B(x,r)$

open ball of radius r around x in a metric space X; $B(x,r) := \{y \in X \mid d(x,y) < r\}$.

$\dot{B}(x,r)$

punctured open ball of radius r around x in a metric space X, i.e. the open ball without origin; $\dot{B}(x,r) := \{y \in X \mid d(x,y) < r\} \setminus \{x\}$.

B^n

(topological) solid n-ball.

D

usually denotes the developing map, see definition 1.3.6.

$d,\ d_{\mathrm{eukl}},\ d_{\mathbb{R}^m}, d_M$

d generally denotes a metric, d_{eukl} and $d_{\mathbb{R}^m}$ is the Euclidean metric, and d_M is the metric on space M.

∂A

boundary of a set A in a topological space.

$D_U(\Sigma)$

image of the singularities or the boundary of U, see definition 2.3.7.

$D_U^{\mathrm{f}}(\Sigma)$

image of the singularities of U, see definition 2.3.2.

$f|_A$

restriction of a function/1-form/… f to a subset A.

F_n

free group of rank n.

$G *_U H$

free product of the groups G and H with amalgamation over U.

$\mathfrak{I}(z)$

the imaginary part of a complex number/function/1-form/… z

$\operatorname{int} A$

interior of a set A in a topological space.

$\ell(\gamma)$

path length of a path γ.

M'

for a cubic translation manifold M: same translation manifold as M but considered as cubic translation manifold where the cubes are subdivided into $3^{\dim M}$ smaller cubes, i.e. the edge length is divided by 3; see definition 3.2.9.

$N_M(A)$

cubic neighbourhood of a face A of a cubic translation manifold M; see definition 3.1.9.

$\pi_1(X,x),\ \pi_1(X)$

first fundamental group of a topological space X with respect to the basepoint x. The basepoint is often ommitted.

$\mathfrak{R}(z)$

the real part of a complex number/function/1-form/… z

$\operatorname{relint}(A)$

the relative interior of a k-face A of a cubic translation manifold; $\operatorname{relint}(A)$ is A without $(k-1)$-faces; see definition 3.1.6

Σ

: set of singularities of a translation manifold, $\Sigma = \overline{M} \setminus M$, see definition 1.1.8.

S^n

: (topological) n-sphere.

$S_U^{\mathrm{f}}(\Sigma)$

: shadows of the singularities of U, see definition 2.3.2.

$S_U(\Sigma)$

: shadows of the singularities or the boundary of U, see definition 2.3.7.

$\mathrm{Deck}(\tilde{X}/X)$

: the Deck transformation group of the covering map $\tilde{X} \to X$.

T^*

: torus of dimension m as translation manifold without its codimension two skeleton.

T^m

: solid torus of dimension m; $T^m = \mathbb{R}^m/\mathbb{Z}^m$.

$X^{\mathbb{N}}$

: set of all sequences $(x_n)_{n\in\mathbb{N}}$ in X.

$[x, y]$

: commutator of the group elements x and y; $[x, y] = xyx^{-1}y^{-1}$.

1 Fundamentals on Translation Manifolds

In this chapter we provide the ground work for the theory of translation manifolds.

We start off with a section giving not one but three different definitions of a translation manifold. This happens in section 1.1. The first definition we give is the simplest and most general one, and uses only the language of a manifold and its transition maps. Less general but a concrete way to construct examples is the second definition which glues polytopes together along their sides to form a translation manifold. The last definition relies on differentiable 1-forms and uses language and results from differential geometry. At the end, in section 1.1.4, we give a comparison of these definitions together with their advantages and uses.

A translation manifold contains special – or better noteworthy – points, the so called *singularities*. We define them in this chapter but they are of such rich structure that they merit a chapter of their own (chapter 2).

We move on to quickly introduce the concept of a translation covering in section 1.2, which will be one of our main focus later on, before we discuss and construct the *developing map* of a translation manifold in section 1.3. We provide the definition of a developing map for a slightly more general type of manifolds – so called (G, X)-manifolds. This comes at no additional cost as the proofs are the same whether we consider a (G, X)-manifold or a translation manifold.

At the end, in section 1.3.3, we will discuss properties of the developing map on a translation manifold. It comes to no surprise that due to the translation nature the developing map of a translation manifold has more nice properties than a generic developing map.

1.1 Definition of a Translation Manifold

Translation manifolds can be defined in different ways. In this section, we will discuss three different definitions. These definitions are generalisations of the three definitions of a translation surface to higher dimensions. It is interesting to note that while the definitions are equivalent for (finite) translation surfaces, they are no longer equivalent in higher dimensions.

1.1.1 Definition as a Manifold

Let us start with the definition which justifies both parts of its name: *translation* and *manifold*.

Definition 1.1.1 (Translation Manifold). An m-dimensional *translation manifold* is an m-dimensional manifold M with atlas $\mathcal{A}$ where the changes of coordinates (also called transition maps) are translations, i.e. they are locally of the form

$$x \mapsto x + c \tag{1.1}$$

for some vector $c \in \mathbb{R}^m$.

Such an atlas is called a *translation atlas.* A maximal translation atlas is called *translation structure.*

Remark 1.1.2. We usually do not make the distinction between a translation atlas and a translation structure and assume without loss of generality that all atlases are maximal.

Remark 1.1.3. Because the changes of coordinates are translations, they are in particular smooth, affine, etc. So a translation structure also induces a smooth structure, affine structure, etc. on M.

Example 1.1.4.

a) $M = \mathbb{R}^m$ with the identity as a global chart $\varphi = \mathrm{id}\colon M \to \mathbb{R}^m$.

b) $M = B(0,1) \subseteq \mathbb{R}^m$ with the inclusion as global chart $\varphi\colon M \hookrightarrow \mathbb{R}^m$.

c) $M = \dot{B}(0,1) \subseteq \mathbb{R}^m$ with the inclusion as global chart $\varphi\colon M \hookrightarrow \mathbb{R}^m$.

d) $M = \mathbb{R}^m/\mathbb{Z}^m$ the m-torus. A chart around a point $p \in M$ is the local inverse of the covering map $\pi\colon \mathbb{R}^m \to M$ on a sheet.

e) Let T_1 and T_2 be two m-tori with the translation structure as above. Embed a solid $(m-1)$-hyperball C in the torus T_1 and in the same way in the torus T_2. Cutting along the hyperball C does not separate the torus and yields two opposed $(m-1)$-hypersurfaces inside each torus (see figure 1.1). Each hypersurface can be identified by a translation with the hypersurface of the opposite side in the other torus because C is embedded in the same way in both tori. This yields a translation structure on $(T_1 \setminus \partial C) \cup (T_2 \setminus \partial C)$ because the gluing was a translation. This is very similar to the gluing appearing in definition 1.1.14.

   This construction can be generalised to more sophistically embedded surfaces embedded in any translation manifold, not just tori, see for example figure 1.2.

We will see more examples of translation manifolds as we go on.

Since the changes of coordinates are translations, the pullback of the Euclidean metric of $\mathbb{R}^m$ via the charts of a translation atlas is locally well-defined. We can use this to define a metric on the translation manifold:

Definition 1.1.5 (Flat metric). The *flat metric* d on a translation manifold M is the metric induced by the path length:

$$d(x,y) := \inf\{\ell(\gamma) \mid \gamma\colon [0,1] \to M \text{ path with } \gamma(0) = x \text{ and } \gamma(1) = y\} \tag{1.2}$$

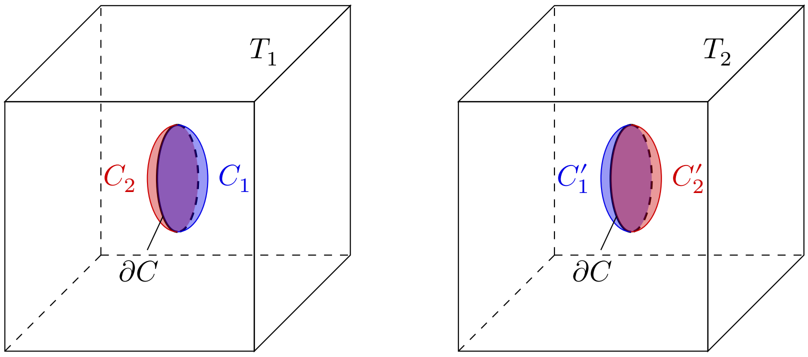

Figure 1.1: The cut along the embedded disc C in each of the tori yields two surfaces, C_1 and C_2, in the first torus and two surfaces, C_1' and C_2', in the second torus T_2. For better visualisation the surfaces are drawn with a curve. The gluing of C_1 with C_1' and of C_2 with C_2' results in a translation structure on $(T_1 \setminus \partial C) \cup (T_2 \setminus \partial C)$. The outer sides of the cube representing the torus T_1 are glued with their respective opposite side in T_1 and likewise for T_2.

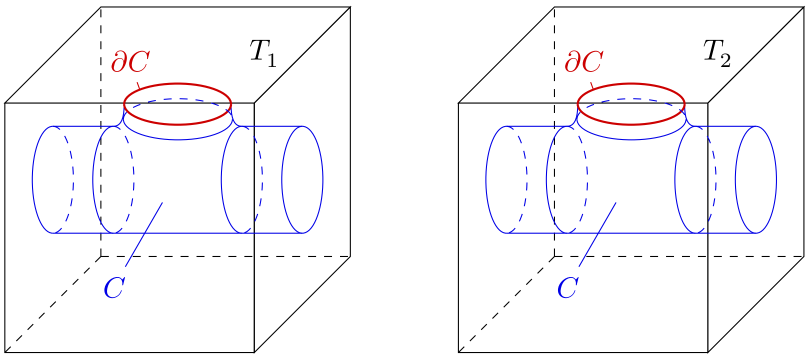

Figure 1.2: Example of a more sophisticated gluing of two tori along an embedded surface C. The surface C has a boundary component ∂C but is not retractable to a disc-surface like in figure 1.1. Again, after the gluing, the translation structure is given on all of the two tori except on the boundary ∂C.

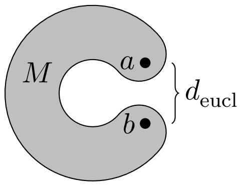

Figure 1.3

for $x, y \in M$. The length $\ell(\gamma)$ of a path γ is measured locally by the pullback of the Euclidean metric, i.e. locally $\ell(\gamma) = \ell(\varphi(\gamma))$ for a chart φ. So the path length of a path γ on M is

$$\ell(\gamma) := \sum_{i=1}^{n} \ell(\varphi_i(\gamma\big|_{[t_{i-1},t_i]})), \tag{1.3}$$

where $0 = t_0 < t_1 < \dots < t_n = 1$ is a subdivision of $[0,1]$ such that $\gamma\big|_{[t_{i-1},t_i]}$ is contained in a chart $\varphi_i \colon U_i \to \mathbb{R}^m$ and the $\ell(\cdot)$ on the right-hand side denotes the usual path length in $\mathbb{R}^m$.

We have that – per definitionem – a translation manifold is a length metric space (also called interior space). It is worth mentioning that the flat metric defined above coincides locally with the pullback of the Euclidean metric and is therefore indeed a metric. Moreover, it immediately follows:

Lemma 1.1.6. *Charts of a translation manifold are local isometries.*

We will later explore in corollary 1.3.11 how local this isometric property is.

Remark 1.1.7. Note that the metric defined via the path length does not need to coincide globally with the (possibly) induced Euclidean metric by pullback via the charts. Consider the C-shaped set in $\mathbb{R}^2$ depicted in figure 1.3 with the atlas consisting of the inclusion $M \hookrightarrow \mathbb{R}^2$. Here the Euclidean distance d_{eucl} between a and b is different than the distance given by the flat metric induced by the path length, which has to 'go around' and is thus longer.

Definition 1.1.8 (Singularities). The *singularities* Σ of a translation manifold M are defined as

$$\Sigma := \overline{M} \setminus M \tag{1.4}$$

where $\overline{M}$ is the metric completion of M with respect to the flat metric.

A singularity $\sigma \in \Sigma$ is called *removable*, if there is an open neighbourhood $U \subseteq \overline{M}$ of σ and a map $\varphi \colon U \to \mathbb{R}^n$ such that $\mathcal{A} \cup \{\varphi\}$ is a translation atlas for $M \cup U$. Note that U might contain more singularities than just σ, which then are also removable using the same chart φ.

Remark 1.1.9. For a flat surface the singularities are (mostly) isolated points, so speaking of *a* singularity is unambiguous. In higher dimensions the singular points themselves can form structures, e.g. all points on an edge of a cube can be singularities. In this case we also refer to the whole edge as *a* singularity although it consists of many individual singular points.

Example 1.1.10. a) $M = \mathbb{R}^m$ with the identity as a global chart $\varphi = \mathrm{id}\colon M \to \mathbb{R}^m$ has no singularities since $\mathbb{R}^m$ is already metrically complete.

b) $M = B(0,1) \subseteq \mathbb{R}^m$ with the inclusion as a global chart $\varphi\colon M \hookrightarrow \mathbb{R}^m$. The metric completion of M is the closed ball $\overline{M} = \bar{B}(0,1)$ and the singularities are/can be thought of as the boundary of M in $\mathbb{R}^m$.

c) $M = \dot{B}(0,1) \subseteq \mathbb{R}^m$ with the inclusion as a global chart $\varphi\colon M \hookrightarrow \mathbb{R}^m$ for $m \geq 2$. Similar to the previous example it has the boundary as singularities but also the centre point 0. We can include the singular centre point into M and still have a translation structure on $M \cup \{0\}$ (which is the previous example). The center is therefore a *removable* singularity. We will discuss this type of isolated singularity in depth in section 2.2.2.

d) The m-torus $M = \mathbb{R}^m/\mathbb{Z}^m$ has no singularities. It is a compact space and thus metrically complete.

e) Gluing of two m-tori along an $(m-1)$-hyperball C as in example 1.1.4 e) yields more interesting singularities. The metric completion is the two m-tori glued along the closed $(m-1)$-hyperball $\overline{C}$ via a translation. The singularities are therefore the embedded boundary of C in each of the tori which get identified. So the set of singularities is homeomorphic to (a single copy of) the boundary of C – see figure 1.4.

In two dimensions there is the well-studied concept of a *finite translation surface* often just referred to as *translation surface*.

Definition 1.1.11 (Finite Translation Surface I). A finite translation surface is a two-dimensional translation manifold with the additional conditions that the metric completion is a closed surface without boundary and the number of singular points is finite.

A finite translation surface can also be defined in two very different but equivalent ways. These different aspects give rise to the beautiful and rich theory of flat surfaces. A summary of these definitions follows, for more details we refer to the higher dimensional definitions (definitions 1.1.14 and 1.1.15).

Definition 1.1.12 (Finite Translation Surface II). A finite translation surface is the object obtained by gluing finitely many polygons in the plane in the following way: each polygon is endowed with an orientation, and two sides of opposite orientation but equal length are glued via a translation in $\mathbb{R}^2$. The metric structure is inherited from $\mathbb{R}^2$ and the singularities are the corners of the polygons.

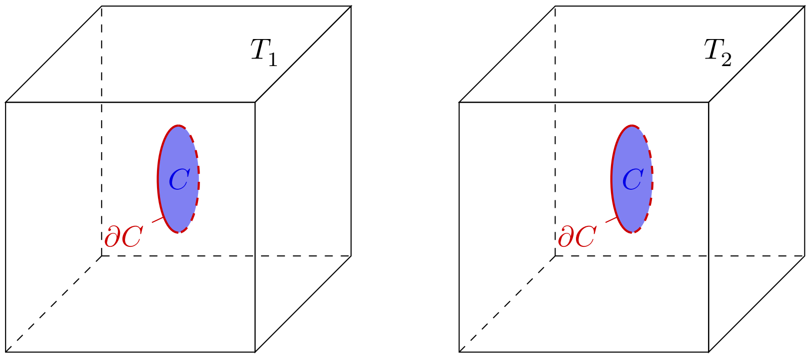

Figure 1.4: Gluing of two tori along an embedded disc C results in the boundary ∂C of C becoming singular. If we walk around a point on ∂C starting in T_1, we pass through C and after a full 360° turn we land at the same point we started but in the other torus. After walking another 360° we come back to the point we started at in T_1. During the walk the surface C works like a portal teleporting us from one torus to the other and back.
The fact that we need to walk 720° to come back to the starting point implies that we cannot find a neighbourhood for that singular point which is locally isometric to the Euclidean space because in the latter we would always only need to walk 360° to return to the starting point. Thus, ∂C consists entirely of non-removable singularities.

Definition 1.1.13 (Finite Translation Surface III). A finite translation surface is a Riemann surface, i.e. a closed surface with a complex structure and without boundary, together with a non-zero holomorphic 1-form on it. The holomorphic 1-form is also called *Abelian differential*. The charts for the translation manifold are obtained by integrating the holomorphic 1-form along paths and the singularities are the zeros of the 1-form.

Both of these definitions can be generalised to higher dimensions. The latter, however, needs to be given some additional consideration as a complex structure cannot be generalised to odd dimensions over $\mathbb{R}$.

1.1.2 Definition with Polytopes

The second definition constructs translation manifolds by gluing polytopes along their sides in the way you would imagine. The precise definition is a little technical as we have to pay attention to the orientation and the singularities.

Definition 1.1.14 (Translation Manifold II). Let $P_i \subseteq \mathbb{R}^m$ be (not necessarily convex) polytopes with $i \in I$ for some index set I. Pick an orientation on each polytope. The sides of the polytopes are partitioned into pairs such that each pair consists of sides with opposite orientation and which are isometric via a translation. Denote by P_i^{m-2}

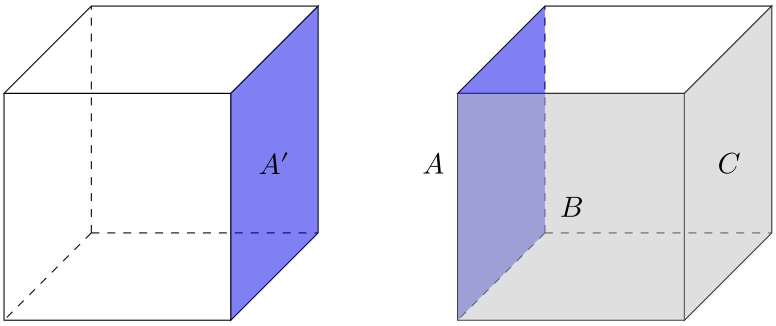

Figure 1.5: The side A' of the first cube can be glued with the side A of the second cube because they are isometric via a translation and have opposite orientations (the interior of the cubes are on opposite sides after identifying A with A'). The side A' cannot be glued with side B even though they are isometric since they are not isometric via a translation.
The side A' can also not be glued with the side C even though they are isometric via translation because they have the same orientation (the interior of the cubes would end up being on the same side after identifying A' with C).

the codimension two skeleton of the polytope P_i. The *translation manifold* M is the topological space

$$M = (\coprod_{i \in I} (P_i \setminus P_i^{m-2}))/_{\sim}, \tag{1.5}$$

where the equivalence relation is induced by identifying the relative interiors of the pairs of sides via the translation. The set of singularities is the union of the codimension two skeletons

$$\Sigma = (\coprod_{i \in I} P_i^{m-2})/_{\sim} \tag{1.6}$$

again identified via the translations for the pairs. The metric structure on M is obtained by locally identifying the polytope with its embedding in $\mathbb{R}^m$. Note that due to the choice of opposite orientation, two solid hemispheres which are glued by an identification of sides match up perfectly to build isometrically a solid ball in $\mathbb{R}^m$ (cf. figure 1.5).

1.1.3 Definition using 1-Forms

The last definition uses 1-forms and the language of differential geometry.

Recall that a 1-form ω on an m-dimensional, differentiable manifold M is a smooth section of the cotangent bundle T^*M, i.e. $\omega\colon M \to T^*M$ with $\omega(p) \in T_p^*M$ for all $p \in M$. Locally on a chart $\varphi\colon U \to \mathbb{R}^m$, any 1-form ω can be written as

$$\omega|_U = \sum_{i=1}^{m} f_i \, \mathrm{d}x_i, \tag{1.7}$$

where $f_i\colon U \to \mathbb{R}$ are smooth functions. A 1-form is called *closed* iff $\mathrm{d}\omega = 0$ and called *exact* iff $\omega = \mathrm{d}f$ for some 0-form f, i.e. for a smooth function $f\colon M \to \mathbb{R}$. Here d denotes the (exterior) derivative operator mapping k-forms to $(k+1)$-forms.

Definition 1.1.15 (Translation Manifold III). A *translation manifold* of dimension m is an m-dimensional, differentiable manifold M together with m closed 1-forms $\omega_1, \ldots, \omega_m$ such that $\omega_1 \wedge \cdots \wedge \omega_m$ has no zeros on M.

The condition for $\omega_1 \wedge \cdots \wedge \omega_m$ to have no zeros on M can be rephrased to: $\omega_1, \ldots, \omega_m$ are pointwise linearly independent, that is $\omega_1(p), \ldots, \omega_m(p)$ are linearly independent in the vector space $T_p^* M$ for all $p \in M$.

When comparing this with the definition of a translation surface, we see that we have replaced one complex-valued 1-form with two real-valued 1-forms. Their connection being – as expected – the mapping of the complex 1-form to its real and imaginary part. The condition for the 1-form of being holomorphic is replaced by being closed. The characterisation of the singularities as zeros is replaced by the characterisation via linearly dependence.

Remark 1.1.16. There are two viewpoints to this definition.

(VP1) The first one is to consider the translation manifold M to be the set of regular points, i.e. does not contain singularities. The singularities are later 'added' via the metric completion: $\Sigma = \overline{M} \setminus M$. This is the point of view we used in definition 1.1.15 and is in line with definition 1.1.1. For our previous definition, this means that we require the 1-forms to be linearly independent on *all* of M.

(VP2) The second point of view is to consider the translation manifold N to be the set of all points (singular and regular). The set Σ of singularities is then defined as a subset of N. That is N corresponds to $\overline{M}$ above and $N \setminus \Sigma$ to M. This viewpoint is used in definition 1.1.13 and usually requires the 1-forms to be defined on all of N, i.e. on regular points and singularities. The singularities are then defined as the points where the 1-forms are not linearly independent.

Note that the two definitions resulting from these points of view are slightly different for corner cases as the latter requires the 1-forms to be extendable to the closure $\overline{M}$.

Example 1.1.17. Let $T := \mathbb{R}^2/\mathbb{Z}^2$ be the torus. Consider the two closed 1-forms $\omega_1 := \sin(2\pi x)\,\mathrm{d}x$ and $\omega_2 := \mathrm{d}y$ on T.

With the second viewpoint (VP2) we have $N := T$ and the set of singularities Σ_2 are all the points where ω_1 and ω_2 are linearly dependent. In this example we have $\Sigma_2 = \{(x, y) \in T \mid x = k\pi \text{ for some } k \in \mathbb{Z}\}$. Hence the translation surface $T^* := T \setminus \Sigma_2$ decomposes into two connected components.

With the first point of view (VP1), we start with $M := T^*$ and ω_1 and ω_2 are only defined on T^*. The set of singularities is then $\Sigma_1 := \overline{M} \setminus M$, where the metric is induced by the charts given by ω_1 and ω_2 (see theorem 1.1.23 below).

Be aware that $\Sigma_1 \neq \Sigma_2$ in this example. This is because T^* has two connected components and thus their metric completions do not overlap. With the second viewpoint (VP2), however, the connected components are still tied together by the underlying N resulting in fewer singularities Σ_2 because Σ_2 is a subset of N .

Remark 1.1.18 (A Translation Surface is a two-dimensional Translation Manifold). Given a translation surface (X, ω) as a closed Riemann surface X together with a non-zero holomorphic 1-form ω, i.e. the classic definition (definition 1.1.13), we get a two-dimensional translation manifold by considering X as a smooth real two-dimensional manifold and the real part $\Re\omega$ and imaginary part $\Im\omega$ of ω as real-valued 1-forms.

Indeed, $\Re\omega$ and $\Im\omega$ are linearly independent at point p if and only if ω is non-zero at p. This can be checked locally: Locally on a chart $U \subseteq X$ we can describe ω as $\omega = f\,\mathrm{d}z = (\Re f + i\Im f)(\mathrm{d}x + i\,\mathrm{d}y)$ for some holomorphic function $f\colon U \to \mathbb{C}$. We have

$$\Re\omega = \Re f\,\mathrm{d}x - \Im f\,\mathrm{d}y, \qquad \Im\omega = \Im f\,\mathrm{d}x + \Re f\,\mathrm{d}y \tag{1.8}$$

and thus

$$\Re\omega \wedge \Im\omega = ((\Re f)^2 + (\Im f)^2)(\mathrm{d}x \wedge \mathrm{d}y). \tag{1.9}$$

Therefore, $(\Re\omega \wedge \Im\omega)(p) = 0$ if and only if $f(p) = 0$, in other words if and only if $\omega(p) = 0$.

So $M := X \setminus \{p \in X \mid \omega(p) = 0\}$ together with $\Re\omega$ and $\Im\omega$ is a (two-dimensional) translation manifold according to definition 1.1.15. Because X is closed and the zeros of ω are discrete, the set of singularities given by the metric completion is precisely the set $\{p \in X \,|\, \omega(p) = 0\}$. Thus, the translation manifold $(M, \Re\omega, \Im\omega)$ is exactly the translation surface (X, ω) and definition 1.1.15 is indeed a generalisation of definition 1.1.13.

The above remark shows that if ω_1 and ω_2 are the real and imaginary part of a holomorphic 1-form, then they become linearly dependent at point p if and only if both vanish at point p. However, if ω_1 and ω_2 are not induced by a holomorphic 1-form, then linearly dependent in point p does not imply that both 1-forms must vanish. It can be that only one of the 1-forms vanishes or that none of them do, see example 1.1.25 for an illustration of this.

Remark 1.1.19. The more correct analogue of holomorphic for higher dimensions is a *harmonic* 1-form. However, harmonic in m dimensions can only be defined on manifolds with a Riemannian metric. This is a slightly stronger prerequisite on our manifold than just differentiable. For the construction of a translation atlas, which we will present in a moment, a closed 1-form is enough. Furthermore, theorem 1.1.24 states that the 1-forms ω_i become harmonic with respect to the translation structure they induce, so in the end we have not lost anything.

Before we begin formulating the translation charts for the above definition of a translation manifold with 1-forms, we need to recall (a consequence of) the Poincaré Lemma [Lee13, Theorem 11.49 and Corollary 11.50, pp. 296–297]:

Lemma 1.1.20 (Poincaré Lemma). *On a simply connected manifold closed k-forms are exact.*

This Poincaré Lemma is the k-form analogue for a holomorphic function to have a primitive on a simply connected domain.

With this tool equipped we can construct (well-defined) charts on M as follows:

Lemma 1.1.21. *Let M be an m-dimensional, differentiable manifold, $p \in M$, and ω_1,…, ω_m closed 1-forms.*

If $\omega_1(p)$,…, $\omega_m(p)$ are linearly independent in T_p^*M, then there is an open neighbourhood $U \subseteq M$ of p such that

$$\varphi: U \to \mathbb{R}^m, \qquad q \mapsto \left(\int_p^q \omega_1, \dots, \int_p^q \omega_m \right) \tag{1.10}$$

is a homeomorphism. Here $\int_p^q \omega_i$ denotes the integral over ω_i along any path from p to q which lies inside U.

Proof. We have three things to check: 1. φ is well-defined, 2. φ is bijective, and 3. φ and φ^{-1} are continuous.

Since linearly independence is an open condition, we find an open neighbourhood $U \subseteq M$ around p on which $\omega_1(p)$,…, $\omega_m(p)$ are linearly independent. Possibly making U smaller, we may further assume that U is simply connected. Then (the restrictions of) the closed 1-forms ω_i are exact on U. Therefore, the integral

$$\int_p^q \omega_i \tag{1.11}$$

for $q \in U$ is independent of the path taken between p and q inside of U [Lee13, Theorem 16.26]. Thus, the map φ of the statement is well-defined.

Since ω_i are exact on U, there are continuous differentiable functions $f_i: U \to \mathbb{R}$ with $\mathrm{d}f_i = \omega_i$ on U. We have [Lee13, Theorem 11.39, p. 291]

$$\int_p^q \omega_i = f_i(q) - f_i(p). \tag{1.12}$$

By altering f_i by a constant we may assume that $f_i(p) = 0$ and this does not affect $\mathrm{d}f_i = \omega_i$. With this choice our chart in spe φ becomes

$$\varphi(q) = (f_1(q), \dots, f_m(q)). \tag{1.13}$$

To show that φ is bijective and a homeomorphism we use the inverse function theorem (or the implicit function theorem). However, instead of proving bijectivity and homeomorphism for φ, we show it for $\varphi \circ \psi^{-1}$, where $\psi: U \to \mathbb{R}^m$ is a chart of the manifold M around p. Indeed, this suffices since ψ is a homeomorphism. By making U smaller if necessary, we find a chart $\psi: U \to \mathbb{R}^m$ around p. In this chart the 1-form ω_i is given by

$$\omega_i(p) = \sum_{j=1}^m \frac{\partial (f_i \circ \psi^{-1})}{\partial x_j}(\psi(p))\, \mathrm{d}_p x_j \tag{1.14}$$

for $p \in U$. The derivative of $\varphi \circ \psi^{-1}: \mathbb{R}^m \supseteq \psi(U) \to \mathbb{R}^m$ at $\psi(p) \in \mathbb{R}^m$ is

$$D(\varphi \circ \psi^{-1})(\psi(p)) = \begin{pmatrix} \dfrac{\partial (f_1 \circ \psi^{-1})}{\partial x_1}(\psi(p)) & \dots & \dfrac{\partial (f_1 \circ \psi^{-1})}{\partial x_m}(\psi(p)) \\ \vdots & & \vdots \\ \dfrac{\partial (f_m \circ \psi^{-1})}{\partial x_1}(\psi(p)) & \dots & \dfrac{\partial (f_m \circ \psi^{-1})}{\partial x_m}(\psi(p)) \end{pmatrix}. \tag{1.15}$$

Note that the entries of a row are exactly the coefficients of $\omega_i(p) = \mathrm{d}_p f_i \in T_p^* M$ with respect to the basis $\mathrm{d}_p x_1,\ldots, \mathrm{d}_p x_m$. Because $\omega_1(p),\ldots, \omega_m(p)$ are linearly independent, the derivative $D(\varphi \circ \psi^{-1})$ is invertible. Hence, by the implicit function theorem we have that on a possible even smaller neighbourhood U around p the map$\varphi \circ \psi^{-1}$ is invertible and a homeomorphism. ∎

Remark 1.1.22. The map φ of lemma 1.1.21 is in fact a chart for M fitting into the already existing (differentiable) atlas. This can be seen in the previous proof: the inverse function theorem states that the transition map $\varphi \circ \psi^{-1}$ has the same differential properties as $f_i \circ \psi^{-1}$.

Now that we have a chart around each point p, we can show that these charts can be combined into a translation atlas:

Theorem 1.1.23. *Let M be an m-dimensional, differentiable manifold and $\omega_1,\ldots, \omega_m$ closed 1-forms.*

Then $\{\varphi_p \mid p \in M, \omega_1(p) \wedge \cdots \wedge \omega_m(p) \neq 0\}$ is a translation atlas for $M \setminus \{p \in M \mid \omega_1(p) \wedge \cdots \wedge \omega_m(p) = 0\}$ where

$$\varphi_p \colon U_p \to \mathbb{R}^m, \qquad q \mapsto \left(\int_p^q \omega_1, \ldots, \int_p^q \omega_m \right) \tag{1.16}$$

is the chart described in lemma 1.1.21 around the point p.

Proof. Let φ_p and φ_q be two charts with $U_p \cap U_q \neq \emptyset$. We must show that the transition map $\varphi_p \circ \varphi_q^{-1}$ is locally a translation. Because this is a local property, we may consider a suitable small neighbourhood. To this end, let $r \in U_p \cap U_q$ and $U_r \subseteq U_p \cap U_q$ a connected neighbourhood of r.

For any point $x \in U_r$ and any closed 1-form ω_i, we have

$$\int_p^x \omega_i = \int_p^r \omega_i + \int_r^x \omega_i \tag{1.17}$$

and all the integrals are well-defined, i.e. independent of the choice of path in U_p connecting the end points, by the construction of the chart φ_p, namely because U_p is simply connected. Moreover, this stays true, when interpreting the last integral $\int_r^x \omega_i$ in U_r, i.e. only paths in $U_r \subseteq U_p$ connecting the end points are allowed.

Similarly, we have

$$\int_r^x \omega_i = \int_r^q \omega_i + \int_q^x \omega_i \tag{1.18}$$

where the integrals are well-defined for any path in U_q connecting the end points. Again, we can reinterpret the left hand side to only consider paths in $U_r \subseteq U_q$.

Plugging equation (1.18) in equation (1.17), we get

$$\int_p^x \omega_i = \int_p^r \omega_i + \int_r^q \omega_i + \int_q^x \omega_i = \overbrace{\underbrace{\int_p^r \omega_i}_{\text{in } U_p} + \underbrace{\int_r^q \omega_i}_{\text{in } U_q}}^{=:c(p,q,r,\omega_i)=:c_i} + \underbrace{\int_q^x \omega_i}_{\text{in } U_q} = c_i + \int_q^x \omega_i. \tag{1.19}$$

The number $c_i = c(p, q, r, \omega_i)$ is a constant independent of x and only depending on the points p, q, r and the 1-form ω_i.

At last, we verify that the transition map $\varphi_p \circ \varphi_q^{-1}$ is a translation on $U_r \subseteq U_p \cap U_q$. Let $y \in \varphi_q(U_r) \subseteq \mathbb{R}^m$, then we have for the i^{th} component of the transition map

$$(\varphi_p \circ \varphi_q^{-1})_i(y) = \int_p^{\varphi_q^{-1}(y)} \omega_i \overset{(1.19)}{=} c_i + \int_q^{\varphi_q^{-1}(y)} \omega_i = c_i + (\varphi_q)_i(\varphi_q^{-1}(y)) = c_i + y_i. \tag{1.20}$$

Thus, $\varphi_p \circ \varphi_q^{-1}$ is locally around $\varphi_q(r)$ (namely on $\varphi_q(U_r)$) a translation on $\mathbb{R}^m$ by the vector with entries $c_i = c(p, q, r, \omega_i)$. ■

Theorem 1.1.24. *The closed* 1*-forms* $\omega_1, \dots, \omega_m$ *are harmonic with respect to the Riemannian metric induced by the translation structure on* $M \setminus \{p \in M | \omega_1(p) \wedge \dots \wedge \omega_m(p) = 0\}$.

Proof. On charts constructed with the 1-forms $\omega_1, \dots, \omega_m$ the 1-form ω_i is just $\mathrm{d}x_i$, thus harmonic. ■

The above proof also shows that for a two-dimensional translation manifold $(M; \omega_1, \omega_2)$, the 1-form $\omega_1 + i\omega_2$ is a holomorphic on $M \setminus \{p \in M \mid \omega_1(p) \wedge \omega_2(p) = 0\}$. Generally we cannot expect that this also holds on the singular points as the following example shows:

Example 1.1.25. Let $M = \mathbb{R}^2$ with $\mathrm{id}\colon \mathbb{R}^2 \to \mathbb{R}^2$ as global chart. Consider the 1-forms

$$\omega_1 = (x^2 + y^2)\,\mathrm{d}x + (2xy + 1)\,\mathrm{d}y \qquad\qquad \omega_2 = \mathrm{d}y. \tag{1.21}$$

Both forms are closed. We note that all the following integrals are well-defined because ω_1 and ω_2 are closed on $\mathbb{R}^2$ and $\mathbb{R}^2$ is simply connected.

We have $(\omega_1 \wedge \omega_2)(x, y) = (x^2 + y^2)(\mathrm{d}x \wedge \mathrm{d}y) = 0$ if and only if $(x, y) = (0, 0)$. We point out that neither form is zero in $(0, 0)$, nevertheless they become linearly dependent in $(0, 0)$.

Let $M^* := M \setminus \{(0, 0)\}$. The charts around a point $p \in M^*$ as described by theorem 1.1.23 are

$$\tilde{\varphi}_p\colon U_p \to \mathbb{R}^2, \qquad q \mapsto \left(\int_p^q \omega_1, \int_p^q \omega_2\right). \tag{1.22}$$

with U_p a sufficiently small neighbourhood around p. However, to avoid unnecessary computation, we post-compose each chart φ_p with the translation

$$t_p\colon x \mapsto x + \left(\int_0^p \omega_1, \int_0^p \omega_2\right). \tag{1.23}$$

Thus, all charts look like

$$\varphi_p = t_p \circ \tilde{\varphi}_p\colon U_p \to \mathbb{R}^2, \qquad q \mapsto \left(\int_0^q \omega_1, \int_0^q \omega_2\right). \tag{1.24}$$

In particular they can be combined into a single map φ (but not necessarily a chart) on all of $\mathbb{R}^2$.

Calculating the integrals yields

$$\varphi_p\colon U_p \to \mathbb{R}^2, \qquad (x,y) \mapsto \left(\frac{1}{3}x^3 + xy^2 + y, y\right). \tag{1.25}$$

From this concrete description we can derive that φ, which is just φ_p but defined on all of $\mathbb{R}^2$, is bijective on $\mathbb{R}^2$. Thus, φ is a global chart on M^*.

Walking around the origin and using that φ is a bijective chart on M^*, we see that $(0,0)$ is a removable singularity.

By the note above $\omega_1 + i\omega_2$ is holomorphic on $M^* \subseteq \mathbb{R}^2 \cong \mathbb{C}$. However, $\omega_1 + i\omega_2$ is not holomorphic on all of M. Because if it were, then at all points p of M either both ω_1 and ω_2 must vanish at p, or they must be linearly independent at p (cf. remark 1.1.18). In this example neither is the case for $p = (0,0)$.

The explanation for this is that on M^* the charts id and φ are compatible, i.e. they define the same differential and complex structure. However, on all of M the two charts are not compatible any more because φ is not a diffeomorphism of $\mathbb{R}^2$. Therefore, the atlases $\{\mathrm{id}\}$ and $\{\varphi\}$ induce different differential and complex structures on M. Be aware that the structures of M are different but they are still isomorphic (via the map φ as map between M and M).

With respect to the structure induced by φ on M, the holomorphic 1-form $\omega_1 + i\omega_2$ on M^* can be extended to a holomorphic 1-form on all of M. However, the same form cannot be extended on M with respect to the structure induced by id.

The same re-interpretation process can also be done with any holomorphic 1-form and other charts where the singularities cover the points at which the charts are not compatible.

1.1.4 Comparison of the Definitions

The obvious question which arises is whether these definitions are equivalent when some restrictions like a finite number of polytopes are imposed.

General Case

For the general case, the definition given via a translation atlas (definition 1.1.1) and via 1-forms (definition 1.1.15) are equivalent. We have already seen how 1-forms give rise to a translation structure, for the other way around we need to realise that the local 1-forms $\mathrm{d}x_1, \dots, \mathrm{d}x_m$ induced by a chart are closed as well as linearly independent on their chart and that they can be combined to a global closed 1-form on all of M because the transition maps between charts are translations, i.e. have derivative 1.

The definition via polytopes (definition 1.1.14) is strictly weaker as for this definition singularities are always 'straight' because they are part of the codimension two skeleton of the polytopes.

Finite Case

For the definition via polytopes (definition 1.1.14) a good finiteness condition seems to be to require only finitely many polytopes to be glued. Again for the same reason as before, this definition is different than the other two as we still only have 'straight' singularities.

For a finite version of the definition via 1-forms (definition 1.1.15) we at least would require that the 1-forms can be extended to the metric completion of M and therefore that the metric completion has a structure as a (differentiable) manifold. With this requirement not all finite translation manifolds glued by polytopes yield a finite translation manifold in the 1-form sense because not all glueings yield a manifold (see example 1.1.26).

This means that the two concepts – 1-form and polytope – which coincide in dimension two, split in higher dimensions into distinct ones, which have some intersection but none is a subset of the other.

Providing a finite equivalent for the definition with a translation atlas (definition 1.1.1) is difficult because it is unclear how to imitate isolated (point) singularities. We could require that the set of singularities Σ has only finitely many connected components. However, even a single component, which can now be a line or similar, can easily accumulate by itself. A better analogue in higher dimensions would be that locally around a singular point, Σ has only a single connected component plus the requirement that the metric completion is a manifold.

The same generalisation problem arises for the points of linear independence in the 1-form definition (definition 1.1.15). In the two-dimensional case the holomorphic nature of the differential guarantees by way of the identity theorem that only finitely many singularities can exist. If the 1-form is only smooth or even only differentiable, there is no such restriction anymore. Thus, we have to impose a similar requirement as in the previous paragraph for the points of linear dependence.

All in all we can say that it is (again) a lucky low dimensional phenomenon that for a translation surface the three definitions of a finite translation surface coincide.

Example 1.1.26 (Three-dimensional L-shaped Translation Manifold).
The manifold depicted in figure 1.6 is called the *three-dimensional L-shaped* translation

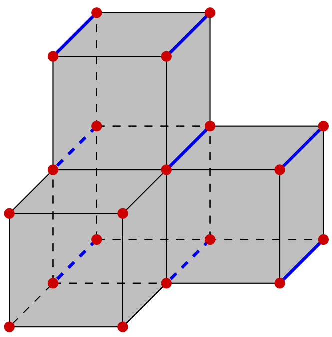

Figure 1.6: The *three-dimensional L-shaped* translation manifold. Opposite sides are glued by a translation.

manifold. It is glued out of four cubes as shown in the image and the remaining (outer) sides are glued with their opposite by a translation in x-, y-, or z-direction.

After the identification the space consists of 4 cubes, 12 sides, 6 edges (one being visualised in the figure), and 1 vertex (also visualised in the figure). The translation manifold itself is the gluing except for the edges and vertices. The metric completion is the gluing including the (glued) edges and vertices.

The Euler characteristic of the metric completion is $1 - 6 + 12 - 4 = 3$. By [Thu97, Proposition 3.2.8, p. 122] a three-dimensional gluing is a manifold if and only if the Euler characteristic is zero, thus the metric completion of this translation manifold is not a manifold.

1.2 Translation Coverings

Translation coverings are a useful tool to describe and compare translation manifolds. Basically, they are topological coverings which interact nicely with the translation structure.

Definition 1.2.1 (Translation Covering). Let $(M, \mathcal{A})$ and $(N, \mathcal{B})$ be two translation manifolds of dimension m with translation atlases $\mathcal{A}$ and $\mathcal{B}$, respectively. A map $p\colon (M, \mathcal{A}) \to (N, \mathcal{B})$ is called *translation covering* if and only if

i) p is a covering map of topological spaces,

ii) the translation structures given by $\mathcal{A}$ and the pullback of $\mathcal{B}$ along p are the same.

The second condition means that for every $x \in M$ and every neighbourhood U of x such that $p\big|_U\colon U \to p(U)$ is a homeomorphism and such that there is a chart $\varphi\colon p(U) \to \mathbb{R}^m$ in $\mathcal{B}$, the map $\varphi \circ p\colon U \to \mathbb{R}^m$ is a chart for M which is compatible with $\mathcal{A}$.

Remark 1.2.2. We can also reverse this definition: Given any (topological) covering $p\colon M \to N$ between a topological space M and a translation manifold $(N, \mathcal{B})$, then p induces a translation structure on M, namely the pullback $p^*(\mathcal{B})$ of $\mathcal{B}$, and with this structure $p\colon (M, p^*(\mathcal{B})) \to (N, \mathcal{B})$ becomes a translation covering.

Lemma 1.2.3. *Let $p\colon (M, \mathcal{A}) \to (N, \mathcal{B})$ be a translation covering of m-dimensional translation manifolds. If $U \subseteq M$ is an open set such that $p\big|_U\colon U \to p(U)$ is a homeomorphism, then $p\big|_U\colon U \to p(U)$ is a translation, i.e. for all suitable charts $\varphi \in \mathcal{A}$ and $\psi \in \mathcal{B}$ the map $\psi \circ p\big|_U \circ \varphi^{-1}$ between subsets of $\mathbb{R}^m$ is a translation $x \mapsto x + c$.*

The following lemma shows that for a translation covering, the group of topological Deck transformations is the same as the group of translation Deck transformations.

Lemma 1.2.4. *Let $p\colon (M, \mathcal{A}) \to (N, \mathcal{B})$ be translation covering between translation manifolds. Furthermore, let* $\mathrm{Deck}(M/N)$ *be the group of topological Deck transformations, i.e.*

$$\mathrm{Deck}(M/N) = \{f\colon M \to M \mid f \in \mathrm{Aut}(M), p \circ f = p\}. \tag{1.26}$$

Then every $f \in \mathrm{Deck}(M/N)$ is locally a translation.

Proof. Let $f\colon M \to N$ be a topological Deck transformation. Choose charts (φ, U) and (ψ, V) of M such that $p|_U\colon U \to p(U)$ and $p|_V\colon V \to p(V)$ are homeomorphisms and $f(U) \subseteq V$. Since $p \circ f = p$ and $p|_V$ is invertible on $p(V)$, we have $f|_U = p|_V^{-1} \circ p|_U$. Because p and hence $p|_V^{-1}$ are locally translations (cf. lemma 1.2.3), so is f. ■

$$\begin{tikzcd} U \supseteq M \arrow[rr, "f"] \arrow[dr, "p"'] & & M \subseteq V \arrow[dl, "p"] \\ & p(U) \subseteq p(V) \subseteq N & \end{tikzcd}$$

1.3 Developing Map

An essential tool for dealing with translation manifolds is the *developing map*. Basically the developing map is glued together out of charts whenever two charts overlap yielding a chart-like map on all of M. This naive construction can have some ambiguity to it so that instead of a map $D\colon M \to \mathbb{R}^m$ we get a map $D\colon \tilde{M} \to \mathbb{R}^m$ from the universal cover $\tilde{M}$ of M.

This idea and concept of a developing map not only applies to translation manifolds but can be done in the more general framework of (G, X)-manifolds.

The concept of a (G, X)-manifold is described in the literature, e.g. [Thu97, chapter 3]. Here we will repeat the definitions and statements we need and fill in some blanks.

1.3.1 (G, X)-Manifolds

A (G, X)-manifold is a generalisation of the concept that the transition maps are translations, which form a group.

Definition 1.3.1 ((G, X)-manifold). Let X be a connected manifold and G a group acting on X via homeomorphisms.

A (G, X)*-manifold* is a topological space M with an (G, X)-atlas. A (G, X)*-atlas* is a collection of maps $\varphi_i\colon U_i \to X$, called *charts*, such that

i) U_i is an open subset of M and all U_i cover M, i.e. $\bigcup_i U_i = M$,

ii) φ_i is an open embedding, i.e. homeomorphism onto an open subset of X,

iii) whenever two open sets U_i and U_j intersect, the *transition map* (also called *change of coordinates*) $\gamma_{ij} := \varphi_i \circ \varphi_j^{-1}\colon \varphi_j(U_i \cap U_j) \to \varphi_i(U_i \cap U_j)$ agrees locally with elements of G, i.e. the domain of γ_{ij} can be covered by open sets $V_k \subseteq X$ such that for each set V_k there is an element $g_k \in G$ with $\gamma_{ij}|_{V_k} = g_k|_{V_k}$.

Example 1.3.2. a) If $X = \mathbb{R}^m$ and G is the group of isometries of the Euclidean space, a (G, X)-manifold is called *Euclidean* or *flat manifold*.

b) If $X = S^m$ is the sphere and $G = O(m+1)$ is the orthogonal group acting on S^m, a (G, S^m)-manifold is called *spherical* or *elliptic manifold*.

c) If $X = \mathbb{H}^m$ is the hyperbolic space and G the group of hyperbolic isometries on $\mathbb{H}^m$, a $(G, \mathbb{H}^m)$-manifold is called *hyperbolic manifold*.

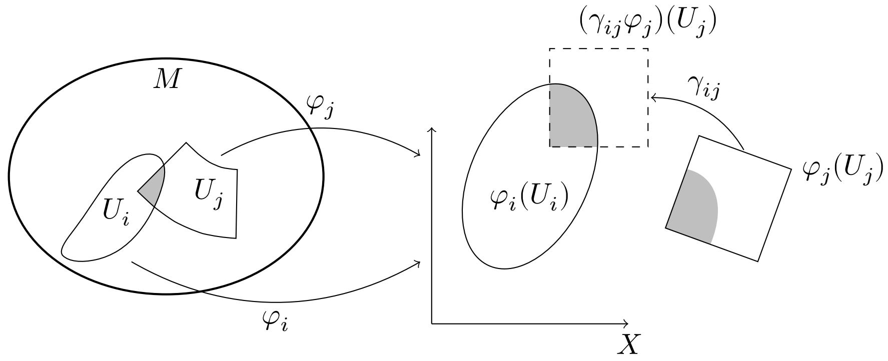

Figure 1.7: Adjusting the image $\varphi_j(U_j)$ by an element $\gamma_{ij} \in G$ so that it matches up with $\varphi_i(U_i)$. After the adjusting we can combine the charts to get a well-defined map $U_i \cup U_j \to X$.

d) If $X = \mathbb{R}^m$ and G is the group of affine maps on $\mathbb{R}^m$, a $(G, \mathbb{R}^m)$-manifold is called *affine manifold*.

e) If $X = \mathbb{R}^m$ and G is the group of translations on $\mathbb{R}^m$, a $(G, \mathbb{R}^m)$-manifold is called *translation manifold*, cf. definition 1.1.1.

1.3.2 Developing Map

To be able to construct the developing map, we need a slightly stronger assumption regarding the action of our group G on X.

Definition 1.3.3 (Analytic Action). A group G acts *analytically* on X when the following condition is fulfilled: If $g, h \in G$ and there is an open set $U \subseteq X$ on which g and h coincide, i.e. $g\big|_U = h\big|_U$, then $g = h$ on each connected component of X intersecting U.

The basic idea of the developing map is to extend a chart $\varphi\colon U \to X$ by combining it with another chart which has some overlap with U. We can achieve this by using an element of G to make the codomains of the charts match-up.

However, we have to be careful to get a well-defined map since combining with different charts in different orders might result in different extensions of the initial chart. It transpires that we can resolve this ambiguity by extending the chart along paths and then ascending to the universal cover of the (G, X)-manifold.

Here is the construction in detail:

Combining two charts

We first describe the gluing construction for two charts: Let M be a (G, X)-manifold with X a connected manifold and G a group acting analytically on X via homeomorphisms.

Consider two charts $\varphi_i: U_i \to X$ and $\varphi_j: U_j \to X$, which overlap, i.e. $U_i \cap U_j \neq \emptyset$. With regard to the definition of a (G, X)-manifold, the transition map

$$\tilde{\gamma}_{ij} := \varphi_i \circ \varphi_j^{-1}: \varphi_j(U_i \cap U_j) \to \varphi_i(U_i \cap U_j) \tag{1.27}$$

locally coincides with some elements $g_k \in G$. Because G acts analytically, all the g_k are equal on the individual connected components of $\varphi_j(U_i \cap U_j)$. In other words $\tilde{\gamma}_{ij}$ is represented by a single group element on each connected component. Phrased differently, we have a locally constant map

$$\gamma_{ij}: U_i \cap U_j \to G \tag{1.28}$$

with $\gamma_{ij}(x) = g$ where g is the element representing $\tilde{\gamma}_{ij}$ on the connected component of $\varphi_j(x)$ in $\varphi_j(U_i \cap U_j)$. Here we also used the homeomorphism φ_j to identify $U_i \cap U_j$ and $\varphi_j(U_i \cap U_j)$ to ease the burden of notation later on.

This constructed γ_{ij} has the following useful property: If we modify φ_j with it, we get a chart which we can attach to φ_i without requiring any further modifications (cf. figure 1.7). More precisely: If $x \in U_i \cap U_j$, then the charts $\gamma_{ij}(x) \circ \varphi_j$ and φ_i coincide around x. In fact, if $U_i \cap U_j$ is connected, then $\gamma_{ij}(x)$ is independent of the choice of x and $\gamma_{ij}(x) \circ \varphi_j$ and φ_i coincide on the whole of $U_i \cap U_j$. In this case, we can extend φ_i onto $U_i \cup U_j$ using $\gamma_{ij}(x) \circ \varphi_j$ resulting in a map $U_i \cup U_j \to X$. This extended map shares many properties with φ_i and φ_j but is not necessarily a chart. For example injectivity might get lost.

We can repeat this process extending φ_i further and further, however, we might run into trouble when extending the map too far (cf. figure 1.8). The way to avoid this is to pass to the universal cover of M.

Extending along paths

Fix a basepoint $x_0 \in M$ and an initial chart $\varphi_0: U_0 \to M$ around that basepoint. We denote by $\pi: \tilde{M} \to M$ the universal covering map of M and identify the universal cover $\tilde{M}$ with the space of homotopy classes of paths starting at x_0. Let $[\alpha] \in \tilde{M}$ be a point in the universal cover represented by a path $\alpha: [0, 1] \to M$. We will extend φ_0 along α.

To this end, we subdivide α at times $t_0 = 0$, t_1,..., $t_n = 1$ such that each subpath $\alpha\big|_{[t_i, t_{i+1}]}$ is contained in a chart $\varphi_i: U_i \to X$ for $0 \leq i \leq n - 1$ (φ_0 being the initial chart chosen above). As we go along α, we successively adjust the chart φ_i at each $\alpha(t_i)$ so that it can be used to extend the previously adjusted φ_{i-1} in a neighbourhood of $\alpha(t_i) \in U_{i-1} \cap U_i$. At step i the newly adjusted chart is

$$\gamma_{01}(\alpha(t_1)) \circ \gamma_{12}(\alpha(t_2)) \circ \cdots \circ \gamma_{i-1,i}(\alpha(t_i)) \circ \varphi_i. \tag{1.29}$$

Note that combining all of these adjusted charts into a single map might not yield a well-defined map on $\bigcup_i U_i$.

Lemma 1.3.4. *Denote by φ_0^α the adjusted map of the last step which contains $\alpha(1)$, i.e.*

$$\varphi_0^\alpha := \gamma_{01}(\alpha(t_1)) \circ \gamma_{12}(\alpha(t_2)) \circ \cdots \circ \gamma_{n-2,n-1}(\alpha(t_{n-1})) \circ \varphi_{n-1}. \tag{1.30}$$

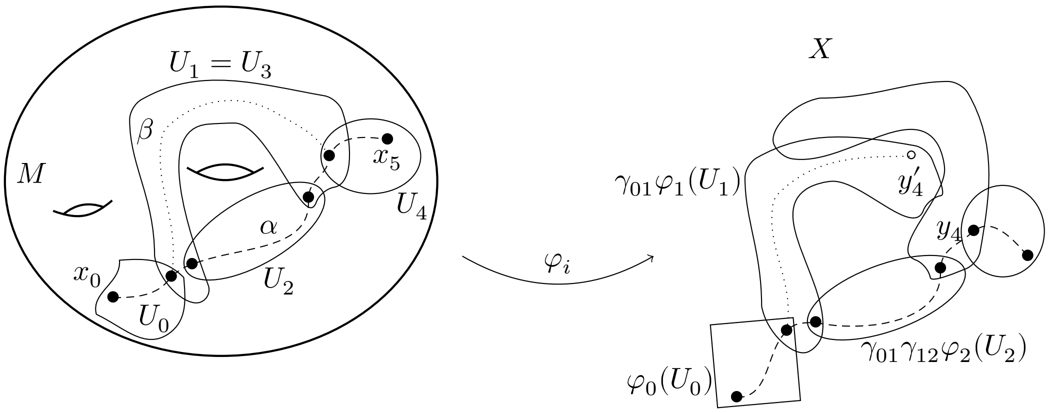

Figure 1.8: Adjusting the charts $\varphi_i: U_i \to \varphi_i(U_i)$ along the path α. Here $x_i := \alpha(t_i)$. Note that the same chart might be used multiple times – in this figure we have $\varphi_1 = \varphi_3$ – and also might have non-connected intersections – here U_1 and U_2. If there are multiple connected components, then the connected component with the current x_i is the one to consider.
This figure also illustrates that different paths can lead to different extensions: Here the image of x_4 can either be y_4 or y_4' depending on whether we adjust along path α or path β.

Then the germ of φ_0^α at $\alpha(1)$ is independent of the choices of charts φ_i, $1 \le i \le n-1$ and the subdivision of $[0,1]$, and only depends on the choice of the homotopy class of α, the basepoint x_0 and the initial chart φ_0. In other words, if $[\alpha] = [\beta]$, then $\varphi_0^\alpha = \varphi_0^\beta$ in a neighbourhood of the end point $\alpha(1) = \beta(1)$.

Proof. We prove the statement in three steps:

1. the independence of the choice of the charts,
2. the independence of the choice of the subdivision, and
3. the independence of the choice of the representative of the homotopy class.

Independence of the choice of charts

Let $t_0, \dots, t_n$ be a subdivision of the interval $[0,1]$ and $\alpha: [0,1] \to M$ a representative of the homotopy class. Let $\varphi_i: U_i \to X$ be charts containing the path segment $\alpha\big|_{[t_i, t_{i+1}]}$ for $0 \le i \le n-1$. Further, let $\psi_i: V_i \to X$ be a second set of charts containing the path segment $\alpha\big|_{[t_i, t_{i+1}]}$ for $0 \le i \le n-1$ and with the initial chart $\psi_0 = \varphi_0$.

We proceed by induction. For $n = 0$ we have $\varphi_0^\alpha = \varphi_0 = \psi_0$ and there is nothing to prove. Let the statement be true for subdivisions consisting of up to n intervals, i.e. we have

$$\begin{aligned} &\gamma_{01}(\alpha(t_1)) \circ \gamma_{12}(\alpha(t_2)) \circ \dots \circ \gamma_{n-2,n-1}(\alpha(t_{n-1})) \circ \varphi_{n-1} \\ = \;&\delta_{01}(\alpha(t_1)) \circ \delta_{12}(\alpha(t_2)) \circ \dots \circ \delta_{n-2,n-1}(\alpha(t_{n-1})) \circ \psi_{n-1} \end{aligned} \tag{1.31}$$

on a neighbourhood of $\alpha(t_n)$, where $\gamma_{ij}(\alpha(t_j)), \delta_{ij}(\alpha(t_j)) \in G$ are the adjustments needed for φ_i and ψ_i, respectively. For a subdivision of $n+1$ intervals we have to show that

$$\begin{aligned} &\gamma_{01}(\alpha(t_1)) \circ \gamma_{12}(\alpha(t_2)) \circ \cdots \circ \gamma_{n-1,n}(\alpha(t_n)) \circ \varphi_n \\ &= \delta_{01}(\alpha(t_1)) \circ \delta_{12}(\alpha(t_2)) \circ \cdots \circ \delta_{n-1,n}(\alpha(t_n)) \circ \psi_n \end{aligned} \tag{1.32}$$

on a neighbourhood of $\alpha(t_{n+1})$.

By the construction of $\gamma_{n-1,n}(\alpha(t_n))$ we have

$$\begin{aligned} &\gamma_{01}(\alpha(t_1)) \circ \gamma_{12}(\alpha(t_2)) \circ \cdots \circ \gamma_{n-2,n-1}(\alpha(t_{n-1})) \circ \varphi_{n-1} \\ &= \gamma_{01}(\alpha(t_1)) \circ \gamma_{12}(\alpha(t_2)) \circ \cdots \circ \gamma_{n-1,n}(\alpha(t_n)) \circ \varphi_n \end{aligned} \tag{1.33}$$

on a neighbourhood of $\alpha(t_n)$ and similarly for ψ_i with δ_{ij}. Therefore, equation (1.32) is true on a neighbourhood of $\alpha(t_n)$. To extend equation (1.32) to a neighbourhood of $\alpha(t_{n+1})$ note that $\alpha(t_n)$ and $\alpha(t_{n+1})$ are connected by the path $\alpha\big|_{[t_n,t_{n+1}]}$ and thus lie in the same connected component of $U_n \cap V_n$. Because G acts analytically being equal on an open set extends to the whole connected component, in particular equation (1.32) holds for a neighbourhood of $\alpha(t_{n+1})$.

Independence of the choice of the subdivision

Let $\alpha\colon [0,1] \to M$ be a representative of the homotopy class. Let $t_0, \dots, t_m$ be a subdivision of the interval $[0,1]$ with suitable charts $\varphi_i\colon U_i \to X$ containing the path segment $\alpha\big|_{[t_i,t_{i+1}]}$ for $0 \le i \le m-1$. Further let $s_0, \dots, s_n$ be a another subdivision of the interval $[0,1]$ with suitable charts $\psi_i\colon V_i \to X$ containing the path segment $\alpha\big|_{[s_i,s_{i+1}]}$ for $0 \le i \le n-1$ and with the initial chart $\psi_0 = \varphi_0$.

Consider the refined subdivision $\{t_1, \dots, t_m\} \cup \{s_0, \dots, s_n\}$ of $[0,1]$. The charts φ_i are suitable for this subdivision by repeating a chart (possibly multiple times) when an s_j occurs. The transition map between a repeated map and itself is obviously the identity. Therefore, φ_0^α is the same whether we use $\{t_0, \dots, t_m\}$ or $\{t_1, \dots, t_m\} \cup \{s_0, \dots, s_n\}$ with the charts φ_i.

The same is true when using the charts ψ_j. As we have already established that the choice of charts does not matter for the subdivision $\{t_1, \dots, t_m\} \cup \{s_0, \dots, s_n\}$, we can conclude that the germ of φ_0^α at $\alpha(1)$ is independent of the choice of the subdivision of the path.

Independence of the representative of the homotopy class

Let $\alpha\colon [0,1] \to M$ and $\beta\colon [0,1] \to M$ be two paths with $\alpha(0) = \beta(0)$ and $\alpha(1) = \beta(1)$ and let $H\colon [0,1] \times [0,1] \to M$ be a homotopy between them which fixes their endpoints. We need to show that $\varphi_0^\alpha = \varphi_0^\beta$ on a neighbourhood around the end point.

If $H([0,1]^2) \subseteq U$, for some chart $\varphi\colon U \to X$, then clearly $\varphi_0^\alpha = \varphi_0^\beta = \varphi$ on U and in particular on a neighbourhood of $\alpha(1) = \beta(1)$. This is also true if the homotopy only happens within a segment $[t_{i-1}, t_i]$ of the subdivision (i.e. $\alpha(t) = \beta(t)$ for $t \notin [t_{i-1}, t_i]$) as only $\gamma_{i-1,i}(\alpha(t_i)) = \gamma_{i-1,i}(\beta(t_i))$ is needed from that step.

Otherwise, because $[0,1]^2$ is compact, we can partition $[0,1]^2$ into finitely many rectangles $[s_i, s_{i+1}] \times [t_j, t_{j+1}]$ with $H([s_i, s_{i+1}] \times [t_j, t_{j+1}]) \subseteq U_{ij}$ for some chart $\varphi_{ij}\colon U_{ij} \to X$.

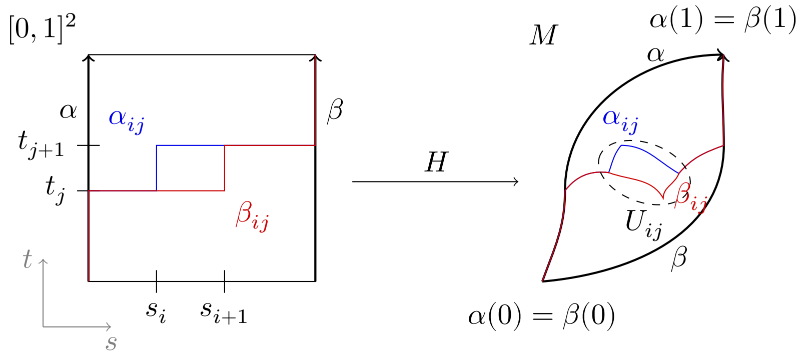

Figure 1.9: The homotopy between α_{ij} and β_{ij} only happens inside the chart domain U_{ij}.

The curves (cf. figure 1.9)

$$\alpha_{ij} = \alpha\big|_{[0,t_j]} H(\,\cdot\,,t_j)\big|_{[0,s_i]} \underbrace{H(s_i,\,\cdot\,)\big|_{[t_j,t_{j+1}]} H(\,\cdot\,,t_{j+1})\big|_{[s_i,s_{i+1}]}}_{\text{homotopy happens here}} H(\,\cdot\,,t_{j+1})\big|_{[s_{i+1},1]} \beta\big|_{[t_{j+1},1]}, \tag{1.34a}$$

$$\beta_{ij} = \alpha\big|_{[0,t_j]} H(\,\cdot\,,t_j)\big|_{[0,s_i]} \overbrace{H(\,\cdot\,,t_j)\big|_{[s_i,s_{i+1}]} H(s_{i+1},\,\cdot\,)\big|_{[t_j,t_{j+1}]}}^{\text{homotopy happens here}} H(\,\cdot\,,t_{j+1})\big|_{[s_{i+1},1]} \beta\big|_{[t_{j+1},1]} \tag{1.34b}$$

are homotopic, have the same fixed start- and endpoints and the homotopy only happens inside of U_{ij}. Thus, $\varphi_0^{\alpha_{ij}}$ and $\varphi_0^{\beta_{ij}}$ coincide around the end point $\alpha(1) = \alpha_{ij}(1) = \beta_{ij}(1) = \beta(1)$. The curves are chosen such that $\beta_{ij} = \alpha_{i+1,j}$ and $\alpha_{0,j-1} = \beta_{i_{\max},j}$. Therefore, we can combine all these small steps leaving the germ at the end point invariant and conclude that φ_0^α and φ_0^β are equal around $\alpha(1) = \beta(1)$. ∎

Remark 1.3.5. It is not true that extending only works on simply connected domains. See example 1.3.9 below for a manifold where the extension of charts is well-defined despite not being simply connected.

Defining the Developing Map

We have seen that an adjusted map around a point only depends on the homotopy class of the path taken. This allows us to define the developing map on the universal cover, which again we regard as the space of homotopy classes of paths.

Definition 1.3.6 (Developing Map). For an initial chart $\varphi_0 \colon U \to X$ and a basepoint $x_0 \in U$, the *developing map* of a (G,X)-manifold M is the map

$$D \colon \tilde{M} \to X, \qquad [\alpha] \mapsto \varphi_0^\alpha(\alpha(1)) \tag{1.35}$$

from the universal cover $\tilde{M}$ to X. In other words, we go along the path α adjusting and extending the charts as we go, taking the value at its endpoint as the value of D.

Because the germ of φ_0^α at $\alpha(1)$ only depends on the initial chart and the basepoint, we can describe the developing map locally around $[\alpha] \in \tilde{M}$ by

$$D = \varphi_0^\alpha \circ \pi, \tag{1.36}$$

where $\pi\colon \tilde{M} \to M$ is the universal covering map. From this description we immediately see that the developing map has locally the same properties as the charts of M.

Remark 1.3.7. Choosing different initial data, i.e. basepoint and initial chart, alters the developing map D by post-composition with an element of G.

The notation of the developing map does not reflect this dependency. Most of the time these are implicit or not relevant.

Remark 1.3.8. Whenever M is simply connected, we can identify M and its universal cover $\tilde{M}$ so that we can regard the developing map to be defined on M. This yields a sort of global chart for M.

Example 1.3.9. If the (G, X)-structure is 'rigid' enough, then the developing map on the universal cover $\tilde{M}$ might descends to a well-defined map on M via the covering map $\pi\colon \tilde{M} \to M$. This is often the case for translation manifolds.

i) The punctured disc $M = \dot{B}(0, 1)$ in $\mathbb{R}^2$ with the inclusion $\varphi_0\colon \dot{B}(0, 1) \to \mathbb{R}^2$ as global chart is a translation manifold, i.e. a $(G, \mathbb{R}^2)$-manifold where G is the group of translations. It is not simply connected and taking φ_0 as initial chart yields a developing map D from the universal cover, which in this case can be thought of as a helix.

   Because adjusting charts only uses translations, it does not change angles and lengths. Therefore, the developing map is the same as the covering map and it descends to a well-defined map on $\dot{B}(0, 1)$, which in this case is φ_0 and the induced map is a proper chart for M.

ii) Let M be the two sheeted connected cover over $\dot{B}(0, 1)$ considered as a translation manifold similar to above. In the same way as before, the developing map descends to a well-defined map on M. This time the induced map is the two-to-one covering map from M to $\dot{B}(0, 1)$ which is not a chart for M.

1.3.3 The Developing Map of a Translation Manifold

As mentioned before, the developing map has locally the same properties as the charts it is made of. This means that for a translation manifold, the developing map is a local isometry and the transition map between the developing map and a chart is a translation. When doing calculations it is essential to know how local the 'local isometry' is and whether we have a uniform bound on it. The following lemma explores this:

Lemma 1.3.10 (D is an isometry on balls). *Let M be an m-dimensional translation manifold, $\Sigma = \overline{M} \setminus M$ the singularities, $x \in M$ and $r > 0$. Let the ball $B(x, r)$ in $\overline{M}$*

be such that $B(x,r) \subseteq M$, i.e. it does not contain any singularities, and such that the developing map descends to a map $D\colon B(x,r) \to \mathbb{R}^m$.

Then D is a homeomorphism from $B(x,r) \subseteq M$ to $B(D(x),r) \subseteq \mathbb{R}^m$ and an isometry from $B(x,\frac{r}{2}) \subseteq M$ to $B(D(x),\frac{r}{2}) \subseteq \mathbb{R}^m$.

The important prerequisite is that the r-ball does not contain any singularities. We would hope that we can avoid reducing the radius for the isometry statement, however, remark 2.1.2 gives an example showing that the developing map is not an isometry on the full ball; in fact the above bound is sharp.

Because charts are the building blocks of the developing map, we have the following corollary:

Corollary 1.3.11. *Let M be an m-dimensional translation manifold, $\varphi\colon U \to \mathbb{R}^m$ a chart, $x \in U$ and $r > 0$ such that the ball $B(x,r)$ in $\overline{M}$ is contained completely in U, i.e. contains no singularities.*

Then φ is a homeomorphism from $B(x,r) \subseteq U$ to $B(\varphi(x),r) \subseteq \mathbb{R}^m$ and an isometry from $B(x,\frac{r}{2}) \subseteq U$ to $B(\varphi(x),\frac{r}{2}) \subseteq \mathbb{R}^m$.

The proof uses a variant of the Hopf-Rinow Theorem. The classic Hopf-Rinow Theorem cannot be used as our metric space in question $B(x,r)$ fails to be geodesically complete. However, there is a generalised version for length metric spaces (which are also called interior metric spaces), i.e. metric spaces where the metric is given by the infimum of path lengths. However, the version for length metric spaces as for example stated by Bridson-Haefliger[BH99, Proposition 3.7, p. 35] cannot be used either because M is a length metric space but $B(x,r)$ is not necessarily one (because geodesics might (and often do) leave the ball). The variant of Hopf-Rinow we use is stated below and is from W. Ballmann[Bal95, Theorem 2.4, pp. 12] and originates in Cohn-Vossen[Coh35; Coh36].

Theorem 1.3.12 (generalised Hopf-Rinow (local version), [Bal95]). *Let X be a locally compact and interior [metric space], and let $x \in X$ and $R > 0$. Then the following are equivalent:*

(i) any geodesic $\gamma\colon [0,1[\to X$ with $\gamma(0) = x$ and $\ell(\gamma) < R$ can be extended to the closed interval $[0,1]$;

(ii) any minimizing geodesic $\gamma\colon [0,1[\to X$ with $\gamma(0) = x$ and $\ell(\gamma) < R$ can be extended to the closed interval $[0,1]$;

(iii) $\bar{B}(x,r)$ is compact for $0 \leq r < R$.

Each of these implies that for any pair y, z of points in $B(x,R)$ with $d(x,y)+d(x,z) < R$ there is a minimizing geodesic from y to z [and that geodesic is contained in $B(x,R)$].

Remark 1.3.13. Since in our case $B(x,r) \subseteq M$ is very close to being a complete Riemannian manifold, even the proof of the classical Hopf-Rinow Theorem as given in [Boo86, Lemma 7.8, p. 347] goes through although the prerequisites required in the statement are not fulfilled.

Proof (of lemma 1.3.10). Before we start note that for a length metric space the closed ball coincides with the closure of the open ball: $\bar{B}(x,r) = \overline{B(x,r)}$.

First, we show that all balls $\bar{B}(x,r')$ are compact for $0 \le r' < r$. Given any geodesic $\gamma\colon [0,1[\to B(x,r)$ with $\gamma(0) = x$ and $\ell(\gamma) < r$, the geodesic γ can be extended to $[0,1]$ since the metric completion is the closure and is contained in M: $\overline{B(x,\ell(\gamma))} \subseteq B(x,r) \subseteq M$. Thus, by theorem 1.3.12 (the generalised Theorem of Hopf-Rinow (local version)) the closed balls $\bar{B}(x,r')$ for $0 \le r' < r$ are all compact.

Second, there is a unique geodesic from x to $y \in B(x,r)$. The existence of the geodesic is given by theorem 1.3.12 (the generalised Theorem of Hopf-Rinow (local version)). For the uniqueness, recall that any geodesic in M is straight – or in other words a Euclidean line segment – because M is a translation manifold. If there are two different geodesics from x to y, then the geodesics must start in different directions at x, otherwise they would be the same as they have the same length. Since x is a regular point, i.e. there is $0 < \varepsilon \le r$ such that $B(x,\varepsilon)$ is isometric to $B(D(x),\varepsilon)$ via D, we can identify this direction with a direction in $\mathbb{R}^m$. Now D maps geodesics to straight line segments. If the geodesics start in different directions, then their images under D start in different directions and yield distinct end points in $\mathbb{R}^m$. Thus, there cannot be two different geodesics from x to y.

Third, we show that the developing map D is a homeomorphism between $B(x,r) \subseteq M$ and $B(D(x),r) \subseteq \mathbb{R}^m$. According to the previous argument, we can identify any point $y \in B(x,r)$ with its geodesic connecting it with x. This geodesic is uniquely specified by its length and the direction it starts in x. Since x is a regular point, we can identify any direction via D with a direction in $\mathbb{R}^m$. Since a direction and length characterise a point in a ball in $\mathbb{R}^m$, we have identified $B(x,r) \subseteq M$ with $B(D(x),r) \subseteq \mathbb{R}^m$ and also vice versa. Moreover, this identification is exactly the developing map D because the process described is the same as developing along the geodesic from x to y. Since D is a local homeomorphism and is bijective on $B(x,r)$, we have that D is a homeomorphism from $B(x,r) \subseteq M$ to $B(D(x),r) \subseteq \mathbb{R}^m$.

Lastly, D is an isometry on $B(x,r/2)$. Let $y,y' \in B(x,r/2)$. In the first paragraph we have shown property (i) of theorem 1.3.12 (the generalised Theorem of Hopf-Rinow (local version)) and hence can apply its conclusion. Thus, we find a geodesic between y and y', which is contained in $B(x,r)$. Because D maps geodesics to straight lines of the same length, we have $d(y,y') = d(D(x),D(y))$ and D is an isometry on $B(x,r/2)$. In particular, $B(x,r/2)$ (with the induced metric of M) is a length metric space of its own. ∎

Remark 1.3.14. The above statement is also true for an open set $U \subseteq M$ which is geodesically convex, that is for every two points y, y' in U the geodesic between these exists and lies in U. The set on which it is an isometry is obtained by shrinking the image of U in $\mathbb{R}^m$ by a factor of 2 (with any image point of U as centre) and than taking the pre-image under D.

2 Singularities

In this chapter we investigate singularities of translation manifolds. Singularities are the points which are outside of the domain where we have the translation structure, i.e. outside of where it looks like the Euclidean space. At such singular points a wide variety of things can happen, from simple things like being a border point of the space, to tame things like being a conic singularity with an angle of $2\pi k$ ($k \in \mathbb{Z}$), to crazy things like a wild singularity on the Chamanara surface[Cha04].

Optimally a singularity is removable. This means that the singularity is no 'real' singularity but merely artificially introduced. So we can add the singularity to the translation manifold and still have a translation manifold. We are particularly interested in understanding under what conditions singularities are removable.

This chapter consists of four parts. The first part is very short and introduces the core theorem (theorem 2.1.1) on which the rest builds. The idea is to use a covering to test whether singularities can be removed.

Applying this theorem to the situation where we have either a tame isolated singularity or an isolated singular point which has a neighbourhood which is a manifold, is done in the second part. Most work is put into proving the existing of the covering in this case. The result is summarised in theorem 2.2.8.

After that the third part generalises the notion of singularity to also take *shadows* of singularities (cf. definitions 2.3.2 and 2.3.7) into account. They are the main obstacles to the developing map having the covering property. Using these concept allows us to prove the existence of a covering map in more situations, thus making theorem 2.1.1 easier applicable. The main results of this part are corollary 2.3.19 and theorem 2.3.21.

The fourth part takes a closer look at the situation in dimension 2 and 3. In these low dimensions some prerequisites of previous theorems are always fulfilled so that the resulting statements become slightly stronger. The statements are summarized in theorem 2.4.6.

2.1 Removing Singularities using a Covering

We have a translation manifold M of dimension m with its singularities $\Sigma = \overline{M} \setminus M$. For a singularity to be removable we must be able to extend the translation structure of M to the singularity. Since a translation structure implies also a metric structure and a manifold structure, we must also be able to extend those, which gives a useful check for when a singularity is not removable.

We have seen that singularities are not necessarily isolated points but can be bended curves, wobbly surfaces, … This makes it more challenging because in contrast to the surface case, where singularities are only (isolated) points, finding a nice, canonical, isometric neighbourhood to classify singularities in tame and wild is impossible. Therefore, we will use translation coverings to compare singularities with each other.

We motivate this with the following observation: Assume that Σ is a set of removable singularities and we have a global chart $\varphi\colon M \to \varphi(M) \subseteq \mathbb{R}^m$. We can then identify M with a subset of $\mathbb{R}^m$. Moreover, the map φ is a covering map of degree 1 and the singularities of $\varphi(M) \subseteq \mathbb{R}^m$, which are removable in $\mathbb{R}^m$, are obviously removable in M and vice versa. We can generalise this observation to arbitrary coverings and also to non-removable singularities:

Theorem 2.1.1. *Let $(M, \mathcal{A})$ and $(N, \mathcal{B})$ be translation manifolds. Let $U \subseteq M$ be an open set. If*

(i) U is connected,

(ii) there exists a (surjective) translation covering $p\colon (U, \mathcal{A}|_U) \to (N, \mathcal{B})$, and

(iii) N is connected and simply connected,

then $p\colon U \to N$ is a homeomorphism and local isometry.

In particular the 'type' of a singularity in N is the same as the 'type' of the corresponding singularity in U. Moreover, all singularities which are removable in N are removable in U (and hence in M).

The crucial part of this theorem is the existence of the translation covering p. This is usually the hardest precondition to meet and can be quite difficult to prove. Note that for a finite translation surface the existence of a (finite) covering of the punctured ball around a singularity is backed directly into the definition of a finite translation surface.

Remark 2.1.2. We cannot expect p to be a global isometry as two points of U can be close in M but have a long path distance inside of U. For example, consider the torus $M = \mathbb{R}^2/\mathbb{Z}^2$, $U =]0.1, 0.9[^2 \subseteq M$ and $N =]0.1, 0.9[^2 \subseteq \mathbb{R}^2$, where N is endowed with the usual Euclidean metric and U has the induced metric from the torus.

The distance of the points $x = (0.2, 0.2)$ and $y = (0.2, 0.8)$ on M is 0.4 with the connecting geodesic γ being a path partially outside of U, see figure 2.1, their images in N, however, the have distance 0.6. This technicality can usually be avoided by making U smaller. Then p becomes a (global) isometry between U and N.

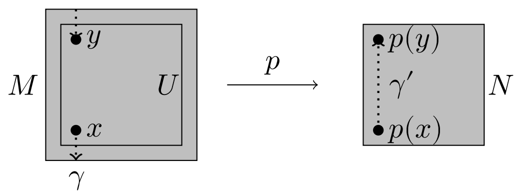

Figure 2.1: Example of a covering p which increases the distance between some points.

Proof (of Theorem 2.1.1). Since N is simply connected, it is its own universal cover. Regarding the property of the universal cover, there exists a covering $q\colon N \to U$, which is also surjective as U is connected, with $p \circ q = \mathrm{id}$. Hence, q is injective. As a covering q is surjective and thus bijective. Therefore, $p = q^{-1}$ and p and q are homeomorphisms which are locally translations. Thus, p is a homeomorphism that is locally an isometry. ■

$$\begin{tikzcd} N \arrow[d, "q"] \arrow[dd, bend right=60, "\mathrm{id}"'] \\ U \arrow[d, "p"] \\ N \end{tikzcd}$$

2.2 Removing Isolated Singularities

The simplest scenario for a singularity we can think of is a single singular point with no other singularities nearby. This situation is quite well-known in the two-dimensional case as nearly all singularities are isolated points.

From the theory of translation surfaces we also know that we can classify singularities broadly into two groups: *tame* singularities and *wild* singularities. The former is well understood while the latter still amazes.

In the first section we discuss the generalisation of tame isolated singularities and prove that they are always removable in higher dimension.

In the second section we take a look at a situation which does not fall under the tame regime: a possibly wild singularity but with a neighbourhood which is a manifold. Although a priori not tame, we are able to show that these singularities are also removable.

Before we start looking into these two situation, let us give a formal definition what an isolated singularity is:

Definition 2.2.1 (Isolated Singularity). Let M be a translation manifold. We call a singularity $\sigma \in \overline{M} \setminus M$ *isolated* iff there is an open neighbourhood $U \subseteq \overline{M}$ of σ such that $U \setminus \{\sigma\} \subseteq M$.

In other words, we have a neighbourhood of σ whose only singular point is σ itself.

2.2.1 Removing Tame Isolated Singularities

For a translation surface the singularities can be categorised into two groups: *tame singularities* and *wild singularities.* The former are completely understood as well as classified and singularities occurring on a finite translation surface are always of this type. Wild singularities, on the other hand, can only be observed on infinite translation surfaces. A precise definition is:

Definition 2.2.2 (Singularities of a Translation Surface). For a translation surface M a singularity σ is called *tame* iff it admits a neighbourhood which is a branched translation covering of a disc. That is, there exists $\varepsilon > 0$ such that $M \supseteq \dot{B}(\sigma, \varepsilon) \to \dot{B}(0, \varepsilon) \subseteq \mathbb{R}^2$ is a translation covering. These tame singularities can be fully classified:

(i) Either it is a *cone angle singularity of multiplicity k*, which means that it has an angle of $2\pi k$ for some $k \in \mathbb{N}$ and the above covering is a k-cyclic translation covering. These are the only singularities which can be found on finite translation surfaces.

(ii) Or it is an *infinite angle singularity* also called *cone angle singularity of multiplicity* ∞, which means that the above covering is an infinite cyclic translation covering.

A singularity σ is called *wild* iff it is not tame.

We can generalise this definition of a tame singularity to translation manifolds in the following way:

Definition 2.2.3 (Tame Isolated Singularity). Let M be a translation manifold of dimension m and let $\sigma \in \overline{M} \setminus M$ be an isolated singularity. We call σ *tame* iff σ has an open neighbourhood $U \subseteq \overline{M}$ such that there is $\varepsilon > 0$ and a translation covering $U \setminus \{\sigma\} \to \dot{B}(0,\varepsilon) = B(0,\varepsilon) \setminus \{0\} \subseteq \mathbb{R}^m$ to the open punctured Euclidean ball.

With this definition we can immediately show that tame isolated singularities are always removable starting from dimension 3; or in other words: *(real) tame isolated singularities do not exist in higher dimensions.*

Theorem 2.2.4. *Let $(M, \mathcal{A})$ be a translation manifold of dimension $m \geq 3$. If $\sigma \in \overline{M} \setminus M$ is a tame isolated singularity, then σ is removable. That is there is a neighbourhood $U \subseteq \overline{M}$ of σ which is isometric to the Euclidean ball $B(0,\varepsilon) \subseteq \mathbb{R}^m$ for some $\varepsilon > 0$.*

Proof. For dimension greater or equal than three the punctured ball is simply connected, i.e. $\pi_1(\dot{B}(0,\varepsilon)) = \{0\}$. Since coverings are classified by the subgroups of the fundamental group $\pi_1(\dot{B}(0,\varepsilon))$, there is only one covering of the punctured ball: the identity. If σ is a tame isolated singularity, then its corresponding covering is the identity and by theorem 2.1.1 σ is removable because 0 is a removable singularity in $\dot{B}(0,\varepsilon) \subseteq \mathbb{R}^m$. ∎

2.2.2 Removing Isolated Singularities on a Manifold

From the surface world we also know that some singular points are wilder than others.

Thus, a good starting point to use the previous results on new ground is to look at a translation manifold of dimension greater than two which has an isolated singularity and the neighbourhood of that singularity is a manifold. The last condition should rule out wild behaviour which we already observe in two-dimensions. We could call such singularities *almost tame* singularities.

The main goal of this section is to prove that any such singularity is removable (thus tame). In other words: *almost tame isolated singularities do not exist in higher dimensions.*

The proof as presented here relies heavily on the symmetry of the situation in particular of the neighbourhood around the singularity and is therefore not easily adaptable to non-point singularities. Nevertheless it lays the groundwork for the more general concepts and techniques outlined in the next sections.

Definition 2.2.5. Let M be a translation manifold and $\sigma \in \overline{M} \setminus M$ a singularity. We say $\overline{M}$ *is a manifold around* σ iff there exists an open neighbourhood $U \subseteq \overline{M}$ of σ such that $M \cup U$ admits the structure of a topological manifold which extends the manifold M. In other words, there exists a chart for U and this chart is compatible with the charts of M as a topological manifold, i.e. changes of coordinates between these charts are homeomorphisms (but not necessarily translations). Note that U is allowed to contain additional singularities other than σ.

Sometimes we require additional properties from the manifold $M \cup U$, e.g. being a smooth manifold, then the changes of coordinates must satisfy these additional requirements, too.

Proposition 2.2.6. *Given a translation manifold M of dimension $m \geq 3$, let $\sigma \in \overline{M} \setminus M$ be an isolated singularity. If the metric completion $\overline{M}$ is a manifold around σ, then there exists an $\varepsilon > 0$ such that*

(i) the developing map $D\colon M \supseteq \dot{B}(\sigma, \varepsilon) \to \mathbb{R}^m$ exists, and

(ii) D is a translation covering onto its image $D(\dot{B}(\sigma, \varepsilon)) = \dot{B}(0, \varepsilon)$.

Remark 2.2.7. We need at least dimension three, as for dimension two the punctured ball is not simply connected and thus the developing map is not well-defined on $\dot{B}(\sigma, \varepsilon)$. However, the statement remains true when talking about the map induced by the developing map.

The condition that $\overline{M}$ is a manifold cannot be dropped without replacement as there are examples of isolated singularities in which every neighbourhood is not simply connected. A better known example might be the isolated singularity of the Chamanara surface[Cha04].

The above proposition has the following theorem as a consequence:

Theorem 2.2.8. *Given a translation manifold M of dimension $m \geq 3$, let $\sigma \in \overline{M} \setminus M$ be an isolated singularity. If the metric completion $\overline{M}$ is manifold around σ, then σ is removable.*

Proof. In dimension three and above the punctured ball $\dot{B}(0, \varepsilon) \subseteq \mathbb{R}^m$ is simply connected and 0 is a removable singularity of it. Thus, theorem 2.1.1 in combination with proposition 2.2.6 imply that σ is a removable singularity. ■

The rest of this section is dedicated to the proof of proposition 2.2.6. The crucial point in the proof is the point symmetry of the neighbourhood around σ which we use throughout. The overall structure is this:

First, we find a round ball-like neighbourhood around σ and show that we can define the developing map on it (to this end we use that M is a manifold around σ). Next, we show that D preserves the radial symmetry, namely D preserves the distance to the centre. Using this we can prove that D is surjective on the Euclidean punctured ball $\dot{B}(0, \varepsilon)$. Lastly, we show that D has the covering property by exploiting the radial symmetry. To verify the covering property, we need some uniform neighbourhoods around every point.

Sadly, the point symmetry does not provide such uniform neighbourhoods, however, it allows us to find neighbourhoods uniform with respect to the angle (but not the radius) which in this case is enough to prove the covering property.

Before we start, a short reminder of a consequence of the Hopf-Rinow Theorem: Since $\overline{M}$ is a manifold and metrically complete around σ, it is also geodesically complete, i.e. any two points near σ can be connected by a geodesic.

Lemma 2.2.9. *There exists $\varepsilon > 0$ such that M admits a developing map*

$$D\colon \dot{B}(\sigma,\varepsilon) \to \mathbb{R}^m \tag{2.1}$$

and this map is not expanding, i.e. $d_{\mathbb{R}^m}(D(x),D(y)) \leq d_M(x,y)$ for all $x,y \in \dot{B}(\sigma,\varepsilon)$. In particular the image of D is contained in $B(0,\varepsilon) \subseteq \mathbb{R}^m$.

Proof. Since $\overline{M}$ is a manifold around σ there is a chart and $\varepsilon > 0$ such that $B(\sigma,\varepsilon)$ is homeomorphic to $B(0,1) \subseteq \mathbb{R}^m$. In particular, for dimension $m \geq 3$ the punctured ball $\dot{B}(\sigma,\varepsilon)$ is simply connected, i.e. it is its own universal cover. Therefore, the developing map

$$D\colon \dot{B}(\sigma,\varepsilon) \to \mathbb{R}^m \tag{2.2}$$

exists and is well-defined (cf. remark 1.3.8). Because the developing map is constructed out of charts of a translation manifold M, it is a local isometry.

Ideally, at this point we would like to apply lemma 1.3.10. However, the ball of interest contains the singularity σ so the lemma is not applicable.

We wish for D to be not expanding. However, for the same reason as pointed out in remark 2.1.2 the shortest path between two points might leave the ball $B(\sigma,\varepsilon)$. To force all geodesics to lie within the ball, it is enough to reduce the radius, e.g. to a quarter (in fact any fraction smaller than or equal to one-third will do).

In this situation D is (globally) not expanding. Indeed, given points x and y in $\dot{B}(\sigma,\varepsilon)$, let γ_n be short paths in M between x and y such that $d_M(x,y) = \lim_{n\to\infty} \ell(\gamma_n)$. Note that we cannot assume γ_n to be a geodesic between x and y as a geodesic might go through σ. However, we can approximate any geodesic through σ with such γ_n. With this choice of γ_n it follows that $D(\gamma_n)$ is a path between $D(x)$ and $D(y)$, which has the same length as γ_n since D is a local isometry. Thus, $d_{\mathbb{R}^m}(D(x),D(y)) \leq \ell(D(\gamma_n)) = \ell(\gamma_n) \xrightarrow{n\to\infty} d_M(x,y)$ and D is not a expanding. ∎

Next we define the image of σ under D (cf. also definition 2.3.2 for a generalisation). Let $(x_n)_{n\in\mathbb{N}}$ be a Cauchy sequence with $x_n \to \sigma$ for $n \to \infty$. Since D is not expanding, the image sequence $(D(x_n))_{n\in\mathbb{N}}$ is a Cauchy sequence in $\mathbb{R}^m$ and converges to some point $0 \in \mathbb{R}^m$. We call 0 the image of σ under D and we will see in a moment that we can indeed extend D to a continuous map on $B(\sigma,\varepsilon)$. Because D is constructed from translation charts, we may assume without loss of generality, that 0 is indeed the origin in $\mathbb{R}^m$ by concatenating the developing map with a translation.

Lemma 2.2.10. *The developing map D preserves the distance to the centre of the respective balls. More precisely, for all $x \in \dot{B}(\sigma,\varepsilon)$*

$$d_{\overline{M}}(x,\sigma) = d_{\mathbb{R}^m}(D(x),0). \tag{2.3}$$

Proof. We already know that D is not expanding on $\dot{B}(\sigma,\varepsilon)$, i.e. for all $x, y \in \dot{B}(\sigma,\varepsilon)$

$$d_{\mathbb{R}^m}(D(x), D(y)) \leq d_M(x,y). \tag{2.4}$$

Let $(x_n)_{n\in\mathbb{N}}$ be a sequence with $x_n \to \sigma$ as $n \to \infty$. Without loss of generality we may assume that $d_{\overline{M}}(x_n,\sigma) < \frac{1}{n}$ and $d_{\mathbb{R}^m}(D(x_n),0) < \frac{1}{n}$.

First, we extend inequality (2.4) to still hold for $y = \sigma$:

$$\begin{aligned} d_{\mathbb{R}^m}(D(x),0) &\leq d_{\mathbb{R}^m}(D(x),D(x_n)) + d_{\mathbb{R}^m}(D(x_n),0) \\ &\leq d_M(x,x_n) + \frac{1}{n} \\ &\leq d_{\overline{M}}(x,\sigma) + d_{\overline{M}}(\sigma,x_n) + \frac{1}{n} \\ &\leq d_{\overline{M}}(x,\sigma) + \frac{1}{n} + \frac{1}{n} \\ &\xrightarrow{n\to\infty} d_{\overline{M}}(x,\sigma) \end{aligned} \tag{2.5}$$

Thus, $d_{\mathbb{R}^m}(D(x),0) \leq d_{\overline{M}}(x,\sigma)$.

Next, we remind ourselves that – since D is a local isometry – if γ is a (local) geodesic in M, then $D(\gamma)$ is a (local) geodesic, too, and $\ell(\gamma) = \ell(D(\gamma))$.

Let $x \in \dot{B}(\sigma,\varepsilon) \subseteq M$, let $d = d_{\overline{M}}(\sigma,x)$ and let $\gamma\colon [0,d] \to \overline{M}$ be a geodesic from σ to x, which exists due to $B(\sigma,\varepsilon)$ being geodesically complete.

Clearly, $\gamma(]0,d]) \subseteq \dot{B}(\sigma,\varepsilon)$ since otherwise $\gamma(t) \in \overline{M} \setminus \dot{B}(\sigma,\varepsilon)$ for some $t > 0$ and as γ does not leave $B(\sigma,\varepsilon)$, we would have $\gamma(t) = \sigma$, a contradiction to γ being a geodesic.

Let $(t_n)_{n\in\mathbb{N}}$ with $t_n \to 0$ as $n \to \infty$. Then we have $\lim_{n\to\infty}\gamma(t_n) = \gamma(0) = \sigma$. Using inequality (2.5) yields

$$d(D(\gamma(t_n)),0) \leq d(\gamma(t_n),\sigma) \xrightarrow{n\to\infty} 0 \tag{2.6}$$

i.e. $\lim_{n\to\infty} D(\gamma(t_n)) = 0$ and we can deduce

$$\begin{aligned} d_{\overline{M}}(\sigma,x) &= d_M(\lim_{n\to\infty}\gamma(t_n),\gamma(d)) = \lim_{n\to\infty} d_M(\gamma(t_n),\gamma(d)) \\ &= \lim_{n\to\infty}\ell(\gamma|_{[t_n,d]}) = \lim_{n\to\infty}\ell(D(\gamma|_{[t_n,d]})) = \lim_{n\to\infty} d_{\mathbb{R}^m}(D(\gamma(t_n)),D(\gamma(d))) \\ &= d_{\mathbb{R}^m}(\lim_{n\to\infty} D(\gamma(t_n)),D(\gamma(d))) = d_{\mathbb{R}^m}(0,D(x)), \end{aligned} \tag{2.7}$$

where we used the continuity of the metrics and that γ is a geodesic. ∎

By the above lemma the developing map preserves the distance to the centre. In particular the pre-image of an open ball around 0 is an open ball around σ. This shows that the extension $\bar{D}$ of D to the unpunctured $B(\sigma,\varepsilon)$ with $\bar{D}(\sigma) = 0$ is continuous. It is, however, not necessarily a developing map for $\overline{M}$.

Lemma 2.2.11. *Let $x \in \dot{B}(\sigma,\varepsilon)$ and $d \coloneqq d(x,\sigma)$. Then, D is an isometry on $B(x,\delta)$ where $\delta \coloneqq \frac{1}{2}\min\{d,\varepsilon - d\}$. In particular, the radius only depends on the distance of x to σ but not on x itself.*

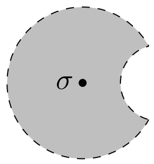

(a) The set $B(0,\varepsilon) \setminus B((1,0,\dots,0),\varepsilon/2) \subseteq \mathbb{R}^m$ is an example of a metric ε-ball which is not isometric to an Euclidean one but homeomorphic to one. This example here is not a 'counter example' as the boundary of the dent would consist of singularities, which are then in the ε-neighbourhood of σ.

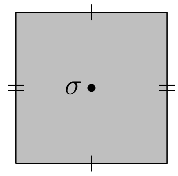

(b) Not every metric ε-ball has a point at a given distance. For example the ball of radius 2 and radius 4 on the standard torus around σ are the same (they are the whole torus). This is also not a 'counter example', as the ball is not homeomorphic to the Euclidean ball in $\mathbb{R}^m$.

Figure 2.2: topological ε-balls

That D is uniformly a local isometry (with respect to the distance from σ) is important. A priori our neighbourhood $B(\sigma,\varepsilon)$ is a metric ball and does *not* need to look like an Euclidean ball at all, although it is homeomorphic to one (see figure 2.2). With this lemma, however, we know that we have a uniform isometry on a sphere around σ, i.e. independent of the direction.

PROOF. The open metric ball $B(x,\delta)$ in $\overline{M}$ is contained in $\dot{B}(\sigma,\varepsilon) \subseteq M$. In particular, it does not contain any singularity of M. This allows us to use lemma 1.3.10 and the claim follows. ■

Lemma 2.2.12. *The image of the developing map D is $\dot{B}(0,\varepsilon) \subseteq \mathbb{R}^m$.*

PROOF. Let $d \in \,]0,\varepsilon[$. Since by lemma 2.2.10 D preserves the distance to the centre, we can see that $D(\partial B(\sigma,d)) \subseteq \partial B(0,d)$.

The developing map is a local isometry and so is $D\big|_{\partial B(\sigma,d)}$, in particular it is an open map. As $\partial B(\sigma,d)$ is open in itself, its image under $D\big|_{\partial B(\sigma,d)}$ is also open. On the other hand $\partial B(\sigma,d)$ is compact (as it is closed and bounded in a set homeomorphic to a subset of $\mathbb{R}^m$), thus its image is also compact. Hence, the image of $\partial B(\sigma,d)$ is open and compact. Thus, it is all of $\partial B(0,d) \subseteq \mathbb{R}^m$ or empty.

It is left to show that this is not the empty set. There is a point $x \in \dot{B}(\sigma,\varepsilon)$ and the geodesic from σ to x exists in $\overline{M}$ by the Hopf-Rinow theorem as $\overline{M}$ is a manifold. Clearly $\partial B(0,r) \neq \emptyset$ for all $0 \leq r < d(\sigma,x)$ since it contains at least the point of the geodesic.

We may now without loss of generality shrink ε to $d(\sigma,x)$ and the statement follows. ■

Lemma 2.2.13. *$D\colon \dot{B}(\sigma,\varepsilon) \to \dot{B}(0,\varepsilon)$ is a covering.*

PROOF. The idea to prove the covering property is based on the last argument in the proof of [Thu97, Proposition 3.4.10, pp. 144]. Thurston's proof, however, requires D to be an isometry on all δ-balls. In general we cannot guarantee this when the balls

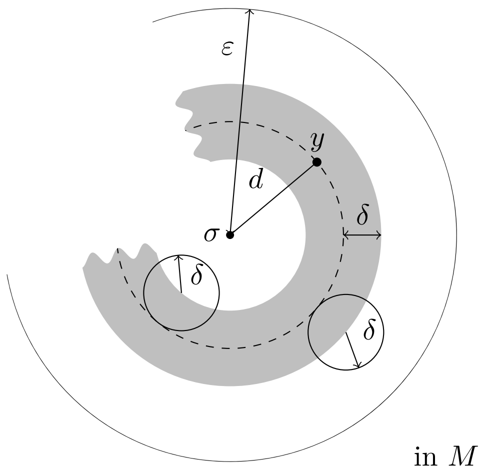

Figure 2.3: The choice of δ in the proof of lemma 2.2.13 must be such that a ball of radius δ fits around every point of the annulus.

approach the singularity σ or the boundary of $B(\sigma, \varepsilon)$ but lemma 2.2.11 gives us a bound uniform enough to make it work nevertheless.

Let $y \in D(\dot{B}(\sigma, \varepsilon)) \subseteq \dot{B}(0, \varepsilon) \subseteq \mathbb{R}^m$ and $d := d(0, y)$. Choose $0 < \delta < d$ such that for any point x in the metric annulus $A := B(\sigma, d+\delta) \setminus B(\sigma, d-\delta)$ the ball $B(x, \delta)$ is contained in $\dot{B}(\sigma, \varepsilon)$ and such that $D\big|_{B(x,\delta)}$ is an isometry. Such a δ exists by lemma 2.2.11, cf. figure 2.3, e.g. choose

$$\delta := \frac{1}{8} \min\{d, \varepsilon - d\}. \tag{2.8}$$

Take $x \in D^{-1}(B(y, \delta/2))$. Because D preserves the distance to the centre (lemma 2.2.10), x is guaranteed to be within the annulus. Figure 2.4 depicts the situation. From here on we follow Thurston's argument. The ball $B(x, \delta) \subseteq M$ maps isometrically and thus must properly contain a homeomorphic copy of $B(y, \delta/2) \subseteq \mathbb{R}^m$. The entire inverse image $D^{-1}(B(y, \delta/2)) \subseteq M$ is then a disjoint union of such homeomorphic copies, since if not, two $\delta/2$-balls would intersect violating $D\big|_{B(x,\delta)}$ being an isometry. Therefore, D evenly covers $D(\dot{B}(\sigma, \varepsilon)) = \dot{B}(0, \varepsilon)$, so it is a covering projection $D\colon \dot{B}(\sigma, \varepsilon) \to \dot{B}(0, \varepsilon)$. ■

This finishes the proof of proposition 2.2.6.

2.3 Images and Shadows of Singularities

The previous discussion of an isolated singularity has shown how we can utilise coverings to describe singularities and in particular to show that isolated singularities are removable. For an isolated singularity there is only a single point we have to take care of and we can find a nice symmetric neighbourhood on which we can build upon. While singularities

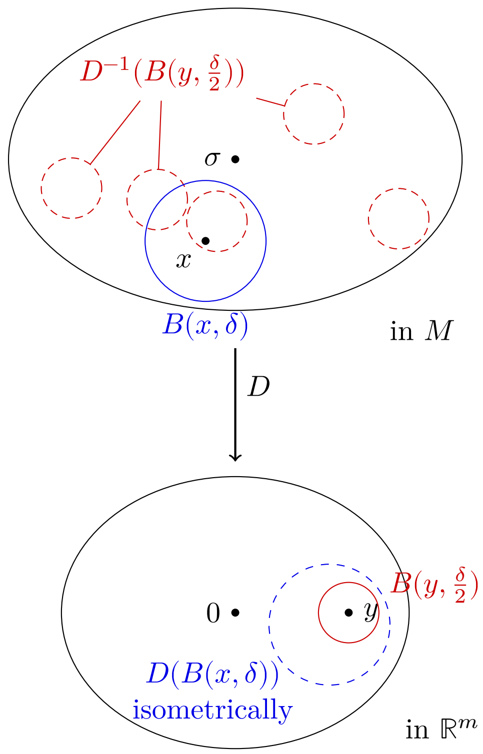

Figure 2.4: The balls involved in the proof of lemma 2.2.13. The depicted situation cannot happen, as D maps $B(x, \delta)$ isometrically but the $\delta/2$-ball has two partial pre-images.

in dimension two are (only) points, this is no longer true in higher dimensions and singularities are usually not isolated at all.

To understand the nature of the developing map in these cases, we must identify the points where complications arise. Obviously, singularities are points which need our attention. However, their influence reaches a little further as they shimmer through to other sheets of the (yet to be proven) covering by the developing map. This yields the concept of images and shadows of singularities which allows us to describe problematic points.

Having taken care of these points, we prove in theorem 2.3.16 under some mild condition (corresponding to the tameness of the singularity) that the developing map is a covering and use that to remove singularities with theorem 2.3.21.

2.3.1 Definition of Shadows and Images

Example 2.3.1. First, let us consider a simple example. Let M be two glued sliced punctured discs where one disc contains a singularity σ but the other does not (see figure 2.5). The developing map D – or more precisely the in this case well-defined map induced by the developing map as M is not simply connected – maps onto the punctured open ball $\dot{B}(0,\varepsilon) \subseteq \mathbb{R}^2$. ($M$ has to be non simply connected because of theorem 2.3.21.) In general we would like to say that D is a covering on $\dot{B}(0,\varepsilon)$. However, D fails to be covering over the point $D(\sigma)$, which can be thought of the image of σ under D, because the pre-image of a small disc around $D(\sigma)$ is in one sheet a punctured disc and in the other a filled disc.

This situation is unpleasant. For a concrete example with knowledge about the singularity in question, the situation can often be resolved easily by choosing a better suited neighbourhood for it. In the above example, we could reduce the radius, so σ is no longer in the neighbourhood we are interested in and thus not a problem any more.

For the generic case, this is much harder to do as singularities are not only isolated points and can accumulate in the points we are interested in. Furthermore, proving the existence of a 'good' neighbourhood can be quite challenging for an abstract space.

To tackle this problem, we introduce *images* and *shadows*. Simply put, we declare the points at which D fails to be a covering map also to be singularities. Then D is a

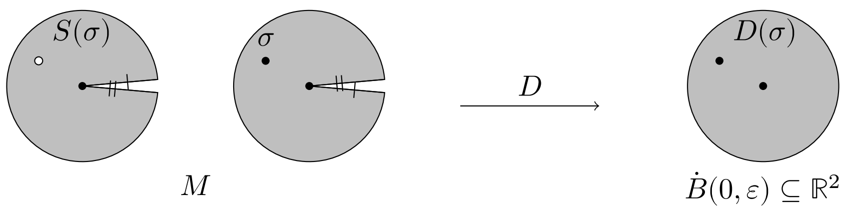

Figure 2.5: The translation manifold M consists of two glued punctured discs. It has two singularities: the centre of the discs (present in both discs) and the singularity σ (only present in one disc).

translation covering again. In the above example we would declare $S(\sigma)$ to also be a singularity. However, singularities without any restriction can be quite wild and expose complicated and unintuitive behaviour so we have to be careful when doing this.

Here is a first definition which we will revise later as it does not cover all situations which can occur. Nevertheless it is useful in its own right especially when we have some control or knowledge about the neighbourhood of the singularities we are interested in.

Definition 2.3.2 (Images and Shadows). Let M be a translation manifold of dimension m, $\overline{M}$ its metric completion and $\Sigma = \overline{M} \setminus M$ the set of singularities. Let $U' \subseteq \overline{M}$ be an open neighbourhood such that the development map on the universal cover of $U := U' \cap M$ descends to a well-defined map on U. The *images* of the singularities of U contained in U' in $\mathbb{R}^m$ are defined as

$$D^{\mathrm{f}}_U(\Sigma) := \bigcup_{(x_n)_{n\in\mathbb{N}}\in\mathcal{C}} \operatorname{Acc}((D(x_n))_{n\in\mathbb{N}}) \subseteq \overline{D(U)}, \tag{2.9a}$$

where

$$\mathcal{C} := \left\{ (x_n)_{n\in\mathbb{N}} \;\middle|\; (x_n)_{n\in\mathbb{N}} \text{ is a Cauchy sequence in } U \text{ with } \lim_{n\to\infty} x_n \in \Sigma \cap U' \right\} \tag{2.9b}$$

and Acc is the set of accumulation points of a sequence. The *shadows* of the singularities $\Sigma \cap U'$ in U are

$$S^{\mathrm{f}}_U(\Sigma) := (\Sigma \cap U') \cup D^{-1}(D^{\mathrm{f}}_U(\Sigma)). \tag{2.10}$$

The ‘f’ in our notation stands for ‘focused’; for an explanation for this name choice see remark 2.3.3 point 5 and remark 2.3.8 point 3.

Remark 2.3.3. A couple of remarks regarding this definition.

1. Images of singularities often do not lie in the image of U under D but in its closure $\overline{D(U)}$, e.g. the punctured ball in $\mathbb{R}^m$ considered as translation manifold.

2. Note that not every image yields a shadow as it might not lie in the image of D, e.g. in the example 2.3.1 the origin 0 is an image but it is not contained in the image $D(M) = \dot{B}(0, \varepsilon)$ and thus has no pre-images, i.e. no corresponding shadow.

3. We include the set of singularities Σ in the set of shadows. This is for convenience but also to honour the spirit of the definition as shadows should be regarded as a generalisation of singularities and thus should be a superset.

4. The images of singularities in $\mathbb{R}^m$ depend on the choice of D and are well-defined up to a translation.

5. The notation is slightly misleading as the singularities depend on U' and not just on U. However, this notation is more in line with our upcoming definition 2.3.7 and the idea that we look at a set in $U \subseteq M$ with developing map D and are interested in the singularities related to U. Having the possibility of choosing U' gives also a more fine-grained control to pick the singularities of interest.

The usual choice for U' is $\operatorname{int}\overline{U}$ but it might be worth noting that in general $U \neq \operatorname{int}\overline{U} \cap M$, so the map induced by the developing map might not be defined for $\operatorname{int}\overline{U} \cap M$. In concrete examples, however, it is clear what U' is supposed to be.

6. The above definition can be adopted for any map D between two translation manifolds which is a local isometry or more generally locally not expanding.

Lemma 2.3.4. *In the situation of definition 2.3.2, we have*

$$D_U^{\mathrm{f}}(\Sigma) = \Big\{ \lim_{n\to\infty} D(x_n) \,\Big|\, (x_n)_{n\in\mathbb{N}} \text{ is a Cauchy sequence in } U \text{ with } \lim_{n\to\infty} x_n \in \Sigma \cap U' \Big\}. \tag{2.11}$$

Proof. If the limit of the sequence $(D(x_n))_{n\in\mathbb{N}}$ exists, then the set of accumulation points $\operatorname{Acc}((D(x_n))_{n\in\mathbb{N}})$ is the singleton set consisting only of the limit point of the sequence $(D(x_n))_{n\in\mathbb{N}}$ and the statement follows. Thus, it suffices to show that taking the limit $\lim_{n\to\infty} D(x_n)$ is well defined.

Let $(x_n)_{n\in\mathbb{N}}$ be a Cauchy sequence in U with $x := \lim_{n\to\infty} x_n \in \Sigma \cap U'$. If D is (globally) not expanding, then $(D(x_n))_{n\in\mathbb{N}}$ would be a Cauchy sequence, too, but this is in general not the case. Even worse, under D a short distance in M can become a large distance in $\mathbb{R}^m$, cf. remark 2.1.2. However, locally D is an isometry and this will suffice as the following shows.

Let $\varepsilon > 0$ and $B(x,\varepsilon) \subseteq U'$ be an open ball in $\overline{M}$ around the limit of the Cauchy sequence. Then for all $y, y' \in B(x, \frac{\varepsilon}{3}) \cap U$, we have $d(y,y') < \frac{2\varepsilon}{3}$ and all the short paths between y and y' are contained in $B(x,\varepsilon) \cap U \subseteq M$. We use the words 'short paths' as the geodesic between y and y' (which has length less than $\frac{2\varepsilon}{3}$) might go through a singularity. By the definition of the metric in M there are short paths whose lengths are arbitrary close to the distance between y and y'. Since D – as a local isometry – preserves the length of paths, we have $d(D(y), D(y')) < \frac{2\varepsilon}{3}$. Although D might be expanding it at least preserves short distances around x, i.e. if y and y' are close to x and thus are close to each other, the distance between their images under D might increase but the images are still close to each other. In particular, D maps a Cauchy sequence converging to x to a Cauchy sequence.

Therefore, the limit of $(D(x_n))_{n\in\mathbb{N}} \in \mathbb{R}^{\mathbb{N}}$ exists and is well defined. ■

Example 2.3.5. Figure 2.6 shows an example of a translation surface where the Cauchy sequence $(x_n)_{n\in\mathbb{N}}$ converges to x but D is not non-expanding on U. Indeed, D increases the distances between x_{2n} and x_{2n+1} (cf. remark 2.1.2). However, if $d_M(x_{2n}, x_{2n+1})$ is small, then $d_{\mathbb{R}^m}(D(x_{2n}), D(x_{2n+1}))$ is small, too.

As outlined above the idea behind the concept of images and shadows is to gather all the problematic points and show that outside of these (the map induced by) the developing map is a covering, in other words

$$D\colon U \setminus S_U^{\mathrm{f}}(\Sigma) \to D(U) \setminus D_U^{\mathrm{f}}(\Sigma) \tag{2.12}$$

is a translation covering.

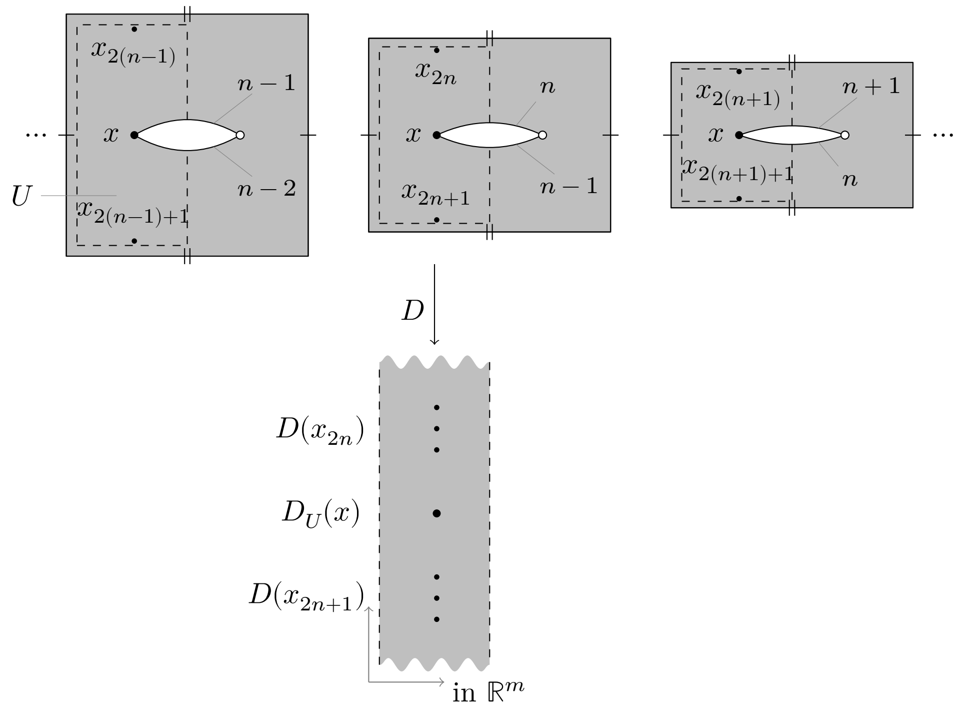

Figure 2.6: Example of a translation surface with not non-expanding developing map.

In example 2.3.1 we have seen that singularities are such problematic points and by extension also their shadows and these capture indeed all the points preventing D from being a covering. While it works out in that situation, generally this is not the case and there is a second kind of points which need our attention: boundary points.

Here is an example illustrating that points of the boundary can lead to covering problems:

Example 2.3.6. This is a slightly modified version of example 2.3.1 but instead of glueing two sliced discs we glue one sliced disc and one sliced square. The situation is depicted in figure 2.7. The developing map D – or more precisely the (in this case

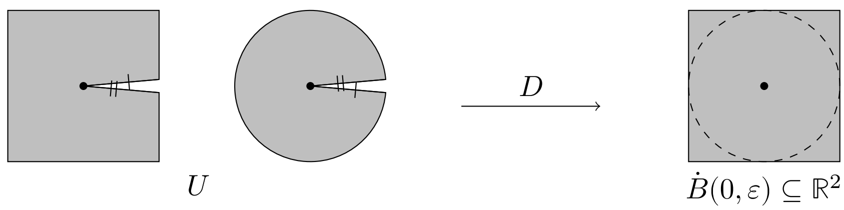

Figure 2.7: The neighbourhood U of the translation manifold M consists of a punctured disc glued with a punctured square. The map D is not a covering above the dashed points.

well-defined) map induced by the developing map as U is not simply connected – maps onto the punctured square in $\mathbb{R}^2$. D is almost everywhere a covering except on the dashed points which resembles the 'image of the boundary of the circle'.

The set of singularities Σ for this translation manifold is only the centre point (a singularity of order two). For definition 2.3.2 the neighbourhood U' is U (the open circle and open square) together with the centre point. Thus, the images $D^{\mathrm{f}}_U(\Sigma)$ are still only the centre point and do not contain the dashed points.

In the previous example, the obstacle which prevents D from being a covering is the boundary of the neighbourhood U. Thus, we need to extend the definition of images and shadows to also take into account the boundaries. The revised definition reads like this:

Definition 2.3.7. Let M be a translation manifold of dimension m, $\overline{M}$ its metric completion and $\Sigma = \overline{M} \setminus M$ the set of singularities. Let $U \subseteq M$ be a connected, open neighbourhood such that the developing map induces a well-defined map $D\colon U \to \mathbb{R}^m$ on U. The *images* of singularities or the boundary of U in $\mathbb{R}^m$ are defined as

$$D_U(\Sigma) := \bigcup_{(x_n)_{n\in\mathbb{N}}\in\mathcal{C}} \operatorname{Acc}((D(x_n))_{n\in\mathbb{N}}) \subseteq \overline{D(U)}, \tag{2.13a}$$

where

$$\mathcal{C} := \left\{(x_n)_{n\in\mathbb{N}} \,\middle|\, (x_n)_{n\in\mathbb{N}} \text{ is a Cauchy sequence in } U \text{ with } \lim_{n\to\infty} x_n \in \partial U\right\}. \tag{2.13b}$$

The *shadows* of the singularities or the boundary of U in U are

$$S_U(\Sigma) := \partial U \cup D^{-1}(D_U(\Sigma)). \tag{2.14}$$

Remark 2.3.8.

1. Lemma 2.3.4 does not hold anymore. To see this consider the torus M where U is everything except the glued boundary, cf. figure 2.8.
2. The above definition does not mention singularities. They are hidden in the boundary: the singularities related to U are a subset of the boundary of U.

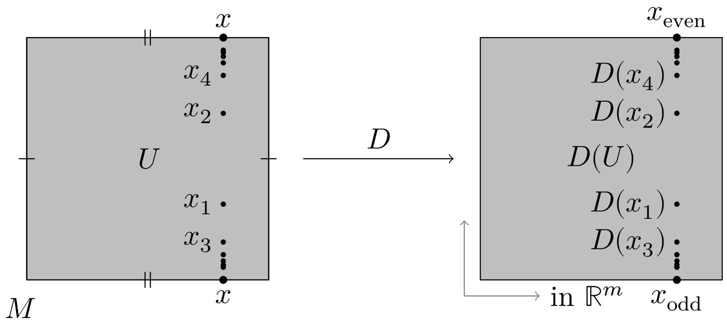

Figure 2.8: The sequence $(x_n)_{n\in\mathbb{N}}$ converges in the translation manifold to a single point x. However, its image under the development map D does not converge but has two accumulation points x_{odd} and x_{even}.

3. In comparison to definition 2.3.2 we only need one open set $U \subseteq M$. This has its merits but also drawbacks:

   The second definition (definition 2.3.7) covers all problematic points allowing us to prove the theorems in the remainder of this section.

   In the first definition (definition 2.3.2) with U and U', the U' can be thought of as a blind or filter which admits only the singularities we care about to cast a shadow. This lets us focus on a particular singularity at the cost that we need additional information about the neighbourhood (like the radial symmetry in example 2.3.1).

4. Note that $U \setminus S_U(\Sigma)$ might not be connected anymore, e.g. example 2.3.6 describes such a situation.

Remark 2.3.9 (Relation between Definition 2.3.2 and Definition 2.3.7).
That we have to include the boundary is not that surprising since if we consider the neighbourhood U to be a translation manifold in its own right, then the set of singularities of U is $\overline{U} \setminus U$, which contains the boundary of U.

We can make this more precise: Let M be a translation manifold, $\overline{M}$ its metric completion and $\Sigma = \overline{M} \setminus M$ its set of singularities. Furthermore, let $U \subseteq M$ an open, connected neighbourhood such that the developing map induces a well-defined map $D\colon U \to \mathbb{R}^m$. Denote by N the metric space U but with the intrinsic metric, i.e. the metric induced by the path lengths of paths in U (cf. definition 1.1.5). The inclusion $\iota\colon N \to M$ is a non-expanding embedding and can be extended to the metric completions $\bar{\iota}\colon \overline{N} \to \overline{M}$, which is still non-expanding. Note that $\bar{\iota}$ might not be injective any more, e.g. figure 2.8 with $\iota = D^{-1}$).

We have that N (and hence U) is a translation manifold with singularities $\Sigma_N = \overline{N} \setminus N$ and we can project $\overline{N}$ onto the closure $\overline{U}$ of U as a subset of $\overline{M}$. Furthermore, definition 2.3.7 can be reduced to definition 2.3.2 by choosing $U' = \overline{N}$ *for the translation manifold* N:

$$S_U(\Sigma) = \bar{\iota}(S_N^{\mathrm{f}}(\Sigma_N)). \tag{2.15}$$

Be aware that U' is open in $\overline{N}$ but $\iota(U') \subseteq \overline{M}$ is usually not open in $\overline{M}$.

If $\bar{\iota}$ is injective, then $\overline{U}$ and $\overline{N}$ can be identified and the above equation can be shortened to $S_U(\Sigma) = S_U^{\mathrm{f}}(\overline{U} \setminus U)$ obfuscating that on the right-hand side U is regarded as a translation manifold on its own and not as a subset of M.

2.3.2 Covering without Shadows

Although the developing map $D\colon U \to \mathbb{R}^m$ might not be a covering map one might think it has the path lifting property, i.e. a path $\gamma\colon [0,1] \to D(U)$ can be lifted to a path $\tilde{\gamma}\colon [0,1] \to U$ with $D \circ \tilde{\gamma} = \gamma$. This is incorrect as a lifted point in the image might be a singularity in one layer but not a singularity in another, see example 2.3.1, which prevents lifting some paths. To make the lifting statement true we have to account for the images of singularities and also for the boundaries as the following example shows.

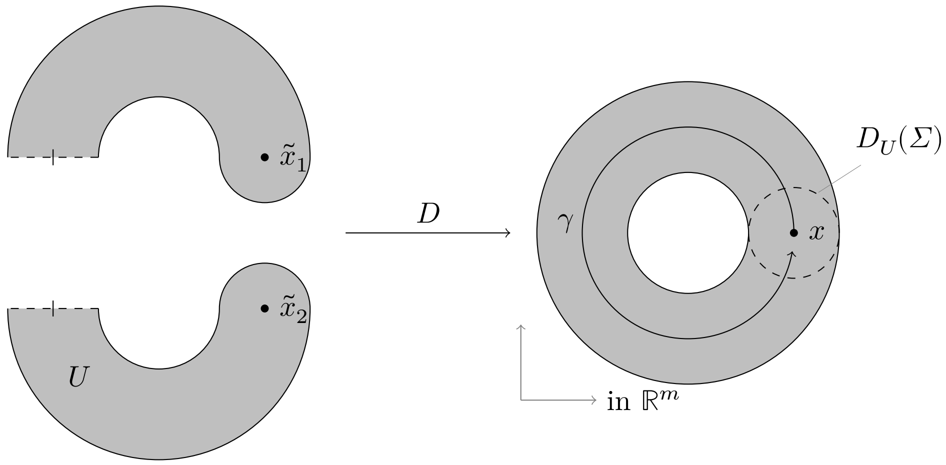

Figure 2.9: The path γ can only be lifted from $\tilde{x}_1$ but not from $\tilde{x}_2$.

Example 2.3.10. For the neighbourhood U depicted in figure 2.9, we have to take the boundary into account when lifting the path γ. The path γ, which crosses $D_U(\Sigma)$, can be lifted from $\tilde{x}_1$ but not from $\tilde{x}_2$.

Lemma 2.3.11 (Path Lifting Property). *Let M be an m-dimensional translation manifold and $\Sigma = \overline{M} \setminus M$ the set of singularities. Let $U \subseteq M$ be an open set such that the developing map descends to a map $D\colon U \to \mathbb{R}^m$.*

Then $D\colon U \setminus S_U(\Sigma) \to D(U) \setminus D_U(\Sigma)$ has the path lifting property. This means that for a path $\gamma\colon [0,1] \to D(U) \setminus D_U(\Sigma)$ or ray $\gamma\colon [0,1[\to D(U) \setminus D_U(\Sigma)$ and a pre-image $\tilde{x} \in D^{-1}(\gamma(0))$ of the starting point, there exists a unique lift $\tilde{\gamma}\colon [0,1] \to U \setminus S_U(\Sigma)$ or $\tilde{\gamma}\colon [0,1[\to U \setminus S_U(\Sigma)$, respectively, that is a path or ray with $D \circ \tilde{\gamma} = \gamma$. For the former, for each choice of $\tilde{x}$, there is an $\tilde{x}_1$ which is the pre-image of the end point $\gamma(1)$.

Remark 2.3.12. The lifting of rays is in particularly useful to aim for singularities, i.e. paths in M which end in a singularity.

Proof. Recall that D is a local isometry.

Let $U_{\tilde{x}}$ be a neighbourhood around the starting point such that $D|_{U_{\tilde{x}}}$ is an isometry. Lift the beginning of the path γ to M by defining $\tilde{\gamma} := (D|_{U_{\tilde{x}}})^{-1} \circ \gamma$. Note that this way we can only lift open parts of γ, i.e. $\gamma|_{[0,t[}$ for some t.

Assume we have already lifted the beginning $\gamma|_{[0,t_1[}$ of γ for some $t_1 \in [0,1[$ to M where the last lifting was done via a local restriction of D to U_{x_0} on which it is an isometry. We extend the lift to t_1. Let $\tilde{x}_1 := \lim_{t \to t_1} \tilde{\gamma}(t)$. This limit is well-defined as $\gamma(t)$ is a Cauchy sequence and the lifting in the last part of $\tilde{\gamma}$ was done via an isometry. We have $\tilde{x}_1 \in U \setminus S_U(\Sigma)$ as otherwise

$$\gamma(t_1) = \lim_{t \to t_1} \gamma(t) = \lim_{t \to t_1} (D \circ \tilde{\gamma})(t) \in D_U(\Sigma), \tag{2.16}$$

a contradiction. Thus, we can pick a neighbourhood $U_{\tilde{x}_1} \subseteq U$ such that D is an isometry around $\tilde{x}_1$ and extend $\tilde{\gamma}$ via $(D|_{U_{\tilde{x}_1}})^{-1} \circ \gamma$, which matches up with $\tilde{\gamma}$ as U_{x_0} and U_{x_1}

intersect. This process can be repeated countably often, thus allowing lifting rays $\gamma\colon [0,1[\to D(U) \setminus D_U(\Sigma)$.

To see that also paths $\gamma\colon [0,1] \to D(U) \setminus D_U(\Sigma)$ can be lifted, first lift its restriction on $[0,1[$, then do the extension to $t = 1$ in the same way as described above for t_1. ■

The developing map also has the homotopy lifting property, the proof of this is similar to the proof of the path lifting property.

Remark 2.3.13 (Homotopy Lifting Property). Let M be an m-dimensional translation manifold, $\Sigma = \overline{M} \setminus M$ the singularities. Let $U \subseteq M$ be an open set such that the developing map descends to a map $D\colon U \to \mathbb{R}^m$.

Then $D\colon U \setminus S_U(\Sigma) \to D(U) \setminus D_U(\Sigma)$ has the homotopy lifting property. This means, given a map $f\colon Y \times I \to D(U) \setminus D_U(\Sigma)$ and map $\tilde{f}\colon Y \times \{0\} \to U$ lifting $f|_{Y\times I}$, i.e. $D \circ \tilde{f} = f$, then there exists a unique lift $\tilde{f}\colon Y \times I \to U \setminus S_U(\Sigma)$ lifting f and restricting to $\tilde{f}$ on $Y \times \{0\}$. Here I is either the closed unit interval $[0,1]$ or the half-open unit interval $[0,1[$.

In particular, paths $\gamma\colon [0,1] \to D(U) \setminus D_U(\Sigma)$ and rays $\gamma\colon [0,1[\to D(U) \setminus D_U(\Sigma)$ can be lifted. For the former, for each choice of $\tilde{x} \in D^{-1}(\gamma(0))$, there is a $\tilde{x}_1$ which is the pre-image of the end point $\gamma(1)$.

Since the set of singularities $\Sigma = \overline{M} \setminus M$ is by construction closed, one might guess that the images of the singularities are also closed in $\mathbb{R}^m$. However, this is not the case as the following example shows.

Example 2.3.14. The translation manifold M as depicted in figure 2.10 consists of an infinite spiral staircase, where at each level $n > 0$ there is a singularity in distance $\ell + \frac{1}{n}$. The images of these singularities converge to a point with distance ℓ but neither of its pre-images is a singularity.

Requiring finiteness in the fibres of D (even uniform finiteness) does not salvage the situation:

Example 2.3.15. Consider the neighbourhood $U = M$ depicted in figure 2.11. It consists of a half-annulus which has spikes at the angles $\frac{\pi}{2^n}$ which reach to the centre up to distance $\frac{1}{2^n}$ combined with a circle. The centre x is not in the image of the singularities or the boundary of U but x is an accumulation point of the image. In this construction the number of pre-images under D is bounded by 2.

A variation of this example uses a singularity of finite order near the tips of the spikes yielding the same result even when not considering the boundary.

Theorem 2.3.16. *Let M be an m-dimensional translation manifold and $\Sigma = \overline{M} \setminus M$ the singularities. Let $U \subseteq M$ be such that the developing map descends to a map $D\colon U \to \mathbb{R}^m$. Furthermore, assume that $D(U) \setminus D_U(\Sigma)$ is open in $D(U)$. Then*

i) $U \setminus S_U(\Sigma)$ is open in U, and

ii) $D\colon U \setminus S_U(\Sigma) \to D(U) \setminus D_U(\Sigma)$ is a translation covering map.

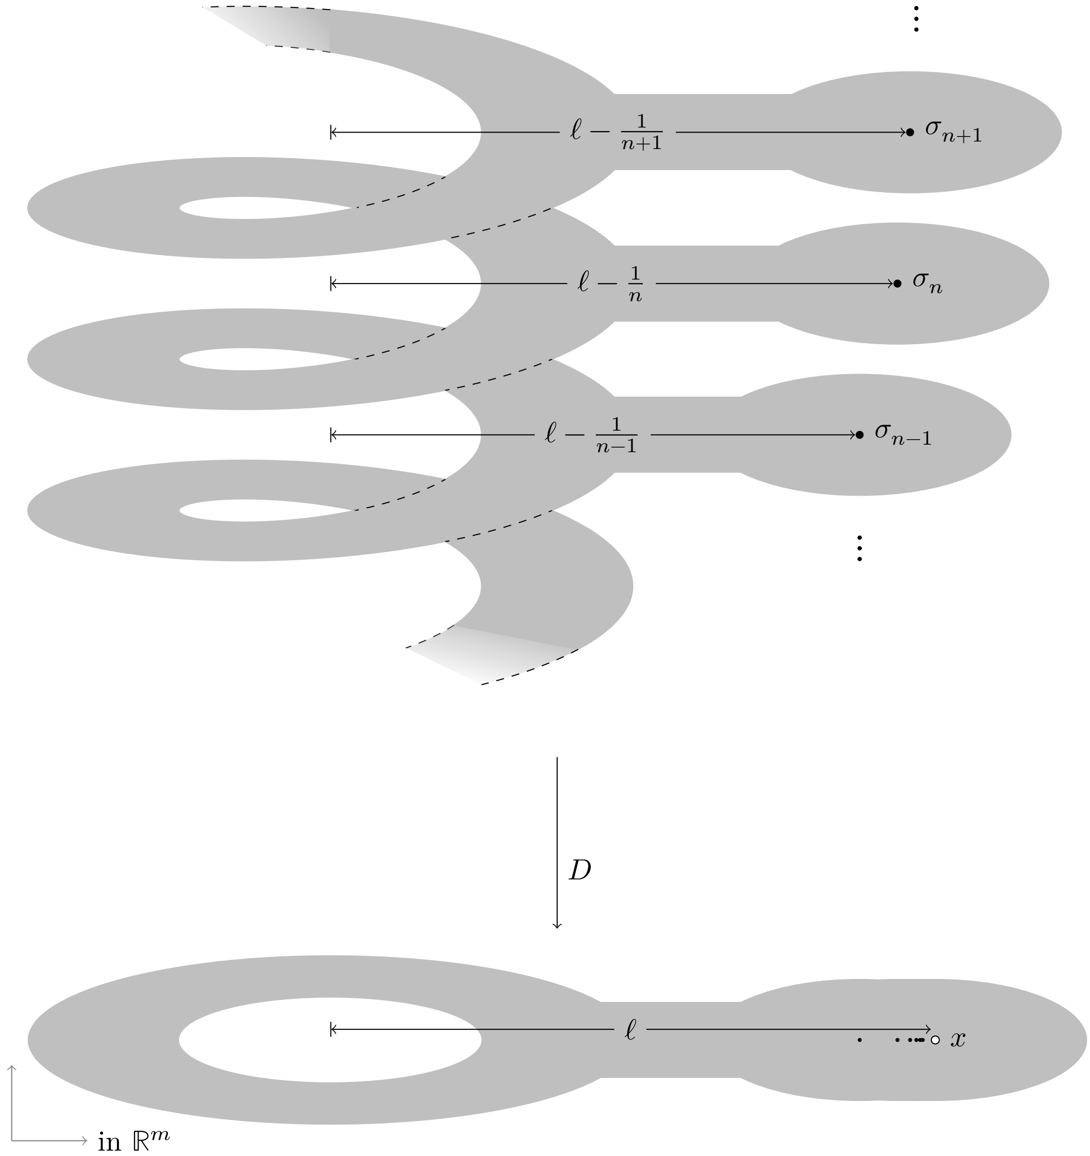

Figure 2.10: Infinite spiral staircase with a singularity σ_n on each level. The images of the singularities on each level converge to x but x itself is not an image of a singularity.

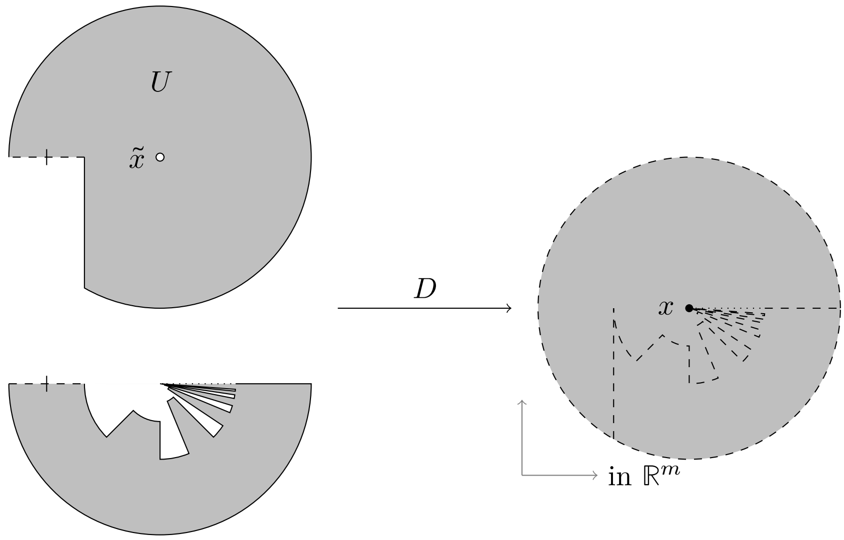

Figure 2.11: In this example the covering is finite but the images of the tips of the spikes accumulate in x although x itself is not an image of a singularity nor the boundary.

Remark 2.3.17. The above can be seen as an indicator or sufficient prerequisite for tame singularities. The first condition is that the images of the singularities are closed in $D(U)$. The second condition is more hidden and is the requirement for the developing map to descend to a map on the translation manifold.

Isolated wild singularities usually fail the second condition.

Remark 2.3.18. Please note that neither $U \setminus S_U(\Sigma)$ nor $D(U) \setminus D_U(\Sigma)$ must be connected. In particular, D can be a covering with different degrees on the individual connected components. For an example see figure 2.12.

Proof. Since $U \setminus S_U(\Sigma) = U \setminus D^{-1}(D_U(\Sigma))$ and D is a local homeomorphism, $D(U) \setminus D_U(\Sigma)$ open implies $U \setminus S_U(\Sigma)$ open.

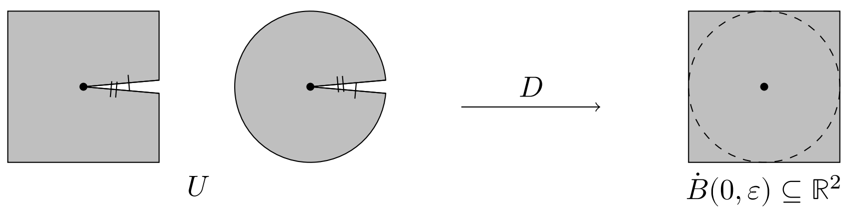

Figure 2.12: The neighbourhood U consists of the punctured disc glued with a punctured square. The image of the boundary of U separates $D(U)$ into five connected components: the four corners and the central punctured disc. Over each of the four corner components D has degree 1, and over the central punctured disc it has degree 2.

Let $y \in D(U) \setminus D_U(\Sigma)$. To show that D is a covering map we again need a uniform bound for the radii of balls on which D is an isometry. More precisely we need $D\big|_{B(\tilde{x},\varepsilon)}$ to be an isometry for all $\tilde{x} \in D^{-1}(B(y, \varepsilon/2))$.

Choose $\varepsilon > 0$ such that $B(y, 5\varepsilon/2) \subseteq D(U) \setminus D_U(\Sigma)$. We claim that with this choice $D\big|_{B(\tilde{x},\varepsilon)} \colon B(\tilde{x}, \varepsilon) \to D(B(\tilde{x}, \varepsilon))$ is an isometry for all $\tilde{x} \in D^{-1}(B(y, \varepsilon/2)) \subseteq U$.

Let $\tilde{x} \in D^{-1}(\{x\})$ be a pre-image of $x \in B(y, \varepsilon/2))$. First, we show that $B(\tilde{x}, 2\varepsilon) \cap S_U(\Sigma) = \emptyset$, in particular $B(\tilde{x}, 2\varepsilon) \subseteq M$. We have $D(B(\tilde{x}, 2\varepsilon)) \subseteq B(x, 2\varepsilon) \subseteq B(y, \frac{5}{2}\varepsilon)$ and since the latter has no images of singularities or the boundary, its pre-image has no shadows of singularities or the boundary. Next, by lemma 1.3.10 the map D restricted to $B(x, \varepsilon)$ is an isometry. Lastly, by the same argument as in lemma 2.2.13 we can conclude that

$$D \colon U \setminus S_U(\Sigma) \to D(U) \setminus D_U(\Sigma) \tag{2.17}$$

is a covering map. It is also a translation covering map as D is glued out of charts, i.e. respects the translation structure. ■

Corollary 2.3.19. *Let M be an m-dimensional translation manifold and $\Sigma = \overline{M} \setminus M$ be the singularities. Let $U \subseteq M$ be an open set such that the developing map descends to a map $D \colon U \to \mathbb{R}^m$. If*

i) $U \setminus S_U(\Sigma)$ is connected,

ii) $D(U) \setminus D_U(\Sigma)$ is open, and

iii) $D(U) \setminus D_U(\Sigma)$ is simply connected,

then the singularities removable in $D(U) \setminus D_U(\Sigma)$ are removable in $U \setminus S_U(\Sigma)$.

Proof. By theorem 2.3.16 the developing map D is a translation covering, thus the statements follows from theorem 2.1.1. ■

In a next step we try to reduce the prerequisites of corollary 2.3.19. A first attempt would be to hope that $U \setminus S_U(\Sigma)$ simply connected would be enough to imply that all singularities inside of U are removable. This is unfortunately not the case as the next example shows:

Example 2.3.20. The universal covering U of the punctured unit disc (considered as translation manifold) is simply connected but the singularity is not removable.

However, for a finite covering we have the following statement:

Theorem 2.3.21. *Let M be an m-dimensional translation manifold and $\Sigma = \overline{M} \setminus M$ be the singularities. Furthermore, let $U \subseteq M$ be such that the developing map descends to a well-defined map $D \colon U \to \mathbb{R}^m$. If*

i) $U \setminus S_U(\Sigma)$ is connected and simply connected,

ii) $\pi_1(D(U) \setminus D_U(\Sigma))$ is torsion free, and *iii) either*

a) $D\colon U \setminus S_U(\Sigma) \to D(U) \setminus D_U(\Sigma)$ *is a finite covering,*

or

b) $D^{-1}(x)$ *is finite for all* $x \in D(U) \setminus D_U(\Sigma)$*, and*

c) $D(U) \setminus D_U(\Sigma)$ *is open,*

then the singularities which are removable in $D(U) \setminus D_U(\Sigma)$ *are removable in* $U \setminus S_U(\Sigma)$*.*

Remark 2.3.22. The interesting fact about the preconditions in this theorem in particular with condition b) and c) is that all conditions (except maybe the finiteness condition) are purely topological properties.

Proof. Firstly, conditions b) and c) imply a) by theorem 2.3.16.

We want to apply theorem 2.1.1. To this end we must show that $D(U) \setminus D_U(\Sigma)$ is simply connected, i.e. $\pi_1(D(U) \setminus D_U(\Sigma)) = 0$. But this follows directly from conditions i), ii) and a) as a torsion free group is either trivial or infinite and the latter cannot happen because of a) and the simply connectedness of $U \setminus S_U(\Sigma)$. ∎

Remark 2.3.23. A condition similar to $U \setminus S_U(\Sigma)$ being simply connected cannot be omitted, for example the classic two disc-covering would fulfill everything of the above except simply connectedness but the singularity is not removable.

2.4 Singularities in Dimension 2 and 3

A question which immediately arises in the situation of theorem 2.3.21 is: Can the fundamental group of the set $D(U) \setminus D_U(\Sigma)$ have torsion? Or more generally, do open sets in $\mathbb{R}^m$ allow torsion?

For $\mathbb{R}^2$ the answer is given by the following theorem:

Theorem 2.4.1. *The fundamental group of an open subset of* $\mathbb{R}^2$ *is free, in particular torsion-free.*

This theorem follows from the fact that an open set in dimension two can be triangulated [Rad26, Hilfssatz 2] in combination with either [Ahl16, Section 44A, p. 102], which proves the above statement for triangulated surfaces or with [Whi61, Lemma 2.1], which implies that the triangulation of an open set in $\mathbb{R}^2$ has a subcomplex of its 1-skeleton to which it is homotopic combined with the fact that a 1-simplicial complex is a graph and the fundamental group of a graph is free.

In the case $\mathbb{R}^3$ we can prove that the fundamental group is torsion-free as follows. The proof uses some knowledge about the classification and composition of 3-manifolds as well as of the 3-sphere. So let us recall them first.

The 3-sphere S^3 has the property that any embedding $S^2 \hookrightarrow S^3$ of a 2-sphere bounds two 3-balls – one on the 'inside' and one on the 'outside' – i.e. we have

$$S^3 \setminus S^2 = B^3 \amalg B^3. \tag{2.18}$$

(a) Loop γ if N is part of the 'inside' of the torus.

(b) Loop γ if N is part of the 'outer surroundings' of the torus.

Figure 2.13: The loop γ, which arises from a torus boundary component of N, generates an infinite group.

This fact is (a corollary of) Alexander's Theorem[Hat07, Theorem 1.1] or a corollary of the generalised Schoenflies Theorem[Bro60; Maz59]. We note that in the literature, the property that every embedded 2-sphere bounds a 3-ball is called *irreducible*.

The *connected sum* $M\#N$ of two m-manifolds M and N is the manifold obtained by cutting an embedded m-ball out of each of the two manifolds and gluing the two resulting boundary $(m-1)$-spheres via a homeomorphism. The connected sum is an associative and commutative operation with S^m being the identity, $M\#S^m = M$ (cf. Alexander's Theorem for $m = 3$).

Lemma 2.4.2. *The fundamental group of an open set in $\mathbb{R}^3$ is torsion free.*

Proof. Let $U \subseteq \mathbb{R}^3$ be an open set and in particular a 3-manifold.

Assume $\pi_1(U)$ has torsion. Then $\pi_1(U)$ has a subgroup of finite order. By [Hem76, Theorem 9.8] the set U decomposes into $U = N\#U'$ such that N is compact and $\pi_1(N)$ is finite and non-trivial; Note that we can consider N not only as a subset of $\mathbb{R}^3$ but also as a subset of S^3.

If N has no boundary, then we can directly skip to the last step. Otherwise we eliminate the boundary as follows:

The boundary of N is a surface. We claim that the boundary of N is a disjoint union of spheres:
$$\partial N = \coprod S^2. \tag{2.19}$$
Clearly, the boundary is a disjoint union of surfaces. If a surface S_g of it has genus $g \geq 1$, then N contains a non-contractable loop along S_g (cf. figure 2.13) and $\mathbb{Z} \subseteq \pi_1(N)$ which is a contradiction to $\pi_1(N)$ being finite. Thus, the boundary of N is a disjoint union of spheres.

According to van Kampen's theorem[Hat15, Theorem 1.20, p. 43] gluing N with a 3-ball B^3 along an S^2 boundary component does not change the first fundamental group of N because
$$\pi_1(N \amalg_{\partial S^2} B^3) \cong \pi_1(N) *_{\underbrace{\pi_1(S^2)}_{=0}} \underbrace{\pi_1(B^3)}_{=0} \cong \pi_1(N). \tag{2.20}$$
Therefore, $\pi_1(N \amalg_{\partial N} \coprod B^3) \cong \pi_1(N)$ and $N \amalg_{\partial N} \coprod B^3$ has no boundary.

Recall that $N \subseteq S^3$. Because $S^3 \setminus S^2 = B^3 \amalg B^3$, equation (2.18), we can choose the 3-ball being glued to be the B^3 of this union which does not contain N. This allows us to consider the gluing of N with the balls as taking unions of subsets of S^3.

Hence, $N \amalg_{\partial N} \coprod B^3$ is a subset of S^3, which is closed and has no boundary, thus is equal to S^3. Therefore,

$$\pi_1(N) \cong \pi_1(N \amalg_{\partial N} \coprod B^3) = \pi_1(S^3) \tag{2.21}$$

but the left side is a finite, non-trivial group and the right side is the trivial group, a contradiction. Therefore, our premise must be false and $\pi_1(U)$ is torsion free. ■

Remark 2.4.3. The answer for the question of a torsion free fundamental group is negative for $\mathbb{R}^m$ with $m \geq 4$ by the observation that a (thickened) projective plane can be embedded into $\mathbb{R}^4$; the thickening is needed as we are talking about open sets. Let us elaborate this in more detail.

We use the representation of the projective plane as the 2-sphere in $\mathbb{R}^3$ with antipodal points identified, i.e. $\mathbb{P}^2(\mathbb{R}) := \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}/_\sim$ where $(x, y, z) \sim (-x, -y, -z)$. The local inverse of the quotient map gives rise to a smooth structure on $\mathbb{P}^2(\mathbb{R})$ and it is thus a smooth manifold.

The smooth map $S^2 \to \mathbb{R}^4$, $(x, y, z) \mapsto (xy, yz, zx, ax^2 + by^2 + cz^2)$ with distinct constants $a, b, c \in \mathbb{R}$ descends to this model of the projective plane and gives rise to an injective immersion of $\mathbb{P}^2(\mathbb{R})$ in $\mathbb{R}^4$, in other words an embedding. Thus, we can regard $\mathbb{P}^2(\mathbb{R})$ as a smooth submanifold of $\mathbb{R}^4$.

By [Lee13, Theorem 6.24, p. 139] this smooth submanifold of $\mathbb{R}^4$ has an open tubular neighbourhood in $\mathbb{R}^4$ and by [Lee13, Theorem 6.25, p. 140] it is a retract of the embedded projective plane. In particular it has the same fundamental group.

To summarise, the tubular neighbourhood of an embedded projective plane in $\mathbb{R}^4$ is an open set whose fundamental group is $\mathbb{Z}/2\mathbb{Z}$, i.e. is not torsion-free.

The above can be generalised from $\mathbb{R}^4$ to any $\mathbb{R}^m$ with $m \geq 4$. Therefore we have the following lemma:

Lemma 2.4.4. *There exist open sets in $\mathbb{R}^m$, $m \geq 4$, with the first fundamental group being of finite order. In particular there exist open sets the first fundamental group of which have torsion.*

These sets also really do occur in our setting with translation manifolds:

Example 2.4.5. Let $M \subseteq \mathbb{R}^m$, $m \geq 4$, be an open subset with finite first fundamental group. This set is a translation manifold with the identity as atlas. Its universal cover $p\colon \tilde{M} \to M$ is finite and $\tilde{M}$ is a translation manifold and p is a translation covering. The developing map $D\colon \tilde{M} \to \mathbb{R}^m$ is the covering map p. The singularities of $\tilde{M}$ are only the boundary points and the shadows do not produce any points outside of it, i.e. $S^{\mathrm{f}}_M(\partial M) = \partial M$.

This is therefore an example of a translation manifold which fulfills every precondition of theorem 2.3.21 except being torsion free.

We can combine these results into the following theorem:

Theorem 2.4.6. *Let M be a 2- or 3-dimensional translation manifold and $\Sigma = \overline{M} \setminus M$ be the singularities. Furthermore, let $U \subseteq M$ be such that the developing map descends to a well-defined map $D\colon U \to \mathbb{R}^m$. If*

i) $U \setminus S_U(\Sigma)$ *is connected and simply connected, and*

ii) either

  a) $D\colon U \setminus S_U(\Sigma) \to D(U) \setminus D_U(\Sigma)$ *is a finite covering,*

  or

  b) $D^{-1}(x)$ *is finite for all* $x \in D(U) \setminus D_U(\Sigma)$*, and*

  c) $D(U) \setminus D_U(\Sigma)$ *is open,*

then the singularities which are removable in $D(U) \setminus D_U(\Sigma)$ *are removable in* $U \setminus S_U(\Sigma)$*.*

Proof. The statement follows from theorems 2.3.21 and 2.4.1 and lemma 2.4.2 ∎

Morally the above theorem states that a singularity in three dimensions is removable if and only if it has a neighbourhood which is simply connected.

2.5 Conclusion

For the ansatz using coverings to describe singularities, not only the singularities themselves have to be considered but also the shadows they cast on other leafs of the covering. This concept of this kind was expected as a comparison of square-tiled surfaces over the punctured torus shows.

In particular, in the last section we have seen that the crucial point for a singularity to be removable is the simply connectedness of a neighbourhood. The simply connectedness implies two things: first the existence of a well-defined developing map on the neighbourhood, and second the removability of the singularity itself.

Morally, the simply connectedness also states that a singularity has to be big-enough in terms of its ‘dimension’ because ‘low-dimensional’ subsets only yield simply connected sets, cf. the punctured ball in 2 vs. 3 dimension. For arbitrary topological spaces the term dimension has to be filled with a proper definition but in all the examples we have seen so far its meaning should be obvious.

Theorem 2.1.1 covers almost all if not all situations which we would call tame singularities (or finite translation manifold). Thus, only ‘codimension two’ singularities exist in finite translation manifolds.

3 Cubic Translation Manifolds

In this chapter we investigate a special type of translation manifold. The type of translation manifold we want to examine are cubic ones, that is a translation manifold which is glued from cubes. In the case of dimension two they are known as *square tiled surfaces* or also as *orgamis* and are from major interest as they provide a rich structure combined with the fact that they are dense in the set of all finite translation surfaces. Furthermore, they give important examples and counterexamples – we have already seen that in example 1.1.26. Moreover, in the future they will guide us the way to understand singularities in more depth.

Because a cubic translation manifold has only one building block – a cube – it admits a covering over the torus. This in turn allows us to give some explicit description of this manifold and to construct suitable and well-behaved neighbourhoods together with translation coverings around any singularity of our interest. Because of this, we will be able to classify their singularities and give a concrete description of them.

In section 3.1 we give the definition of a cubic translation manifold as well as some notation and background on cubes. Most importantly we define what a cubic neighbourhood is (definition 3.1.9), which is an adapted neighbourhood which allows to exercise control over the singularity resulting in a complete classification of them in the next section.

After that, in section 3.2 we take a closer look at the singularities of a cubic translation manifold. Using the cubic neighbourhoods, we start with the classification of singularities of codimension two. Then we show in theorem 3.2.7 that singularities of codimension greater than two are always removable. Finally, we give a complete classification of the singularities and their intersection behaviour in theorem 3.2.29.

3.1 Definition of a Cubic Translation Manifold

Throughout this chapter we denote by T^* the translation manifold built out of a single m-dimensional unit cube in $\mathbb{R}^m$ where opposite faces are identified but without its $(m-2)$-skeleton. The $(m-2)$-skeleton are the (removable) singularities of T^*. We can use the standard identification of the torus with $\mathbb{R}^m/\mathbb{Z}^m$ to obtain a concrete model for T^*. Because the torus is homogenous we can – or must, depending on the point of view – choose also the skeleton. In this model we regard the points with integer coordinates as the vertices of the unit cube. Thus, a point of the $(m-2)$-skeleton has at least two integer entries in its coordinates. Furthermore, we denote by c the centre of T^* – we could use any point here and choose the centre purely because its more memorable. In our concrete model the centre c would be $c := (0.5, ..., 0.5) + \mathbb{Z}^m \in T^*$. Sometimes we will shorten the edge length of our cubes, in this case we also shrink T^* by the same factor.

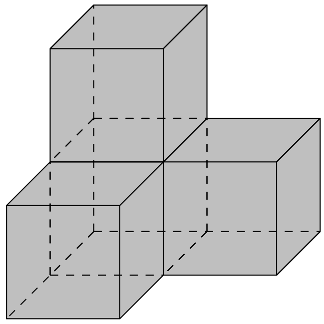

Figure 3.1: The *three-dimensional L-shaped* translation manifold. It is glued out of four cubes and the remaining sides are glued with their opposite by a translation in x-, y-, or z-direction. All the drawn edges are not part of the cubic translation manifold and form the singularities.

Definition 3.1.1 (Cubic Translation Manifold I). A *cubic translation manifold* of dimension m is a translation manifold M glued out of m-dimensional cubes, all of the same size, as described in definition 1.1.14.

We say M is *finite* iff it is glued out of finitely many cubes.

Example 3.1.2. a) The simplest example is the torus T^* itself, which is the unique translation manifold built out of a single cube.

b) $\mathbb{R}^m$ is an infinite cubic translation manifold. To strictly obey our definition, we must exclude its codimension 2 skeleton but we do not reflect that in our notation and write $\mathbb{R}^m$ for the cubic translation manifold with and without its (removable) singularities.

c) Another example glued out of four cubes can be seen in figure 3.1. We have already encountered this translation manifold in example 1.1.26, where we give some more details about its singularities. See also figure 3.3 were we discuss the singularities of it.

d) If M is a cubic translation manifold of dimension m, then $M \times T^n$ is a cubic translation manifold of dimension $m + n$, where we already added some removable singularities. Here T^n denotes the n-dimensional torus.

Because a cubic translation manifold M consists only of cubes of the same size and shape, we have a map $p\colon M \to T^*$. This map is well-defined not only in the interior of the cubes but also on the faces and is a translation covering map.

Since coverings are connected to the first fundamental group, it is good to know what the fundamental group of the torus is:

Lemma 3.1.3. *The fundamental group of T^* is the free group in m generators:*

$$\pi_1(T^*) = F_m. \tag{3.1}$$

The loops which go through c in positive direction x_i, $i = 1, ..., m$, form a basis for $\pi_1(T^*)$. We will denote the elements of this basis also by x_i.

Proof. Thickening the codimension two skeleton and shrinking the interior of the cube, yields a bouquet of m circles, one circle for each dimension. ■

Each cube in M can be identified with its centre and the set of centres is exactly the fibre $p^{-1}(\{c\})$. Therefore, we can identify the cubes and the fibre of the covering $p\colon M \to T^*$. The usual operation of the generators x_i on the fibre becomes an operation on the cubes. This operation encodes the gluing pattern of the cubic translation manifold because the generator x_i, which is the path in direction x_i maps a cube C to the cube which is glued to the face of C in positive x_i-direction.

Similarly, if we have a translation covering $p\colon M \to T^*$, then we can use p to partition M into cubes and for these cubes we get a gluing pattern by lifting the generators $x_i \in \pi_1(T^*)$ to M. Thus, M is glued out of cubes and hence a cubic translation manifold.

With this we have just shown that the following definition and definition 3.1.1 are equivalent:

Definition 3.1.4 (Cubic Translation Manifold II). A *cubic translation manifold* of dimension m is a translation manifold M together with a translation covering $p\colon M \to T^*$.

M is a *finite* cubic translation manifold iff the covering map is finite, i.e. if each fibre $p^{-1}(\{x\})$ for $x \in T^*$ is finite.

Remark 3.1.5. Instead of starting with a translation covering between translation manifolds, we can also start with a topological covering map $p\colon M \to T^*$ between a topological space M and the torus and then endow M with the translation structure of T^* lifted via p, cf. remark 1.2.2.

3.1.1 Notation

From now on we assume that M is connected.

Covering Map

Results of the covering map theory allows us to identify the elements in the fibre $p^{-1}(\{c\})$ with the elements in the quotient $p_*(\pi_1(M))\backslash\pi_1(T^*)$. Note that the quotient is only a group if the covering is normal, i.e. if $p_*(\pi_1(M))$ is a normal subgroup in $\pi_1(T^*)$.

We introduce the following short-hand notations: For a cube C and an element $g \in \pi_1(T^*)$, we denote by Cg the cube under the permutation of the fibre (and hence the cubes) induced by g. Under the identification of cubes and cosets the permutation corresponds to the usual right multiplication.

Since the fundamental group and coverings are sensitive to the basepoint, we will fix some cube C_1 of M as base and its centre c_1 as basepoint. Using C_1 we can map each element $g \in \pi_1(T^*)$ to the cube $C_1 g$. Under the identification of cubes in M with elements in the fibre $p^{-1}(\{c\})$ which in turn can be identified with elements in $p_*(\pi_1(M))\backslash\pi_1(T^*)$, the base cube C_1 corresponds to the coset $p_*(\pi_1(M, c_1))$. We picked the notation C_1

instead of C_0 because we are writing the concatenation of group elements in $\pi_1(T^*)$ multiplicatively.

To ease our notation we will suppress the injective map p_* and directly regard $\pi_1(M, c_1)$ as a subgroup of $\pi_1(T^*, c)$, e.g. C_1 becomes $\pi_1(M, c_1)$.

Cubes

The number of k-faces, i.e. faces of dimension k, in an m-dimensional cube is

$$2^{m-k}\binom{m}{k}. \tag{3.2}$$

We can see this by realising that a point in a k-face in $\mathbb{R}^m$ has $m-k$ entries which are fixed to an integer value and k entries without that restriction. The entry in an integer coordinate can, in the representation $[0,1]^m \subseteq \mathbb{R}^m$, either be 0 or 1. Since the entries of the integer coordinates can be chosen freely, there are $\binom{m}{k}$ different k-faces which are not translates of each other. Moreover, a k-face in $\mathbb{R}^m$ is contained in exactly 2^{m-k} different cubes.

Cubic neighbourhood

In the upcoming discussion we will need some special neighbourhoods for a k-face. Since a cubic translation manifold is built out of cubes we can utilise this to define a nice neighbourhood which is also built out of cubes.

Definition 3.1.6 (Relative Interior). Let M be a cubic translation manifold and A a k-face in M. The *relative interior* of A, written $\operatorname{relint}(A)$, is A without its $(k-1)$-faces.

Remark 3.1.7. The name relative interior is motivated by the following observation. Let A be a face in M of a cube C. Then we can also think of A as a face $A_{\mathbb{R}}$ of a cube $C_{\mathbb{R}}$ in $\mathbb{R}^m$ because the (unglued) cube C can be embedded into $\mathbb{R}^m$. The relative interior of A (in M) is the image of the relative interior of $A_{\mathbb{R}}$ in $\mathbb{R}^m$ (in the meaning of affine geometry of $\mathbb{R}^m$) under this identification. Recall that for a subset $X \subseteq \mathbb{R}^m$ the relative interior of X is defined as $\operatorname{relint}(X) := \{x \in X \mid \exists \varepsilon > 0 \colon B(x, \varepsilon) \subseteq X \cap \operatorname{aff}(X)\}$, where $\operatorname{aff}(X)$ is the affine hull of X in $\mathbb{R}^m$.

Example 3.1.8. The relative interior of a vertex, i.e. a point, is the vertex itself.

For the three-dimensional torus T^* and an edge A the relative interior is the edge without the vertex of the torus although the edge A is 'glued at its ends', i.e. is topologically the circle S^1.

Definition 3.1.9 (Cubic Neighbourhood). Let M be a cubic translation manifold of dimension m and let A be a k-face of M.

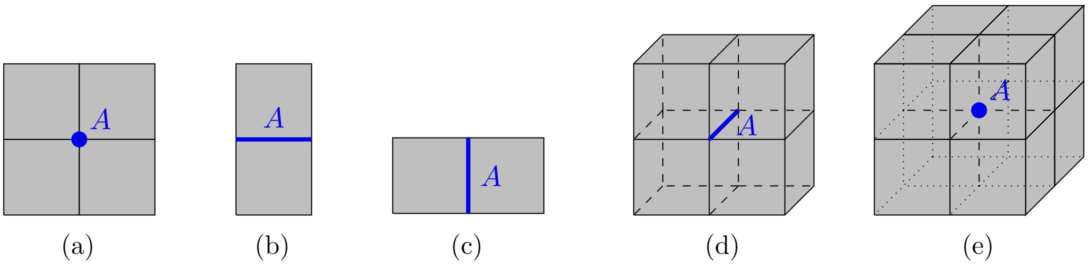

Figure 3.2: Cubic neighbourhoods in $\mathbb{R}^2$ and $\mathbb{R}^3$

The *cubic neighbourhood* $N_M(A)$ of A is the interior of all cubes touching A with the codimension 2 skeleton removed. More precisely

$$N_M(A) := \operatorname{int}\left(\bigcup_{C \in \mathcal{C}(A)} C\right) \setminus \text{codim 2 skeleton} \tag{3.3}$$

where $\mathcal{C}(A) := \{C \text{ closed face of } M \,|\, C \cap \operatorname{relint}(A) \neq \emptyset\}$. This neighbourhood is open and path-connected.

Remark 3.1.10. Taking the interior of the cubes is strictly speaking not necessary. Doing so ensures that we have an open neighbourhood, which has its merits. Using the closed cubes has the advantage that also all faces of the cubes of the cubic neighbourhood are in the neighbourhood except those of codimension 2 and greater, i.e. the blocks we build the closed cubic neighbourhood with are precisely the building blocks which we used for gluing the cubic translation manifold to begin with.

Example 3.1.11. We consider $\mathbb{R}^m$ as a cubic translation manifold (cf. example 3.1.2 b)). Figure 3.2 depicts different cubic neighbourhoods $N_{\mathbb{R}^m}(A)$ for different dimensions and faces: 3.2(a) is the cubic neighbourhood of a 0-face (point) in two dimensions; (b) and (c) are the two different cubic neighbourhoods of a 1-face (edge) in two dimensions; (d) is the cubic neighbourhood of a 1-face (edge) in three dimensions; (e) is the cubic neighbourhood of a 0-face (point) in three dimensions; in (d) and (e) are the edges not part of the neighbourhood.

For an example of a cubic neighbourhood not in $\mathbb{R}^m$ see figure 3.5.

Example 3.1.12. Let T^* be the three dimensional torus without the edges and let A be an edge. Then $N_{T^*}(A) = T^*$. In particular $N_{T^*}(A)$ consists only of one cube where we would have expected four cubes and A touches this one cube in four edges.

We will see later that this poses a problem (which we can circumvent) as $N_{T^*}(A)$ cannot be embedded into $\mathbb{R}^3$.

3.2 Singularities of Cubic Translation Manifolds

As described in the construction in definition 1.1.14, the singularities of a cubic translation manifold are the union of the (glued) $(m-2)$-skeletons of the cubes. Again we want

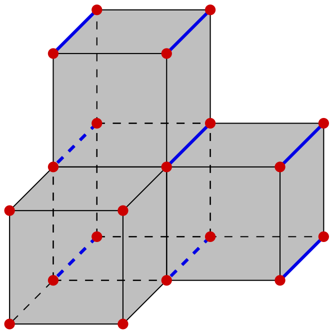

Figure 3.3: The *three-dimensional L-shaped* translation manifold. After the identification (cf. figure 3.1) the translation manifolds has a single vertex and six edges (one of which is highlighted). The non-removable singularities consist of the single vertex and the three edges which are all part of the centre cube. The edge singularities have order 3 each (see proposition 3.2.2).

to classify and understand the singularities, in particular we want to find criteria for a singularity to be removable.

Example 3.2.1. a) The set of singularities of the torus T^* is the $(m-2)$-skeleton and all singularities are removable.

b) $\mathbb{R}^m$ as cubic translation manifold (cf. example 3.1.2 b)) has only removable singularities.

c) Figure 3.3 shows the L-shaped surface and its singularities.

d) For the construction $M \times T^n$ of example 3.1.2 d), the non-removable singularities are $\Sigma \times T^n$, where Σ is the set of non-removable singularities of M. A singularity $\{\sigma\} \times T^n$ is removable if and only if σ is removable in M.

3.2.1 Faces of Codimension 2

Let us start by looking at the codimension 2 faces of a cubic translation manifold – in three dimensions they would correspond to edges. Denote by E the codimension 2 face of C_1 which is perpendicular to x_1 and x_2 and lies in positive x_1 and positive x_2 direction, cf. figure 3.4.

If the relative interior of E only consists of removable singularities, then walking around that face in the x_1-x_2-plane with an angle of 2π will end in the same cube C_1 and at the same point were we started at. Walking around in the described way corresponds to first walking one cube in x_1-direction, then one cube in x_2-direction, then one cube in negative x_1-direction and finally one cube in negative x_2-direction. Note that the cubes we are walking in might all be the same cube (as it is the case for the torus T^*) but entered from different directions. As described earlier in definition 3.1.4 and section 3.1.1, we can express the walk with the group operation of the fundamental

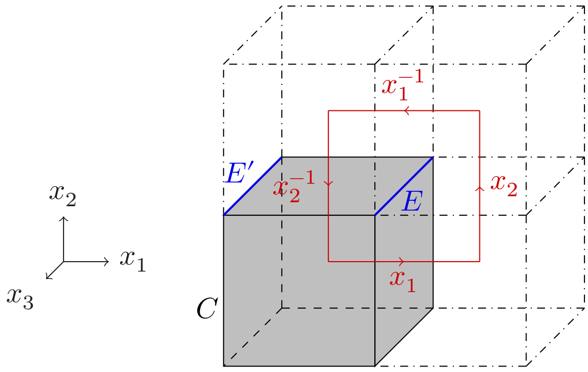

Figure 3.4: The edge E is the codimension two face of the cube C perpendicular to the x_1-x_2-plane and in $(+x_1,+x_2)$-direction. The edge E' is also perpendicular to the x_1-x_2-plane but in $(-x_1,+x_2)$-direction.
If E is removable, then the path $x_1x_2x_1^{-1}x_2^{-1}$ is a closed path starting and ending in C.

group. Therefore, E is removable if and only if $C_1 = C_1x_1x_2x_1^{-1}x_2^{-1} = C_1[x_1,x_2]$, where $[x_1,x_2] \in \pi_1(T^*,c)$ denotes the commutator of x_1 and x_2 considered as elements in $\pi_1(T^*,c)$. If we express the above equation in terms of coset, where $C_1 = \pi_1(M,c_1)$, we get $\pi_1(M,c_1)x_1x_2x_1^{-1}x_2^{-1} = \pi_1(M,c_1)$ in $\pi_1(T^*,c)$, which is equivalent to $[x_1,x_2] \in \pi_1(M,c_1) \subseteq \pi_1(T^*,c)$ (recall that we consider $\pi_1(M,c_1)$ as a subgroup of $\pi_1(T^*,c)$). This is a nice group theoretic description of a removable codimension 2 singularity.

Now let E be a codimension 2 face of any cube C, not necessarily the base cube, but still being in positive x_1 and positive x_2 direction. We can express C as C_1g with $g \in \pi_1(T^*,c)$ a path whose lift to M connects C_1 with C. The above reasoning still applies and we obtain that E is removable if and only if $C[x_1,x_2] = C$ if and only if $g[x_1,x_2]g^{-1} \in \pi_1(M,c_1) \subseteq \pi_1(T^*,c)$ The last equation corresponds to starting in c_1, then first walking along g to the cube C, then following the commutator around E and eventually walking back along g^{-1} to c_1 yielding a closed path.

If the codimension two face E is not removable, then, in order to come back to the starting point, we have to walk around E with an angle of $2\pi k$ for a certain k, which corresponds to using the k^{th} power of the commutator.

Summarising the findings above, we have showed the following proposition:

Proposition 3.2.2. *Let M be a cubic translation manifold, with basepoint $c_1 \in M$ and let $C = C_1g = \pi_1(M,c_1)g$ be a cube in M, where $g \in \pi_1(T^*,c)$ and C_1 is the cube containing the basepoint c_1.*

For the codimension two face E of the cube C perpendicular to the x_i-x_j-plane and located in the $(\pm x_i, \pm x_j)$-direction the following is equivalent:

(i) $C[x_i^{\pm}, x_j^{\pm}]^k = C$

(ii) $g[x_i^{\pm 1}, x_j^{\pm 1}]^k g^{-1} \in \pi_1(M,c_1)$.

In this case we say that E has order k*.*

If M is a normal covering over T^*, then we have the following statement:

Corollary 3.2.3. *In the situation of proposition 3.2.2, if $p\colon M \to T^*$ is a normal covering map, then the following is equivalent:*

(i) E has order $k \in \mathbb{Z}$,

(ii) $g[x_i^{\pm 1}, x_j^{\mp 1}]^k g^{-1} = 1$ in the quotient $\pi_1(M, c_1)\backslash\pi_1(T^*, c) \cong \mathrm{Deck}(M/T^*)$.

Remark 3.2.4. The second entry in proposition 3.2.2 has two interpretations:
First, the interpretation as (lifted) paths in M, which asks the question whether the concatenation of lifted paths yields a closed path in M.
Second, the interpretation as elements in the fundamental group of the torus T^*, where the question becomes whether the product of paths lies in a certain subgroup, viz. $\pi_1(M, c_1)$.

Remark 3.2.5. Forming a local tubular neighbourhood around the interior of a codimension two face E and orthogonally projecting this neighbourhood to the plane perpendicular to E yields a translation surface, where E corresponds to the singularity in the centre and that singularity has the same order as E.

Example 3.2.6. In the construction of examples 3.1.2 d) and 3.2.1 d), if $\sigma \subseteq \Sigma$ is a codimension two face in M, then $\sigma \times T^n$ is a codimension two face in $M \times T^n$ and the order of $\sigma \times T^n$ is equal to the order of σ.
This also shows that codimension two faces of any order exist: It is well-known (or easily constructable) that cubic translation manifolds of dimension two – better known as square-tiled surfaces – can have singularities of any order $k \in \mathbb{N}$. The above construction then yields singularities of order k in any dimension m.

3.2.2 Faces of Codimension 3 and higher

Next we are looking at the lower dimensional parts of the skeleton of the cubic translation manifold. In three dimension this would be the vertices/corners of the cubes.

Theorem 3.2.7. *Let M be a cubic translation manifold of dimension m and σ a point in the skeleton of codimension k with $k \geq 3$. If σ is not part of a codimension two face which is non-removable, i.e. not part of the relative boundary of a codimension two face with order greater than one, then σ is removable.*

In other words, in a cubic translation manifold the singularities must have codimension two unless they are removable.

Proof. In dimension $m = 2$ there is nothing to prove. Let $m \geq 3$ and σ be a singularity in the skeleton of codimension k. It suffices to show that there are no singularities in the 0-skeleton, i.e. the statement that vertices are regular except when part of a singular

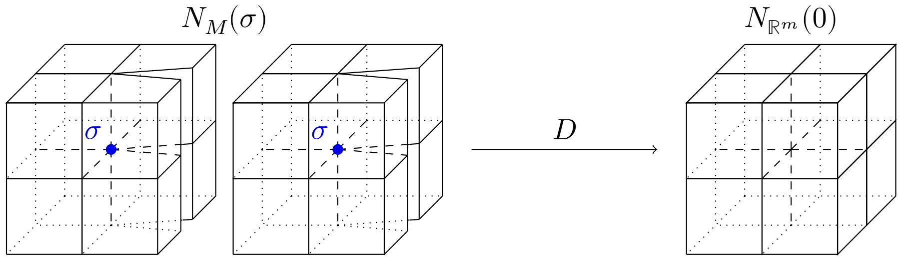

Figure 3.5: A cube of dimension m has 2^m codimension three faces in $\mathbb{R}^m$, therefore σ touches 2^m different 'kinds' of cubes in M. Here each type occurs twice. Note that the outer sides are *not* glued.

codimension two face. Indeed, if one point of a k-face is singular, then all points of that face are singular and since a k-face contains its $(k-1)$-skeleton, $(k-2)$-skeleton, … it contains a point of the 0-skeleton. So let σ be a singular vertex (i.e. 0-face) which is not contained in a codimension two face of order greater than one.

We subdivide the cubes of M into smaller cubes with a third of the edge length and call this refined cubic translation manifold M'. Note that a point p is removable in M if and only if it is removable in M' because the chart used for removing p can be used in M and M'. We take the neighbourhood $N_{M'}(\sigma)$ of σ, which consists of all the refined cubes having σ as one of their vertices (cf. definition 3.1.9 and figures 3.2(a) and 3.2(e)). A cube in dimension m has 2^m vertices, so one vertex is contained in 2^m cubes and because we split the edge length these cubes are distinct cubes in M'.

The developing map for this neighbourhood $N_{M'}(\sigma)$ descends to a well-defined map $D\colon N_{M'}(\sigma) \to N_{\mathbb{R}^m}(0)$ because the neighbourhood is glued out of equally sized cubes and the position of the vertex σ in the cube uniquely determines one of the 2^m cubes in the image $N_{\mathbb{R}^m}(0)$, see figure 3.5 and since the 'outer sides' of $N_{M'}(\sigma)$ are not glued to any other side of $N_{M'}(\sigma)$.

Because all codimension two faces consist of removable singularities, we can remove them yielding a new slightly larger neighbourhood U and extend D to it. Then the image of U under D in $\mathbb{R}^m$ consists of 2^m glued cubes where only the central codimension three faces are omitted. In particular the image $D(U) \subseteq \mathbb{R}^m$ is simply connected. Thus, corollary 2.3.19 is applicable and since the codimension three part in the image is removable in $\mathbb{R}^m$, so is the singularity σ. ∎

Remark 3.2.8. Here is a sketch for proving theorem 3.2.7 in a different way using the fundamental group of $N_{M'}(\sigma)$. The induced map D is for the same reason as above a covering map. Looking at the generators of $\pi_1(N_{\mathbb{R}^m}(0))$, i.e. the paths around the codimension two faces in $\mathbb{R}^m$, we can find pre-images under D_* in $N_{M'}(\sigma)$, which shows that $D_*\colon \pi_1(N_{M'}(\sigma)) \to \pi_1(N_{\mathbb{R}^m}(0))$ is surjective. Because D is a covering map, D_* is also injective, hence $\pi_1(N_{M'}(\sigma)) \cong \pi_1(N_{\mathbb{R}^m}(0))$ and D is the identity covering.

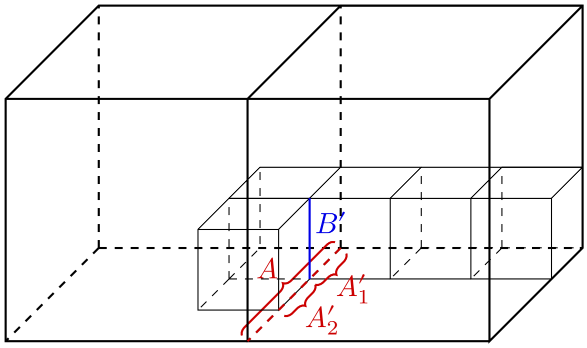

Figure 3.6: When subdividing a single cubes gets split into 3^m smaller cubes. The edge A gets split into three subparts A'_1, A'_2 and A'_3 (the last one is not named in the figure).
A side effect is that new singularities arise for example the vertical edge B'. However, these new singularities are all removable, cf. lemma 3.2.11.

3.2.3 Subdividing Cubic Translation Manifolds

In the proof of theorem 3.2.7 we have subdivided the cubes of the cubic translation manifold M into smaller cubes yielding a new cubic translation manifold M' which is the same as M except being built out of smaller cubes. We did this to ensure that the cubes we use to construct $N_{M'}(\sigma)$ are distinct and yield a covering of (a subset of) $\mathbb{R}^m$.

In this subsection we will discuss the process of subdividing showing that it is sadly necessary but merrily a technical detail imposed by our choice of cubic neighbourhood and not a feature of cubic translation manifolds or their singularities themselves. We could define the cubic neighbourhood differently but this leads to other difficulties which can be surmounted by dealing with a lot of corner cases resulting in more complex definitions which in turn need more thorough argumentations later on. In this work we do not follow this alternative path.

Definition 3.2.9. Let M be a cubic translation manifold of dimension m. We denote by M' the cubic translation manifold which is essentially M but each cube is subdivided into 3^m equally sized smaller cubes. In other words, we divided the edge length by 3. As we are considering M' to be a cubic translation manifold this also means that we remove all codimension 2 faces with respect to the smaller cubes as these are now considered singularities.

Definition 3.2.10. Let M be a cubic translation manifold, M' its subdivision, $k \in \mathbb{N}_0$ and A a k-face of M. We call a ℓ-face A' of M' a *subpart* of A iff $k = \ell$ and $A' \subseteq A$, cf. figure 3.6.

The additional singularities are a technical artifact as the following lemma shows:

Lemma 3.2.11. *Let M be a cubic translation manifold and M' be its subdivision. All new singularities which are created by the subdividing process, i.e. all singularities of M' which are not singularities of M, are removable (cf. figure 3.6).*

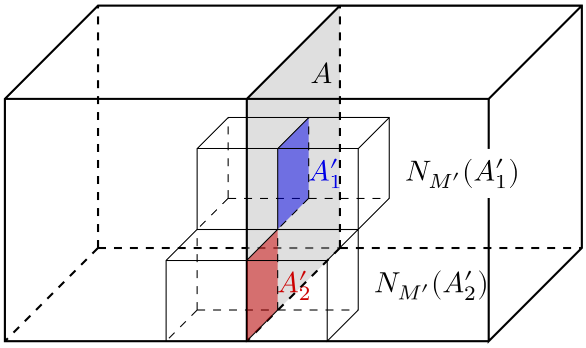

Figure 3.7: The face A in a translation manifold M together with two subparts A_1' and A_2' and their cubic neighbourhoods $N_{M'}(A_1')$ and $N_{M'}(A_2')$, respectively. The subparts, and thus their cubic neighbourhoods in M', are isometric by a translation along A.

Proof. A visual argument for the statement is given by realizing what subdivision does: It merily renames some interior points of a cube to be now part of the codimension 2 skeleton, i.e. to be singularities.

A more thorough argument is this: The metric completion $\overline{M'}$ of M' can be identified with the metric completion $\overline{M}$ of M because M' is dense in M. If σ is a singularity of M', i.e. $\sigma \in \overline{M'} \setminus M'$, which is not a singularity of M, i.e. $\sigma \notin \overline{M} \setminus M$, then $\sigma \in M \setminus M'$ and σ was a regular point of M. Hence, σ has a neighbourhood and a chart compatible with the translation atlas of M. The same chart is also compatible with the translation atlas of M' as the latter is a subset of the atlas of M. Thus, σ is removable. ∎

Remark 3.2.12. Let M be a cubic translation manifold built out of unit cubes. Subdividing M and considering it to be glued out of smaller cubes, is equivalent to scaling M up and considering it to be built out 3^m times as many unit cubes as before.

The following lemma emphasizes that the subdivision is more a technical artifact and does not change the description of a singular face, namely that they have isomorphic neighbourhoods:

Lemma 3.2.13. *Let M be a translation manifold, M' its subdivision and A a k-face of M. If A' and B' are k-faces of M' with $A', B' \subseteq A$, then $N_{M'}(A')$ and $N_{M'}(B')$ are isometric as translation manifolds.*

Proof. In a cube C of M the parts A' and B' of the face A are related by a translation. Applying this translation cubewise in M to all cubes of the neighbourhood $N_{M'}(A')$ gives an isometry between $N_{M'}(A')$ and $N_{M'}(B')$, cf. figure 3.7. ∎

Lemma 3.2.14. *Consider $\mathbb{R}^m$ as cubic translation manifold and let A be a k-face of $\mathbb{R}^m$ and A' a subpart of A, i.e. a k-face in $(\mathbb{R}^m)'$ with $A' \subseteq A$. Then,*

$$\pi_1(N_{\mathbb{R}^m}(A)) \cong \pi_1(N_{(\mathbb{R}^m)'}(A')). \tag{3.4}$$

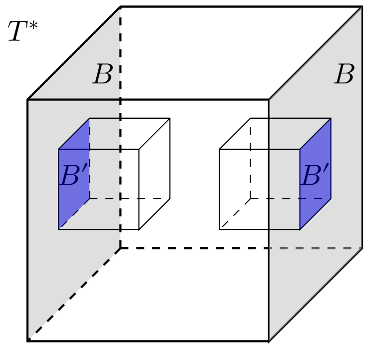

Figure 3.8: The three-dimensional torus T^* with (glued) face B. The cubic neighbourhood of B consists of only one cube, viz. T^*. Hence, cannot be a covering of a cubic neighbourhood in $\mathbb{R}^3$.
After subdividing, the cubic neighbourhood of the subpart B' consists of two cubes. Hence, we obtain a covering over a cubic neighbourhood in $\mathbb{R}^3$.

Proof. In the case of $\mathbb{R}^m$ the subdivided cubic translation manifold $(\mathbb{R}^m)'$ is just a scaled version of $\mathbb{R}^m$. ∎

The following example shows that subdividing the cubes is necessary if we want to use the cubic neighbourhood for a covering like in the proof of theorem 3.2.7.

Remark 3.2.15. Consider the torus T^* (cf. section 3.1). To keep things simple, we consider the three-dimensional torus, though other dimensions exhibit the same problems. Let A be an edge of T^*, i.e. an edge of the cube the torus is built of.

(i) If we do not subdivide the torus, then we have for the above definition of cubic neighbourhood $N_{T^*}(A) = T^*$. Later we want that this cubic neighbourhood covers a cubic neighbourhood $N_{\mathbb{R}^3}(B)$ of an edge B of $\mathbb{R}^3$ (as cubic translation manifold). But $N_{T^*}(A) = T^*$ cannot cover $N_{\mathbb{R}^3}(B)$ as the former consists of one cube and the latter consists of four cubes, see also figure 3.8.

   If we subdivide first, then the one cube of the torus splits into multiple cubes and $N_{(T^*)'}(A)$ has enough cubes to cover $N_{(\mathbb{R}^3)'}(B)$ like we want.

   We can also see the problem by using covering theory: The fundamental group of $N_{\mathbb{R}^3}(B)$ is $\mathbb{Z}$ which does not have the free group F_2 as a subgroup, which is the fundamental group of $N_{T^*}(A) = T^*$ by lemma 3.1.3.

(ii) If we want to use $1/3 \times 1/3 \times 1$ cuboids as the building blocks for the cubic neighbourhood of A (now better named cuboidal neighbourhood), we must break the cycle in the direction of the edge A; see figure 3.9. This can be achieved by removing the facets on the ends of the edge, i.e. those which intersect with A in a single vertex, but not the facets which have A as an edge. This works but has the drawback that it is difficult to describe and that it cannot be isometrically embedded into $\mathbb{R}^3$.

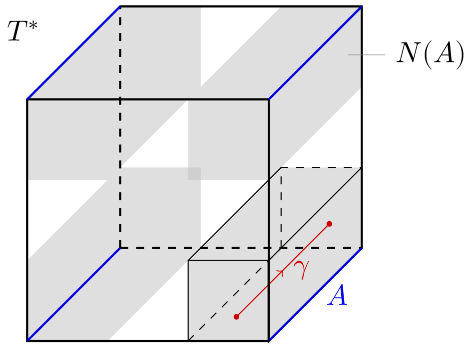

Figure 3.9: Cuboidal neighbourhood $N(A)$ of the edge A in the torus T^*. This is not a covering over a cuboidal neighbourhood of $\mathbb{R}^3$ because the fundamental group of $N(A)$ also contains the non-trivial loop γ which is not present in $\mathbb{R}^3$ because the front and back are not glued as it is the case in the torus.

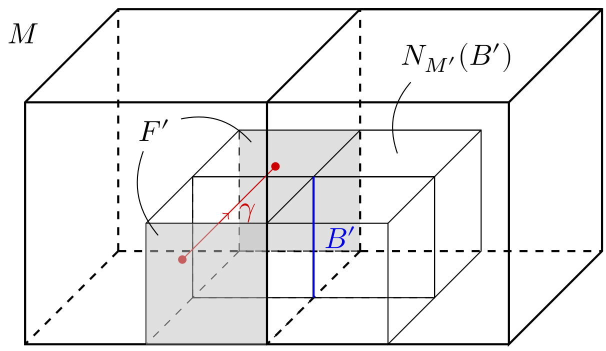

Figure 3.10: M are two cubes as depicted with opposite sides glued (a cuboidal torus). After the subdivision by the factor of 2, the neighbourhood $N_{M'}(B')$ of the newly created edge B' is not a covering of a cubic neighbourhood of $\mathbb{R}^3$ because the fundamental group of $N_{M'}(B')$ has the additional loop γ.

Remark 3.2.16. The reason that we subdivide by a factor of 3 instead of a factor of 2 is more another technicality. We have to decide which faces A' of M' are allowed when considering cubic neighbourhoods $N_{M'}(A')$. If we only allow k-faces A' that are a subpart of a k-face A of M *for the same* k, then factor 2 would be enough. This restriction, for example, rules out the vertex point which appears in the middle of an edge.

However, if we allow all k-faces A' of M', then we need a factor making the edge length strictly smaller than $1/2$ otherwise we have the following problem:

For the three dimensional torus T^* and edge A, let B' in $(T^*)'$ be the newly created edge orthogonal to A, similar to figure 3.10. The cubic neighbourhood $N_{(T^*)'}(B')$ of B' are four cubes which are also glued on the outer sides facing away in directions orthogonal to B. This neighbourhood deformation retracts to $\bigvee_{i=1}^{3} S^1$, three circles glued together on a single point and its fundamental group is the free group in three generators F_3. As above, this neighbourhood cannot cover a cubic neighbourhood in $\mathbb{R}^3$.

3.2.4 Classification of Singularities

We have seen in proposition 3.2.2 that the order k of a codimension two face classifies it completely. What is missing is the classification of the places where two such singularities meet, i.e. the codimension k faces with $k \geq 3$. To classify those we again use the map D induced by the developing map on a neighbourhood similar to the one in the proof of theorem 3.2.7.

As promoted earlier, in this section we will encounter the point where we see that we need to subdivide the cubes into smaller cubes in order to obtain a covering map for our cubic neighbourhoods (definition 3.1.9). As mentioned before this can be avoided by defining a suitable neighbourhood for faces differently but this in turn would lead to more complex argumentations elsewhere. More philosophically, by choosing to subdivide we make the neighbourhoods smaller which is morally the better choice because order and type of a singularity should be regarded as – and indeed are – local properties of that point.

Remark 3.2.17. Throughout this section we formulate and prove all the statements with respect to the original (i.e. not subdivided) cubic translation manifold M. This increases the length of the statements and increases the complexity of the arguments.

To get the gist it is therefore recommended to consider M already subdivided and then read all the statements for $M = M'$ ignoring the subdivision and choosing of subparts altogether.

For the classification of singularities we first need to define when two singularities are considered equal. There are two kinds of equality for translation manifolds which pop into mind: via translations and via isometries.

Definition 3.2.18 (Isomorphism for Singularities/Faces). Let M and N be cubic translation manifolds and denote by M' and N' their respective subdivisions. Let A be a k-face and B be an ℓ-face of M and N, respectively.

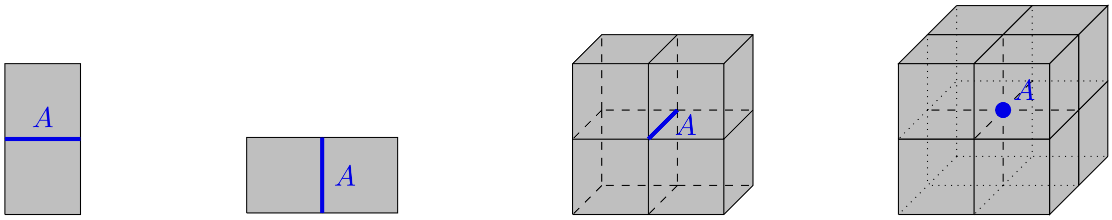

(a) (b) (c) (d)

Figure 3.11: The faces A (as a singularity) in (a) and (b) are isometric but not translation-isomorphic (because they are rotated by $\pi/2$). Neither of them is isometric to the faces of (c) or (d) and thus also not translation-isomorphic. Note that the faces A in (a) and (c) are isometric when considered as metric spaces, i.e. as a line and *without* their neighbourhood, but they are not isometric as a singularity of a translation manifold.

(i) We say that A and B (as a singularity) are *translation-isomorphic* or *isomorphic via a translation* iff there exists a k-face A' of M' which is a subpart of A, i.e. $A' \subseteq A$, and a ℓ-face B' of N' which is a subpart of B, i.e. $B' \subseteq B$, such that $N_{M'}(A')$ and $N_{N'}(B')$ are isomorphic as translation manifolds.

In this case we write $A \cong B$ and denote the equivalence class of a face A as $[A]_{\cong}$.

(ii) We say that A and B (as a singularity) are *isometric* or *isometric-isomorphic* or *isomorphic via an isometry* iff there exists a k-face A' of M' which is a subpart of A, i.e. $A' \subseteq A$, and a ℓ-face B' of N' which is a subpart of A, i.e. $B' \subseteq B$, such that $N_{M'}(A')$ and $N_{N'}(B')$ are isometric, i.e. isomorphic as Riemannian manifolds.

In this case we write $A \approxeq B$ and denote the equivalence class of a face A as $[A]_{\approxeq}$.

Some examples are given in figure 3.11.

Remark 3.2.19. By lemma 3.2.13 all the neighbourhoods of subparts of a face A are isomorphic as translation surfaces. Thus, the above definition is well-defined with respect to the choices of A' and B'.

Remark 3.2.20. The subdividing of the cubic translation manifolds is necessary as otherwise glueings might change the metric structure. Thus, the neighbourhoods might not be isometric any more. For example, we want to consider the edge of the torus T^* and the edge of $\mathbb{R}^m$ translation-isomorphic. However, the edge in the torus is glued to a S^1 and is thus not isometric to an edge in the Euclidean space.

Remark 3.2.21. We may skip subdividing the cubic translation manifold given that no cube is glued to itself (not even in a vertex) and any two cubes are only glued in a single face. By glued in a single face we mean that the two cubes are glued along a face A and all other faces in which they touch each other are contained in that face A. The fine print is necessary because if two cubes are glued along an edge then they are also glued in the vertices of that edge which are faces in their own right.

Remark 3.2.22. The isomorphism up to translation is the canonic definition for a translation manifold where nearly everything is considered up to translation. However, when we are purely interested in the 'look' of a singularity we often do not care whether it is a line along the x-axis or along the y-axis, we only care that it is a line. In this case the second definition is the more useful one.

For cubic translation manifolds this is enough as we have a nice cubic lattice fixing all faces and hence singularities into discrete places and directions. In contrast, for a generic translation manifold this should be relaxed even a little further, e.g. such that a straight line and a bent line 'with the same order' are considered isomorphic singularities. However, these considerations are not scope of this work.

Lemma 3.2.23. *Let A be a k-face of M and B be an ℓ-face of N with M and N cubic translation manifolds not necessarily of the same dimension. If $k \neq \ell$, then $A \not\cong B$ and $A \not\approxeq B$; or equivalently $A \approxeq B \Rightarrow k = \ell$.*

Proof. If M and N have different dimension, then the neighbourhoods of A and B cannot be homeomorphic. Thus, the cubic neighbourhoods of A and B consist of cubes of the same dimension. We now prove the contraposition of the statement above: $A \approxeq B \Rightarrow k = \ell$.

If the cubic neighbourhoods of A and B are isometric, then surely, A and B, which are in the centre of the neighbourhoods, are isometric. Thus, they have the same dimension. ∎

Lemma 3.2.24 (Equivalence classes of faces in $\mathbb{R}^m$). *Consider $\mathbb{R}^m$ as cubic translation manifold and let $k \in \mathbb{N}_0$.*

(i) All k-faces are isomorphic up to isometry, i.e. there is only one equivalence class up to isometry. Stated differently, for a k-face A and a ℓ-face B we have $A \approxeq B \Leftrightarrow k = \ell$.

(ii) There are $\binom{m}{k}$ different k-faces up to translation.

Proof. If two faces A and B of two cubes have the same dimension k, then – because we are in $\mathbb{R}^m$ – we can use an isometry to map one to the other. Thus, their cubic neighbourhoods are isometric and there is only one equivalence class up to isometry.

The second statement follows from counting the k-faces, which we have already done in section 3.1.1. ∎

Example 3.2.25. In a three-dimensional cube there are $2^2\binom{3}{1} = 12$ different 1-faces but only three different cubic neighbourhoods up to translation; see figure 3.12. The three different neighbourhoods correspond to the three directions, which in turn correspond to the $\binom{m}{k}$ factor. The position of the edge within the cube (corresponding to the 2^{m-k} factor) does not matter for the neighbourhood as we can use translations to switch to another cube where the edge is located at a different position relative to the cube.

Before we state the classification theorem, we need one last things: the shape of the singularity itself, i.e. how it looks without its neighbourhood.

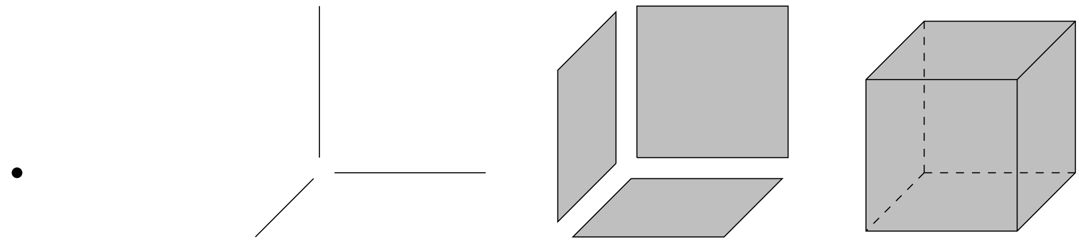

$k = 0$ $k = 1$ $k = 2$ $k = 3$

Figure 3.12: Representatives of all equivalence classes up to translation of k-faces of $\mathbb{R}^3$.

Definition 3.2.26 (Shape of a Face/Singularity). We consider $\mathbb{R}^m$ to be a translation manifold glued out of cubes (cf. example 3.1.2 b)). Let M be a cubic translation manifold of dimension m and A a k-face of it. The face A belongs to a cube and that cube can be identified with a cube of $\mathbb{R}^m$ ignoring possible glueings of faces. The *shape* of the face A is the corresponding face of that cube in $\mathbb{R}^m$, which we denote by A_{shape}. Note that A_{shape} depends on the cube choosen in $\mathbb{R}^m$, however, the shape becomes unique when considered up to translation, i.e. we get a well-defined map $A \mapsto [A_{\text{shape}}]_{\cong}$. Similar, the shape up to isometry is the equivalence class $[A_{\text{shape}}]_{\cong}$ and is also well-defined.

Because the choice of the cube in $\mathbb{R}^m$ is not important, we will call any representative of the equivalence class $[A_{\text{shape}}]_{\cong}$ the *shape* of A.

Remark 3.2.27. Often it is not important whether a k-face of a cube is part of an (abstract) cubic translation manifold M or part of $\mathbb{R}^m$, in this case we denote both, the k-face A in M and the k-face A_{shape} in $\mathbb{R}^m$, i.e. its shape, by the symbol A.

Remark 3.2.28. In $\mathbb{R}^m$ there is – by definition – no difference between the face A and its shape A_{shape}. In particular the shape determines the cubic neighbourhood of the face in $\mathbb{R}^m$.

Theorem 3.2.29 (Classifying Singularities of Cubic Translation Manifolds). *Let $m \geq 2$ and $k \in \{0, \dots, m-2\}$. Further, fix a k-face F of $\mathbb{R}^m$ as cubic translation manifold – this determines the shape of the singularities we classify.*

The k-faces (up to isometry) of cubic translation manifolds which have the same shape as F up to isometry are classified by the conjugacy classes of the subgroups of $\pi_1(N_{\mathbb{R}^m}(F))$, i.e. we have a bijection

$$\left\{ [A]_{\cong} \;\middle|\; A \textit{ face with } [A_{\text{shape}}]_{\cong} = [F]_{\cong} \right\} \\ 1:1 \updownarrow \\ \left\{ \textit{conjugacy classes of subgroups of } \pi_1(N_{\mathbb{R}^m}(F)) \right\} \tag{3.5}$$

Likewise, we have a classification of the singularities up to translation: The k-faces (up to translation) of cubic translation manifolds which have the same shape as F up to *translation are classified by the conjugacy classes of the subgroups of* $\pi_1(N_{\mathbb{R}^m}(F))$*, i.e. we have a bijection*

$$\begin{gathered}\Big\{[A]_{\cong} \;\Big|\; A \text{ face with } [A_{\text{shape}}]_{\cong} = [F]_{\cong}\Big\} \\ 1:1 \updownarrow \\ \Big\{\text{conjugacy classes of subgroups of } \pi_1(N_{\mathbb{R}^m}(F))\Big\}\end{gathered} \tag{3.6}$$

Remark 3.2.30. The first bijection classifies a singular face up to isometry, i.e. it classifies all singularities which 'look' like the reference face F: point, line, plane, …

The second bijection is a little more fine-grained as it also fixes the orientation of the singularity in addition to its 'look': point, line in x-direction, x-y-plane, …

The reason that the lower set is equal in both cases, is that we take slightly different equivalence classes of the faces in the upper set. The equivalence relation chosen reflects what we are interested in: if we only want the 'look' without orientation then there is no need to use the translation-isomorphic relation.

Proof. Let M be a cubic translation manifold of dimension m and A a face of dimension k we want to classify. Denote by M' the subdivided cubic translation manifold and pick a k-face A' of M' with $A' \subseteq A$, i.e. A' is a subpart of A. We argue with this subpart as isomorphism for edges is defined via subparts (see definition 3.2.18).

Because we subdivided the cubes, all cubes of $N_{M'}(A')$ touch A' only in one face (counted without the identification by gluing). Thus, the developing map induces a well-defined map $D\colon N_{M'}(A') \to N_{(\mathbb{R}^m)'}(A'_{\text{shape}})$ and D is a covering map, which we can check easily as we are mapping cubes of $N_{M'}$ to cubes of $N_{(\mathbb{R}^m)'}$. We want to emphasise that the whole point of subdividing the cubes was to ensure that this delevoping map becomes a covering (in a 'neighbourhood' of A), i.e. 'has enough cubes in its domain', cf. remark 3.2.15

Note that for a different representative of $[A_{\text{shape}}]_{\cong}$ the cubic neighbourhood of a subpart of it is isometric to $N_{\mathbb{R}^m}(A'_{\text{shape}})$, thus in particular isometric to $N_{\mathbb{R}^m}(F')$. Similar, if B is a side of another cubic translation manifold N and A and B are isometric-isomorphic, then their neighbourhoods of subparts A' and B' are isometric. Moreover, they also yield isomorphic coverings of $N_{(\mathbb{R}^m)'}(A'_{\text{shape}})$.

From covering theory we know that the covering spaces of $N_{(\mathbb{R}^m)'}(A'_{\text{shape}})$ are classified by the conjugacy classes of subgroups of its fundamental group $\pi_1(N_{(\mathbb{R}^m)'}(A'_{\text{shape}}))$. Since the choice of the subpart A' of A does not matter (lemma 3.2.13) and because $A_{\text{shape}} \cong F$, this becomes a bijection to the conjugacy classes of subgroups of $\pi_1(N_{(\mathbb{R}^m)'}(F'))$. Finally, by lemma 3.2.14 we can replace $\pi_1(N_{(\mathbb{R}^m)'}(F'))$ with $\pi_1(N_{\mathbb{R}^m}(F))$ yielding the desired bijection between the faces up to isometry and the subgroups of $\pi_1(N_{\mathbb{R}^m}(F))$ up to conjugation.

Likewise, if instead of isometry-isomorphic we take isomorphic via translation, then the above reasoning still applies. However, we have restricted us to those coverings where the orientation of the face also matches the face F. ∎

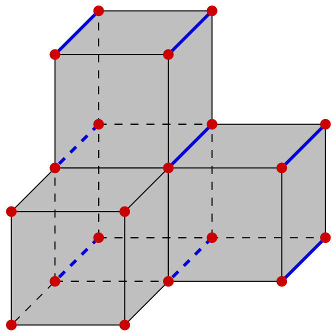

Figure 3.13: We have already seen that the edges around the central cube have order three. The vertex has order four when considered as a 0-face on its own because the subdivided cubic neighbourhood consists of 32 small cubes.

From the above classification we can derive two invariants for singularities:

Definition 3.2.31 (Type and Order of a Singularity). Let M be a cubic translation manifold of dimension m and A a k-face of the singular set $\Sigma = \overline{M} \setminus M$, i.e. a face in the codimension two skeleton of M. We denote by D the well-defined map $D\colon N_{M'}(A') \to N_{(\mathbb{R}^m)'}(A'_{\text{shape}})$ induced by the developing map on the subdivided manifolds of the proof of theorem 3.2.29 where $A' \subseteq A$ is a k-face of M', i.e. a subpart of A.

The *type* of A is given by the conjugacy class of the subgroup $D_*(\pi_1(N_{M'}(A')))$ in $\pi_1(N_{(\mathbb{R}^m)'}(A'_{\text{shape}}))$.

The *order* of A is given by the index $[\pi_1(N_{(\mathbb{R}^m)'}(A'_{\text{shape}})) : D_*(\pi_1(N_{M'}(A')))]$ of the subgroup $D_*(\pi_1(N_{M'}(A')))$ in $\pi_1(N_{(\mathbb{R}^m)'}(A'_{\text{shape}}))$.

Remark 3.2.32. By lemma 3.2.13 the choice of A' does not matter and the above definitions are well-defined. Moreover, by lemma 3.2.14 we could even replace $\pi_1(N_{(\mathbb{R}^m)'}(A'_{\text{shape}}))$ with $\pi_1(N_{\mathbb{R}^m}(A_{\text{shape}}))$ (but sadly not $\pi_1(N_{M'}(A'))$ with $\pi_1(N_M(A))$).

Remark 3.2.33. Be aware that order and type is defined for faces not for individual points of M which can be part of multiple faces! We can define the order and type of a point $\sigma \in M$ to be the order and type of the face containing σ with the lowest dimension, see figure 3.13.

Remark 3.2.34. Let A be a face of codimension 2 in a cubic translation manifold. If we compare the above result of theorem 3.2.29 with the investigation for faces of codimension 2 of section 3.2.1, then we see that $N_{(\mathbb{R}^m)'}(A'_{\text{shape}})$ consists of $2^2 = 4$ glued cubes and (a scaled version of) $[x_i, x_j]$ is precisely the generator of $\pi_1(N_{(\mathbb{R}^m)'}(A'_{\text{shape}}))$. The order of the generator in $\pi_1(N_{(\mathbb{R}^m)'}(A'_{\text{shape}}))/D_*(\pi_1(N_{M'}(A')))$ corresponds to the order of A as defined earlier in proposition 3.2.2. Thus, the new definition coincides with our old one and is also applicable to faces other than codimension 2.

Corollary 3.2.35 (Classification of Singularities). *The shape $F \subseteq \mathbb{R}^m$ (as metric subspace of $\mathbb{R}^m$) and type $G \subseteq \pi_1(N_{\mathbb{R}^m}(F))$ (as subgroup) completely determine a face/singularity up to translation or isometry.*

Proof. Follows directly from theorem 3.2.29 as this is merely a rephrasing of that statement with the new terms. ∎

Corollary 3.2.36 (Classification of Singularities up to Isometry). *The numbers $m, k \in \mathbb{N}_0$ with $0 \leq k \leq m$ and type $G \subseteq \pi_1(N_{\mathbb{R}^m}(A))$ (as subgroup) of an arbitrary k-face A in $\mathbb{R}^m$ completely determine a face/singularity up to isometry.*

Proof. By lemma 3.2.24, $\mathbb{R}^m$ has only one k-face up to isometry. The claim follows now from corollary 3.2.35. ∎

Remark 3.2.37. The type as abstract group does not determine the singularity uniquely. For example, a neighbourhood of a vertex σ of $\mathbb{R}^2$ (as cubic translation manifold) has fundamental group $\pi_1(N_{\mathbb{R}^2}(\sigma)) = \mathbb{Z}$. All non-trivial subgroups of $\mathbb{Z}$ are of the form $n\mathbb{Z}$ for some $n \in \mathbb{Z} \setminus \{0\}$ and thus are all isomorphic as abstract groups.

Likewise, the type as (abstract) subgroup of an (abstract) group alone also does not determine the singularity uniquely. For example, a vertex of $\mathbb{R}^2$ and an edge of $\mathbb{R}^3$ (both as cubic translation manifold) have the same type, namely $\mathbb{Z} \subseteq \mathbb{Z}$, but are obviously not isomorphic.

4 Conclusion and Prospects

We have seen how the concepts of (finite) translation surfaces can be generalised to higher dimensions. An important observation was that the different definitions which are equivalent in dimensions two, split into distinct concepts in higher dimensions. The consequence is that results proven for one definition do not necessarily apply to all translation manifolds. However, they provide a good indication of what to expect.

Our research on the singularities shows that the developing map and translation coverings are useful tools to describe and compare singularities and their nature. We have seen criteria when singularities are removable using these methods. We explicitly gave a construction for how to remove isolated singularities and generalised the idea to make it applicable to arbitrary singularities.

All the examples we have seen show that all the non-removable singularities are of codimension two, i.e. if singularities are too 'low-dimensional', then they become removable. For arbitrary topological spaces the definition of dimension is difficult. However, corollary 2.3.19 and theorem 2.3.21 give criteria for removability using simply connected neighbourhoods and this matches up with the observation of the dimensions: If the singularities have codimension two, e.g. a point in a plane or a curve in $\mathbb{R}^3$, then they induce a non-trivial fundamental group. On the other hand, if they are of lower dimension, then the fundamental group is trivial, e.g. a point in $\mathbb{R}^3$ or a curve in $\mathbb{R}^4$.

Based on our investigation and the examples we have seen, in particular in the chapter about cubic translation manifolds, we propose the following definition:

Definition 4.0.1. A singularity $\sigma \in \overline{M} \setminus M$ is *tame* if and only if there exists a neighbourhood $U' \subseteq \overline{M}$ of σ such that the developing map of $U := U' \cap M$ descends to a well defined map $D\colon U \to \mathbb{R}^m$.

In this case we define the *type* and *order* of the singularity as:
The *type* is the group $D_*(\pi_1(U))$ as subgroup of $\pi_1(D(U))$.
The *order* is the index $[\pi_1(D(U)) : D_*(\pi_1(U))]$ of the subgroup $D_*(\pi_1(U))$ in $\pi_1(D(U))$.

This definition coincides with the classification of singularities in two dimensions: For a conic singularity with angle $2\pi k$ the type would be $k\mathbb{Z}$ as subgroup of $\pi_1(\dot{B}(0,1)) \cong \mathbb{Z}$ and the order is $[\mathbb{Z} : k\mathbb{Z}] = k$. It also coincides with the definitions for cubic translation manifolds (see chapter 3).

Please note that we did not require D to be a covering. We believe that the requirement of theorem 2.3.16 that $D(U) \setminus D_U(\Sigma)$ is open in $D(U)$ can be satisfied by choosing U appropriately. Then, theorem 2.3.16 is applicable and D is a covering map on $U \setminus S_U(\Sigma)$. In this case the type of the singularity is the subgroup of the fundamental group corresponding to that covering.

Bibliography

[Ahl16] Lars Valerian Ahlfors. *Riemann Surfaces*. Ed. by Leo Sario. Princeton, NJ, 2016. URL: https://doi.org/10.1515/9781400874538.

[Bal95] Werner Ballmann. *Lectures on spaces of nonpositive curvature : with an appendix by Misha Brin Ergodicity of geodesic flows*. DMV-Seminar ; 25. Basel: Birkhäuser, 1995. ISBN: 3-7643-5242-6 ; 0-8176-5242-6. URL: https://people.mpim-bonn.mpg.de/hwbllmnn/archiv/NPC0606.pdf.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen ; 319. Berlin: Springer, 1999. ISBN: 3-540-64324-9.

[Boo86] William M. Boothby. *An introduction to differentiable manifolds and Riemannian geometry*. 2. ed. Pure and applied mathematics ; 120. Orlando, Fla.: Academic Pr., 1986. ISBN: 0-12-116052-1 ; 0-12-116053-X.

[Bro60] Morton Brown. 'A proof of the generalized Schoenflies theorem'. In: *Bulletin of the American Mathematical Society* 66 (1960), pp. 74–76. ISSN: 0002-9904. DOI: https://doi.org/10.1090/S0002-9904-1960-10400-4.

[Cha04] Reza Chamanara. 'Affine automorphism groups of surfaces of infinite type'. In: *Contemporary mathematics* 355 (2004): *In the tradition of Ahlfors and Bers, III*. Ed. by William Abikoff and Andrew Haas, pp. 123–145.

[Coh35] Stefan Cohn-Vossen. 'Existenz kürzester Wege'. de. In: *Doklady SSSR* 8 (1935), pp. 339–342.

[Coh36] Stefan Cohn-Vossen. 'Existenz kürzester Wege'. de. In: *Compositio Mathematica* 3 (1936), pp. 441–452. URL: http://www.numdam.org/item/CM_1936__3__441_0.

[Hat07] Allen Hatcher. *Notes on Basic 3-Manifold Topology*. 2007. URL: http://pi.math.cornell.edu/~hatcher/3M/3Mdownloads.html.

[Hat15] Allen Hatcher. *Algebraic topology*. 14th printing 2015. Cambridge: Cambridge University Press, 2015. ISBN: 9780521795401 ; 9780521791601. URL: http://pi.math.cornell.edu/~hatcher/AT/AT.pdf.

[Hem76] John Hempel. *3-manifolds*. Annals of mathematics studies ; 86. Princeton, N.J.: Princeton Univ. Press, 1976. ISBN: 0691081786; 0691081832.

[Lee13] John M. Lee. *Introduction to smooth manifolds*. 2. ed. Graduate texts in mathematics ; 218. New York: Springer, 2013. ISBN: 978-1-4419-9981-8.

[Maz59] Berry Mazur. 'On embeddings of spheres'. In: *Bulletin of the American Mathematical Society* 65 (1959), pp. 59–65. DOI: https://doi.org/10.1090/S0002-9904-1959-10274-3.

[Rad26] Tibor Radó. 'Über den Begriff der Riemannschen Flache'. In: *Acta Scientiarum Mathematicarum* 2.2-2 (1924-26), pp. 101–121.

[Thu97] William P. Thurston. *Three-dimensional geometry and topology.* Vol. 1: Princeton mathematical series ; 35. Princeton, NJ: Princeton Univ. Press, 1997. ISBN: 0691083045; 9780691083049.

[Whi61] J. H. C. Whitehead. 'The Immersion of an Open 3-Manifold in Euclidean 3-Space'. In: *Proceedings of the London Mathematical Society* s3-11.1 (1961), pp. 81–90. DOI: 10.1112/plms/s3-11.1.81. eprint: https://londmathsoc.onlinelibrary.wiley.com/doi/pdf/10.1112/plms/s3-11.1.81. URL: https://londmathsoc.onlinelibrary.wiley.com/doi/abs/10.1112/plms/s3-11.1.81.

Index