

Astérisque

THIERRY LÉVY

Two-dimensional Markovian holonomy fields

Astérisque, tome 329 (2010)

http://www.numdam.org/item?id=AST_2010__329__R1_0

© Société mathématique de France, 2010, tous droits réservés.

L'accès aux archives de la collection « Astérisque » (<http://smf4.emath.fr/Publications/Asterisque/>) implique l'accord avec les conditions générales d'utilisation (<http://www.numdam.org/conditions>). Toute utilisation commerciale ou impression systématique est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

NUMDAM

Article numérisé dans le cadre du programme
Numérisation de documents anciens mathématiques

<http://www.numdam.org/>

329

ASTÉRISQUE

2010

**TWO-DIMENSIONAL MARKOVIAN
HOLONOMY FIELDS**

Thierry LÉVY

SOCIÉTÉ MATHÉMATIQUE DE FRANCE

Publié avec le concours du CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

Thierry Lévy

CNRS, École Normale Supérieure, 45, rue d'Ulm, 75005 Paris (France).

E-mail : levy@dma.ens.fr

Url : <http://www.dma.ens.fr/~levy/>

2000 Mathematics Subject Classification. — 60G60, 60J99, 60G51, 60B15, 57M20, 57M12, 81T13, 81T27.

Key words and phrases. — Random field, Markov process, Lévy process, compact Lie group, graph on a surface, rectifiable path, ramified covering, Yang-Mills, lattice gauge theory, continuous limit.

TWO-DIMENSIONAL MARKOVIAN HOLONOMY FIELDS

Thierry Lévy

Abstract. — This text defines and studies a class of stochastic processes indexed by curves drawn on a compact surface and taking their values in a compact Lie group. We call these processes two-dimensional Markovian holonomy fields. The prototype of these processes, and the only one to have been constructed before the present work, is the canonical process under the Yang-Mills measure, first defined by Ambar Sengupta and later by the author. The Yang-Mills measure sits in the class of Markovian holonomy fields very much like the Brownian motion in the class of Lévy processes. We prove that every regular Markovian holonomy field determines a Lévy process of a certain class on the Lie group in which it takes its values, and we construct, for each Lévy process in this class, a Markovian holonomy field to which it is associated. When the Lie group is in fact a finite group, we give an alternative construction of this Markovian holonomy field as the monodromy of a random ramified principal bundle. Heuristically, this agrees with the physical origin of the Yang-Mills measure as the holonomy of a random connection on a principal bundle.

Résumé (Champs d'holonomie markoviens bidimensionnels). — Ce travail est consacré à la définition et à l'étude d'une classe de processus stochastiques indexés par des chemins tracés sur une surface, qui prennent leurs valeurs dans un groupe de Lie compact et qui satisfont une propriété d'indépendance conditionnelle analogue à la propriété de Markov. Nous appelons ces processus des champs d'holonomie markoviens bidimensionnels. L'exemple fondamental de cette sorte de processus est le processus canonique sous la mesure de Yang-Mills, qui a été construite d'abord par Ambar Sengupta puis plus tard par l'auteur. C'est aussi le seul champ d'holonomie markovien qui ait été construit avant ce travail. Le processus canonique sous la mesure de Yang-Mills est assez exactement aux champs d'holonomie markoviens ce que le mouvement brownien est aux processus de Lévy. Deux de nos principaux résultats affirment qu'à tout champ d'holonomie markovien suffisamment régulier est associé un processus de Lévy d'une certaine classe sur le groupe de Lie dans lequel il prend ses valeurs et réciproquement que pour tout processus de Lévy dans cette classe il existe un champ

d'holonomie markovien auquel il est associé. Dans le cas particulier où le groupe de Lie considéré est un groupe fini, nous parvenons à réaliser ce champ d'holonomie markovien comme la monodromie d'un fibré principal ramifié aléatoire. Ceci nous rapproche de l'interprétation originelle de la mesure de Yang-Mills, issue de la théorie quantique des champs, comme mesure de probabilités sur l'espace des connexions sur un fibré principal.

CONTENTS

Introduction	1
0.1. A 1-dimensional analogue	1
0.2. A first definition	4
0.3. Transition kernels	6
0.4. Two fundamental examples	9
0.5. Markovian holonomy fields and gauge fields	12
0.6. Finite groups, gauge fields and ramified coverings	14
0.7. Markovian holonomy fields and the Yang-Mills measure	15
0.8. Topological quantum field theories	16
0.9. Structure	20
0.10. References	21
0.11. Acknowledgements	22
1. Surfaces and graphs	23
1.1. Surfaces	23
1.1.1. Classification of surfaces	23
1.1.2. Surgery of surfaces	26
1.2. Curves and paths	30
1.2.1. Definitions	30
1.2.2. The complete metric space of rectifiable paths	34
1.3. Graphs	41
1.3.1. Graphs and the sewing of patterns	41
1.3.2. The boundary of a face	48
1.3.3. Adjunction of edges	59
1.3.4. The group of loops in a graph	63
1.3.5. Graphs with one face	67
1.4. Riemannian metrics	69
1.4.1. Measured and Riemannian marked surfaces	69
1.4.2. Partially ordered sets of graphs	72
1.4.3. Approximation of graphs	74
2. Multiplicative processes indexed by paths	79
2.1. Multiplicative functions	79
2.2. Multiplicative families of random variables	83

2.3. Uniform multiplicative functions on a graph	86
2.3.1. Constraints on marked surfaces	86
2.3.2. Uniform measures	87
2.3.3. Surgery of uniform measures	90
2.4. Tame generators of the group of reduced loops	94
3. Markovian holonomy fields	107
3.1. Definition	107
3.2. Discrete Markovian holonomy fields	110
3.3. An abstract extension theorem	114
3.4. Extension of discrete holonomy fields	120
4. Lévy processes and Markovian holonomy fields	125
4.1. The partition functions of a Markovian holonomy field	125
4.2. The Lévy process associated to a Markovian holonomy field	134
4.3. A Markovian holonomy field for each Lévy process	140
4.3.1. A discrete Markovian holonomy field	141
4.3.2. A Markovian holonomy field	146
5. Random ramified coverings	151
5.1. Ramified G -bundles	151
5.2. Monodromy of ramified G -bundles	153
5.3. Measured spaces of ramified G -bundles	155
5.4. The monodromy process as a Markovian holonomy field	159
Index	165
Bibliography	169

INTRODUCTION

The elementary theory of Markov processes establishes a correspondence between several types of objects among which transition semigroups and stochastic processes. These stochastic processes can take their values in fairly general spaces, but they are usually indexed by a subset of the real numbers, for the Markov property relies on the distinction between past and future. In the present work, we investigate a correspondence between certain transition semigroups and another kind of stochastic processes, where the notions of past and future are replaced by the notions of inside and outside. The processes that we consider are indexed by curves, or rather loops, drawn on a surface, and they take their values in a compact Lie group. We call them (2-dimensional) *Markovian holonomy fields*. They are Markovian in the following sense: if some piece of a surface is bounded by a finite collection of loops, then the values of the process on loops located inside this piece and outside this piece are independent given the value of the process on the finite collection of loops which bound this piece.

0.1. A 1-dimensional analogue

Let us start by discussing the 1-dimensional analogues of Markovian holonomy fields, which are just Markov processes looked at from a slightly unusual point of view. Let us choose a transition semigroup $P = (P_t)_{t \geq 0}$ on some state space \mathcal{X} . For each $t \geq 0$, $P_t(x, dy)$ is a transition kernel on $\mathcal{X} \times \mathcal{X}$. Under suitable assumptions, we can associate to P a homogeneous Markov process with values in \mathcal{X} , which we denote by X . This Markov process is not really a single stochastic process, it is rather a collection of processes, essentially one for each initial condition at a specific time. In fact, if we consider X restricted to segments, we can say that to each segment $[a, b] \subset \mathbb{R}$ and each initial condition $x \in \mathcal{X}$ we associate a process $(X_t)_{t \in [a, b]}$ with values in \mathcal{X} such that $X_a = x$ almost surely. Within the structure implied by the fact that $[a, b]$ is a subset of \mathbb{R} , what we really use is the topological structure of this interval, its orientation and our ability to measure the distance between any two of its points. Of course, in the present 1-dimensional setting, this structure suffices to characterise

the interval up to translation, and the last sentence may seem pointless. Its content should however become clearer in the 2-dimensional setting.

Let us push the abstraction a little further and try to define, for all compact 1-dimensional manifold M , a process $(X_t)_{t \in M}$ with values in \mathcal{X} . As we have just observed, we need an orientation of M and a way to measure distances. If M is not connected, let us agree that the restrictions of our process to the various connected components of M will be independent. So, let M be a connected oriented compact Riemannian 1-dimensional manifold. There are not so many options: M is either homeomorphic to a segment or to a circle, it has a certain positive total length, and this information characterises it completely up to orientation-preserving isometry. If M is a segment of length L , it is isometric to $[0, L]$ and there is no difficulty in defining the process $(X_t)_{t \in M}$ given an initial condition. Before turning to the case of the circle, let us interpret the Markov property of X in terms of these 1-dimensional manifolds.

Let M_1 and M_2 be two manifolds as above, isometric to segments. Let $M_1 \cdot M_2$ denote the manifold obtained by identifying the final point of M_1 with the initial point of M_2 . It is still homeomorphic to a segment. Choose an initial condition $x \in \mathcal{X}$. We are able to construct two stochastic processes indexed by $M_1 \cdot M_2$. On one hand, we can take $M_1 \cdot M_2$ as a segment on its own and simply consider the process $(X_t)_{t \in M_1 \cdot M_2}$ with initial condition x . On the other hand, we can also proceed as follows. For all segment M and all $x \in \mathcal{X}$, let $\mathcal{L}(x, M)$ denote the distribution of the process $(X_t)_{t \in M}$ with initial condition x . Let us also denote by $\mathcal{L}(x, M, dy)$ the disintegration of $\mathcal{L}(x, M)$ with respect to the value of X at the final point of M . Then the probability measure $\int_{\mathcal{X}} \mathcal{L}(x, M_1, dy) \otimes \mathcal{L}(y, M_2)$ is the distribution of a process indexed by the disjoint union of M_1 and M_2 which takes the same value at the final point of M_1 and the initial point of M_2 . It can thus be identified with the distribution of a process indexed by $M_1 \cdot M_2$. It is exactly the content of the Markov property of X that the two measures that we have considered are equal:

$$(1) \quad \int_{\mathcal{X}} \mathcal{L}(x, M_1, dy) \otimes \mathcal{L}(y, M_2) = \mathcal{L}(x, M_1 \cdot M_2).$$

This example illustrates in the simplest possible way the fact that the Markov property can be nicely formulated in terms of surgery of manifolds, in this case in terms of concatenation of intervals. Manifolds of dimension 1 undergo another kind of surgery, when the two endpoints of a single interval are glued together (see Figure 1). If we try to mimic (1), we are tempted to define the distribution of a process indexed by the circle S^1 of length L , seen as the interval $[0, L]$ of which the endpoints have been

identified, by

$$(2) \quad \mathcal{L}(S^1) = \int_{\mathcal{X}} \mathcal{L}(x, [0, L], dx),$$

an expression which unfortunately is meaningless. Still, this formula is consistent with the fact that a circle has no boundary, so that there is no initial condition to specify. What we are attempting to define here is a bridge, that is to say a probability measure on closed trajectories in \mathcal{X} , from the transition semigroup $(P_t)_{t \geq 0}$.




FIGURE 1. The surgery of 1-dimensional manifolds.

Without aiming at the greatest possible generality, let us describe a situation in which this is possible. We will assume that \mathcal{X} is a smooth finite-dimensional manifold, for example a vector space or a Lie group, which carries a Borel probability measure μ which is stationary for the semigroup P . We will also assume that for all $t > 0$ and all $x \in \mathcal{X}$, the measure $P_t(x, dy)$ has a continuous density with respect to μ , which we denote by $y \mapsto Q_t(x, y)$. In this situation, it is possible to define bridges of the Markov process X between any two points of \mathcal{X} . Hence, for each segment $[0, L]$, it is possible to define the finite measure $\mathcal{L}(x, [0, L], y)$ which is the conditional distribution of $(X_t)_{t \in [0, L]}$ given $X_0 = x$ and $X_L = y$, multiplied by the real number $Q_L(x, y)$. With this definition, $\mathcal{L}(x, [0, L], y)$ is not a probability measure in general but the relation $\mathcal{L}(x, [0, L], dy) = \mathcal{L}(x, [0, L], y)\mu(dy)$ holds. We can then define a measure on the set of trajectories in \mathcal{X} indexed by a circle of length L by setting

$$(3) \quad \mathcal{L}(S^1) = \int_{\mathcal{X}} \mathcal{L}(x, [0, L], x)\mu(dx).$$

The identification of a trajectory indexed by the circle S^1 with a process indexed by $[0, L]$ requires the choice of a base point on S^1 but the stationarity of μ implies that the resulting definition of $\mathcal{L}(S^1)$ is independent of this choice.

The measure $\mathcal{L}(S^1)$ defined by (3) is not a probability measure in general and, in order to construct a stochastic process indexed by S^1 , one should consider the associated normalised measure. Finite measures, as opposed to probability measures, will play an important role throughout this work.

Let us summarise this discussion of the 1-dimensional case. Starting with a Markovian transition semigroup on \mathcal{X} with good properties, we have been able to associate

to each compact oriented 1-dimensional Riemannian manifold, endowed with an initial condition for each connected component homeomorphic to a segment, a stochastic process with values in \mathcal{X} indexed by the points of this manifold. This collection of processes exhibits a Markovian behaviour with respect to the operations of cutting or concatenation of 1-manifolds.

0.2. A first definition

A 2-dimensional Markovian holonomy field is a 2-dimensional analogue of the object that we have just described. It is a collection of stochastic processes, one for each compact surface endowed with boundary conditions and some way of measuring areas. For each such surface, the process is indexed by a set of *loops* drawn on this surface and the boundary conditions specify the value of the process on the loops which bound the surface. Moreover, this collection of processes satisfies Markovian properties with respect to the operations of cutting surfaces along curves or gluing them along boundary components.

In contrast with the 1-dimensional case, the index sets of our processes now carry a partially defined internal composition law, namely the concatenation of loops. We are interested in stochastic processes which satisfy a property of additivity with respect to the concatenation of loops. This requires that they take their values in a group and since this group will usually not be assumed to be Abelian, and denoted multiplicatively, we will rather call this a property of *multiplicativity*. Let us give a precise definition. If M is a 2-dimensional manifold and m is a point of M , let $L_m(M)$ be a set of loops on M based at m . We will discuss later which loops exactly we wish to consider (see Section 0.4). To each loop $l \in L_m(M)$ we can associate the inverse loop l^{-1} , which is simply l traced backwards. Also, to each pair of loops $l_1, l_2 \in L_m(M)$ we can associate their concatenation which we denote by $l_1 l_2$. Let G be a group, which plays the role of the space \mathcal{X} above. A stochastic process $(H_l)_{l \in L_m(M)}$ with values in G is said to be multiplicative if

$$(4) \quad \forall l \in L_m(M), \quad H_{l^{-1}} = H_l^{-1} \text{ a.s.},$$

$$(5) \quad \forall l_1, l_2 \in L_m(M), \quad H_{l_1 l_2} = H_{l_2} H_{l_1} \text{ a.s.}$$

We will explain later, in Section 0.5, why we choose to reverse the order on the right-hand side of this equality. For the time being, suffice it to say that the processes which constitute a Markovian holonomy field will be multiplicative in the sense just defined.

Let us now describe further the transitions of a Markovian holonomy field. Just as an interval has two extremities, a surface has a boundary which is homeomorphic to a disjoint union of circles. The only difference is that the number of connected

components of the boundary of a connected surface can be any natural number. Another specificity of the 2-dimensional case, as opposed to the 1-dimensional one, is that even when a surface is oriented, its boundary components are indistinguishable from a topological point of view: any permutation of the boundary components can be realised by an orientation-preserving homeomorphism. So, in order to see our surface as realizing a cobordism, that is, a topological transition, between two sets of circles (see Figure 2), we need to arbitrarily declare that some of these components are incoming and the others are outgoing. If the surface is oriented, then we orient incoming boundary components negatively and outgoing boundary components positively. This matters because, according to (4), the boundary conditions on the incoming components are associated with oriented loops: to each oriented incoming boundary component we associate an element of G , and to the same boundary component with the opposite orientation is associated the inverse element of G .




FIGURE 2. A surface seen as a topological transition between two collections of circles

The comparison with the 1-dimensional case reveals that we need to endow our surfaces with a structure analogous to the metric structure of intervals and circles, and the role of which is to incorporate a notion of time into the picture. It turns out, in the present context, that the correct analogue of the ability to measure distances is the ability to measure areas. On a surface, this requires much less than a Riemannian metric, just a Borel measure which is smooth enough to be the Riemannian volume of a Riemannian metric.

We can now give a loose definition of a Markovian holonomy field. The actual definition unfortunately requires a great deal too many prerequisites to be given in this introduction but the curious reader is welcome look it up on page 108.

Definition 0.2.1 (Markovian holonomy fields - first approach). — *Let G be a group. A 2-dimensional G -valued Markovian holonomy field is a collection of G -valued stochastic processes, one for each compact surface endowed with a measure of area and a choice of*

boundary conditions along its incoming boundary components. For each such surface, the G -valued stochastic process is indexed by the set of loops on the surface and it is multiplicative. Moreover, as a whole, the collection of stochastic processes behaves with respect to the surgery of surfaces in a way which is governed by the Markov property.

0.3. Transition kernels

The transition kernels of a Markovian holonomy field describe the distribution of the process on the outgoing circles of a surface given its values on the incoming ones. There is a much greater variety of situations than in the 1-dimensional case, but fortunately for us, this variety has been very well understood for about a century and is easy to describe. Up to homeomorphism, a connected compact surface is characterised by the fact that it is orientable or not, by the number of connected components of its boundary, and by a single other topological invariant called its genus which can be any non-negative integer if the surface is orientable and any positive integer if it is not. Moreover, according to a theorem of Moser, the only invariant of a smooth measure of area under diffeomorphisms is its total area. It turns out that, when one deals with orientable and non-orientable surfaces at the same time, the genus is not the most convenient way to label the possible topological types of surfaces. We prefer to work with a slightly modified notion of genus, which is simply the genus if the surface is not orientable, and twice the classical genus if it is orientable. The main advantage of this genus is that it is additive with respect to the operation of connected sum.

For all integers p and q such that $0 \leq q \leq p$, let us denote by $\mathcal{M}(G^{p-q})$ the space of probability measures on G^{p-q} . Then the transition kernels of a Markovian holonomy field consist of a collection of mappings

$$(6) \quad P_{p,g,t}^{\pm}(x_1, \dots, x_q, dx_{q+1}, \dots, dx_p) : G^q \longrightarrow \mathcal{M}(G^{p-q}),$$

where the sign indicates the orientability of the surface, g is its genus, t its total area and p the number of its boundary components, of which q are incoming and $p - q$ outgoing.

If $q = 0$, then the transition kernel is a measure on G^p , if $q = p$ it is a function on G^p and if $p = 0$, it is a real number. In order to avoid the problems that we have encountered with the circle in the 1-dimensional case, we will make fairly strong assumptions. Firstly, we will assume that G is a compact Lie group. Such a group carries a unique probability measure invariant by translations, the Haar measure, which we denote by dx . We will also assume that the transition kernels of the Markovian

holonomy field that we consider can be put under the form

$$(7) \quad \begin{aligned} P_{p,g,t}^\pm(x_1, \dots, x_q, dx_{q+1}, \dots, dx_p) \\ = Z_{p,g,t}^\pm(x_1^{-1}, \dots, x_q^{-1}, x_{q+1}, \dots, x_p) dx_{q+1} \cdots dx_p \end{aligned}$$

for some functions $Z_{p,g,t}^\pm$ which we will call the *partition functions* of the holonomy field. The exponents that we have introduced take care of the issue of orientation. They restore the symmetry between the boundary components of a surface, so that the partition functions are invariant by permutation of their arguments.

The possibility of gluing together boundary components of one or two surfaces leads to infinitely many relations analogous to (1) between the transition kernels and to just as many relations between the partition functions. For instance, consider a cylinder of area s with one incoming circle and one outgoing circle. A cylinder has genus 0 and the corresponding transition kernel is $P_{2,0,s}^+(x, dy)$. By gluing the incoming circle of this cylinder along an outgoing circle of an arbitrary surface, we do not change this surface up to homeomorphism, we only increase its area by s . Hence, we have the following relation between transition kernels:

$$(8) \quad \int_G P_{p,g,t}^\pm(x_1, \dots, x_q, dx_{q+1}, \dots, dx_{p-1}, dx) P_{2,0,s}^+(x, dx_p) \\ = P_{p,g,t+s}^\pm(x_1, \dots, x_q, dx_{q+1}, \dots, dx_p).$$




FIGURE 3. Gluing a cylinder does not affect a surface up to homeomorphism but it increases its area.

As another example, let us consider an orientable surface with at least one incoming and one outgoing circle. If we glue these circles one along the other in such a way that the result is still orientable, then we obtain a surface with two less boundary components, a classical genus increased by 1, hence a reduced genus increased by 2, and the same area. This example is reminiscent of the situation where we obtained a circle by identifying the two endpoints of a segment. It is thus not surprising that in this case, the Markov property is best expressed in terms of the partition functions, rather than the transition kernels. It reads

$$(9) \quad \int_G Z_{p,g,t}^+(x_1, \dots, x_{p-2}, x, x^{-1}) dx = Z_{p-2,g+2,t}^+(x_1, \dots, x_{p-2}).$$




FIGURE 4. Gluing together two boundary components of a connected surface.

If we identify an outgoing circle with an incoming one in a non-orientable surface, then the reduced genus of the surface is also increased by 2. In this case, the Kolmogorov-Chapman equation is

$$(10) \quad \int_G Z_{p,g,t}^-(x_1, \dots, x_{p-2}, x, x^{-1}) dx = Z_{p-2,g+2,t}^-(x_1, \dots, x_{p-2}).$$

A more unusual topological operation consists in identifying a boundary component of a surface with itself by a fixed point free orientation-preserving involution (see Figure 2 page 28). Up to homeomorphism, this operation is equivalent to gluing a Möbius band along this boundary component. The resulting surface is always non-orientable and its genus is increased by one. The corresponding relation for partition functions is

$$(11) \quad \int_G Z_{p,g,t}^\pm(x_1, \dots, x_{p-1}, x^2) dx = Z_{p-1,g+1,t}^-(x_1, \dots, x_{p-1}).$$

The equalities (8), (9), (10) and (11) essentially generate all the relations which hold between the partition functions. Using these relations, we will prove that the whole set of partition functions of a Markovian holonomy field is determined by just those associated to a disk and a three-holed sphere of arbitrary areas. Indeed, any surface can be built from these three elementary bricks by a finite number of gluings. In fact, if the Markovian holonomy field is regular enough, then we will prove that its transition kernels are all determined by the sole transition kernels associated to disks. These transition kernels $(P_{1,0,t}^+(dx))_{t>0}$, also written $(Z_{1,0,t}^+(x)dx)_{t>0}$ in terms of the partition functions, constitute a one-parameter family of probability measures on G which will turn out to form a continuous convolution semigroup, hence the collection of 1-dimensional marginals of a classical Lévy process on G , a process indexed by \mathbb{R}_+ and with independent and stationary multiplicative increments.

We are now able to give an idea of two of our main results. After giving an axiomatic definition of a 2-dimensional Markovian holonomy field with values in a compact Lie group G (Definition 3.1.2), we will prove, under a suitable regularity assumption, that there is a classical Lévy process with values in G associated with each Markovian holonomy field (Proposition 4.2.1). This Lévy process is characterised by the fact

that its 1-dimensional distributions are the transition kernels of the holonomy field associated to disks. Moreover, all partition functions of the holonomy field can be expressed in terms of these 1-dimensional distributions (Proposition 4.1.10). Then, we will prove that for all Lévy process which satisfies some regularity properties, there exists a Markovian holonomy field to which this Lévy process is associated (Theorem 4.3.1).

We do not tackle in this work the question of the unicity of a Markovian holonomy field with a given associated Lévy process. When the group G is Abelian, I believe that this unicity holds and that the proof should be straightforward. However, when G is not Abelian, the question seems to be more delicate and the answer is still unclear to me.

0.4. Two fundamental examples

In this paragraph, we present two examples of stochastic processes which do not exactly fit into our definition of Markovian holonomy fields, but are in a sense simpler and should absolutely be kept in mind as fundamental examples. We will also take the second example as an opportunity to discuss the role of loops with finite length.

Consider the plane \mathbb{R}^2 endowed with the Lebesgue measure denoted by dx . Let Π denote a Poisson point process on \mathbb{R}^2 with intensity dx . Let $l : [0, 1] \rightarrow \mathbb{R}^2$ be a continuous loop. For each point $x \in \mathbb{R}^2 \setminus l([0, 1])$, the topological index of l with respect to x is an integer denoted by $n_l(x)$ and defined, if we identify \mathbb{R}^2 with \mathbb{C} , by

$$(12) \quad n_l(x) = \frac{1}{2i\pi} \oint_{\tilde{l}} \frac{dz}{z - x},$$

where \tilde{l} is any piecewise smooth loop which is uniformly close enough to l , for example closer than the distance of x to the range of l . Provided the range of l is negligible, we may define

$$N_l = \int_{\mathbb{R}^2} n_l(x) \Pi(dx) = \sum_{x \in \Pi} n_l(x) dx.$$

Definition 0.4.1. — *The stochastic process $\{N_l : l \text{ has negligible range}\}$ is called the Poisson process indexed by loops on \mathbb{R}^2 .*

Let us define the Brownian motion indexed by loops in parallel to the Poisson process. Let $W : L^2(\mathbb{R}^2, dx) \rightarrow \mathcal{G}$ be a white noise, that is, an isometry into a Hilbert space of centred real Gaussian random variables.

Let us say that a continuous loop $l : [0, 1] \rightarrow \mathbb{R}^2$ *winds gently* if its range is negligible and if $n_l \in L^2(\mathbb{R}^2, dx)$. In this case, it is legitimate to define

$$(13) \quad B_l = W(n_l).$$

Definition 0.4.2. — *The stochastic process $\{B_l : l \text{ winds gently}\}$ is called the Brownian motion indexed by loops on \mathbb{R}^2 .*

To each subset D of the plane, we may associate two σ -fields \mathcal{F}_D^N and \mathcal{F}_D^B , which we define by

$$\begin{aligned}\mathcal{F}_D^N &= \sigma(N_l : l \text{ has negligible range, } l([0, 1]) \subset D), \\ \mathcal{F}_D^B &= \sigma(B_l : l \text{ winds gently, } l([0, 1]) \subset D).\end{aligned}$$

If J is a Jordan curve in \mathbb{R}^2 , then we denote respectively by $\text{int}(J)$ and $\text{ext}(J)$ the bounded and unbounded connected components of $\mathbb{R}^2 \setminus J$, which we call the interior and exterior of J . A basic property of the Poisson point process Π (resp. the white noise W) is that it associates independent random variables to disjoint subsets of the plane (resp. to functions with disjoint supports). This implies the following property.

Proposition 0.4.3. — *Let J_1 and J_2 be two Jordan curves with disjoint interiors. Then the σ -fields $\mathcal{F}_{\text{int}(J_1)}^N$ and $\mathcal{F}_{\text{int}(J_2)}^N$ (resp. $\mathcal{F}_{\text{int}(J_1)}^B$ and $\mathcal{F}_{\text{int}(J_2)}^B$) are independent.*

Proof. — If l is a loop whose range is contained in $\text{int}(J_1)$, then the support of n_l is also contained in $\text{int}(J_1)$. The same holds for J_2 and the result follows. \square

Let us now prove a Markov property.

Proposition 0.4.4. — *Let J be a Jordan curve with negligible range. The σ -fields $\mathcal{F}_{\text{int}(J)}^N$ and $\mathcal{F}_{\text{ext}(J)}^N$ (resp. $\mathcal{F}_{\text{int}(J)}^B$ and $\mathcal{F}_{\text{ext}(J)}^B$) are independent conditionally on N_J (resp. B_J).*

Proof. — We prove the result for the Brownian motion. By the Jordan curve theorem, a Jordan curve with negligible range winds gently. Let L be the line in $L^2(\mathbb{R}, dx)$ generated by $n_J = \mathbf{1}_{\text{int}(J)}$. Let H_{in} (resp. H_{out}) be the closed linear subspace of $L^2(\mathbb{R}, dx)$ spanned by n_l for l which winds gently and $l([0, 1]) \subset \overline{\text{int}(J)}$ (resp. $l([0, 1]) \subset \overline{\text{ext}(J)}$). The inclusions $H_{\text{in}} \subset \{f \in L^2(\mathbb{R}, dx) : \text{supp}(f) \subset \text{int}(J)\}$ and $H_{\text{out}} \subset \{f \in L^2(\mathbb{R}, dx) : f \text{ is constant on } \text{int}(J)\}$ are straightforward. They are actually equalities but we do not need this fact. In particular, $H_{\text{in}} \cap H_{\text{out}} = L$. Moreover, with the notation \ominus for the orthogonal complement, we have

$$H_{\text{in}} \ominus L \subset \{f \in L^2(\mathbb{R}, dx) : \text{supp}(f) \subset \text{int}(J) \text{ and } \int_{\text{int}(J)} f(x) dx = 0\},$$

$$H_{\text{out}} \ominus L \subset \{f \in L^2(\mathbb{R}, dx) : f = 0 \text{ on } \text{int}(J)\}.$$

In particular, the orthogonality relation $H_{\text{in}} \ominus L \perp H_{\text{out}} \ominus L$ holds and the result follows. \square

It is illuminating to discuss the role of loops of finite length with the example of the Brownian motion indexed by loops in mind. A loop with finite length admits a Lipschitz continuous parametrization. Hence, its range has Hausdorff dimension 1, unless it is constant. In any case, its range is negligible. The fact that its topological index is square-integrable is not at all obvious. It is granted by a generalization of the isoperimetric inequality discovered by T. Banchoff and W. Pohl. We denote the length of l by $\ell(l)$.

Theorem 0.4.5 (Banchoff-Pohl). — *Let $l : [0, 1] \rightarrow \mathbb{R}^2$ be a Lipschitz continuous loop. Then*

$$4\pi \int_{\mathbb{R}^2} n_l(x)^2 dx \leq \ell(l)^2.$$

The original reference for this theorem is the article [6] of T. Banchoff and W. Pohl. They prove the inequality for a loop of class C^2 . An elementary proof of the inequality for rectifiable curves can be found in a paper by A. Vogt [37].

Of course, there are many loops with infinite length whose topological index is square-integrable, for instance fractal Jordan curves or simply loops of infinite length whose range is contained in a line. It would probably be difficult to characterise the set of loops which wind gently in a way which significantly differs from its definition. Nevertheless, on the scale of roughness given by the p -variation, as defined by L.C. Young [42], the space of rectifiable loops is the largest which contains only loops which wind gently. Recall that the p -variation of a loop $l : [0, 1] \rightarrow \mathbb{R}^2$ is defined as the supremum over all subdivisions $\{t_0 \leq \dots \leq t_r\}$ of $[0, 1]$ of the quantity $(\sum_i \|l(t_{i+1}) - l(t_i)\|^p)^{\frac{1}{p}}$. A loop has finite length if and only if it has finite 1-variation. A loop with finite p -variation for $p < 2$ has negligible range.

Proposition 0.4.6. — *There exists a loop $l : [0, 1] \rightarrow \mathbb{R}^2$ such that l has finite p -variation for all $p > 1$ and $\int_{\mathbb{R}^2} n_l(x)^2 dx = +\infty$.*




FIGURE 5. This loop has finite p -variation for all $p > 1$ but its topological index is not square-integrable.

Proof. — For each $n \geq 1$, let l_n be the loop based at the origin which goes once along the circle of radius $\frac{1}{n}$ through the origin, tangent to the horizontal axis and contained in the upper half-plane. Assume that l_n is parametrised at constant speed by an interval of length 2^{-n} . Let l be the uniform limit of the finite concatenations $l_1 \cdots l_n$ as n tends to infinity. This limit exists because the radii of the circles l_n tend to 0 as n tends to infinity. For all $p > 1$, the p -variation of l raised to the power p is equal, up to some constant, to $\sum_{n \geq 1} n^{-p}$, hence it is finite. On the other hand, the squared L^2 norm of l is $\pi \sum_{n \geq 1} n^2 \left(\frac{1}{n^2} - \frac{1}{(n+1)^2} \right) = +\infty$. \square

0.5. Markovian holonomy fields and gauge fields

The original motivation for the study of processes indexed by loops is issued from theoretical physics, indeed from quantum field theory and more precisely from quantum gauge theories. Let us explain this on the example of electrodynamics. The classical theory of electrodynamics, as established in the second half of the nineteenth century, is summarised by the Lorentz law and Maxwell's equations. Maxwell's equations relate the electric and magnetic fields to the density of electric charge and the density of electric current in space.

In order to derive these equations from a principle of least action, which is usually the first step in the procedure of quantization of a physical theory, it is convenient to express the electric and magnetic fields in terms of a scalar potential and a vector potential. These potentials are not uniquely defined by the fields and this indeterminacy is called the gauge symmetry of the theory. It turns out that the geometric nature of the pair formed by the scalar and vector potentials is that of a connection on a principal bundle with structure group $U(1)$ over the space-time. At the level of rigour of this discussion, we do not make a serious mistake by identifying this object with a differential 1-form on the space-time. This 1-form is usually denoted by A and called the gauge field. The exterior differential of the gauge field is a mixture of the electric and magnetic fields called the electromagnetic field and denoted by F . The relation $F = dA$ implies the equality $dF = 0$, which is equivalent to the two homogeneous Maxwell equations. The two inhomogeneous equations can be put under the form $*d * F = J$, where $*$ is the Hodge operator on space-time associated with the Minkowski metric, and J is a differential 1-form built from the densities of charge and current.

The main fact for us in this discussion is that the gauge field, the object in terms of which the classical electrodynamics is best described, is a differential 1-form on space-time. The most natural way to evaluate a 1-form is to integrate it along paths. In the case of electrodynamics, the gauge symmetry of the theory implies that any two gauge fields which differ by a total differential describe the same physics. This

indicates that the natural gauge-invariant functionals of the gauge field are in fact its integrals along loops.

If instead of doing quantum mechanics we prefer to do statistical mechanics, then we turn the Minkowski space-time into a Euclidean space-time and put on the space of gauge fields the Gibbs measure corresponding to the action which gives back Maxwell's equations through the least action principle. In an empty space-time, this action is called the Yang-Mills action and it is essentially the squared L^2 norm of the electromagnetic field. Just as in the deterministic case, an appropriate way of studying a random gauge field is again to consider the random variables determined by its integrals along loops, sometimes called Wilson loops, and this constitutes indeed a stochastic process indexed by loops. The Markov property of such a process, in this physical context, reflects the following property of locality of the Yang-Mills action: if the space-time is partitioned into several regions, then each region contributes to the action by a quantity which can be computed from the values of the field inside this region only.

In general, the random object that we are studying is thus an analogue of the electromagnetic field, or rather of the gauge field formed by the scalar potential of the electric field and the vector potential of the magnetic field. Let us suggest an idea of the physical meaning of the gauge field A . This field interacts with particles which carry an analogue of electric charge. In fact, the 1-form A takes its values in the Lie algebra of a Lie group G and the charge of a particle is, mathematically, a linear action of G on some vector space in which the wave function of the particle takes its values. For example, in the case of electrodynamics, the group is $U(1)$, the wave functions take their values in \mathbb{C} and for a particle of charge ne , where $-e$ is the charge of the electron, the group $U(1)$ acts on \mathbb{C} by the representation $e^{i\theta} \cdot z = e^{in\theta} z$. The exponential of the integral of the gauge field along a certain loop, as an element of $U(1)$, describes the modification of the phase of the wave function of a particle which travels along l . More generally, let $l : [0, 1] \rightarrow M$ be a loop. Assume that G is a Lie group and that A is a differential 1-form on M with values in the Lie algebra of G . Then, under fairly weak regularity assumptions, the differential equation

$$(14) \quad h_0 = 1, \quad \dot{h}_t h_t^{-1} = -A(\dot{l}_t), \quad t \in [0, 1]$$

has a unique solution $h : [0, 1] \rightarrow G$. The element h_1 determined by (14) is called the holonomy of A along l and this is why we call our processes holonomy fields. The action of this element h_1 of G on the left on the vector space in which the wave function of a particle takes its values determines how the state of a particle is modified when it travels along l . The reversed order in the right-hand side of (5) is due to the fact that the state of a particle travelling along the concatenation of l_1 and l_2 is

modified first by the transformation due to the displacement along l_1 and then by the transformation due to the displacement along l_2 .

0.6. Finite groups, gauge fields and ramified coverings

When the group G is finite, the heuristic idea that a Markovian holonomy field is a reflection of a probability measure on a space of connections or differential 1-forms with values in the Lie algebra of G becomes awkward. Indeed, this Lie algebra is the null vector space. Topologically, a principal bundle with finite structure group is a covering and it carries a unique connection, which is flat. If M is simply connected, then the holonomy along any loop is the unit element of G . On the other hand, there exist Markovian holonomy fields with values in finite groups, which are non-trivial processes indexed by loops, even on the sphere S^2 .

When G is finite, the correct geometric picture is the following: a Markovian holonomy field with values in G is the monodromy of a random ramified G -bundle. By a ramified G -bundle, we mean a ramified covering whose regular fibres are endowed with a free transitive action of G , or equivalently, a principal G -bundle over the complement of a finite set in the surface. There is still a unique connection on a ramified principal bundle and this connection is flat at each point which is not a ramification point, but each ramification point acts like a macroscopic amount of curvature concentrated at a single point. For instance, if $n \geq 2$ is an integer, then the mapping $z \mapsto z^n$ from \mathbb{C} to itself is naturally a ramified $\mathbb{Z}/n\mathbb{Z}$ -bundle. The group $\mathbb{Z}/n\mathbb{Z} = \{e^{2ik\pi/n} : k \in \{0, \dots, n-1\}\}$ acts by multiplication on \mathbb{C}^* , freely and transitively on the fibres of the covering map $z \mapsto z^n$. Any loop which goes once positively around a disk which contains 0 has monodromy $e^{\frac{2i\pi}{n}}$, no matter how small this disk. This is consistent with the picture of a concentration of curvature at 0.

The idea that a Markovian holonomy field with values in a connected Lie group is a probability measure on a space of connections looked at through its holonomy is only a guide for the intuition and has in general no firm rigorous ground. On the contrary, the fact that a large class of Markovian holonomy fields with values in a finite group G can be realised as the monodromies of random ramified G -bundles is a theorem that we will prove. In the intuitive picture of the Yang-Mills measure, the curvature of the random connection is supposed to have the distribution of a white noise. The correct distribution of the ramified G -bundles which correspond with a given Markovian holonomy field can be roughly described as follows: first choose the ramification locus by throwing a Poisson point process on the surface with intensity the measure of area, then give a weight to every ramified G -bundle with this ramification locus which depends on its monodromy at each ramification point and the Lévy process associated to the Markovian holonomy field.

0.7. Markovian holonomy fields and the Yang-Mills measure

The Yang-Mills measure is to Markovian holonomy fields what the Brownian motion is to Lévy processes. It is indeed a Markovian holonomy field, whose associated Lévy process is the Brownian motion on G . It is also the only Markovian holonomy field to have been constructed before the present work. The Yang-Mills measure has been the object of mathematical work since around 1990. It was first constructed on an arbitrary compact surface by Ambar Sengupta [35]. The author gave later a different construction of essentially the same measure in [24]. One of the by-products of the present work is to provide another construction of the Yang-Mills measure, which really is very close to that given in [24]. Yet, we would like to emphasise several aspects in which it differs from the previous ones.

The first difference is a slight shift of point of view which we have already illustrated. Instead of considering a specific surface, choosing boundary conditions and constructing a process indexed by some class of loops on this surface, we now consider as one single object a whole collection of processes indexed by loops on all possible surfaces. The advantage of this view is that it leads quite naturally to an axiomatic characterization of Markovian holonomy fields. The one which we propose is inspired by the classical definition of a Markov process and by the axiomatic definition of a topological quantum field theory (see the next section).

Another difference lies in the class of loops that we consider on a surface. Indeed, on a surface M endowed with boundary conditions, a measure of the area, and a base point m , the Yang-Mills measure produces a multiplicative G -valued stochastic process $(H_l)_{l \in \mathcal{L}_m(M)}$ for some class of loops $\mathcal{L}_m(M)$. In the two previous constructions of the Yang-Mills measure mentioned above, the loops to which one was able to attach a random variable were finite concatenations of very special curves, for example segments of submanifolds in the author's construction. On the other hand, it was proved in [24] that the mapping $l \mapsto H_l$ is continuous in L^1 norm on the class of loops considered there and endowed with the topology of uniform convergence and convergence of the length. This suggested that it should be possible to associate a random variable at least to each loop of finite length. This is indeed what we achieve in the present construction, thus increasing the coherence between the regularity property of the stochastic process and the set of loops on which it is defined. Let us point out that it is not necessary to be able to measure lengths in order to decide if a loop has finite length. Indeed, a diffeomorphism does not alter the fact that a curve has finite or infinite length. Thus, the definition of $\mathcal{L}_m(M)$ as the set of loops on M with finite length based at m does not require the choice of a Riemannian metric on M , a smooth structure is more than enough. It is natural to wonder if one could define the Yang-Mills measure or any other Markovian holonomy field for a

significantly larger class of loops. As the discussion at the end of Section 0.4 suggests, we believe that this is not possible with only the techniques used in this work, and probably very difficult anyway.

The third important difference between this work and the existing constructions of the Yang-Mills measure concerns the group in which the stochastic process takes its values, which we have already denoted by G . Since the Yang-Mills measure is associated with the Brownian motion on G , it was natural to assume that it was a connected group. In the present work, the Brownian motion is replaced by a Lévy process, thus a process with jumps. This opens the possibility of considering non-connected compact Lie groups and in particular finite groups. In the case of finite groups, the Lévy processes have a very simple structure and we are able to give a completely geometrical picture of Markovian holonomy fields, as we have already briefly explained in the previous section.

0.8. Topological quantum field theories

So far, our presentation of Markovian holonomy fields has been based mainly on the analogy with usual Markov processes indexed by intervals of time. Before closing this introduction, we would like to indicate another analogy, between the notions of Markovian holonomy field and topological quantum field theory (TQFT). We will not pursue this analogy in this work, but still we would like to acknowledge it as an important source of inspiration.

Let us start by recalling briefly what a TQFT is. Let $n \geq 0$ be an integer. There is a category, denoted by Cob_n , whose objects are the compact oriented smooth manifolds without boundary of dimension n and whose morphisms are given by the cobordisms: if N_1 and N_2 are two oriented n -manifolds, then $\text{Hom}(N_1, N_2)$ is the set of oriented $(n + 1)$ -manifolds whose boundary is endowed with an orientation-preserving diffeomorphism with the disjoint union of N_1^* and N_2 , where N_1^* denotes the manifold obtained from N_1 by reversing the orientation. On the other hand, there exists a perhaps much more familiar category, denoted by Vect , whose objects are complex linear spaces and whose morphisms are linear mappings.

A $(n + 1)$ -dimensional TQFT is a covariant functor from Cob_n into Vect . This functor is usually denoted by Z , in reference to the fact that this definition has initially been given as a general framework for the partition functions of certain statistical or quantum field theories. Thus, by definition, Z associates to each n -manifold N a linear space $Z(N)$ and to each $(n + 1)$ -manifold M endowed with a diffeomorphism between its boundary and the disjoint union $N_1^* \sqcup N_2$ of two n -manifolds, a linear map $Z(M) : Z(N_1) \rightarrow Z(N_2)$.

This functor is required to satisfy a certain number of natural properties.




FIGURE 6. Two examples of cobordisms, in dimensions $n + 1 = 1$ et $n + 1 = 2$.

1) To be multiplicative: for all n -manifolds N_1 and N_2 , we insist that

$$Z(N_1 \sqcup N_2) = Z(N_1) \otimes Z(N_2)$$

(the tensor product of linear spaces) and, for any two $(n + 1)$ -manifolds M_1 and M_2 , that

$$Z(M_1 \sqcup M_2) = Z(M_1) \otimes Z(M_2)$$

(the tensor product of linear maps).

2) To be involutive: for all n -manifold N , we insist that

$$Z(N^*) = Z(N)^*$$

(the dual linear space) and, for all $(n + 1)$ -manifold, that

$$Z(M^*) = Z(M)^*$$

(the adjoint linear map).

3) To behave well with respect to the composition of cobordisms: if M_1 realises a cobordism between N_1 and N_2 , and M_2 a cobordism between N_2 and N_3 , and if M is formed by gluing M_1 and M_2 along N_2 , then we insist that

$$Z(M) = Z(M_2) \circ Z(M_1).$$

This last axiom can be reformulated in a more general way by slightly shifting our point of view. If we identify the space of linear maps from $Z(N_1)$ into $Z(N_2)$ with $Z(N_1)^* \otimes Z(N_2) = Z(N_1^* \sqcup N_2)$, we can say that the linear map $Z(M)$ is an element of the linear space $Z(\partial M)$. In these terms, the axiom becomes the following.

3') If the boundary of M can be written as $N_1 \sqcup N_1^* \sqcup N_2$ and if M' is formed by gluing the two copies of N_1 , then we insist that

$$Z(M') = \kappa(Z(M)),$$

where $\kappa : Z(N_1)^* \otimes Z(N_1) \otimes Z(N_2) \rightarrow Z(N_2)$ is the natural contraction of the first two factors.




FIGURE 7. Axiom 3: composition of two cobordisms in dimension 2.




FIGURE 8. Axiom 3': gluing two components of the boundary of a surface.

4) Finally, we insist that Z behaves well with respect to orientation-preserving diffeomorphisms of manifolds of dimension n and $n + 1$, that is to say that to diffeomorphic n -manifolds it associates isomorphic linear spaces and to diffeomorphic cobordisms it associates linear maps conjugated by the linear isomorphisms associated to the diffeomorphisms induced on the boundaries.

If \emptyset designates the empty n -manifold, then the multiplicativity condition implies for example that $Z(\emptyset) = Z(\emptyset) \otimes Z(\emptyset)$. Thus, $Z(\emptyset)$ is either the null linear space or \mathbb{C} . In order to avoid trivial situations, we will assume that $Z(\emptyset) = \mathbb{C}$. Thus, for all $(n + 1)$ -manifold M without boundary, $Z(M)$ is a linear map from \mathbb{C} into itself, that is, a complex number.

If N is an n -manifold, then $Z(N \times [0, 1])$ is an endomorphism of $Z(N)$. Since the cylinder $N \times [0, 1]$ is diffeomorphic to two copies of itself glued together along two copies of N , one has the identity

$$Z(N \times [0, 1]) = Z(N \times [0, 2]) = Z(N \times [0, 1]) \circ Z(N \times [0, 1]).$$

Thus, $Z(N \times [0, 1])$ is a projection in the linear space $Z(N)$ and, since a surface is not modified up to homeomorphism when one glues a cylinder along one of its boundary components, one loses no generality in assuming that it is the identity. Then, gluing the two ends of this cylinder and using the axiom 3', one finds that $Z(N \times S^1)$, where S^1

is the circle, is the natural contraction of the identity of $Z(N)$, that is, the trace of this linear map, which is nothing but the dimension of $Z(N)$.

In the context of Markovian holonomy fields, we should modify slightly the definition of a TQFT in order to incorporate in it the notion of volume, or area when $n + 1 = 2$. We just need to replace $(n + 1)$ -manifolds by manifolds endowed with a notion of $(n + 1)$ -volume, of which the only invariant under diffeomorphisms is, according to a theorem of Moser, the total volume. In doing this, we modify only the morphisms in the category Cob_n . The axioms 1 to 4 are unchanged. The object thus defined is called an *area-dependent topological quantum field theory* (ad-TQFT). For all $(n + 1)$ -manifold M , we will denote by $Z_t(M)$ the element of $Z(\partial M)$ associated to M when it is endowed with a measure of $(n + 1)$ -volume of total area t .

Let us investigate a 1-dimensional ad-TQFT, which we denote by Z . There is only one connected 0-manifold, the point, which we denote by pt . It can have two orientations, denoted by $+$ and $-$. There is only one linear space involved, namely $Z(\text{pt}_+)$, which we denote by V . Then, there are only two oriented connected 1-manifolds, the interval and the circle. It is convenient to represent intervals as intervals of \mathbb{R} endowed with their natural length. Thus, for all $s < t$ reals, $Z([s, t])$ is an endomorphism of V , which depends only on the length of $[s, t]$ and which we denote by P_{t-s} . Moreover, the axiom 3 implies that the relation $Z([s, u]) = Z([t, u]) \circ Z([s, t])$ holds for all $s < t < u$, so that $(P_t)_{t>0}$ is a semigroup of endomorphisms of V . Finally, according to the general argument which we have presented above, we have for all $t > 0$ the relation $Z_t(S^1) = \text{Tr}(P_t)$.

Finally, a 1-dimensional ad-TQFT is the same thing as a semigroup of endomorphisms of a linear space. In contrast with the purely topological case, nothing here prevents this linear space from being infinite-dimensional. For instance, in a Hilbertian framework, it suffices that the semigroup should be of trace class.

A sufficiently regular classical Markov process, indexed by intervals of time, gives rise, through its transition kernels, to a 1-dimensional ad-TQFT. Consider for instance the Brownian motion X on a compact Lie group G . For all $t > 0$, let Q_t denote the density of the distribution of X_t with respect to the Haar measure on G . Define $Z(\text{pt})$ as the space $L^2(G)^G$ of square-integrable functions on G which moreover are invariant by conjugation. Then, for all $s < t$, we can define $Z([s, t])$ as the mapping $(x, y) \mapsto Q_{t-s}(x^{-1}y)$, seen either as the integral kernel of a linear map from $L^2(G)^G$ into itself, or equivalently as an element of $L^2(G^2)^{G^2} \simeq L^2(G)^G \otimes L^2(G)^G$.

Our definition of the 2-dimensional Markovian holonomy fields contains the definition of an ad-TQFT, given by the partition functions of the field. In the 2-dimensional setting, there is also a unique connected 1-manifold, which is the circle S^1 , so that there is also a unique linear space involved. We set $Z(S^1) = L^2(G)^G$. Note that this

space is finite-dimensional when G is finite. We will prove (see Proposition 4.1.4), that the partition functions of a Markovian holonomy field are indeed invariant by conjugation and square-integrable with respect to all their arguments. They are moreover symmetric. Each function $Z_{p,g,t}^\pm$ can thus be seen as a symmetric tensor of $(L^2(G)^G)^{\otimes p}$. Let us summarise this whole section in the following result.

Proposition 0.8.1. — *Let G be a compact Lie group. Let $(Z_{p,g,t}^+)_{p,g \geq 0, t > 0}$ be the partition functions of a Markovian holonomy field, associated to oriented surfaces. Set $Z(S^1) = L^2(G)^G$, the space of square-integrable functions on G invariant by conjugation. For all surface M with genus g , total area t and such that ∂M has p connected components, set $Z(M) = Z_{p,g,t}^+$. Then the functor Z is a 2-dimensional ad-TQFT.*

0.9. Structure

The present work consists of five chapters. In Chapters 1 and 2, we develop the tools and prove most of the technical results that we use in our study of Markovian holonomy fields. The first chapter covers the topology of surfaces and their surgery, the topological space of paths on a surface, the fundamental notion of graph on a surface, which we treat both topologically and combinatorially, and finishes with a discussion of Riemannian metrics. The second chapter introduces the space of multiplicative functions of paths with its measurable structures and its uniform measure. The last section is devoted to a study of the free group of loops in a graph in relation with this uniform measure.

In Chapter 3, we define Markovian holonomy fields and their discrete analogues. We prove the first central result of this work (Theorem 3.2.9) which encapsulates in an abstract way and extends the procedure which allowed us, in our previous construction of the Yang-Mills measure, to take the continuous limit of a discrete gauge theory. In our present language, we prove that every regular discrete Markovian holonomy field can be extended in a unique way to a regular Markovian holonomy field.

In Chapter 4, we prove that a regular Markovian holonomy field with values in a compact Lie group G determines a classical Lévy process in G , which in turn determines completely the partition functions of the holonomy field (Propositions 4.2.1 and 4.1.10). We then prove that each Lévy process of a wide class can be obtained in this way (Theorem 4.3.1). Whether or not two distinct Markovian holonomy fields can have the same associated Lévy process is a natural question which we do not settle here.

In Chapter 5, we prove that when the group G is finite, the Markovian holonomy field constructed in the previous chapter is the monodromy process of a random

ramified covering (Theorem 5.4.2). In fact, most of the chapter is devoted to the construction of this random ramified covering.

The choice that we have made of concentrating to the extent possible the technical results in the first two chapters has the obvious drawback that the results exposed there often lack their real motivation, and that a linear reading of these two chapters may not be very rewarding. We hope that this is compensated by the fact that the study of Markovian holonomy fields themselves is much more straightforward than it would be if one had to constantly interrupt the exposition to prove technical results. In order to allow as much as possible the reader to jump from a section to another, we have included an index of notation which should be helpful in locating the first occurrence of a notation or a symbol.

0.10. References

The original paper of C.N. Yang and R. L. Mills where they introduced non-Abelian gauge theories appeared in 1954 [41]. An explicit description of the Yang-Mills measure on a lattice was given by A. Migdal in 1975 [31]. The importance of the 2-dimensional Yang-Mills theory in relation with the geometry of the moduli space of flat connections was emphasised by E. Witten in two papers [39] and [40] published in the early 1990's.

An introduction to some aspects of the physics of gauge theories accessible to a mathematician can be found in the book by J. Baez and J. Muniain [5]. The book by D. Bleecker [7] is a mathematical exposition of the theory of connections on a principal bundle with a strong bias towards gauge theories. The mathematical reference on connections is the book by S. Kobayashi and K. Nomizu [22]. Topological quantum field theories were introduced by M. Atiyah in [4].

The mathematical study of the Yang-Mills measure was developed around 1985 by S. Albeverio, R. Høegh-Krohn and H. Holden [1], [2], L. Gross [18], [19], C. King and A. Sengupta [20], B. Driver [10], [11]. The idea that one could construct a variant of the Yang-Mills measure associated to a Lévy process which is not necessarily the Brownian motion was already present, twenty years ago, in the work of Albeverio et al. [2]. A. Sengupta gave the first construction of the measure on the sphere S^2 in 1992 [34] and on an arbitrary compact surface in 1996 [35]. The author gave another construction on an arbitrary surface in 2003 [24], later refined to take non-trivial principal bundles into account [25]. In a different direction, around 1990, D. Fine investigated the Yang-Mills measure in [13], [14] with a focus on the geometrical structure of the space of connections on a principal bundle.

The fact that gauge theories with finite structure groups are related to random ramified coverings was already observed by A. D'Adda and P. Provero [8]. Ramified coverings are also important in relation with the large N limit of the 2-dimensional $U(N)$ Yang-Mills theory, for a reason which is not directly related to the content of the present work. This relation was revealed by D. Gross, in collaboration with W. Taylor [17] and A. Matytsin [16]. It has recently been investigated by the author [26].

0.11. Acknowledgements

This work has greatly benefitted from many discussions with Ambar Sengupta, as well as from his continual interest and encouragement. Comments which he has made on an earlier version of the manuscript have also helped me to improve it significantly.

Guy David has kindly provided me with appropriate references and technical help for the proofs of Section 1.3.3.

CHAPTER 1

SURFACES AND GRAPHS

In this chapter, we introduce the tools of topology and geometry of surfaces that we use in the rest of this work. We set up the notation, collect the necessary classical results and prove less classical ones. After a short review of compact surfaces, we describe their surgery and study in some detail the paths and graphs drawn on them. In particular, we describe carefully the boundary of a face of a graph. Then we define the group of reduced loops based at a point in a graph and recall why it is free. In the next chapter, we will prove the existence of sets of generators of this group with specific properties. Finally, we discuss Riemannian metrics on surfaces in relation with our problem.

1.1. Surfaces

1.1.1. Classification of surfaces. — Let us start by recalling the definition of a surface.

Definition 1.1.1. — A topological compact surface is a Hausdorff compact topological space in which every point admits a neighbourhood homeomorphic to \mathbb{R}^2 or homeomorphic to $\mathbb{R}_+ \times \mathbb{R}$.

A smooth compact surface, or simply a compact surface, is a topological compact surface equipped with a structure of smooth 2-dimensional manifold with boundary.

The distinction between topological and smooth surfaces is not essential, as the following well-known result shows (see [33]).

Theorem 1.1.2. — Any topological compact surface is homeomorphic to a smooth compact surface. Moreover, two smooth compact surfaces are diffeomorphic if and only if they are homeomorphic.

The classification theorem for compact surfaces is thus the same for smooth and topological surfaces. We will now describe this theorem, using a slightly unorthodox convention about the genus of a surface.

Recall that the connected sum of two surfaces is the surface obtained by removing a small disk from both surfaces and gluing the two resulting surfaces along the boundaries of these disks. Of course, the connected sum of two surfaces is defined up to homeomorphism only.

Let us describe two infinite series of surfaces. The first series is built from the torus, which is the Cartesian product of two circles. For each even integer $g \geq 0$ and each integer $p \geq 0$, let $\Sigma_{p,g}^+$ be the surface obtained by removing p pairwise disjoint open disks from the connected sum of $\frac{1}{2}g$ tori. For $g = 0$, the surface $\Sigma_{p,0}^+$ is a sphere with p holes. The second series is built from the projective plane, which is the quotient of the unit sphere of \mathbb{R}^3 by the group of isometries $\{\text{id}, -\text{id}\}$. For each integer $g \geq 1$ and each $p \geq 0$, let $\Sigma_{p,g}^-$ be the surface obtained by removing p pairwise disjoint open disks from the connected sum of g projective planes.

Recall that a smooth compact surface is *orientable* if it carries a non-vanishing differential 2-form. We say that a topological compact surface is orientable if a smooth compact surface to which it is homeomorphic is orientable. The following theorem is proved for example in [30].

Theorem 1.1.3. — *Any connected orientable topological compact surface is homeomorphic to one and exactly one of the surfaces $\{\Sigma_{p,g}^+ : p, g \geq 0, g \text{ even}\}$. Any connected non-orientable compact surface is homeomorphic to one and exactly one of the surfaces $\{\Sigma_{p,g}^- : p \geq 0, g \geq 1\}$. Any oriented smooth compact surface admits an orientation-reversing diffeomorphism.*

We call the integer g which appears in this classification the *genus* of a surface. For orientable surfaces, it is twice the number which is usually called the genus. The advantage of our convention is illustrated by Proposition 1.1.4. We denote the genus of a surface M by $\mathbf{g}(M)$ and the number of connected components of its boundary by $\mathbf{p}(M)$

With this notation, we have $\mathbf{g}(\Sigma_{p,g}^\pm) = g$ and $\mathbf{p}(\Sigma_{p,g}^\pm) = p$. Let us define a binary operation \wedge on $\{+, -\}$ by setting

$$+ \wedge + = + \quad \text{and} \quad + \wedge - = - \wedge + = - \wedge - = -.$$

If M_1 and M_2 are two compact topological surfaces, we denote by $M_1 \# M_2$ the connected sum of M_1 and M_2 .

Proposition 1.1.4. — *Let $\Sigma_{p,g}^\varepsilon$ and $\Sigma_{p',g'}^{\varepsilon'}$ be two surfaces of the list described above. Then*

$$(15) \quad \Sigma_{p,g}^\varepsilon \# \Sigma_{p',g'}^{\varepsilon'} = \Sigma_{p+p',g+g'}^{\varepsilon \wedge \varepsilon'}.$$

Proof. — It is clear that $\Sigma_{p,g}^\varepsilon \# \Sigma_{p',g'}^{\varepsilon'}$ has $p + p'$ boundary components. Moreover, the connected sum of two manifolds is orientable if and only if both manifolds are. The only non-trivial point is that the value of the reduced genus is correct when exactly one of the two surfaces $\Sigma_{p,g}^\varepsilon$ and $\Sigma_{p',g'}^{\varepsilon'}$ is orientable. In this case, since the operation of connected sum is commutative and associative, this boils down to the fact that the connected sum of a projective plane and a torus is homeomorphic to the connected sum of three projective planes. This is a classical result, proved as Lemma 7.1 in [30]. \square

It is useful to keep in mind that if M is a non-orientable compact topological surface, then the connected sum of M with the torus $\Sigma_{0,2}^+$ is homeomorphic to the connected sum of M with the Klein bottle $\Sigma_{0,2}^-$.

The fundamental groups of surfaces are most easily described by generators and relations. We denote by

$$\langle x_1, \dots, x_n \mid r_1, \dots, r_m \rangle$$

the group generated by x_1, \dots, x_n subject to the relations r_1, \dots, r_m . If a and b are two elements of a group, we denote their commutator by $[a, b] = aba^{-1}b^{-1}$.

Theorem 1.1.5. — *The fundamental groups of compact surfaces are, up to isomorphism, the following.*

1) $\pi_1(\Sigma_{0,0}^+) = \{1\}$ and for all $g \geq 1$,

$$\pi_1(\Sigma_{0,2g}^+) = \langle a_1, b_1, \dots, a_g, b_g \mid [a_1, b_1] \cdots [a_g, b_g] = 1 \rangle.$$

2) For all $g \geq 0$ and all $p \geq 1$, $\pi_1(\Sigma_{p,2g}^+)$ is free of rank $2g + p - 1$.

3) For all $g \geq 1$, $\pi_1(\Sigma_{0,g}^-) = \langle a_1, \dots, a_g \mid a_1^2 \cdots a_g^2 = 1 \rangle$.

4) For all $g \geq 0$ and all $p \geq 1$, $\pi_1(\Sigma_{p,g}^-)$ is free of rank $g + p - 1$.

It follows from this theorem that a compact surface is not characterised up to homeomorphism by its fundamental group. For example, a sphere with three holes and a torus with one hole both have a fundamental group which is free of rank 3. However, *closed* surfaces, that is, surfaces without boundary, are indeed characterised by their fundamental group.

If M is a closed compact surface, then the genus of M is the minimal number of generators in a presentation of the fundamental group of M . On the other hand, if M has a non-empty boundary, then its fundamental group is free of rank $g(M) + p(M) - 1$.

The boundary of a compact surface is a compact 1-dimensional manifold without boundary, hence a finite union of circles. If a surface is oriented, then every connected component of ∂M carries an induced orientation, such that the surface stays on the left of a person walking along the boundary in the positive direction.

Definition 1.1.6. — *Let M be a compact surface. We denote by $\mathcal{B}(M)$ the set of connected components of the boundary of M , each taken twice, once with each orientation. If M is oriented, we denote by $\mathcal{B}^+(M)$ the subset of $\mathcal{B}(M)$ formed by the oriented connected components of ∂M which bound M positively.*

Any diffeomorphism of a compact surface induces a diffeomorphism of its boundary. We need to know which diffeomorphisms of the boundary can be obtained in this way. For this, observe that if the boundary of an oriented surface M has p connected components, then, among the 2^p distinct orientations of ∂M , the two distinct orientations of M determine two preferred orientations. We say that a diffeomorphism of ∂M is orientation-preserving if it preserves these orientations and orientation-reversing if it exchanges them. Of course, if $p \geq 2$, then there exist diffeomorphisms of ∂M which are neither orientation-preserving nor orientation-reversing.

Theorem 1.1.7. — *Let M be a smooth compact surface. If M is non-orientable, then any diffeomorphism of ∂M can be extended to a diffeomorphism of M . If M is orientable, then any orientation-preserving (resp. orientation-reversing) diffeomorphism of ∂M can be extended to an orientation-preserving (resp. orientation-reversing) diffeomorphism of M .*

In order to prove this theorem, it suffices to prove the following two facts: that on a disk with two holes (that is, a three-holed sphere) there exists a diffeomorphism which fixes the outer boundary and exchanges the two holes ; and that on a Möbius band with one hole (that is, a two-holed projective plane) there exists a diffeomorphism which fixes the outer boundary and reverses the orientation of the boundary of the hole. In both cases, these diffeomorphisms can be constructed by hand.

1.1.2. Surgery of surfaces. — Surfaces undergo natural operations of surgery such as cutting along a curve or gluing one or two boundary components. When one performs gluings and wants to keep track of where they have occurred, one ends up with surfaces which carry marks. On the other hand, when one cuts a surface along one or several curves, a convenient way of keeping track of what has been done is to maintain an involution of the set of boundary components of the current surface.

Definition 1.1.8. — *A marked surface is a pair (M, \mathcal{C}) , where M is a smooth compact surface and \mathcal{C} is a finite collection of oriented smooth 1-dimensional submanifolds of the interior of M , which are pairwise either disjoint or equal up to orientation,*

and such that an oriented curve belongs to \mathcal{C} if and only if the same curve with the opposite orientation belongs to \mathcal{C} . The elements of \mathcal{C} are called marks.




FIGURE 1. An orientable surface of genus 4 with two holes, marked by three curves.

The marked surface (M, \mathcal{C}) is said to be oriented if every connected component of M is orientable and oriented.

Let us emphasise that on a marked surface, even an oriented one, the marks do not carry a preferred orientation. Recall the notation $\mathcal{B}(M)$ from Definition 1.1.6. The group $\mathbb{Z}/2\mathbb{Z}$ acts on $\mathcal{C} \cup \mathcal{B}(M)$ by reversing the orientation, and we denote this action by $b \mapsto b^{-1}$.

Definition 1.1.9. — A tubular pattern is a triple (M, \mathcal{C}, τ) where (M, \mathcal{C}) is a marked surface and τ is an involution of $\mathcal{B}(M)$ which commutes with the orientation reversal, that is, which satisfies $\tau(b^{-1}) = \tau(b)^{-1}$. If $\mathcal{C} = \emptyset$, the tubular pattern is said to be split.

Let (M, \mathcal{C}, τ) be a tubular pattern. Choose $b \in \mathcal{B}(M)$. If $\tau(b)$ is a component of ∂M distinct from b , then b is meant to be identified with $\tau(b)$ by an orientation-preserving diffeomorphism. If $\tau(b) = b$, then b is not meant to be glued or altered in any way. Finally, if $\tau(b) = b^{-1}$, then b is meant to be glued on itself according to an orientation-preserving involution. Of course, this way of encoding the possible gluing operations is purely conventional.

We define now the basic operation of surgery, which is the operation of gluing.

Definition 1.1.10. — Let (M, \mathcal{C}, τ) and $(M', \mathcal{C}', \tau')$ be two tubular patterns. A smooth mapping $f : M' \rightarrow M$ is called an elementary gluing if one of the following sets of conditions is satisfied.

1) The mapping f is the quotient map which identifies b' with $\tau'(b')$ by an orientation-preserving diffeomorphism for some $b' \in \mathcal{B}(M')$ such that

$$\{b', b'^{-1}\} \neq \{\tau'(b'), \tau'(b')^{-1}\}.$$

Moreover, $f(\mathcal{C}' \cup \{b', b'^{-1}\}) = \mathcal{C}$ and, on the set $\mathcal{B}(M') \setminus \{b', b'^{-1}, \tau'(b'), \tau'(b')^{-1}\}$, $\tau \circ f = f \circ \tau'$. Such a gluing is called binary, and the pair of curves $\{f(b'), f(b'^{-1})\}$ is called its joint.

2) The mapping f is the quotient map which identifies the points of b' by pairs according to an orientation-preserving fixed-point free smooth involution for some $b' \in \mathcal{B}(M')$ such that

$$\tau(b') = b'^{-1}.$$

Moreover, $f(\mathcal{C}' \cup \{b', b'^{-1}\}) = \mathcal{C}$ and, on $\mathcal{B}(M') - \{b', b'^{-1}\}$, $\tau \circ f = f \circ \tau'$. In this case, the gluing is called unary and the pair of curves $\{f(b'), f(b'^{-1})\}$ is called its joint.

A gluing is a map which can be written as the composition of several elementary gluings. A gluing is complete if the involution of the set of boundary components of the target surface is the identity.




FIGURE 2. Binary and unary gluings.

Up to homeomorphism of the underlying surfaces and disregarding the markings, performing a unary gluing along a boundary component is equivalent to gluing a Möbius band along this boundary component. The result of this operation is never orientable.

The other basic surgery operation is that of splitting. It is really the same thing as a gluing, looked at in the other direction.

Proposition 1.1.11. — *Let (M, \mathcal{C}, τ) be a tubular pattern with $\mathcal{C} \neq \emptyset$. Choose $\{l, l^{-1}\} \subset \mathcal{C}$. Then there exists a tubular pattern $(\text{Spl}_l(M), \text{Spl}_l(\mathcal{C}), \text{Spl}_l(\tau))$, and an elementary gluing $f : \text{Spl}_l(M) \rightarrow M$ such that the joint of f is $\{l, l^{-1}\}$. Moreover, this gluing is unique up to isomorphism: if $(M', \mathcal{C}', \tau')$ and $f' : M' \rightarrow M$ satisfy the same properties, then there exists a diffeomorphism $\psi : \text{Spl}_l(M) \rightarrow M'$ such that*

$$\psi(\text{Spl}_l(\mathcal{C})) = \mathcal{C}', \quad \psi \circ \text{Spl}_l(\tau) = \tau' \circ \psi \quad \text{and} \quad f' \circ \psi = f.$$

This result is intuitively obvious and yet it lacks a concise and rigorous proof. Considering that the splitting operation is defined by a picture (Figure 21.1) in E. Moise's reference book [33], we feel excused for not attempting what would be a lengthy and uninformative proof. The surface $\text{Spl}_l(M)$ can be defined as the topological space underlying the universal completion of the metric space $(M \setminus l, d)$, where d is the restriction to $M \setminus l$ of an arbitrary Riemannian distance on M . The universal property of the completion actually provides us with a continuous gluing map.

It is easy to determine from the pair $(M, \{l, l^{-1}\})$ if the gluing is binary or unary. In fact, the gluing is unary if and only if l does not admit an orientable neighborhood, which is equivalent to the fact that for every neighborhood U of l there exists a neighborhood $V \subset U$ of l such that $V \setminus l$ is connected. Another equivalent statement is that l admits a compact neighborhood which is homeomorphic to a Möbius band of which l is an equator.

When these equivalent properties do not hold and the gluing is binary, there is an issue of orientation about the way in which the two boundary components of $\text{Spl}_l(M)$ which are glued together are identified. If M is orientable, then $\text{Spl}_l(M)$ is also orientable, because a gluing performed on a non-orientable surface always results in another non-orientable surface, and the identification must be orientation-reversing. If M is non-orientable, then two situations arise. Either $\text{Spl}_l(M)$ is orientable, in which case the identification must be orientation-preserving, or $\text{Spl}_l(M)$ is non-orientable. In this last case, any identification of the two boundary components is convenient. This is not in contradiction with the uniqueness part of the statement since, thanks to Theorem 1.1.7, any diffeomorphism of the boundary of a non-orientable compact surface can be extended to a diffeomorphism of the whole surface.

Let (M, \mathcal{C}) be a marked surface. By successively applying Proposition 1.1.11 to the marks of M , one eventually gets a tubular pattern with no marks, from which one can reconstruct (M, \mathcal{C}) .

Proposition 1.1.12. — *Let (M, \mathcal{C}) be a marked surface. Let $(M, \mathcal{C}, \text{id})$ be the associated tubular pattern. There exists a tubular pattern (M', \emptyset, τ') and a smooth mapping $f : M' \rightarrow M$ which is a complete gluing in the sense of Definition 1.1.10. This gluing is unique up to isomorphism. The pattern (M', \emptyset, τ') is called a split tubular pattern of (M, \mathcal{C}) .*

The reader may find it an amusing exercise to draw a split tubular pattern of the marked surface depicted in Figure 1. In the next definition, we still use our special convention about the genus of a surface (see Theorem 1.1.3).

Definition 1.1.13. — Let (M, \mathcal{C}) be a marked surface. Let (M', \emptyset, τ') be a split tubular pattern of (M, \mathcal{C}) . Assume that M' is connected. We define the split genus of (M, \mathcal{C}) as the genus of M' . We denote it by $\text{sg}(M, \mathcal{C})$.

1.2. Curves and paths

1.2.1. Definitions. — In the theory of Markovian holonomy fields, curves and paths on a surface play the role of points of a time interval for classical Markov processes. Please note that the words *curve* and *path* are not interchangeable in this work: a path is a curve with finite length.

Definition 1.2.1. — Let M be a topological compact surface.

A parametrised curve on M is a continuous curve $c : [0, 1] \rightarrow M$ which is either constant on $[0, 1]$ or constant on no open sub-interval of $[0, 1]$. The set of parametrised curves is denoted by $\text{PC}(M)$.

Two parametrised curves on M are said to be equivalent if they differ by an increasing homeomorphism of $[0, 1]$. An equivalence class is simply called a curve and the set of curves on M is denoted by $\text{C}(M)$.

A continuous loop is a curve whose endpoints coincide. A continuous loop is said to be simple if it is injective on $[0, 1]$.

If c is a curve, then we denote respectively by \underline{c} and \bar{c} the starting and finishing point of c . We denote its inverse by c^{-1} . It is defined as the class of the parametrised curve $t \mapsto c(1 - t)$, which does not depend on the particular parametrization of c . The concatenation of curves is defined in the usual way. It is only partially defined on $\text{C}(M)$ but associative whenever this makes sense.

The space $\text{C}(M)$ is too large for many of our purposes. Let us define another space of curves which we call *paths*. Let M be a smooth compact surface endowed with a Riemannian metric. Let $c : [0, 1] \rightarrow M$ be a Lipschitz continuous curve. Then the derivative of c is defined almost-everywhere and its norm is bounded above. We are going to consider curves whose speed is also bounded *below* by a positive constant. Since the range of a curve is compact, this notion is independent of the choice of the Riemannian metric, as it would be on any smooth manifold.

Definition 1.2.2. — Let M be a smooth compact surface.

A parametrised path on M is a continuous curve $c : [0, 1] \rightarrow M$ which is either constant or Lipschitz continuous with speed bounded below by a positive constant. The set of parametrised paths is denoted by $\text{PP}(M)$.

Two parametrised paths on M are said to be equivalent if they differ by an increasing bi-Lipschitz homeomorphism of $[0, 1]$. An equivalence class is simply called a path and the set of paths on M is denoted by $P(M)$.

A loop is a path whose endpoints coincide. The set of loops is denoted by $L(M)$. A loop is said to be simple if it is injective on $[0, 1]$.

We use for paths the same notation as for curves. If c is a path, we denote its endpoints by \underline{c} and \bar{c} , and its inverse by c^{-1} . The concatenation of paths is also associative whenever it is defined. When M is endowed with a specific Riemannian metric, we usually identify $P(M)$ with the subset of $PP(M)$ consisting of paths which are parametrised at constant speed.

While the inclusion $PP(M) \subset PC(M)$ does not strictly speaking determine an inclusion $P(M) \subset C(M)$, because we are not using the same equivalence relation on parametrised curves and parametrised paths, it is true that a path, as a set of parametrised curves, is a subset of a unique curve. Moreover, two parametrised paths which are equivalent as parametrised curves are also equivalent as parametrised paths. Hence, there is a natural injection $P(M) \subset C(M)$ which we use without further comment.

Let us define a relation on $L(M)$ by saying that two loops l_1, l_2 are related if and only if there exists $c, d \in P(M)$ such that $l_1 = cd$ and $l_2 = dc$. It is not difficult to check that this is an equivalence relation.

Definition 1.2.3. — Let M be a smooth compact surface. A cycle is an equivalence class of loops for the relation on $L(M)$ just defined. We call non-oriented cycle a pair $\{l, l^{-1}\}$ where l is a cycle. We say that a cycle is simple if one of its representatives (hence all) are simple loops.

A cycle is simply a loop from which one has forgotten the origin. It is important to observe that an oriented 1-dimensional submanifold of M determines a simple cycle. Another definition derived from that of loops and which will be useful is the following.

Definition 1.2.4. — A path $l \in P(M)$ is called a lasso if there exists a path s and a simple loop m such that $l = sms^{-1}$.

Lemma 1.2.5. — Let l be a lasso. There exist a unique path s and a unique simple loop m such that $l = sms^{-1}$. The path s is called the spoke of l and the simple loop m the meander of l .

Proof. — Endow M with a Riemannian metric. Assume that l is parametrised at constant speed by $[0, 1]$. Then the meander of l is the restriction of l to the largest interval of the form $[\frac{1}{2} - t, \frac{1}{2} + t)$ on which l is injective. \square




FIGURE 3. A lasso. The dotted path is the spoke and the solid loop is the meander. Note that the spoke can intersect itself and the meander can intersect the spoke.

We will use the topology of uniform convergence on $C(M)$.

Definition 1.2.6. — Let M be a compact surface endowed with a Riemannian metric, whose Riemannian distance we denote by d . Let c_1, c_2 be two curves of M . We define

$$d_\infty(c_1, c_2) = \inf_{\text{param}} \sup_{t \in [0,1]} d(c_1(t), c_2(t)),$$

where the infimum is taken over all parametrizations of c_1 and c_2 . The distance d_∞ is known as the Fréchet metric.

The distance d_∞ depends on the Riemannian metric chosen on M . However, the topology on $C(M)$ does not.

Lemma 1.2.7. — Let M be a compact surface. The distances on $C(M)$ associated with any two Riemannian metrics on M are equivalent.

Proof. — Since M is compact, the Riemannian distances on M determined by any two Riemannian metrics are equivalent. \square

On $P(M)$, we will use a topology which is stronger than the trace of the uniform topology. We use an analogue of the topology of convergence in 1-variation of Lipschitz continuous paths, for which a sequence of Lipschitz continuous paths in a Euclidean space converges if it converges uniformly and the sequence of the derivatives of the paths converges in L^1 . For the moment, we introduce a metric on $P(M)$ which depends on a Riemannian metric on M and is apparently weaker than the distance in 1-variation.

When c is a path on a Riemannian surface, we denote by $\ell(c)$ its length.

Definition 1.2.8. — Let M be a compact surface endowed with a Riemannian metric. Let c_1, c_2 be two paths on M . We define

$$d_\ell(c_1, c_2) = d_\infty(c_1, c_2) + |\ell(c_1) - \ell(c_2)|.$$

It is not obvious, and we will prove in the next section, that the topology induced by d_ℓ on $P(M)$ does not depend on the Riemannian metric on M . Moreover, we will prove that this topology can be metrised by a distance for which $P(M)$ is a complete metric space.

The topology on $P(M)$ induced by the distance d_ℓ is the only one that we consider in this work and it is always the one to which we refer when we say that a sequence of paths converges. We will often add a condition on the endpoints of the paths which we consider.

Definition 1.2.9. — *Let $(c_n)_{n \geq 0}$ be a sequence of paths on M . Let c be a path on M . We say that $(c_n)_{n \geq 0}$ converges to c with fixed endpoints if*

- 1) $d_\ell(c_n, c) \rightarrow 0$ as $n \rightarrow \infty$,
- 2) for all $n \geq 0$, $\underline{c}_n = \underline{c}$ and $\bar{c}_n = \bar{c}$.

When M is endowed with a Riemannian metric, we will also make use of piecewise geodesic paths.

Definition 1.2.10. — *Let M be a compact surface endowed with a Riemannian metric γ . We define $A_\gamma(M)$ as the subset of $P(M)$ containing the piecewise geodesic paths, that is, the finite concatenations of segments of geodesics.*

The letter A stands for *affine*, instead of the letter G which will be used for many other things. We claim that $A_\gamma(M)$ is dense in $P(M)$. Indeed, there is an obvious way to approximate an arbitrary path by piecewise geodesic ones.

Recall that any two points of a Riemannian manifold which are closer than the injectivity radius are joined by a unique minimizing geodesic.

Definition 1.2.11. — *Let M be a compact surface endowed with a Riemannian metric γ . Consider $c \in P(M)$, identified with a path parametrised at constant speed. Let $n \geq 0$ be an integer. Assume that $2^{-n}\ell(c)$ is smaller than the injectivity radius of M . For each $k \in \{0, \dots, 2^n - 1\}$, let $\sigma_{n,k}$ be the minimizing geodesic which joins $c(k2^{-n})$ to $c((k+1)2^{-n})$. Then define*

$$D_n(c) = \sigma_{n,0} \cdots \sigma_{n,2^n-1}.$$

Proposition 1.2.12. — *Let M be a compact surface endowed with a Riemannian metric γ . For all path $c \in P(M)$, the sequence $(D_n(c))$ defined for n large enough converges to c with fixed endpoints. In particular, the space $A_\gamma(M)$ is dense in $P(M)$ for the convergence with fixed endpoints.*

Proof. — Let n be large enough for the path $D_n(c)$ to be defined. It has the same endpoints as c by construction. Let us parametrise it in such a way that for each

$k \in \{0, \dots, 2^n - 1\}$, the restriction of $D_n(c)$ to $[k2^{-n}, (k+1)2^{-n}]$ is the minimizing geodesic which joins $c(k2^{-n})$ to $c((k+1)2^{-n})$. It is straightforward that

$$\sup_{t \in [0,1]} d(c(t), D_n(c)(t)) \leq 2^{-n+1} \ell(c).$$

Hence, $D_n(c)$ converges uniformly towards c . Since the length is lower semi-continuous with respect to pointwise convergence, this implies that $\liminf \ell(D_n(c)) \geq \ell(c)$. On the other hand, $\ell(D_n(c)) \leq \ell(c)$, hence $\ell(D_n(c))$ converges to $\ell(c)$. \square

1.2.2. The complete metric space of rectifiable paths. — The goal of this section is to prove that the topology that we have introduced on $P(M)$ does not depend on a particular choice of a Riemannian metric on M and can be metrised by a complete distance. Let us start by a negative result.

Lemma 1.2.13. — *Let M be a compact surface endowed with a Riemannian metric. The metric space $(P(M), d_\ell)$ is not complete.*




FIGURE 4. Left: a piecewise geodesic approximation of a path. Right: a Cauchy sequence for d_ℓ which does not converge to its uniform limit, the dotted straight line.

Proof. — Let $c : [-\frac{1}{4}, \frac{5}{4}] \rightarrow M$ be a segment of minimizing geodesic parametrised at constant speed. For each $n \geq 1$, define $c_n : [0, 1] \rightarrow M$ by $c_n(t) = c(t + \frac{1}{n} \sin(2\pi nt))$. For all $n \geq 1$, $d_\infty(c_n, c|_{[0,1]}) = \frac{1}{n} \ell(c|_{[0,1]})$. Moreover, for all $n \geq 1$, $\ell(c_n) = 4\ell(c|_{[0,1]})$. Hence, the sequence $(c_n)_{n \geq 1}$ is a Cauchy sequence for d_ℓ which converges uniformly to $c|_{[0,1]}$. Its only possible limit is $c|_{[0,1]}$, but $\ell(c_n)$ does not converge to $\ell(c|_{[0,1]})$. \square

The main result of this section is the following.

Proposition 1.2.14. — *Let M be a compact surface.*

1) *The topologies induced on $P(M)$ by the distances d_ℓ associated to any two Riemannian metrics on M are the same.*

2) *Endow M with a Riemannian metric. There exists a metric d_1 on $P(M)$ which induces the same topology as d_ℓ and such that $(P(M), d_1)$ is a complete metric space.*

In order to prove this proposition, we define a new distance on $P(M)$, which is the analogue of the distance in 1-variation between Lipschitz continuous paths. Let TM denote the tangent bundle of M . The Levi-Civita connection of γ determines a splitting of the tangent bundle to TM as

$$T(TM) = T^V(TM) \oplus T^H(TM).$$

The vertical part $T^V(TM)$ is the kernel of the derivative of the bundle map $\pi : TM \rightarrow M$. It is canonically identified with TM by associating to $X \in T_mM$ the vector $\frac{d}{dt}|_{t=0}(Y + tX) \in T_Y(T_mM)$. The horizontal part $T^H(TM)$ is mapped isomorphically onto TM by the differential of π and the reciprocal mapping can be defined as follows. Consider $X, Y \in T_mM$. Let $c : (-1, 1) \rightarrow M$ be a smooth curve such that $c(0) = m$ and $\dot{c}(0) = X$. Let $Y(t)$ be the unique vector field along c such that $Y(0) = Y$ and $\nabla_{\dot{c}(t)}Y = 0$ for all t . Then the element of $T_Y^H(T_mM)$ which is sent to X by $T\pi$ is $\dot{Y}(0)$.

Since the tangent space to TM at each vector $X \in T_mM$ splits into the direct sum of two subspaces isomorphic to T_mM , there is a natural Riemannian metric on TM , which we denote by $\gamma \oplus \gamma$. The corresponding Riemannian distance on TM can be described as follows: if m and n are close enough on M to be joined by a unique minimizing geodesic and if $X \in T_mM, Y \in T_nM$, then

$$d_{TM}(X, Y) = (d(m, n)^2 + \|\parallel_{[m,n]}(X) - Y\|^2)^{\frac{1}{2}},$$

where $\parallel_{[m,n]}$ denotes the parallel transport along the unique minimizing geodesic from m to n .

Definition 1.2.15. — *Let M be a compact surface endowed with a Riemannian metric γ . Let c_1, c_2 be two paths on M . We define*

$$d_1(c_1, c_2) = \inf_{\text{param.}} \left(\sup_{t \in [0,1]} d(c_1(t), c_2(t)) + \int_0^1 d_{TM}(\dot{c}_1(t), \dot{c}_2(t)) dt \right),$$

where the infimum is taken over all parametrizations of c_1 and c_2 .

We define also

$$\bar{d}_1(c_1, c_2) = \sup_{t \in [0,1]} d(c_1(t), c_2(t)) + \int_0^1 d_{TM}(\dot{c}_1(t), \dot{c}_2(t)) dt,$$

where c_1 and c_2 are parametrised at constant speed.

It is clear that the inequalities $d_\ell \leq d_1 \leq \bar{d}_1$ hold. We are going to prove that these three metrics induce the same topology on $P(M)$. The main result is the following.

Proposition 1.2.16. — *Let M be a compact surface endowed with a Riemannian metric. Let c be a path on M and $(c_n)_{n \geq 1}$ a sequence of paths such that $d_\ell(c_n, c)$ tends to 0. Then $\bar{d}_1(c_n, c)$ tends to 0.*

Let us start with two preparatory lemmas.

Lemma 1.2.17. — *Let M be a compact surface endowed with a Riemannian metric. Let c be a path on M and $(c_n)_{n \geq 1}$ a sequence of paths such that $d_\ell(c_n, c)$ tends to 0. Then, the paths c_n and c being parametrised at constant speed, the uniform convergence holds:*

$$\sup_{t \in [0,1]} d(c_n(t), c(t)) \xrightarrow{n \rightarrow \infty} 0.$$

Proof. — Let us parametrise c and c_n for all $n \geq 1$ at constant speed. Let us also choose for all n a parametrization \tilde{c}_n of c_n such that the uniform convergence $\sup_{t \in [0,1]} d(\tilde{c}_n(t), c(t)) \rightarrow 0$ holds as n tends to infinity. Consider $t \in [0, 1]$. Since $\tilde{c}_n|_{[0,t]}$ and $\tilde{c}_n|_{[t,1]}$ converge uniformly respectively to $c|_{[0,t]}$ and $c|_{[t,1]}$, we have

$$\liminf \ell(\tilde{c}_n|_{[0,t]}) \geq \ell(c|_{[0,t]}) = t\ell(c) \quad \text{and} \quad \liminf \ell(\tilde{c}_n|_{[t,1]}) \geq \ell(c|_{[t,1]}) = (1-t)\ell(c).$$

Since $\ell(\tilde{c}_n)$ tends to $\ell(c)$, this implies that $\ell(\tilde{c}_n|_{[0,t]})$ tends to $t\ell(c)$ as n tends to infinity. This convergence holds for all $t \in [0, 1]$ and, since the functions $t \mapsto \ell(\tilde{c}_n|_{[0,t]})$ are non-decreasing, a classical result ensures that the convergence is uniform:

$$\sup_{t \in [0,1]} |\ell(\tilde{c}_n|_{[0,t]}) - t\ell(c)| \xrightarrow{n \rightarrow \infty} 0.$$

Now, for all $t \in [0, 1]$, $\tilde{c}_n(t) = c_n(\ell(\tilde{c}_n|_{[0,t]})/\ell(c_n))$. Since c_n is $\ell(c_n)$ -Lipschitz continuous, we have thus

$$d(\tilde{c}_n(t), c_n(t)) \leq |\ell(\tilde{c}_n|_{[0,t]}) - t\ell(c_n)| \leq |\ell(\tilde{c}_n|_{[0,t]}) - t\ell(c)| + t|\ell(c_n) - \ell(c)|.$$

The result follows easily. □

The second lemma is the Euclidean version of Proposition 1.2.16.

Lemma 1.2.18. — *Let $N \geq 1$ be an integer. Consider \mathbb{R}^N endowed with its usual Euclidean structure, with norm $\|\cdot\|$. Let f and $(f_n)_{n \geq 1}$ be measurable functions from $[0, 1]$ to the unit sphere of \mathbb{R}^N . Assume that the primitives of f_n converge uniformly to the primitive of f as n tends to infinity, that is,*

$$\sup_{t \in [0,1]} \left\| \int_0^t f_n(s) ds - \int_0^t f(s) ds \right\| \xrightarrow{n \rightarrow \infty} 0.$$

Then f_n converges in L^1 towards f , that is,

$$\int_0^1 \|f_n(t) - f(t)\| dt \xrightarrow{n \rightarrow \infty} 0.$$

The difficulty of this lemma is that the assumptions do not imply that the sequence (f_n) converges almost-everywhere to f . For example, if f is constant, f_n can be constant except on a small interval, which wanders around $[0, 1]$, and inside which f_n oscillates rapidly around the value of f .

Proof. — Since all functions take their values in a bounded subset of \mathbb{R}^N , it suffices to prove that the sequence $(f_n)_{n \geq 1}$ converges in measure to f , that is, denoting by Leb the Lebesgue measure on $[0, 1]$, to prove that

$$\forall \varepsilon > 0, \quad \lim_{n \rightarrow \infty} \text{Leb}(\{t \in [0, 1] : \|f_n(t) - f(t)\| > \varepsilon\}) = 0.$$

According to Lebesgue's differentiation theorem,

$$(16) \quad \frac{1}{2h} \int_{t-h}^{t+h} f(\tau) d\tau \xrightarrow{h \rightarrow 0} f(t) \quad \text{for a.e. } t \in (0, 1).$$

Let $p, q \geq 1$ be two integers. Set

$$C_{p,q} = \left\{ t \in [0, 1] : \forall h \in \left[0, \frac{1}{p}\right], \left\| \frac{1}{2h} \int_{t-h}^{t+h} f(\tau) d\tau - f(t) \right\| \leq \frac{1}{q} \right\}.$$

The relation (16) is equivalent to the fact that for all $q \geq 1$, $\text{Leb}(\bigcup_{p \geq 1} C_{p,q}) = 1$. Hence, for all $\alpha > 0$ and all $q \geq 1$, there exists $p \geq 1$ such that $\text{Leb}(C_{p,q}) \geq 1 - \alpha$.

Let us fix $\varepsilon > 0$. Then, let us choose two reals $\alpha > 0$, $r > 0$ and an integer $q \geq 1$. Let $p(q, \alpha) \geq 1$ be an integer such that $\text{Leb}(C_{p(q,\alpha),q}) \geq 1 - \alpha$. Set $h = 1/p(q, \alpha)$. Let $n_0(r) \geq 1$ be an integer such that,

$$\forall n \geq n_0(r), \quad \sup_{t \in [0,1]} \left\| \int_0^t f_n(s) ds - \int_0^t f(s) ds \right\| \leq \frac{1}{2r}.$$

Choose $n \geq n_0(r)$ and $t \in C_{p(q,\alpha),q}$. Then

$$\left\| \frac{1}{2h} \int_{t-h}^{t+h} f_n(\tau) d\tau - f(t) \right\| < \frac{1}{q} + \frac{1}{2hr}.$$

Since for all $s \in [0, 1]$, $\|f_n(s)\| = \|f(s)\| = 1$, we have

$$\forall \tau \in [t-h, t+h], \quad \|f_n(\tau) - f(t)\| > \frac{\varepsilon}{2} \implies 1 - \langle f_n(\tau), f(t) \rangle > \frac{\varepsilon^2}{8}.$$

Hence,

$$\begin{aligned} & \frac{1}{2h} \text{Leb}(\{\tau \in [t-h, t+h] : \|f_n(\tau) - f(t)\| > \frac{1}{2}\varepsilon\}) \\ & \leq \frac{8}{\varepsilon^2} \frac{1}{2h} \int_{t-h}^{t+h} \langle f(t) - f_n(\tau), f(t) \rangle d\tau \leq \frac{8}{\varepsilon^2} \left(\frac{1}{q} + \frac{1}{2hr} \right). \end{aligned}$$

The same inequality holds when f_n is replaced by f . Hence,

$$(17) \quad \frac{1}{2h} \text{Leb}(\{\tau \in [t-h, t+h] : \|f_n(\tau) - f(\tau)\| > \varepsilon\}) \leq \frac{16}{\varepsilon^2} \left(\frac{1}{q} + \frac{1}{2hr} \right).$$

This inequality holds for every $t \in C_{p(q,\alpha),q}$. Consider a subset T of $C_{p(q,\alpha),q}$ such that any two distinct points of T are at distance at least h . Take T to be maximal for inclusion among all subsets of $C_{p(q,\alpha),q}$ with this property. Then by the assumption of separation of the points of T , T has less than $1/h + 1$ points and by the maximality of T , the intervals $[t - h, t + h]$ with $t \in T$ cover $C_{p(q,\alpha),q}$. By applying (17) at the points of T , we find

$$\text{Leb}(\{t \in C_{p(q,\alpha),q} : \|f_n(t) - f(t)\| > \varepsilon\}) \leq \frac{32(1+h)}{\varepsilon^2} \left(\frac{1}{q} + \frac{1}{2hr} \right).$$

Since $\text{Leb}(C_{p(q,\alpha),q}) \geq 1 - \alpha$, and since $h = 1/p(q, \alpha) \leq 1$, we have finally proved that for all $\alpha > 0$, $r > 0$ and all $q \geq 1$, there exists $n_0(r)$ such that

$$\forall n \geq n_0(r), \quad \text{Leb}(\{t \in [0, 1] : \|f_n(t) - f(t)\| > \varepsilon\}) \leq \frac{64}{\varepsilon^2} \left(\frac{1}{q} + \frac{p(q, \alpha)}{2r} \right) + \alpha.$$

By choosing α sufficiently small and q sufficiently large, then r such that $p(q, \alpha)/(2r)$ is sufficiently small, this proves that the left-hand side of this inequality can be made arbitrarily small by choosing n sufficiently large. This is exactly the desired convergence. \square

Let us now prove Proposition 1.2.16.

Proof of Proposition 1.2.16. — Let us parametrise $(c_n)_{n \geq 1}$ and c at constant speed. For each $n \geq 1$, set $U_n = \sup_{t \in [0,1]} d(c_n(t), c(t))$. By Lemma 1.2.17, U_n tends to 0 as n tends to infinity. Hence, we need to prove that $\int_0^1 d_{TM}(\dot{c}_n(t), \dot{c}(t)) dt$ tends to 0. Let us choose n large enough for U_n to be smaller than the injectivity radius of M . Then

$$\begin{aligned} \int_0^1 d_{TM}(\dot{c}_n(t), \dot{c}(t)) dt &= \int_0^1 (d(c_n(t), c(t))^2 + \|\|_{[c_n(t), c(t)]} \dot{c}_n(t) - \dot{c}(t)\|^2)^{\frac{1}{2}} dt \\ &\leq U_n + \int_0^1 \|\|_{[c_n(t), c(t)]} \dot{c}_n(t) - \dot{c}(t)\| dt. \end{aligned}$$

Nash's embedding theorem grants the existence of an isometric embedding of M in a Euclidean space. Let $i : M \rightarrow \mathbb{R}^N$ be such an embedding. We denote its differential by di and, using the natural identification $T\mathbb{R}^N \simeq \mathbb{R}^N \times \mathbb{R}^N$, we see di as a map from TM to \mathbb{R}^N . For all $X \in TM$, we have $\|di(X)\|_{\mathbb{R}^N} = \|X\|$. Hence,

$$\begin{aligned} \|\|_{[c_n(t), c(t)]} \dot{c}_n(t) - \dot{c}(t)\| &= \|\|(di \circ \|\|_{[c_n(t), c(t)]})(\dot{c}_n(t)) - di(\dot{c}(t))\|_{\mathbb{R}^N} \\ &\leq \|\|(di \circ \|\|_{[c_n(t), c(t)]})(\dot{c}_n(t)) - di(\dot{c}_n(t))\|_{\mathbb{R}^N} + \|\|di(c_n(t)) - di(c(t))\|_{\mathbb{R}^N}. \end{aligned}$$

In order to estimate the first term, let us observe that, in local coordinates, the linear operator $di \circ \|\|_{[m,n]} - di$ is a rectangular matrix whose entries depend on (m, n)

and vanish when $m = n$. The square of the first term of the right-hand side above is bounded by

$$\mathrm{Tr}((\mathrm{di} \circ //_{[c_n(t), c(t)]} - \mathrm{di})(\mathrm{di} \circ //_{[c_n(t), c(t)]} - \mathrm{di})^*) \|\dot{c}_n(t)\|^2.$$

The trace, as a function of $(c_n(t), c(t))$, is a sum of squares of smooth functions on $M \times M$ which vanish on the diagonal. Since M is compact, any such smooth function is dominated by the Riemannian distance. Hence, there exists a constant $K > 0$ such that for all $m, n \in M$ and all $X \in T_m M$,

$$\|(\mathrm{di} \circ //_{[m, n]})(X) - \mathrm{di}(X)\|_{\mathbb{R}^N} \leq Kd(m, n)\|X\|.$$

Since the lengths of the paths c_n converge, $L = \sup\{\ell(c_n) : n \geq 1\}$ is finite. Hence, for n large enough, we have

$$\int_0^1 d_{TM}(\dot{c}_n(t), \dot{c}(t)) dt \leq (1 + KL)U_n + \int_0^1 \|\mathrm{di}(\dot{c}_n(t)) - \mathrm{di}(\dot{c}(t))\|_{\mathbb{R}^N} dt.$$

It suffices to prove that the last integral converges to 0. As a mapping between metric spaces, i is 1-Lipschitz continuous. Hence, $i(c_n)$ converges uniformly to $i(c)$ as n tends to infinity. Since i is a Riemannian isometry, $i(c_n)$ and $i(c)$ are also parametrised at constant speed for all $n \geq 1$, respectively $\ell(c_n)$ and $\ell(c)$. For all $n \geq 1$, define

$$f_n : [0, 1] \longrightarrow \mathbb{R}^N, \quad f_n = \frac{1}{\ell(c_n)} \mathrm{di}(\dot{c}_n).$$

Define also $f : [0, 1] \rightarrow \mathbb{R}^N$ by $f = 1/\ell(c) \cdot \mathrm{di}(\dot{c})$. These functions take their values in the unit sphere of \mathbb{R}^N . Since $i(c_n)$ converges uniformly to $i(c)$ as n tends to infinity and $\ell(c_n)$ tends to $\ell(c)$, the primitives of $(f_n)_{n \geq 1}$ converge uniformly to the primitive of f . By Lemma 1.2.18, this implies that $(f_n)_{n \geq 1}$ converges in L^1 to f . Using again the fact that $\ell(c_n)$ converges to $\ell(c)$, we find that the derivative of $i(c_n)$ converges in L^1 to the derivative of $i(c)$:

$$(18) \quad \int_0^1 \|\mathrm{di}(\dot{c}_n(t)) - \mathrm{di}(\dot{c}(t))\|_{\mathbb{R}^N} dt \xrightarrow{n \rightarrow \infty} 0.$$

This is the expected convergence. □

Lemma 1.2.19. — *Let M be a compact surface. The topology on $\mathcal{P}(M)$ induced by the distances d_1 associated to any two Riemannian metrics on M are the same.*

Proof. — Consider two Riemannian metrics γ and γ' on M . We will denote with a prime the quantities associated with γ' .

Let c be a path and $(c_n)_{n \geq 1}$ a sequence of paths such that $d_1(c_n, c)$, and thus also $d_\ell(c_n, c)$, tend to 0 as n tends to infinity. Let us parametrise c and each path c_n at constant speed with respect to γ . By Proposition 1.2.16, we have

$$\sup_{t \in [0, 1]} d(c_n(t), c(t)) + \int_0^1 d_{TM}(\dot{c}_n(t), \dot{c}(t)) dt \xrightarrow{n \rightarrow \infty} 0.$$

Set $L = \sup\{\ell(c_n) : n \geq 1\} \geq \ell(c)$. On the compact subset

$$B_L(TM) = \{X \in TM : \|X\|_\gamma \leq L\}$$

of TM , the distances d_{TM} and d'_{TM} are equivalent. Moreover, the distances d and d' on M are also equivalent. It follows that

$$\sup_{t \in [0,1]} d'(c_n(t), c(t)) + \int_0^1 d'_{TM}(\dot{c}_n(t), \dot{c}(t)) dt \xrightarrow{n \rightarrow \infty} 0$$

for some parametrization of c and the paths c_n . Hence, $d'_1(c_n, c)$ tends to 0. \square

Lemma 1.2.20. — *Let M be a compact surface endowed with a Riemannian metric. The metric space $(P(M), \bar{d}_1)$ is complete.*

Proof. — Let $(c_n)_{n \geq 1}$ be a Cauchy sequence of $P(M)$ for the distance \bar{d}_1 . Let us parametrise all these paths at constant speed. They form a Cauchy sequence for the uniform distance between continuous mappings from $[0, 1]$ to M , so they converge uniformly to some continuous mapping $c : [0, 1] \rightarrow M$.

Let us use Nash's theorem again to find an isometric embedding $i : M \rightarrow \mathbb{R}^N$. Since i is 1-Lipschitz continuous, the sequence $(i(c_n))_{n \geq 1}$ of paths in \mathbb{R}^N converges uniformly to $i(c)$.

The sequence $(c_n)_{n \geq 1}$ is in particular Cauchy for the distance d_ℓ , so that the sequence $(\ell(c_n))_{n \geq 1}$ converges to some real l . Set $L = \sup\{\ell(c_n) : n \geq 1\} < +\infty$. The restriction to the compact set $B_L(TM) = \{X \in TM : \|X\| \leq L\}$ of the smooth mapping $di : TM \rightarrow \mathbb{R}^N$ is Lipschitz continuous. Hence, the sequence $(i(c_n))_{n \geq 1}$ of paths in \mathbb{R}^N is also a Cauchy sequence for the L^1 distance of the derivatives. Hence, the derivatives $di(\dot{c}_n)$ converge in L^1 to some function $f : [0, 1] \rightarrow \mathbb{R}^N$ which takes its values in the sphere of radius $\lim_{n \rightarrow \infty} \ell(c_n) = l$. Passing the equality

$$\int_0^t di(\dot{c}_n(s)) = i(c_n(t)) - i(c_n(0))$$

to the limit, we find that f is the derivative of $i(c)$. Hence, $i(c)$ is a Lipschitz continuous path parametrised at constant speed l , and so is c . In particular, $l = \ell(c)$.

Finally, the sequence $(c_n)_{n \geq 1}$ satisfies $d_\ell(c_n, c) \rightarrow 0$ as n tends to infinity. By Proposition 1.2.16, this implies that $\bar{d}_1(c_n, c)$ tends to 0 as n tends to infinity. \square

Let us collect the results that we have proved and deduce Proposition 1.2.14.

Proof of Proposition 1.2.14. — Since $d_\ell \leq d_1 \leq \bar{d}_1$ and by Proposition 1.2.16, the three distances d_ℓ , d_1 and \bar{d}_1 induce the same topology on $P(M)$. By Lemma 1.2.19, this topology does not depend on the Riemannian metric on M . By Lemma 1.2.20, it is the topology of a complete metric space. \square

1.3. Graphs

1.3.1. Graphs and the sewing of patterns. — A graph on a surface is a finite set of paths or curves called *edges* and which satisfy several conditions. For Markovian holonomy fields, these finite sets of paths play the role of the finite sets of points in a time interval along which one considers the finite-dimensional marginals of a Markov process. The fact that most finite collections of paths are not the set of edges of a graph leads to quite a lot of technical complication: Markovian holonomy fields are stochastic processes of which only a small number of finite-dimensional marginals can be described by a simple formula.

Recall the distinction between curves, which are just continuous (see Definition 1.2.1) and paths, which are Lipschitz continuous (see Definition 1.2.2).

Definition 1.3.1. — *Let M be a topological compact surface.*

A curve on M which is injective or a simple continuous loop is called a continuous edge. The set of continuous edges on M is denoted by $\text{CE}(M)$.

Let M be a smooth compact surface. A path on M which is injective or a simple loop is called an edge. The set of edges on M is denoted by $\text{E}(M)$.

When we consider an edge or a continuous edge e , we will often abusively denote the range of e by e instead of $e([0, 1])$.

Definition 1.3.2. — *Let M be a connected compact surface (resp. a topological compact surface).*

A pre-graph on M is a triple $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$, where:

- 1) \mathbb{V} is a finite subset of M ,
- 2) \mathbb{E} is a non-empty finite subset of $\text{E}(M)$ (resp. $\text{CE}(M)$), stable by inversion, such that $\mathbb{V} = \bigcup_{e \in \mathbb{E}} \{e, \bar{e}\}$, and such that two edges of \mathbb{E} which are not each other's inverse meet, if at all, only at some of their endpoints,
- 3) \mathbb{F} is the set of the connected components of $M - \bigcup_{e \in \mathbb{E}} e([0, 1])$. The elements of $\mathbb{V}, \mathbb{E}, \mathbb{F}$ are called the vertices, edges and faces of the pre-graph.

A graph on M is a pre-graph which satisfies the condition

- 4) *Each face of \mathbb{G} is homeomorphic to an open disk of \mathbb{R}^2 .*

The skeleton of a pre-graph \mathbb{G} is the subset of M defined by $\text{Sk}(\mathbb{G}) = \bigcup_{e \in \mathbb{E}} e([0, 1])$. The set of paths (resp. curves) that can be obtained by concatenating edges of \mathbb{G} is denoted by $\text{P}(\mathbb{G})$ (resp. $\text{C}(\mathbb{G})$). The subset of $\text{P}(\mathbb{G})$ (resp. $\text{C}(\mathbb{G})$) consisting of loops is denoted by $\text{L}(\mathbb{G})$ (resp. $\text{CL}(\mathbb{G})$).

If M is homeomorphic to a sphere and m is a point of M , we include the exceptional triple $(\{m\}, \emptyset, \{M - \{m\}\})$ in the set of graphs.

A graph on a non-connected surface is defined as the data of a graph on each connected component of this surface.

Let (M, \mathcal{C}) be a marked surface. Let \mathbb{G} be a graph on M . We say that \mathbb{G} is a graph on (M, \mathcal{C}) if each cycle of \mathcal{C} is represented by a loop of $\mathbb{L}(\mathbb{G})$.

In the terminology of Mohar and Thomassen [32], what we call a graph on a topological surface is a cellular embedding of a combinatorial multigraph. It was proved by Radó in 1925 that every surface can be triangulated. In particular, on every topological compact surface there exists a graph. On a Riemannian surface, the proof of the fact that there exists a triangulation given in [32] is still valid if one uses only piecewise geodesic paths. Hence, a Riemannian surface admits a piecewise geodesic triangulation. This triangulation is a graph and, by adding some vertices, one may assume that the edges of this graph are geodesic. Finally, a Riemannian surface admits a graph with geodesic edges.

The goal of this section is twofold: to provide us with an efficient combinatorial tool for dealing with graphs (Proposition 1.3.8) and to establish a robust criterion which tells us when a pre-graph is a graph (Proposition 1.3.10).

In order to analyze a pre-graph or a graph, an effective method consists in splitting it along some of its edges. This is very similar to the surgery of smooth marked surfaces described in Section 1.1.2. The operations described here are however less regular and best defined in the category of topological surfaces.

Definition 1.3.3. — A pattern is a triple (M, \mathbb{G}, ι) , where M is a topological surface, \mathbb{G} is a pre-graph on M and ι is an involution of the set \mathbb{E} of edges of \mathbb{G} such that for all edge $e \in \mathbb{E}$,

- 1) $\iota(e) \neq e^{-1}$,
- 2) $\iota(e^{-1}) = \iota(e)^{-1}$,
- 3) $e \not\subset \partial M \Rightarrow \iota(e) = e$.

A pattern (M, \mathbb{G}, ι) is split if $\text{Sk}(\mathbb{G}) \subset \partial M$.

Two patterns (M, \mathbb{G}, ι) and $(M', \mathbb{G}', \iota')$ are isomorphic if there exists a homeomorphism $\psi : M \rightarrow M'$ such that $\psi(\mathbb{G}) = \mathbb{G}'$ and $\psi \circ \iota = \iota' \circ \psi$.

A pattern is meant to be sewed according to the identifications determined by its involution. Our convention here is slightly simpler than in the case of tubular patterns, in that we exclude the case $\iota(e) = e^{-1}$ which was the purely conventional encoding of unary gluings. Here, an edge e is always meant to be identified by an orientation-preserving homeomorphism with $\iota(e)$.

When $f : M' \rightarrow M$ is a continuous mapping between two surfaces and e' is a continuous edge on M' , we denote by $f(e')$ the curve $f \circ e'$.

Definition 1.3.4. — Let (M, \mathbb{G}, ι) and $(M', \mathbb{G}', \iota')$ be two patterns.

A continuous mapping $f : M' \rightarrow M$ is called an elementary sewing if it is the quotient map which identifies e' with $\iota'(e')$ by an orientation-preserving homeomorphism for some $e' \in \mathbb{E}'$. Moreover, it is required that $f(\mathbb{E}') = \mathbb{E}$ and $\iota \circ f = f \circ \iota'$ on $\mathbb{E}' - \{e', e'^{-1}, \iota(e'), \iota(e')^{-1}\}$.

The unoriented edge $\{f(e'), f(e')^{-1}\}$ is called the joint of the elementary sewing.

A sewing is a map which can be written as the composition of several elementary sewings. A sewing is complete if the involution of the set of edges of the target surface is the identity.

We have results for sewings which are similar to those we had for gluings. In particular, sewings can always be performed and a surface can always be split along an edge, provided the interior of the edge does not meet the boundary of the surface.

Proposition 1.3.5. — 1) Let $(M', \mathbb{G}', \iota')$ be a pattern. Consider $e' \in \mathbb{E}'$ such that $\iota(e') \neq e'$. There exists a pattern (M, \mathbb{G}, ι) and an elementary sewing $f : M' \rightarrow M$ such that the joint of f is $\{f(e'), f(e')^{-1}\}$. Moreover, this gluing is unique up to isomorphism: if $(M'', \mathbb{G}'', \iota'')$ and $f'' : M' \rightarrow M''$ satisfy the same properties, then there exists an isomorphism $\psi : (M, \mathbb{G}, \iota) \rightarrow (M'', \mathbb{G}'', \iota'')$ such that $\psi \circ f = f''$.

2) Let (M, \mathbb{G}, ι) be a pattern. Choose $\{e, e^{-1}\} \subset \mathbb{E}$ such that $e \cap \partial M \subset \{e, \bar{e}\}$. Then there exists a pattern $(M', \mathbb{G}', \iota')$ and an elementary sewing $f : M' \rightarrow M$ such that the joint of f is $\{e, e^{-1}\}$. Moreover, this sewing is unique up to isomorphism: if $(M'', \mathbb{G}'', \iota'')$ and $f'' : M'' \rightarrow M$ satisfy the same properties, then there exists an isomorphism $\psi : (M', \mathbb{G}', \iota') \rightarrow (M'', \mathbb{G}'', \iota'')$ such that $f'' \circ \psi = f$.

Just as Proposition 1.1.11, this result is obvious at a certain intuitive level but lacks a concise proof. The first assertion relies on the fact that the result of the identification of e with $\iota(e)$ is always a compact surface. This fact is explained in [32], at the beginning of Section 3.1. The second assertion relies on the Jordan curve theorem and on the theorem of Schönfliess, which asserts that the group of homeomorphisms of \mathbb{R}^2 acts transitively on the set of parametrised Jordan curves. A self-contained exposition of the theorems of Jordan and Schönfliess and of results which are very close to the forthcoming Proposition 1.3.8 can be found in the book of B. Mohar and C. Thomassen [32].

Let us only discuss the second assertion when the edge e is a simple loop and the equator of a Möbius band. In this case, the surface M' has one more boundary component than M and this boundary component is covered by two unoriented edges e'_1 and e'_2 , which we may assume to be oriented in such a way that the concatenation $e'_1 e'_2$ makes sense. In this case, $e'_1 e'_2$ is a loop which represents the new boundary component of M' and the involution ι exchanges e'_1 and e'_2 .




FIGURE 5. A three-holed sphere endowed with a graph, the same surface cut along all edges of the graph, and the corresponding split pattern.

Since a gluing is a special case of a sewing, Proposition 1.3.5 implies that a pregraph can be lifted through a splitting, in a way which is unique up to homeomorphism. It also implies the following result.

Lemma 1.3.6. — *Let M be a compact topological surface. The group of homeomorphisms of M acts transitively on the set of injective continuous edges contained in the interior of M .*

Proof. — An injective continuous edge contained in the interior of M determines a pre-graph on M . The associated split pattern is simply M to which a disk has been removed. The boundary of this disk is the concatenation of two edges which are identified with each other's inverse by the involution. Hence, this split pattern does not depend, up to homeomorphism, on the edge. \square

Corollary 1.3.7. — *Let M be a topological compact surface. Let e be an injective continuous edge contained in the interior of M . There exists a graph on M of which e is an edge.*

Proof. — Let \mathbb{G} be a graph on M . If M is a disk and $\text{Sk}(\mathbb{G}) \subset \partial M$, or if M is a sphere and \mathbb{G} is the exceptional graph with no edge, let us add to \mathbb{G} a continuous edge whose interior is contained in the interior of M . In any other case, \mathbb{G} contains an edge whose interior is contained in the interior of M . By adding vertices to \mathbb{G} if necessary, we may assume that it has an edge, say e_1 , contained in the interior of M . The image of \mathbb{G} by a homeomorphism of M which sends e_1 to e is a graph on M of which e is an edge. \square

By successively applying Proposition 1.3.5 in order to split all the edges of a pre-graph which are not located on the boundary, we end up with a split pattern.

Proposition 1.3.8. — *Let M be a topological compact surface. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a pre-graph on M . Assume that each edge of \mathbb{G} is either contained in ∂M or has no interior point on ∂M . Endow \mathbb{E} with the identity involution. There exists a split pattern (M', \mathbb{G}', ι) and a sewing $f : M' \rightarrow M$ such that $f(\mathbb{E}') = \mathbb{E}$. For each face F of \mathbb{G} , $f^{-1}(F)$ is the interior of a connected component of M' which we denote by M'_F . The sewing map f applies $M'_F \setminus (\text{Sk}(\mathbb{G}') \cap M'_F)$ homeomorphically onto F and $\text{Sk}(\mathbb{G}') \cap M'_F$ continuously onto the topological boundary of F .*

We call $(M', \mathbb{G}', \iota, f)$ a split pattern of the pair (M, \mathbb{G}) .

If $(M'', \mathbb{G}'', \iota'')$ is another split pattern and $f'' : M'' \rightarrow M$ a sewing which sends \mathbb{G}'' to \mathbb{G} , then there exists an isomorphism of patterns $\psi : M' \rightarrow M''$ such that $f'' \circ \psi = f$. Finally, if M is oriented, then M' can be oriented and the sewing map can be assumed to be orientation-preserving.

One of the simplest consequences of this result is that a pre-graph has a finite number of faces. Let us identify a simple condition under which the assumption on the edges of pre-graph made in Proposition 1.3.8 are satisfied.

Lemma 1.3.9. — *Let M be a topological compact surface. Let \mathbb{G} be a pre-graph on M . Let c be a subset of M homeomorphic to a circle. Then c is the image of a simple loop of $\text{CL}(\mathbb{G})$ if and only if $c \subset \text{Sk}(\mathbb{G})$. Moreover, if $c \subset \text{Sk}(\mathbb{G})$, then for each edge e of \mathbb{G} , either e is contained in c or e has no interior point on c .*

Proof. — One implication in the first assertion is obvious. Assume that $c \subset \text{Sk}(\mathbb{G})$. The image of $(0, 1)$ by an edge is homeomorphic to $(0, 1)$, hence it cannot contain a subset homeomorphic to a circle. Thus, there is at least one vertex on c . Let us choose a continuous parametrization of c by $[0, 1]$, injective on $(0, 1)$, such that $c(0) = c(1) \in \mathbb{V}$. The set $c \cap \mathbb{V}$ is finite and its complement in c is a finite union of open intervals. Let (a, b) be such an interval. Each point of (a, b) belongs to one single edge of \mathbb{G} . Assume that there exists u, v with $a < u < v < b$, such that u and v do not belong to the same edge. Since for each given edge, the subset of $[0, 1]$ covered by that edge is closed, there should necessarily be a point between u and v which is covered by at least two distinct edges. Hence, (a, b) is covered by a single edge. Both a and b must be vertices of this edge and the result follows.

Let e be an edge which has an interior point on c . Let us choose a parametrization of e and $t \in (0, 1)$ such that $e(t) \in c$. Let $I \subset [0, 1]$ be the largest segment containing t such that $e(I) \subset c$, that is, the connected component of t in $\{s \in [0, 1] : e(s) \in c\}$. Assume first that $I = \{t\}$. In this case, since c is contained in $\text{Sk}(\mathbb{G})$, $e(t)$ belongs to the closure of another edge of \mathbb{G} , hence to another edge, and it is a vertex of \mathbb{G} . This is impossible since $t \notin \{0, 1\}$ by assumption. Let us now assume that $I = [a, b]$ with $a < b$. Since e is an edge, the equality $e(a) = e(b)$ can occur only if $a = 0$ and $b = 1$, in which case e is contained in c . Actually, in this case, e is a simple loop whose range is c . Assume now that $e(a) \neq e(b)$. Then $e(I)$ is a subset of c homeomorphic to a segment. Since $c \subset \text{Sk}(\mathbb{G})$, each endpoint of this segment belongs to the range of another edge of \mathbb{G} . Hence $e(a)$ and $e(b)$ are vertices. This forces $a = 0$, $b = 1$ and in particular the fact that e is contained in c . \square

In our definition of graphs, the focus is put on edges: a graph is a set of edges which satisfies certain properties. It is important that we find a robust criterion which tells us when a pre-graph satisfies the topological condition (labelled 4. in Definition 1.3.2) which makes it a graph. By a robust criterion, we mean a criterion which makes it obvious that a pre-graph whose edges are close to those of a graph is also a graph. Let us apply Proposition 1.3.8 to establish such a criterion.

Proposition 1.3.10. — *Let M be a connected topological surface. Let \mathbb{G} be a pre-graph on M . The following properties are equivalent.*

- 4) *Each face of \mathbb{G} is homeomorphic to an open disk of \mathbb{R}^2 .*
- 4') *The skeleton of \mathbb{G} is connected, contains ∂M , and there exists $v \in \text{Sk}(\mathbb{G})$ such that any loop in M based at v is homotopic to a loop whose image is contained in $\text{Sk}(\mathbb{G})$.*

In particular, if \mathbb{G} is a graph, then each connected component of ∂M is the image of a loop of $\text{CL}(\mathbb{G})$.

Of course, if 4' is satisfied for some $v \in \text{Sk}(\mathbb{G})$, it is satisfied for all such v . In the course of the proof, we use the following lemma.

Lemma 1.3.11. — *Let M be a topological surface. Let \mathbb{G} be a topological pre-graph on M . Assume that $\text{Sk}(\mathbb{G})$ is connected and contains ∂M . Let v be a point of $\text{Sk}(\mathbb{G})$. Consider the quotient topological space $M/\text{Sk}(\mathbb{G})$, in which all the points of $\text{Sk}(\mathbb{G})$ are identified. Then the natural mapping $\pi_1(M, v) \rightarrow \pi_1(M/\text{Sk}(\mathbb{G}), [v])$ is onto.*

Proof. — Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . Since $\text{Sk}(\mathbb{G})$ contains ∂M , we have $\text{Sk}(\mathbb{G}') = \partial M'$. Hence, the ill-defined mapping $f^{-1} : M \rightarrow M'$ descends to a well-defined mapping $f^{-1} : M/\text{Sk}(\mathbb{G}) \rightarrow M'/\partial M'$, which is a homeomorphism. In particular, $[v]$ admits in $M/\text{Sk}(\mathbb{G})$ a neighbourhood homeomorphic to a finite bunch of disks whose centres are identified, thus a contractible neighbourhood. Hence, any loop in $M/\text{Sk}(\mathbb{G})$ based at $[v]$ is homotopic to a finite product of loops based at $[v]$ and whose interior does not visit $[v]$.

Choose a loop l based at $[v]$ on $M/\text{Sk}(\mathbb{G})$. Assume that l is homotopic to $l_1 \cdots l_n$ and the interiors of l_1, \dots, l_n do not visit $[v]$. Each loop l_i corresponds via f^{-1} to a loop in the space $M'_F/\partial M'_F$ for some $F \in \mathbb{F}$. Such a loop can be lifted to a path c_i on M'_F which starts and finishes on $\partial M'_F$ and stays in the interior of M'_F in the meantime. Since $\text{Sk}(\mathbb{G})$ contains ∂M , the paths $f(c_i)$ are paths on M which start and finish in $f(\partial M') = \text{Sk}(\mathbb{G}) \cup \partial M = \text{Sk}(\mathbb{G})$. Since $\text{Sk}(\mathbb{G})$ is connected, it is possible to connect their endpoints inside $\text{Sk}(\mathbb{G})$ and thus to produce a loop in M based at v whose image in the quotient $M/\text{Sk}(\mathbb{G})$ is $l_1 \cdots l_n$. \square

Proof of Proposition 1.3.10. — 4) \Rightarrow 4'). Let us assume that assumption 4 is satisfied. Since the faces of \mathbb{G} are homeomorphic to open disks, they contain no point of ∂M . Hence, the skeleton of \mathbb{G} contains ∂M . Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . By Proposition 1.3.8, the interior of M'_F is homeomorphic to an open disk for each F , so that M'_F is a closed disk. In particular, $\partial M'_F$ is connected. By Proposition 1.3.8 again, it follows that the boundary of each face is connected. If $\text{Sk}(\mathbb{G})$ was not connected, there would exist a face whose boundary meets two distinct connected components of $\text{Sk}(\mathbb{G})$. The boundary of this face would not be connected: this is impossible.

For each face F of \mathbb{G} , choose a point x_F in F . Choose $v \in \text{Sk}(\mathbb{G})$. It is well known that any continuous loop in M based at v is homotopic to a loop which avoids the points $x_F, F \in \mathbb{F}$. To see this, endow M with a Riemannian metric. Then, any two loops which are closer in uniform distance than the convexity radius of M are homotopic to each other. In particular, any loop is homotopic to a piecewise geodesic loop and it is possible to choose this loop such that it avoids the points $x_F, F \in \mathbb{F}$. By Proposition 1.3.8, the skeleton of \mathbb{G} is a retract by deformation of $M - \{x_F : F \in \mathbb{F}\}$.

Hence, any loop based at v is homotopic in M to a loop which stays in the skeleton of \mathbb{G} .

4') \Rightarrow 4). Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . Let $[v]$ denote the class of v in the quotient topological space $M/\text{Sk}(\mathbb{G})$. This class is nothing but $\text{Sk}(\mathbb{G})$. The third part of assumption 4' implies that the homomorphism $\pi_1(M, v) \rightarrow \pi_1(M/\text{Sk}(\mathbb{G}), [v])$ induced by the quotient mapping is trivial. By Lemma 1.3.11 above, this homomorphism is surjective. Hence, assumption 4' implies that $M/\text{Sk}(\mathbb{G})$ is simply connected. Hence, it implies that $M'/\partial M'$ is simply connected. The fundamental group of this space is isomorphic to the free product of the fundamental groups of the spaces $M'_F/\partial M'_F$. Hence, assumption 4' implies that each space $M'_F/\partial M'_F$ is simply connected. Up to homeomorphism, there exist only two connected compact surfaces which, when all their boundary points are identified to a single point, are simply connected: the sphere and the disk. Finally, assumption 4' implies that all surfaces M'_F are homeomorphic to disks, that is, the assertion 4.

The last assertion follows from Lemma 1.3.9 and the fact that the skeleton of a graph covers ∂M . \square

Corollary 1.3.12. — *Let M be a topological compact surface. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on M . For each $n \geq 0$, let $\mathbb{G}_n = (\mathbb{V}, \mathbb{E}_n, \mathbb{F}_n)$ be a pre-graph on M equipped with a bijection $S_n : \mathbb{E} \rightarrow \mathbb{E}_n$ such that for all $e \in \mathbb{E}$, $S_n(e^{-1}) = S_n(e)^{-1}$. Assume that for all $n \geq 0$ and all edge e such that $e \subset \partial M$, $S_n(e) = e$. Assume also that for all $e \in \mathbb{E}$, the sequence $(S_n(e))_{n \geq 0}$ converges uniformly to e with fixed endpoints. Then, for n large enough, \mathbb{G}_n is a graph on M .*

Proof. — It suffices to check that \mathbb{G}_n satisfies the condition 4' of Proposition 1.3.10 for n large enough. Firstly, for all $n \geq 0$, the skeleton of \mathbb{G}_n contains ∂M because the skeleton of \mathbb{G} does and every edge of \mathbb{G} located on ∂M is also an edge of \mathbb{G}_n . Then, for all $n \geq 0$, the skeleton of \mathbb{G}_n is connected. Indeed, let m and m' be two points of $\text{Sk}(\mathbb{G}_n)$. They can be respectively joined inside $\text{Sk}(\mathbb{G}_n)$ to two vertices v and v' , which are also vertices of \mathbb{G} . Since $\text{Sk}(\mathbb{G})$ is connected, there exist a curve $e_1 \cdots e_k$ in $\text{C}(\mathbb{G})$ which joins v to v' . The curve $S_n(e_1) \cdots S_n(e_k)$ joins v to v' inside $\text{Sk}(\mathbb{G}_n)$. Finally, for n large enough and for all $e \in \mathbb{E}$, $S_n(e)$ is homotopic with fixed endpoints to e . By choosing a point of \mathbb{V} as base point, we find that any loop in $\text{Sk}(\mathbb{G})$ is homotopic to a loop in $\text{Sk}(\mathbb{G}_n)$. This finishes the proof. \square

1.3.2. The boundary of a face. — Although a face in a graph is, by definition, homeomorphic to an open disk, its closure needs not be homeomorphic to a closed disk and even when it is the case, the topological boundary of the face may not be homeomorphic to a circle. The boundary of a face of a graph can in fact be defined as a cycle in the graph and this is the notion which matters for us. The appropriate

intuitive picture is that of someone walking in the interior of the face, keeping her right hand on the boundary. If the surface is non-orientable, the boundary of the face is a non-oriented cycle.




FIGURE 6. These are two faces of the graph represented on Figure 5. The closure of the first is not homeomorphic to a closed disk. The closure of the second is homeomorphic to a closed disk, but its topological boundary is not homeomorphic to a circle.

The following definition makes sense thanks to Proposition 1.3.8, in particular the statement of uniqueness.

Definition 1.3.13. — Let M be a topological compact surface. Let \mathbb{G} be a graph on M . Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . Let F be a face of \mathbb{G} . Let M'_F be the connected component of M' such that $f(M'_F) = \bar{F}$.

If M is oriented, then the boundary of F is defined as the cycle $\partial F = f(\partial M'_F)$ in $\text{CL}(\mathbb{G})$. If M is not oriented, we may still orient M'_F and the boundary of F is defined as the unoriented cycle $\partial F = \{f(\partial M'_F), f(\partial M'_F)^{-1}\}$.

A cycle of the form ∂F for some face F is called a facial cycle of (M, \mathbb{G}) .

This definition allows us to make sense of an edge adjacent to a face.

Definition 1.3.14. — Let M be a topological compact surface endowed with a graph \mathbb{G} . Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . Let F be a face of \mathbb{G} and M'_F the corresponding connected component of M' . Let e be an edge of \mathbb{G} .

We say that the unoriented edge $\{e, e^{-1}\}$ is adjacent to F if there exists an edge e' of \mathbb{G}' such that $e' \subset \partial M'_F$ and $f(e') \in \{e, e^{-1}\}$.

If M is oriented and M' is oriented accordingly, we say that e is adjacent to F if there exists an edge e' with the same properties as above and e' bounds M'_F positively.

An unoriented edge is adjacent to a face if and only if it is contained in its topological closure. When M is oriented, it follows from Proposition 1.3.8 that each oriented edge is adjacent to exactly one face. It may however occur that e and e^{-1} are adjacent to the same face.

The content of the next result is that an edge which is adjacent to two distinct faces can be removed from a graph. Recall that a marked surface is a pair (M, \mathcal{C}) where

\mathcal{C} is a set of smooth curves on M (see Definition 1.1.8) and that a graph on (M, \mathcal{C}) is a graph \mathbb{G} on M such that each curve of \mathcal{C} is the range of a cycle of $\mathbb{C}(\mathbb{G})$.

Proposition 1.3.15. — *Let (M, \mathcal{C}) be a marked surface. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on (M, \mathcal{C}) .*

1) *Let e be an edge of \mathbb{G} which is not contained in any curve of \mathcal{C} . Assume that e is adjacent to two distinct faces F_1 and F_2 . Write $\partial F_1 = ce$ and $\partial F_2 = e^{-1}d$ for some $c, d \in \mathbb{P}(\mathbb{G})$.*

Then $\mathbb{E} \setminus \{e, e^{-1}\}$ is the set of edges of a graph on (M, \mathcal{C}) , denoted by $\mathbb{G} \setminus e$, with the same faces as \mathbb{G} , except for the faces F_1 and F_2 which are replaced by $F = F_1 \cup F_2 \cup e((0, 1))$. Moreover, $\partial F = cd$.

2) *Let e be an edge of \mathbb{G} which finishes at a vertex of degree 1, that is, such that the terminal point of e is the terminal point of no other edge of \mathbb{G} . Let F be the unique face adjacent to e . Let $c \in \mathbb{P}(\mathbb{G})$ be such that $\partial F = cee^{-1}$ or $\partial F = ce^{-1}e$.*

Then $\mathbb{E} \setminus \{e, e^{-1}\}$ is the set of edges of a graph on (M, \mathcal{C}) , denoted by $\mathbb{G} \setminus e$, with the same faces as \mathbb{G} , except for the face F which is replaced by $F \cup e((0, 1])$. Moreover, $\partial F = c$.




FIGURE 7. Removing one edge of the graph depicted on Figure 5.

Proof. — The proofs of the two assertions are very similar. We prove only the first one. Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . By suitably orienting M' , we may assume that e is the image by f of an edge e'_1 which bounds M'_{F_1} positively and e^{-1} the image of an edge $e'_2 = \iota(e'_1)$ which bounds M'_{F_2} positively. Let us write $\partial M'_{F_1} = c'e'_1$ and $\partial M'_{F_2} = e'_2d'$, where c' and d' are curves in \mathbb{G}' which satisfy $f(c') = c$ and $f(d') = d$.

Let us assume first that either c' or d' is not the constant curve, that is, that either $\partial M'_{F_1} \neq e'_1$ or $\partial M'_{F_2} \neq e'_2$. In this case, sewing e'_1 and e'_2 results in a new surface M'_F which is still homeomorphic to a closed disk.

By removing the inner edge of this disk, we obtain a new split pattern $(M'', \mathbb{G}'', \iota'')$ with one connected component less than M' . The mapping $f'' : M'' \rightarrow M$ induced by f is a complete sewing of this pattern, so that $(M'', \mathbb{G}'', \iota'', f'')$ is a split pattern of $(M, \mathbb{G} \setminus e)$. It follows on one hand that $\mathbb{G} \setminus e$ is a graph and on the other hand that the boundary of the new face F is $f''(c'd') = f(c')f(d') = cd$.

It remains to check that each curve of \mathcal{C} is represented by a loop of $\mathbb{G} \setminus e$. By Lemma 1.3.9, it suffices to check that each curve of \mathcal{C} is contained in $\text{Sk}(\mathbb{G} \setminus e)$. Consider $l \in \mathcal{C}$. By the second assertion of Lemma 1.3.9, the assumption that e is not contained in any curve of \mathcal{C} ensures that e has at most some of its endpoints on l . Hence, $\text{Sk}(\mathbb{G} \setminus e)$ contains at least the complement of a finite set in l , hence l itself because it is closed. This finishes the proof.

Let us now treat the case where $\partial M'_{F_1} = e'_1$ and $\partial M'_{F_2} = e'_2$. In this case, the image of $M'_{F_1} \cup M'_{F_2}$ by f is a sphere of which e is an equator. This sphere is a connected component of M and, ignoring possible other connected components, we have $\mathbb{E} = \{e, e^{-1}\}$ and $\mathcal{C} = \emptyset$. Hence, $\mathbb{G} \setminus e$ is indeed a graph, namely the exceptional graph with no edge and a single vertex. It has a unique face whose boundary is the constant curve at this vertex. \square

The difficulty with the definition of the boundary of a face given by Definition 1.3.13 is the same as that regarding the topological condition on the faces of a graph and which led us to state Proposition 1.3.10. It is not obvious from this definition that a small deformation of the edges of a graph cannot significantly modify the facial cycles. Since there will come a point in this work at which we will need to compare graphs with close edges, we need to be able to extract in a fairly explicit and robust way the amount of combinatorial structure of a graph which determines its facial cycles.

The content of the theorem of Schönfiess is that there is no local topological invariant of a simple curve in a surface. Hence, in a graph, the only place where some local structure arises is at the vertices. This structure at a given vertex is completely described by the cyclic order of the edges which share this vertex as an endpoint. When the surface is orientable, the information of these cyclic orders is

in fact sufficient to determine completely the facial cycles of the graph, hence, by Proposition 1.3.8, the pair (M, \mathbb{G}) up to homeomorphism. When the surface is not orientable, a small amount of global information is needed to recover the facial cycles. Before explaining this, let us describe precisely what we mean by the cyclic order of the edges at a vertex. The characterization of this order which we establish now will be useful later.

In the next lemma, the surface is equipped with a differentiable structure because we need to consider a Riemannian metric on it. Nevertheless, the graph is allowed to have continuous edges. Also, the result is stated on a surface without boundary. If M has a boundary, then the cyclic order of the edges at a vertex should be computed after a disk has been glued along each boundary component of M .

Lemma 1.3.16. — *Let M be a smooth surface without boundary. Let \mathbb{G} be a graph on the topological surface underlying M . Let $v \in \mathbb{V}$ be a vertex. Let e_1, \dots, e_n be n parametrised curves which represent the edges of \mathbb{G} which share v as their starting point. Let γ be a Riemannian metric on M , whose Riemannian distance is denoted by d . Let R be the injectivity radius of γ . Set*

$$r_0 = \min (\{R\} \cup \{d(v, e_i(\frac{1}{2})) : i \in \{1, \dots, n\}\}).$$

For each $r \in (0, r_0)$ and each $i \in \{1, \dots, n\}$, define $s_i(r), t_i(r) \in [0, \frac{1}{2}]$ by

$$s_i(r) = \inf \{t \in [0, \frac{1}{2}] : d(v, e_i(t)) = r\}, \quad t_i(r) = \sup \{t \in [0, \frac{1}{2}] : d(v, e_i(t)) = r\}.$$

If M is not oriented, choose an orientation of the ball $B(v, r_0)$. For each $r \in (0, r_0)$, let $\omega_{\text{first}}(r)$ be the cyclic permutation of $\{e_1, \dots, e_n\}$ corresponding to the cyclic order of the points $e_1(s_1(r)), \dots, e_n(s_n(r))$ on the circle $C(v, r)$, oriented as the boundary of the ball $B(v, r)$. Similarly, let $\omega_{\text{last}}(r)$ be the cyclic permutation of $\{e_1, \dots, e_n\}$ corresponding to the cyclic order of the points $e_1(t_1(r)), \dots, e_n(t_n(r))$ on the circle $C(v, r)$. Then the following properties hold:

- 1) The cyclic order $\omega_{\text{first}}(r)$ does not depend on $r \in (0, r_0)$. We denote it simply by ω_{first} .
- 2) There exists $r_1 \in (0, r_0)$ such that for all $r \in (0, r_1)$, $\omega_{\text{last}}(r) = \omega_{\text{first}}$.

In the proof of this lemma, we take the following fact (which can be deduced from Proposition 1.3.8) for granted. On the compact cylinder $[0, 1] \times S^1$, consider n injective continuous curves c_1, \dots, c_n which do not meet each other. Assume that each curve starts at a point of $\{0\} \times S^1$ and finishes at a point of $\{1\} \times S^1$. Assume that no point of these curves other than their endpoints is located on the boundary of the cylinder. Then there exists an orientation-preserving homeomorphism of the cylinder onto itself which sends each curve to a set of the form $[0, 1] \times \{z\}$ for some $z \in S^1$.

In particular, the cyclic order of the initial points of c_1, \dots, c_n on the circle $\{0\} \times S^1$ is the same as the cyclic order of their terminal points on the circle $\{1\} \times S^1$.

Proof. — Let us choose $r \in (0, r_0)$ and $r' \in (0, r)$. For each $i \in \{1, \dots, n\}$, let us define

$$u_i(r', r) = \sup \{t \in [0, s_i(r)] : d(v, e_i(t)) = r'\}.$$

Thus, $c_i = e_i([u_i(r', r), s_i(r)])$ is an injective curve which joins $C(v, r')$ to $C(v, r)$ and stays in the annulus $r' \leq d(v, \cdot) \leq r$. Moreover, only the endpoints of c_i lie on the boundary of the annulus. This annulus is homeomorphic to a cylinder and the curves c_1, \dots, c_n do not meet each other. According to the remark made before the proof, the cyclic order of the points $e_1(u_1(r', r)), \dots, e_n(u_n(r', r))$ on the circle $C(v, r')$, which we denote by $\omega_{\text{mixed}}(r', r)$, is the same as the cyclic order of $e_1(s_1(r)), \dots, e_n(s_n(r))$ on $C(v, r)$, which is by definition $\omega_{\text{first}}(r)$.

Set $r_1(r) = \min(\{r\} \cup \{d(v, e_i([s_i(r), \frac{1}{2}])) : i \in \{1, \dots, n\}\})$. Since the edges e_1, \dots, e_n are injective paths, $r_1(r)$ is a positive number and $u_i(r', r) = t_i(r')$, for all $r' < r_1(r)$. Hence, for all $r \in (0, r_0)$ and $r' \in (0, r_1(r))$,

$$\omega_{\text{last}}(r') = \omega_{\text{mixed}}(r', r) = \omega_{\text{first}}(r).$$

Both assertions follow from this equality. □




FIGURE 8. The cyclic order of first exit of the disk bounded by the dotted circle for these six edges is (123456) but their cyclic order of last exit of the same disk is (214365). For a disk smaller than the grey disk, the two cyclic orders coincide.

Let us describe informally the algorithm which one uses to compute the facial cycles of a graph. First, one has to land somewhere on the surface, near the boundary of a face and to grasp the nearest edge with either hand. Then, one walks forward to the next vertex without breaking the contact with the edge. There, one performs two

operations. The first consists in changing the hand which holds the edge in order to turn round. One now has the vertex in one's back. The second operation consists in grasping with one's free hand the only outgoing edge at this vertex that one is not already holding and that is accessible without crossing any edge, and finally releasing the first edge. When one reaches a vertex at which there is only one outgoing edge, one turns around this vertex and walks back along the same edge, on the other side and holding it with the other hand. This process has to be iterated until one comes back to one's initial position.




FIGURE 9. A possible track followed by our explorer. The landing point is denoted by a cross.

Formally, the facial cycles arise as the cycles of a certain permutation on a set which corresponds to the possible ways of our explorer holding an edge. We call this set a *framing* of the graph.

Definition 1.3.17. — Let M be a compact topological surface. Let \mathbb{G} be a graph on M . An orientation of the vertices of \mathbb{G} is a collection $(U_v)_{v \in \mathbb{V}}$ of pairwise disjoint open subsets of M such that for all vertex $v \in \mathbb{V}$, U_v is an orientable and oriented neighbourhood of v .

Given an orientation of the vertices of \mathbb{G} , and for each edge e , we use the orientation of U_e to determine a left and a right of e , at least in the vicinity of e . If e is located on the boundary of M , we say that it bounds M positively if M is on the left of e .

Definition 1.3.18. — Let M be a topological compact surface. Let \mathbb{G} be a graph on M . Let $(U_v)_{v \in \mathbb{V}}$ be an orientation of the vertices of \mathbb{G} . For each $e \in \mathbb{E}$, set

$$\text{fr}(e) = \begin{cases} \{-1, 1\} & \text{if } e \text{ is not contained in } \partial M, \\ \{1\} & \text{if } e \subset \partial M \text{ and } e \text{ bounds } M \text{ positively,} \\ \{-1\} & \text{if } e \subset \partial M \text{ and } e \text{ bounds } M \text{ negatively.} \end{cases}$$

The framing of \mathbb{E} is the subset $\text{fr}(\mathbb{E})$ of $\mathbb{E} \times \{-1, 1\}$ defined by

$$\text{fr}(\mathbb{E}) = \bigcup_{e \in \mathbb{E}} \{e\} \times \text{fr}(e).$$

We have already mentioned that, without the assumption that M is orientable, some amount of global information is needed to determine the facial cycles.

Definition 1.3.19. — Let M be a compact topological surface. Let \mathbb{G} be a graph on M . Let $(U_v)_{v \in \mathbb{V}}$ be an orientation of the vertices of \mathbb{G} . The signature of this orientation is the collection of signs $(\lambda_e)_{e \in \mathbb{E}} \in \{-1, 1\}^{\mathbb{E}}$ defined as follows. For each edge e which is a simple loop, set $\lambda_e = 1$ if e admits an orientable neighbourhood and $\lambda_e = -1$ otherwise. Then, for each edge e such that $\underline{e} \neq \bar{e}$, consider an orientable neighbourhood U_e of e and set $\lambda_e = 1$ if there exists an orientation of U_e compatible with the orientations of $U_{\underline{e}}$ and $U_{\bar{e}}$, and $\lambda_e = -1$ otherwise.

If M is orientable, then it is possible to choose an orientation of the vertices of \mathbb{G} which is induced by an orientation of M . The signature of such an orientation is simply given by $\lambda_e = 1$ for all $e \in \mathbb{E}$.

We are now ready to define the permutation on $\text{fr}(\mathbb{E})$ which determines the facial cycles.

Definition 1.3.20. — Let M be a compact topological surface. Let \mathbb{G} be a graph on M . Let $(U_v)_{v \in \mathbb{V}}$ be an orientation of the vertices of \mathbb{G} . Let $\text{fr}(\mathbb{E})$ be the associated framing of \mathbb{G} . Let $(\lambda_e)_{e \in \mathbb{E}}$ be the signature of this orientation.

The collection of the cyclic orders of the outgoing vertices at each vertex relatively to the orientation specified by the collection $(U_v)_{v \in \mathbb{V}}$ is the set of cycles of a unique permutation of \mathbb{E} which we denote by σ . The involution $e \mapsto e^{-1}$ is another permutation of \mathbb{E} which we denote by α .

We define now three permutations $\bar{\alpha}$, $\bar{\sigma}$ and $\bar{\varphi}$ of $\text{fr}(\mathbb{E})$ as follows. Firstly, we set

$$\forall (e, \varepsilon) \in \text{fr}(\mathbb{E}), \quad \bar{\alpha}(e, \varepsilon) = (e^{-1}, -\lambda_e \varepsilon) \text{ and } \bar{\sigma}(e, \varepsilon) = (\sigma^\varepsilon(e), -\varepsilon).$$

Then, we define $\bar{\varphi}$ by the relation $\bar{\varphi} \circ \bar{\alpha} \circ \bar{\sigma} = \text{id}$. Hence,

$$\forall (e, \varepsilon) \in \text{fr}(\mathbb{E}), \quad \bar{\varphi}(e, \varepsilon) = (\sigma^{-\lambda_e \varepsilon}(e^{-1}), \lambda_e \varepsilon).$$

It is easy to check that $\bar{\sigma}$ and $\bar{\alpha}$, hence $\bar{\varphi}$ take indeed their values in $\text{fr}(\mathbb{E})$. Both permutations $\bar{\alpha}$ and $\bar{\sigma}$ are involutions. They correspond respectively to the first and second operations performed by our explorer after reaching a vertex. The permutation $\bar{\varphi}$ is the one whose cycles give the facial cycles of the graph.

Proposition 1.3.21. — Let M be a connected topological surface. Let \mathbb{G} be a graph on M . Let $(U_v)_{v \in \mathbb{V}}$ be an orientation of the vertices of \mathbb{G} . Let $\text{fr}(\mathbb{E})$ be the associated framing of \mathbb{G} . Let $\bar{\varphi}$ be the permutation of $\text{fr}(\mathbb{E})$ defined in Definition 1.3.20.

The range of the mapping which to each cycle $((e_1, \varepsilon_1) \cdots (e_n, \varepsilon_n))$ of the permutation $\bar{\varphi}$ associates the cycle $e_1 \cdots e_n$ in \mathbb{G} is exactly the set of the facial cycles of \mathbb{G} , taken once with each orientation.

If M is orientable and oriented, and if the orientation of the sets U_v is induced by the orientation of M , then $\bar{\varphi}$ leaves $\text{fr}(\mathbb{E}) \cap (\mathbb{E} \times \{1\})$ globally invariant. Moreover, the set of cycles of the restriction of $\bar{\varphi}$ to $\text{fr}(\mathbb{E}) \cap (\mathbb{E} \times \{1\})$ determines exactly the set of facial cycles which bound positively a face. These cycles are also those of the permutation $\varphi = \sigma^{-1} \circ \alpha^{-1}$ on \mathbb{E} .

The best proof of this result is probably the one which the reader will make for himself by drawing pictures. Another option is to read the section 3.3 of the book by B. Mohar and C. Thomassen [32], although their description of the permutations is slightly less formal than ours. This whole discussion is also a variation on the theme of *ribbon graphs* or *maps*, which are discussed extensively in [23].

Let us apply Proposition 1.3.21 to prove a result in the same vein as Corollary 1.3.12. If A and B are two subsets of a same set, we use the notation

$$A \dot{+} B = (A \cup B) \setminus (A \cap B).$$

Proposition 1.3.22. — *Let M be a connected compact topological surface. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on M . For each $n \geq 0$, let $\mathbb{G}_n = (\mathbb{V}_n, \mathbb{E}_n, \mathbb{F}_n)$ be a graph on M equipped with a bijection $S_n : \mathbb{V} \rightarrow \mathbb{V}_n$ and a bijection $S_n : \mathbb{E} \rightarrow \mathbb{E}_n$ such that for all $e \in \mathbb{E}$, $S_n(e)$ is the starting point of $S_n(e)$ and $S_n(e^{-1}) = S_n(e)^{-1}$. We assume that for all $n \geq 0$ and all edge e such that $e \subset \partial M$, $S_n(e) = e$. We assume also that for all $e \in \mathbb{E}$, the sequence $(S_n(e))_{n \geq 0}$ converges uniformly to e . Finally, we assume that for all $n \geq 0$ and for some orientation $(U_v)_{v \in \mathbb{V}}$ of the vertices of \mathbb{G} such that $S_n(v) \in U_v$ for all $n \geq 0$ and all $v \in \mathbb{V}$, the cyclic order of the outgoing edges at every vertex is preserved by the bijection S_n .*

Then for all $n \geq 0$, there exists a unique bijection $S_n : \mathbb{F} \rightarrow \mathbb{F}_n$ such that for all $F \in \mathbb{F}$, $\partial S_n(F) = S_n(\partial F)$. Moreover, for all $F \in \mathbb{F}$, one has

$$\limsup_{n \rightarrow \infty} (F \dot{+} S_n(F)) = \bigcap_{n \geq 0} \bigcup_{m \geq n} (F \dot{+} S_m(F)) \subset \text{Sk}(\mathbb{G}).$$

We use the following simple lemma.

Lemma 1.3.23. — *Let M be a connected compact surface endowed with a graph \mathbb{G} . If M is non-orientable, then two faces cannot have the same bounding unoriented cycle, and if M is oriented, then two faces cannot have the same oriented bounding cycle.*


More specifically, assume that there exist two faces of \mathbb{G} whose boundaries are equal as unoriented cycles. Then M is homeomorphic to a sphere and $\text{Sk}(\mathbb{G})$ is homeomorphic to a circle. In particular, after choosing an orientation of M , the boundaries of the two faces as oriented cycles are each other's inverse.

Proof. — Let F_1 and F_2 denote two faces of \mathbb{G} which share the same unoriented bounding cycle. Let c be a simple loop which represents this cycle, oriented in an arbitrary way. Consider, in a split pattern $(M', \mathbb{G}', \iota, f)$ of (M, \mathbb{G}) , the two disks M'_1 and M'_2 corresponding to F_1 and F_2 respectively. They are bounded by the same number of edges, which is also the combinatorial length of c . Let $e'_{1,1}, \dots, e'_{1,n}$ and $e'_{2,1}, \dots, e'_{2,n}$ denote respectively the set of edges located on the boundary of M'_1 and M'_2 , in such a way that $\partial M'_1 = e'_{1,1} \cdots e'_{1,n}$ and $\partial M'_2 = e'_{2,1} \cdots e'_{2,n}$. Each edge on the boundary of M'_1 is sent by f to an edge of \mathbb{G} which is also adjacent to F_2 , hence is identified by ι with an edge bounding M'_2 . We may assume that $\iota(e'_{1,1}) = e'_{2,1}$. We may also assume that $f(e'_{1,1}) = f(e'_{2,1})$ is the first edge traversed by c and this characterises fully c among all representatives of the unoriented cycle ∂F_1 . Indeed, c traverses each unoriented edge of \mathbb{G} at most once, for ι does never identify two distinct edges of the boundary of M'_1 , or M'_2 . Hence, $c = f(e'_{1,1}, \dots, e'_{1,n}) = f(e'_{2,1}, \dots, e'_{2,n})$, so that $\iota(e'_{1,i}) = e'_{2,i}$ for all $i \in \{1, \dots, n\}$. The result follows. \square

Proof of Proposition 1.3.22. — The first step of this proof is a rather unpleasant construction whose aim is to provide us with a simple way of telling apart the faces of \mathbb{G} . We achieve this by adding an edge with one loose endpoint in each face, as follows.

In the case where M is a disk and $\text{Sk}(\mathbb{G}) \subset \partial M$, $\mathbb{G}_n = \mathbb{G}$ for all n and the result is true. In any other case, each face of \mathbb{G} is bounded by at least one edge which is not contained in ∂M . For each face F of \mathbb{G} , let us choose a point $m_F \in F$ and a point v_F in the interior of an edge e_F adjacent to F and not contained in ∂M . Let us choose a continuous edge \tilde{f}_F which crosses $\text{Sk}(\mathbb{G})$ exactly once at v_F , and finishes at m_F . We assume that v_F is not the initial point of \tilde{f}_F . We denote by f_F the portion of \tilde{f}_F which joins v_F to m_F . For $n \geq 0$ large enough, \tilde{f}_F meets $\text{Sk}(\mathbb{G}_n)$, more precisely the edge $S_n(e_F)$ and only this edge. For such n , let $v_{F,n}$ be the last exit point of \tilde{f}_F from $\text{Sk}(\mathbb{G}_n)$. It is an interior point of $S_n(e_F)$. Let $S_n(f_F)$ be the portion of \tilde{f}_F which joins $v_{F,n}$ to m_F .

Let us perform this construction for each face F , with the edges \tilde{f}_F chosen to be pairwise disjoint. Let us define \mathbb{G}' as the graph obtained from \mathbb{G} by subdividing the edges e_F at v_F and adding the edges f_F . Also, for all $n \geq 0$, let \mathbb{G}'_n be the graph obtained from \mathbb{G}_n by subdividing the edges $S_n(e_F)$ at $v_{F,n}$ and adding the edges $S_n(f_F)$. We extend $S_n : \mathbb{V} \rightarrow \mathbb{V}_n$ to $S'_n : \mathbb{V}' \rightarrow \mathbb{V}'_n$ by setting $S_n(v_F) = v_{F,n}$, and extend also $S_n : \mathbb{E} \rightarrow \mathbb{E}_n$ to $S'_n : \mathbb{E}' \rightarrow \mathbb{E}'_n$ in the obvious way. It is not difficult to check that $v_{F,n}$ tends to v_F as n tends to infinity, and hence that $S'_n(f_F)$ converges uniformly to f_F . By considering a small ball around v_F , one checks also that the bijections S'_n still preserve the cyclic order at each vertex, including v_F and m_F .

FIGURE 10. The construction of the graphs \mathbb{G}' and \mathbb{G}'_n .

The faces of \mathbb{G} and \mathbb{G}' (respectively \mathbb{G}_n and \mathbb{G}'_n) are obviously in bijective correspondence which we denote simply with a prime. For instance, for all $F \in \mathbb{F}$, the cycle which bounds F' is deduced from ∂F by the insertion of a sequence $f_F f_F^{-1}$. In particular, F and F' have the same topological closure. The boundary of F' is the only facial cycle of \mathbb{G}' which involves the edge f_F and it involves no other edge of \mathbb{G}' which is not contained in the skeleton of \mathbb{G} . Hence, any facial cycle of \mathbb{G}' which involves f_F must be the boundary of F' .

The advantage of this tedious construction is that for all $F \in \mathbb{F}$, F' is the only face of \mathbb{F}' whose closure contains the edge f_F and in particular the vertex m_F .

Consider now $n \geq 1$ and a face F of \mathbb{G} . Since S_n preserves the cyclic order of the edges at each vertex, it follows from Proposition 1.3.21 that the cycle $S_n(\partial F)$ is a facial cycle of \mathbb{G}_n . By Lemma 1.3.23, two faces of \mathbb{G}_n cannot have the same boundary. Hence, $S_n(\partial F)$ is the boundary of a unique face of \mathbb{G}_n which we denote by $S_n(F)$. By construction, the equality $\partial S_n(F) = S_n(\partial F)$ holds for all $F \in \mathbb{F}$.

The same construction provides us with a bijection S'_n between \mathbb{F}' and \mathbb{F}'_n . Now consider $F \in \mathbb{F}$. By definition, $S'_n(F')$ is the face of \mathbb{G}'_n whose boundary is the cycle $S_n(\partial F)$ in which $S'_n(f_F)S'_n(f_F)^{-1}$ has been inserted at the occurrence of the vertex $v_{F,n}$. Hence, m_F belongs to the closure of $S'_n(F')$ and, by the discussion a few lines above, $S'_n(F') = S_n(F)'$.

Now this implies that m_F belongs to the closure of $S_n(F)'$, which is equal to the closure of $S_n(F)$. Since m_F does not lie on the skeleton of \mathbb{G}_n , this implies that $m_F \in S_n(F)$.

For each face F of \mathbb{G} , let us choose a connected open subset U_F of F which contains m_F and such that the closure of U_F is contained in F . The last assumption implies that for n large enough, $\text{Sk}(\mathbb{G}_n)$ is disjoint from U_F . Hence, U_F is contained in a unique face of \mathbb{G}_n , which must be $S_n(F)$.

In particular, $F \dot{+} S_n(F)$ is contained in $(F \cup S_n(F)) \setminus U_F$. Since F nor $S_n(F)$ meet another subset of the form U_{F_1} for some $F_1 \in \mathbb{F}$, we deduce from this inclusion that, for n large enough, $F \dot{+} S_n(F)$ is contained in $M \setminus \bigcup_{F \in \mathbb{F}} U_F$. Hence, $\limsup(F \dot{+} S_n(F))$ is contained in the same set. This inclusion holds for any choice of the sets U_F and the result follows. \square

1.3.3. Adjunction of edges. — Proposition 1.3.15 gives us a way of removing some edges from a graph. We will also need a way of adding edges to a graph. The typical problem is the following: we are given a compact surface M endowed with a graph \mathbb{G} , we consider a face F of this graph, two vertices v_1 and v_2 on the boundary of F and we would like to join them by a new edge whose interior is contained in F .

If we are working in the category of topological surfaces and graphs with continuous edges, then Proposition 1.3.8 suffices to guarantee the existence of a continuous edge with the desired properties. However, if we are working with a graph with rectifiable edges and insist that the new edge be rectifiable too, then we need something more. The problem is a purely local one and we loose nothing by formulating it in the plane. The difficulty is that it seems not to be known whether a rectifiable edge can be straightened, even locally, by a bi-Lipschitz continuous homeomorphism of the plane (see for instance [28]).

In this section and in this section only, we use the symbol ∂ to denote the topological boundary of a set. We denote by \mathcal{H}^1 the 1-dimensional Hausdorff measure on \mathbb{R}^2 .

Proposition 1.3.24. — *Let K be a compact subset of \mathbb{R}^2 . Assume that ∂K is connected and satisfies $\mathcal{H}^1(\partial K) < +\infty$. Let v be a point of ∂K . Let m be a point of $\mathbb{R}^2 \setminus K$. Assume that v is curve-accessible from m , that is, that there exists a continuous curve $c : [0, 1] \rightarrow \mathbb{R}^2$ such that $c(0) = m$, $c(1) = v$ and $c([0, 1]) \cap K = \emptyset$. Then there exists an injective Lipschitz-continuous curve with the same properties as c .*

We start by proving an intermediary result, whose content is that two points of the bounded connected component of the complement of a Jordan curve with finite length can be joined inside this component by a path with finite length controlled by the length of the Jordan curve.

Proposition 1.3.25. — *Let U be a non-empty bounded connected open subset of \mathbb{R}^2 with connected boundary. Assume that $\mathcal{H}^1(\partial U) < +\infty$. For all $a, b \in U$, there exists a rectifiable path c which joins a to b and such that $\ell(c) = \mathcal{H}^1(c) \leq 100\mathcal{H}^1(\partial U)$.*

Proof. — Since U is connected and open, it is arcwise connected. Let γ be a continuous curve which joins a to b inside U . Set

$$\varepsilon = \frac{1}{4} \min(d(\gamma([0, 1]), \partial U), \text{diam}(\partial U)) > 0.$$

Let X be a maximal subset of ∂U such that for all $x, y \in X$ with $x \neq y$, $d(x, y) \geq 2\varepsilon$. The assumptions on U imply that ∂U is compact, so that X is finite. Write $X = \{x_1, \dots, x_n\}$.

Let us say that a finite set of circles are in *generic position* if no two distinct of them are tangent and no three pairwise distinct of them have a common point. We claim that it is possible to choose n positive real number r_1, \dots, r_n in the interval $(2\varepsilon, 4\varepsilon)$ such that the boundaries of the balls $B_i = B(x_i, r_i)$ are in generic position.

Indeed, let us choose $r_1 = 3\varepsilon$. There are only a finite number of values of r_2 (actually, two values) for which ∂B_2 is tangent to ∂B_1 . Thus, we can choose $r_2 \in (2\varepsilon, 4\varepsilon)$ such that ∂B_1 and ∂B_2 are not tangent. Assume that r_1, \dots, r_k have been chosen such that $\partial B_1, \dots, \partial B_k$ are in generic position. Then there are only a finite number of values of r_{k+1} for which $\partial B_1, \dots, \partial B_{k+1}$ would not be in generic position. Hence, we can choose r_{k+1} in $(2\varepsilon, 4\varepsilon)$ such that this does not happen.

Set $K = \bigcup_{i=1}^n \bar{B}(x_i, r_i)$. This is a compact set which contains ∂U , by maximality of X . Moreover, each connected component of K meets ∂U , which is connected by assumption. Hence, K itself is connected. Finally, K does not meet γ , so that a and b are in the same connected component of $\mathbb{R}^2 \setminus K$.

Consider $x \in \partial K$. Then x is on the boundary of one or two of the balls B_1, \dots, B_n . In any case, x admits a neighbourhood in which ∂K is a simple curve composed of one or two arcs of circle. It follows that the boundary of K is a compact set each point of which admits a neighbourhood homeomorphic to \mathbb{R} . Hence, it is homeomorphic to a finite union of pairwise disjoint circles.




FIGURE 11. The compact set K is the union of all the small disks. In this case, its boundary is the union of three Jordan curves. The thicker path is the one of which we control the length by $100\mathcal{H}^1(\partial U)$.

We have proved that K is bounded by a finite collection of pairwise disjoint Jordan curves. Let us call *interior* of a Jordan curve the bounded connected component of its complement. Since K is bounded and connected, ∂K is necessarily the union of a

Jordan curve J_0 such that K is contained in the interior of J_0 , and a certain number of Jordan curves J_1, \dots, J_k whose interiors are disjoint and contained in the interior of J_0 . Since a and b belong to the same bounded connected component of $\mathbb{R}^2 \setminus K$, they both lie in the interior of one of the curves J_1, \dots, J_k , say J_1 .

Let us give a bound on the length of J_1 by bounding the total length of ∂K . Since $\varepsilon < \text{diam}(\partial U)$, none of the balls $B(x_i, \varepsilon)$ contains ∂U . Let us choose $i \in \{1, \dots, n\}$ and set, for $y \in \mathbb{R}^2$, $\pi_i(y) = d(x_i, y)$. Then $\pi_i(\partial U)$ is connected, contains 0 and also contains a real larger than ε . Hence, $\pi_i(\partial U) \supset (0, \varepsilon)$ and in fact $\pi_i(\partial U \cap B(x_i, \varepsilon))$ contains $(0, \varepsilon)$. Since π_i is 1-Lipschitz continuous, it follows that $\mathcal{H}^1(\partial U \cap B(x_i, \varepsilon)) \geq \varepsilon$. Using the fact that the n balls $B(x_i, \varepsilon)$, $i \in \{1, \dots, n\}$ are pairwise distinct, we deduce that $n\varepsilon \leq \mathcal{H}^1(\partial U)$, which combined with inequality $\mathcal{H}^1(\partial K) \leq 8\pi n\varepsilon$ yields $\mathcal{H}^1(\partial K) \leq 8\pi \mathcal{H}^1(\partial U)$.

Let s be the straight path from a to b . Its length is smaller than $\text{diam}(U) \leq \text{diam}(\partial U) \leq \mathcal{H}^1(\partial U)$. If s does not meet K , it does not meet ∂U either and the result is proved. Otherwise, let a' and b' be respectively the first and the last point at which s meets ∂K . Let c be the path obtained by concatenating the straight path from a to a' , then an arc of J_1 from a' to b' and finally the straight path from b' to b . The length of c is bounded by $\ell(c) \leq \ell(s) + \ell(J_1) \leq 100\mathcal{H}^1(\partial U)$ as expected. \square

Proof of Proposition 1.3.24. — Since \mathcal{H}^1 is a σ -additive measure and $\mathcal{H}^1(\partial K)$ is finite, $\mathcal{H}^1(\partial K \cap B(v, r))$ tends to $\mathcal{H}^1(\{v\}) = 0$ as r tends to 0. Let $(r_n)_{n \geq 0}$ be a decreasing sequence of positive reals such that

$$(19) \quad \sum_{n \geq 0} (\mathcal{H}^1(\partial K \cap B(v, 2r_n)) + 4\pi r_n) < +\infty.$$

By shifting the sequence (r_n) if necessary, we may assume that $B(v, 2r_0)$ does not contain K . Since ∂K is connected, this implies that for all $n \geq 0$, $K \cap \partial B(v, r_n) \neq \emptyset$.

Let c be a continuous curve which joins m to v and meets K only at v . For each $n \geq 0$, let m_n be the last point of the curve c which is on $\partial B(v, r_n)$. Choose $n \geq 0$. Consider the compact set $K_n = (K \cap B(v, 2r_n)) \cup \partial B(v, 2r_n)$. It is connected, as the image of the connected set $K \cup \partial B(v, 2r_n)$ by the projection on the closed convex set $B(v, 2r_n)$. The points m_n and m_{n+1} belong to the same connected component of the complement of K_n (see Figure 12). Let U_n denote this connected component.

As a bounded connected component of the complement of a connected compact subset of \mathbb{R}^2 , U_n has a connected boundary (see [38, p. 47], where this property is called the Brouwer property of the sphere). Moreover, ∂U_n is contained in $\partial K_n = (\partial K \cap B(v, 2r_n)) \cup \partial B(v, 2r_n)$. Hence, $\mathcal{H}^1(\partial U_n) \leq \mathcal{H}^1(\partial K \cap B(v, 2r_n)) + 4\pi r_n$. So, by Proposition 1.3.25, there exists a path c_{n+1} which joins m_n to m_{n+1} and has length smaller than $100\mathcal{H}^1(\partial U_n)$.

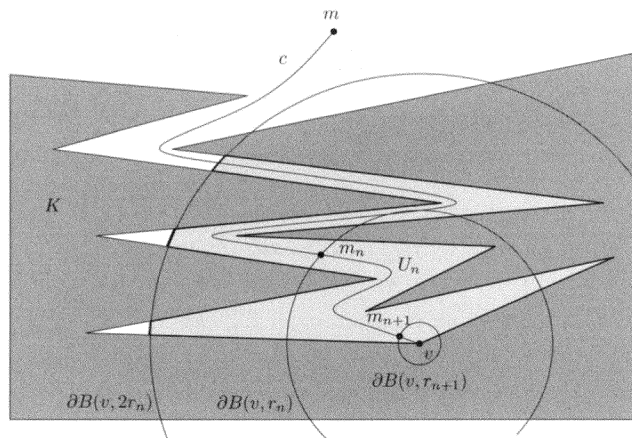


FIGURE 12. The path c_{n+1} will join m_n to m_{n+1} inside U_n .

Finally, let c_0 be a path of finite length which joins m to m_0 outside K . Set $C = \{v\} \cup \bigcup_{n \geq 0} c_n([0, 1])$. By (19), it is the Lipschitz-continuous image of a segment and it joins m to v . The set C needs not be the image of an injective path but since it is compact and connected, by Proposition 14 in [9], there exists an injective Lipschitz-continuous path which joins m to v in C . Such a path is exactly what we were looking for. \square

We will make use of this result under the following form.

Proposition 1.3.26. — *Let M be a compact surface. Let \mathbb{G} be a graph on M . Let F be a face of \mathbb{G} . Let v_1 and v_2 be two vertices which lie on the boundary of F . There exists an edge e such that $\underline{e} = v_1$, $\bar{e} = v_2$ and $e((0, 1)) \subset F$. In particular, $\mathbb{E} \cup \{e, e^{-1}\}$ is the set of edges of a graph on M .*

In this statement, the assumption that the vertices lie on the boundary of the face F can be understood in the topological sense or, as it is equivalent, in the sense that they are traversed by the facial cycle associated to F .

Proof. — Let us endow M with a Riemannian metric. By the same result used at the end of the previous proof ([9], Proposition 14), it suffices to prove that there exists a compact connected subset C of M with finite 1-dimensional Hausdorff measure which contains v_1 and v_2 and such that $C \setminus \{v_1, v_2\} \subset F$. Since F is arcwise connected by paths of finite length (for instance piecewise geodesic paths), it suffices to prove that v_1 , hence v_2 , can be joined to at least one point of F by a curve of finite length which has only its starting point outside F .

For this, choose a point n in F and choose a continuous curve c which joins n to v and has only its finishing point outside F . That such a curve exists is obvious in a split pattern of \mathbb{G} . Choose also $r > 0$ such that the metric ball $B(v, r)$ is diffeomorphic to a disk and each edge starting from v_1 crosses the circle $\partial B(v_1, r)$. Choose a point m on c which c traverses after its last entry time in the ball $B(v_1, r)$. By applying Proposition 1.3.24 to m and $(\text{Sk}(\mathbb{G}) \cap B(v_1, r)) \cup \partial B(v_1, r)$ inside the ball $\bar{B}(v_1, r)$ smoothly identified with a ball in \mathbb{R}^2 , we find the desired curve with finite length. \square

1.3.4. The group of loops in a graph. — The concatenation of paths is not a group operation, even when it is restricted to a set of loops based at the same point, in which case all pairs of loops can be concatenated. The obstruction is the fact that if c is a non-constant path, then there is no path c' such that cc' is constant. However, the path cc^{-1} is equivalent to the constant path for a natural equivalence relation.

Definition 1.3.27. — Let M be a compact topological compact surface. Two curves $c, c' \in \mathcal{C}(M)$ are elementarily equivalent if there exist three curves a, b, d such that $\{c, c'\} = \{ab, add^{-1}b\}$. We say that c and c' are equivalent and write $c \simeq c'$ if there exists a finite chain $c = c_0, \dots, c_n = c'$ of curves such that c_i is elementarily equivalent to c_{i+1} for each $i \in \{0, \dots, n - 1\}$.

This relation is an equivalence relation on $\mathcal{C}(M)$ similar to the equality of words in a free group, with the important difference that for the relation \simeq , there is no analogue of the reduced form of a word, even if we restrict ourselves to rectifiable paths. For example, the class of the rectifiable infinite polygonal path in the complex plane which joins the points $0, e^{i\pi}, 0, 2^{-1}e^{2^{-1}i\pi}, 0, \dots, 0, 2^{-n}e^{2^{-n}i\pi}, 0, \dots$ contains no path of minimal length.



FIGURE 13. The equivalence class of this rectifiable path contains no shortest element.

B. Hambly and T. Lyons have defined in [21] an equivalence relation on rectifiable paths for which the path described above is equivalent to the constant path equal to 0. This relation is strictly less fine than \simeq and each class contains a unique element of minimal length. We plan to investigate in a future work the importance of this equivalence relation in the framework of the present theory.

On a graph however, these subtleties do not arise. Until the end of this section, we work on smooth surfaces and consider paths instead of curves but all our results apply to graphs on topological surfaces.

Lemma 1.3.28. — *Let M be a compact surface endowed with a graph \mathbb{G} . Let c and c' be elements of $\mathcal{P}(\mathbb{G})$. Assume that c and c' are elementarily equivalent. Then there exist a, b, d in $\mathcal{P}(\mathbb{G})$ such that $\{c, c'\} = \{ab, add^{-1}b\}$.*

Proof. — Let a, b, d be given by the definition of the fact that c and c' are elementarily equivalent. Let us assume that $c = ab$ and $c' = add^{-1}b$. Since a, b, d are pieces of a path in \mathbb{G} , it suffices to show that their endpoints are vertices of \mathbb{G} to show that they are themselves paths in \mathbb{G} . This is clear for $\underline{a} = \underline{c}$ and $\bar{b} = \bar{c}$.

Let us say that the curve c *backtracks* at a point $m \in M$ if there exists $t, \varepsilon > 0$ such that $(t - \varepsilon, t + \varepsilon) \subset [0, 1]$ and a parametrization of c such that $c(t) = m$ and for all $h \in [0, \varepsilon)$, $c(t + h) = c(t - h)$. A point at which a path in \mathbb{G} backtracks must be a vertex, hence \bar{d} , which is a point at which c' backtracks, is a vertex. There remains to prove that $m = \bar{a} = \underline{d} = \underline{b}$ is a vertex.

Let \mathbb{G}' be the graph obtained by adding m to the set of vertices of \mathbb{G} and subdividing the edges of \mathbb{G} accordingly. The graph \mathbb{G}' has the same skeleton as \mathbb{G} . We claim that either m is a backtracking point for c or c' , or \mathbb{G}' has at least three distinct outgoing edges at m . In both cases, it follows that m is a vertex of \mathbb{G} .

Let e_a^{-1} denote the last edge of a as a path in \mathbb{G}' , and e_d and e_b the first edges of d and b as paths in \mathbb{G}' . Let us assume that c does not backtrack at m . Then $e_a \neq e_b$. Let us assume that c' does not backtrack at m either. Then ad and $d^{-1}b$ do not backtrack at m , so that $e_a \neq e_d$ and $e_d \neq e_b$. The claim is proved. \square

If c is a path in a graph, we call *combinatorial length* of c and denote by $\text{length}(c)$ the number of edges which constitute c . A path is said to be *reduced* if it contains no sequence ee^{-1} for some $e \in \mathbb{E}$. Equivalently, a path is reduced if it is not elementarily equivalent to a combinatorially shorter path.

Corollary 1.3.29. — *Let M be a compact surface endowed with a graph \mathbb{G} . Every class of equivalence of $\mathcal{P}(\mathbb{G})$ contains a unique element of shortest combinatorial length, which is characterised by the fact that it is reduced.*

Proof. — By Lemma 1.3.28, two paths in a graph which are equivalent differ by a finite number of insertions or erasures of sequences ee^{-1} , where e is an edge of \mathbb{G} .

Let us consider an equivalence class of paths. This class contains paths of minimal combinatorial length. These paths are necessarily reduced. Thus, it suffices to prove that the given class contains only one reduced path. Assume that there are two distinct reduced paths, say c and c' . Consider a chain of paths $c = c_0, \dots, c_n = c'$ obtained by successive erasures and insertions of sequences ee^{-1} where $e \in \mathbb{E}$. Assume that this chain minimises $\max\{\text{length}(c_0), \dots, \text{length}(c_n)\}$ among all chains from c to c' and that, among those minimisers, it also minimises the number of intermediate paths of maximal length. Consider an integer k such that c_k has maximal length

among c_0, \dots, c_n . Since c_0 and c_n are reduced, c_1 is deduced from c_0 by an insertion and c_n from c_{n-1} by an erasure. Thus, $k \in \{1, \dots, n-1\}$. So, c_k is deduced from c_{k-1} by an insertion of, say, ee^{-1} and c_{k+1} from c_k by an erasure of ff^{-1} . Let us assume that $e \notin \{f, f^{-1}\}$. Then the sequence ff^{-1} is already present in c_{k-1} and could have been removed before the insertion of ee^{-1} , thus diminishing the number of intermediate paths of maximal length. By assumption, this is impossible. Hence, $e = f$. Moreover, for the same reason, the sequence ee^{-1} removed between c_k and c_{k+1} is not present in c_{k-1} . It cannot be the sequence ee^{-1} inserted between c_{k-1} and c_k , for then $c_{k-1} = c_{k+1}$ and by removing c_k and c_{k+1} from the chain, we would again diminish the number of intermediate paths of maximal length. Hence, exactly one of the letters inserted in c_k is removed between c_k and c_{k+1} . There are two cases and in both, it appears that $c_{k-1} = c_{k+1}$. This is again impossible. Finally, there is exactly one reduced path in each equivalence class of paths. \square

We can now define the group of reduced loops.

Definition 1.3.30. — *Let M be a smooth surface endowed with a graph \mathbb{G} . Let v be a vertex of \mathbb{G} . We denote by $\text{RL}(\mathbb{G})$ (resp. $\text{RL}_v(\mathbb{G})$) the subset of $\text{L}(\mathbb{G})$ formed by reduced loops (resp. reduced loops based at v).*

The set $\text{RL}_v(\mathbb{G})$ is a group for the operation of concatenation-reduction, which to two loops l_1 and l_2 associates the unique reduced loop equivalent to l_1l_2 .

It is a well-known fact that the group $\text{RL}_v(\mathbb{G})$ is free. At a later stage, we will devote some time to finding families of generators of this group which satisfy special properties. For the moment, let us simply recall why it is a free group, by using spanning trees.

Definition 1.3.31. — *Let M be a compact surface endowed with a graph \mathbb{G} .*

A spanning tree of \mathbb{G} is a subset $T \subset \mathbb{E}$ such that $T = T^{-1}$ and such that by concatenating edges of T , one may construct a path from any vertex to any other but no simple cycle.

For all vertices v_1, v_2 of \mathbb{G} , we denote by $[v_1, v_2]_T$ the unique injective path in T which joins v_1 to v_2 . Let v be a vertex of \mathbb{G} . For each edge e of \mathbb{E} , we define the loop

$$l_{e,T} = [v, \underline{e}]_T e [\bar{e}, v]_T.$$

If $e \in T$, then $l_{e,T}$ is equivalent to the constant loop. Otherwise, it is reduced and in fact a lasso (see Definition 1.2.4). Indeed, the paths $[v, \underline{e}]_T$ and $[v, \bar{e}]_T$ can be written in a unique way as $[v, w]_T [w, \underline{e}]_T$ and $[v, w]_T [w, \bar{e}]_T$ with $[w, \underline{e}]_T \cap [w, \bar{e}]_T = \{w\}$. Then l_e is the lasso with spoke $[v, w]_T$ and meander $[w, \underline{e}]_T e [\bar{e}, w]_T$.

Definition 1.3.32. — *Let M be a compact surface endowed with a graph \mathbb{G} . An orientation of \mathbb{G} is a subset \mathbb{E}^+ of \mathbb{E} such that for all $e \in \mathbb{E}$, exactly one of the two edges*




FIGURE 14. A spanning tree of our favourite graph. On the right, the position of the vertex w at which the injective paths $[\underline{e}, v]_T$ and $[\bar{e}, v]_T$ first meet, for a certain choice of v and e .

e and e^{-1} belongs to \mathbb{E}^+ . If M is oriented, e is an edge which lies on the boundary of M and which bounds M positively, we insist that $e \in \mathbb{E}^+$.

Given an orientation \mathbb{E}^+ of \mathbb{G} and a subset $Q \subset \mathbb{E}$, we use the notation

$$Q^+ = Q \cap \mathbb{E}^+.$$

Given a graph \mathbb{G} , we set

$$v(\mathbb{G}) = \#\mathbb{V}, \quad e(\mathbb{G}) = \frac{1}{2}\#\mathbb{E} \quad \text{and} \quad f(\mathbb{G}) = \#\mathbb{F}.$$

The following lemma is classical.

Lemma 1.3.33. — *Let M be a compact surface endowed with a graph \mathbb{G} . Let v be a vertex of \mathbb{G} . Let $T \subset \mathbb{E}$ be a spanning tree of \mathbb{G} . Let \mathbb{E}^+ be an orientation of \mathbb{G} . The group $\text{RL}_v(\mathbb{G})$ is freely generated by the loops $\{l_{e,T} : e \in (\mathbb{E} \setminus T)^+\}$. In particular, it is free of rank $e(\mathbb{G}) - v(\mathbb{G}) + 1$. Moreover, the natural mapping $\text{L}_v(\mathbb{G}) \rightarrow \pi_1(\text{Sk}(\mathbb{G}), v)$ descends to a group isomorphism $i : \text{RL}_v(\mathbb{G}) \xrightarrow{\sim} \pi_1(\text{Sk}(\mathbb{G}), v)$.*

Proof. — If $l = e_1 \cdots e_n$ belongs to $\text{L}_v(\mathbb{G})$, then $l \simeq l_{e_1,T} \cdots l_{e_n,T}$. Hence, the loops $l_{e,T}$, $e \in \mathbb{E} \setminus T$ generate $\text{RL}_v(\mathbb{G})$. Since $l_{e^{-1},T} = l_{e,T}^{-1}$ for all $e \in \mathbb{E}$, this implies that the loops $l_{e,T}$, $e \in (\mathbb{E} \setminus T)^+$ generate $\text{RL}_v(\mathbb{G})$. Now let X be a group. Let $x = \{x_e : e \in \mathbb{E} \setminus T\}$ be a collection of elements of X such that $x_{e^{-1}} = x_e^{-1}$ for all $e \in \mathbb{E} \setminus T$. Complete the collection x by setting $x_e = 1$ for all $e \in T$. The mapping from $\text{L}_v(\mathbb{G})$ to X which sends the loop $l = e_1 \cdots e_n$ to $x_1 \cdots x_n$ descends to a group homomorphism from $\text{RL}_v(\mathbb{G})$ to X which sends $l_{e,T}$ to x_e for all $e \in \mathbb{E}$. Thus, $\text{RL}_v(\mathbb{G})$ satisfies the universal property which characterises freeness. Finally, since T has $v(\mathbb{G})$ vertices, it has $v(\mathbb{G}) - 1$ unoriented edges. Hence, $\text{RL}_v(\mathbb{G})$ is free of rank $e(\mathbb{G}) - v(\mathbb{G}) + 1$.

It is obvious that two equivalent loops are homotopic in $\text{Sk}(\mathbb{G})$. Hence, the morphism i is well defined. Let us use the letter T to denote the subset $\bigcup_{e \in T} e$ of $\text{Sk}(\mathbb{G})$. This subset is contractible and it is easy to check that $\text{Sk}(\mathbb{G})$ has the same homotopy type as $\text{Sk}(\mathbb{G})/T$, the topological space obtained from $\text{Sk}(\mathbb{G})$ by identifying all

the points of T . This topological space is a bunch of circles, one for each element of $(\mathbb{E} \setminus T)^+$. Moreover, each loop $l_{e,T}$, composed with the continuous projection $\text{Sk}(\mathbb{G}) \rightarrow \text{Sk}(\mathbb{G})/T$, becomes a loop which goes once around the circle corresponding to e . Thus, the composition of i with the isomorphism $\pi_1(\text{Sk}(\mathbb{G}), v) \rightarrow \pi_1(\text{Sk}(\mathbb{G})/T, v)$ is an isomorphism, and i is also an isomorphism. \square

1.3.5. Graphs with one face. — By Proposition 1.3.8, a graph with a single face on a connected surface determines a way of realizing this surface as the quotient of a disk by a suitable identification of its boundary. On the other hand, many non-isomorphic patterns with a single face give rise, when they are completely sewed, to homeomorphic surfaces. In this section, we discuss this fact in relation with the classical proof of the theorem of classification of surfaces (see Theorem 1.1.3) by cut-and-paste.

It is convenient to represent a pattern with one face by a word in a free group. This is what we explain now. For each integer $n \geq 1$, let us call n -gon the split pattern (D, \mathbb{G}_n) formed by a closed disk D and a graph \mathbb{G}_n with n unoriented edges on the boundary of D . This split pattern is unique up to homeomorphism.

Definition 1.3.34. — Consider a set X and let $\langle X \rangle$ denote the free group over X . Let w be an element of $\langle X \rangle$. Write w as a reduced word $x_1 \cdots x_n$ with x_1, \dots, x_n in $X \cup X^{-1}$. We say that w is admissible if w is cyclically reduced, that is, if $x_n \neq x_1^{-1}$, and each letter of X appears at most twice in w , that is, for each $x \in X$,

$$\#\{i \in \{1, \dots, n\} : x_i \in \{x, x^{-1}\}\} \leq 2.$$

We say that an admissible word w is closed if no letter appears exactly once in w .

The fact that a word is admissible is not changed if this word is submitted to a circular permutation of its letters nor if it is replaced by its inverse. Of course, it is not changed either by changing the names of the letters: the set X plays no special role and we identify two words which differ only by relabelling the letters which constitute them. We define now a correspondence between admissible words and graphs with one face.

Definition 1.3.35. — 1) Let M be a compact surface endowed with a graph \mathbb{G} . Assume that \mathbb{G} has a single face and that each vertex of \mathbb{G} is the starting point of at least two distinct edges. Let \mathbb{E}^+ be an orientation of \mathbb{G} . Then each cycle which represents the boundary of the unique face of \mathbb{G} is a cyclically reduced word in the letters of \mathbb{E}^+ , that is, an element of $\langle \mathbb{E}^+ \rangle$. We define

$$W(M, \mathbb{G}) = \{w \in \langle \mathbb{E}^+ \rangle : w \text{ is a facial cycle of } \mathbb{G}\}.$$

2) Let $w = x_1 \cdots x_n$ be an admissible word of length n . Let (D, \mathbb{G}_n) be an n -gon. Write $\mathbb{E}_n = \{e_1^{\pm 1}, \dots, e_n^{\pm 1}\}$ in such a way that $e_1 \cdots e_n$ represents ∂D . Let ι_w be the involution of \mathbb{E}_n defined as follows: for all $i \in \{1, \dots, n\}$,

$$\iota_w(e_i) = \begin{cases} e_j^\varepsilon & \text{if there exists } j \neq i \text{ and } \varepsilon = \pm 1 \text{ such that } x_i = x_j^\varepsilon, \\ e_i & \text{otherwise.} \end{cases}$$

The closed compact surface obtained by completely sewing the pattern $(D, \mathbb{G}_n, \iota_w)$ is said to be associated with w and we denote it by $\Sigma(w)$.

Let M be a compact surface. It follows from the definitions that for all graph \mathbb{G} with a single face on M and for all $w \in W(M, \mathbb{G})$, the surface $\Sigma(w)$ is homeomorphic to M . On the other hand, there are in general many admissible words which are not in $W(M, \mathbb{G})$ whose associated surface is homeomorphic to M . Ignoring the precise set to which the letters of the words that we consider belong, we define a set of words as follows:

$$W(M) = \{w \text{ admissible word} : M(w) \text{ is homeomorphic to } M\}.$$

Each word of $W(M)$ belongs to $W(M, \mathbb{G})$ for some graph \mathbb{G} , for instance the graph constructed by sewing the pattern associated to this word. We use this discussion to prove the following theorem.

Proposition 1.3.36. — *Let M be a compact surface. Let \mathbb{G}_1 and \mathbb{G}_2 be two graphs on M . There exists a homeomorphism $f : M \rightarrow M$ which preserves each connected component of ∂M , and which is orientation-preserving if M is oriented, and a finite sequence of graphs $\mathbb{G}_{1,0}, \dots, \mathbb{G}_{1,r}$ such that $\mathbb{G}_{1,0} = \mathbb{G}_1$ and $\mathbb{G}_{1,r} = f(\mathbb{G}_2)$, and such that for all $i \in \{0, \dots, r-1\}$, $\mathbb{G}_{1,i+1}$ is deduced from $\mathbb{G}_{1,i}$ by erasure of an edge in the sense of Proposition 1.3.15 or by adjunction of an edge in the sense of Proposition 1.3.26.*

Proof. — By erasing enough edges of \mathbb{G}_1 and \mathbb{G}_2 , we can transform them into two graphs with a single face and of which every vertex is the initial point of at least two distinct edges. Such graphs determine two words of $W(M)$ which we denote by w_1 and w_2 .

The theorem of classification of surfaces as it is proved in [30] asserts that, by repeated operations of cutting and pasting, w_1 and w_2 can be put under one of the standard words $[a_1, a_2] \cdots [a_{g-1}, a_g]$ or $a_1^2 \cdots a_g^2$ if M is closed of genus g , or the same words multiplied by a word of the form $d_1 c_1 d_1^{-1} \cdots d_p c_p d_p^{-1}$ if M has a boundary, with c_1, \dots, c_p corresponding to the p boundary components of M .

The general operation of cutting and pasting, described at the level of a split pattern (M', \mathbb{G}') of \mathbb{G} , consists in choosing two vertices v_1 and v_2 on the boundary of M' and a pair of edges $(e, \iota(e))$ which are identified by ι and separated by v_1 and v_2 . One then adds to \mathbb{G}' an edge inside M' which joins v_1 to v_2 , identifies e and $\iota(e)$ and

removes the joint of this identification (see Figure 1.3.5). Seen on M , these operations can be described simply as follows: add an edge to \mathbb{G}_1 joining v_1 to v_2 , thanks to Proposition 1.3.26, thus creating two faces, and remove the edge e , thus retrieving a graph with a single face.




FIGURE 15. The basic operation of cutting and pasting

Thus, by successively adding and erasing edges to \mathbb{G}_1 and \mathbb{G}_2 , we may arrive at a point where $W(M, \mathbb{G}_1) = W(M, \mathbb{G}_2)$. Then, by Proposition 1.3.8, there exists a homeomorphism of M which sends \mathbb{G}_1 to \mathbb{G}_2 . \square

1.4. Riemannian metrics

1.4.1. Measured and Riemannian marked surfaces. — In the setting of Markovian holonomy fields, the scale of time is provided by a means to measure areas on each surface that one considers.

Definition 1.4.1. — *Let M be a smooth compact surface. A measure of area on M is a smooth non-vanishing density on M , that is, a Borel measure which has a smooth positive density with respect to the Lebesgue measure in any coordinate chart.*

A gluing is a diffeomorphism outside a negligible subset. Hence, a measure of area on a surface determines a measure of area on any other surface obtained by splitting this surface along a curve (see Proposition 1.1.11).

Definition 1.4.2. — *Let M be a smooth compact surface endowed with a measure of area denoted by vol . Let l be a mark on M and let $\text{Spl}_l(M)$ be the surface obtained by splitting M along l . Let $f : \text{Spl}_l(M) \rightarrow M$ be the associated gluing. Then $\text{vol} \circ f$ is a measure of area on $\text{Spl}_l(M)$ which we denote by $\text{Spl}_l(\text{vol})$.*

On an oriented surface, a measure of area is also the same thing as a non-vanishing differential 2-form. We are going to work with surfaces endowed with a specific measure of area, on which we will choose Riemannian metrics. We would like these two structures to be compatible. In fact, we have the following result.

Proposition 1.4.3. — *Let (M, \mathcal{C}) be a marked surface. Let vol be a measure of area on M . There exists a Riemannian metric on M whose Riemannian volume is vol and such that the curves of $\mathcal{C} \cup \mathcal{B}(M)$ are closed geodesics.*

Before proving this proposition, let us state a definition.

Definition 1.4.4. — *A marked Riemannian surface is a quadruple $(M, \text{vol}, \gamma, \mathcal{C})$, where $(M, \text{vol}, \mathcal{C})$ is a marked surface endowed with a smooth density and γ is a Riemannian metric on M with Riemannian volume vol and such that the curves of $\mathcal{C} \cup \mathcal{B}(M)$ are closed geodesics.*

Proof. — We denote by vol_γ the Riemannian volume of a Riemannian metric γ on M . Let us first assume that M is orientable and has no boundary. We write $\mathcal{C} = \{l_1^{\pm 1}, \dots, l_q^{\pm 1}\}$.

Let γ_0 be any Riemannian metric on M . Let f be the smooth real function on M such that $\text{vol} = e^f \text{vol}_{\gamma_0}$. Set $\gamma_1 = e^f \gamma_0$. The Riemannian volume of γ_1 is vol .

For each $i \in \{1, \dots, q\}$, let T_i denote a neighbourhood of l_i diffeomorphic to $[-1, 1] \times \mathbb{R}/2\pi\mathbb{Z}$ such that l_i corresponds to $\{0\} \times \mathbb{R}/2\pi\mathbb{Z}$. We assume that T_1, \dots, T_q are pairwise disjoint. For each $i \in \{1, \dots, q\}$, we denote by $(r, \theta)_i$ the natural coordinates in T_i .

Let $\varphi : [-1, 1] \rightarrow [0, 1]$ be a smooth function such that $\varphi([-1, -\frac{3}{4}] \cup [\frac{3}{4}, 1]) = \{0\}$ and $\varphi([-\frac{1}{2}, \frac{1}{2}]) = \{1\}$. Let Φ be the smooth real function on M defined by

$$\Phi(m) = \begin{cases} \varphi(r) & \text{if } m = (r, \theta)_i \text{ for some } i \in \{1, \dots, q\}, \\ 0 & \text{otherwise.} \end{cases}$$

For each $i \in \{1, \dots, q\}$, write $T_i^+ = [0, 1] \times \mathbb{R}/2\pi\mathbb{Z}$, $T_i^- = [-1, 0] \times \mathbb{R}/2\pi\mathbb{Z}$, and define $\delta_i^\pm = \int_{T_i^\pm} \Phi \text{vol}_{\gamma_1}$. Then, choose a smooth function $\Psi : M \rightarrow \mathbb{R}_+$, vanishing outside the union of the T_i 's, such that the following properties are satisfied: for each $i \in \{1, \dots, q\}$, $\Psi = 1$ in a neighbourhood of l_i and $\Psi = 0$ in a neighbourhood of ∂T_i , $\Phi + \Psi > 0$ on T_i and $\int_{T_i^\pm} \Psi \text{vol}_{(dr^2 + d\theta^2)_i} = \delta_i^\pm$. Now consider the Riemannian metric

$$\gamma_2 = (1 - \Phi)\gamma_1 + \Psi \sum_{i=1}^q (dr^2 + d\theta^2)_i.$$

The Riemannian volume of γ_2 coincides with vol on $M \setminus (T_1 \cup \dots \cup T_q)$ and in a neighbourhood of $\partial T_1 \cup \dots \cup \partial T_q$. Moreover, for all $i \in \{1, \dots, q\}$, the total volume of T_i^\pm is the same for vol and vol_{γ_2} . Finally, l_1, \dots, l_q are geodesic for γ_2 .

For each $i \in \{1, \dots, q\}$, and all $r \in [-1, 1]$, set $V_i(r) = \int_0^r \int_0^{2\pi} \text{vol}$. It is understood that $V_i(r) < 0$ when $r < 0$. Similarly, set $V_{i, \gamma_2}(r) = \int_0^r \int_0^{2\pi} \text{vol}_{\gamma_2}$. The functions V_i and V_{i, γ_2} are both smooth, increasing, equal to 0 at $r = 0$ and they coincide on

neighbourhoods of 1 and -1 . Define a diffeomorphism ρ of M by setting

$$\rho(m) = \begin{cases} (V_i^{-1}(V_{i,\gamma_2}(r)), \theta)_i & \text{if } m = (r, \theta)_i \text{ for some } i \in \{1, \dots, q\}, \\ m & \text{otherwise.} \end{cases}$$

The metric $\gamma_3 = (\rho^{-1})^*\gamma_2$ satisfies $\text{vol}_{\gamma_3} = \rho_*\text{vol}_{\gamma_2}$. Hence, the volume of any cylinder $[r, s] \times \mathbb{R}/2\pi\mathbb{Z}$, $r, s \in [-1, 1]$ is the same for vol_{γ_3} and for vol . Moreover, since ρ preserves the curves l_1, \dots, l_q , they are still geodesic for γ_3 .

Let D and D_3 be the two smooth functions defined on $T_1 \cup \dots \cup T_q$ such that $\text{vol} = Ddr \wedge d\theta$ and $\text{vol}_{\gamma_3} = D_3dr \wedge d\theta$. For each $i \in \{1, \dots, q\}$, define

$$A_i(r, \theta) = \int_0^\theta D((r, \xi)_i) d\xi \quad \text{and} \quad A_{i,\gamma_3}(r, \theta) = \int_0^\theta D_3((r, \xi)_i) d\xi.$$

By construction of γ_3 , we have for each $i \in \{1, \dots, q\}$ and all $r \in [-1, 1]$ the equality $A_i(r, 2\pi) = A_{i,\gamma_3}(r, 2\pi)$. It is easy to check that the mapping α from M to itself defined by

$$\alpha(m) = \begin{cases} (r, A_i^{-1}(A_{i,\gamma_3}(\theta)))_i & \text{if } m = (r, \theta)_i \text{ for some } i \in \{1, \dots, q\}, \\ m & \text{otherwise} \end{cases}$$

is a diffeomorphism. Set $\gamma_4 = (\alpha^{-1})^*\gamma_3$. Then by construction, vol_{γ_4} and vol give the same volume to any subset of T_i which is a rectangle in the coordinates $(r, \theta)_i$. Hence, they agree on $T_1 \cup \dots \cup T_q$, hence on M . Since α preserves the curves l_1, \dots, l_q , they are still geodesic for γ_4 . Thus, γ_4 has the desired properties.

Let us explain how the result extends to surfaces with boundary and non-orientable surfaces. Let M be a non-orientable surface without boundary. Then there exists an orientable double of M , that is, an orientable surface $D(M)$ and a smooth mapping $f : D(M) \rightarrow M$ which is a covering of degree 2. This surface $D(M)$ can for instance be constructed as the unitary frame bundle of the real line bundle $\wedge^2 T^*M$ for some Riemannian metric on M . The density vol and the marks of M can be lifted through f . This yields an orientable marked surface $(D(M), D(\mathcal{C}))$ endowed with a density $D(\text{vol})$ whose total area is equal to twice that of vol . The result that we have proved above applied on $D(M)$ yields a Riemannian metric γ with Riemannian volume $D(\text{vol})$ and such that the curves of $D(\mathcal{C})$ are geodesics. Let α be the non-trivial automorphism of the covering $f : D(M) \rightarrow M$, that is, the diffeomorphism of $D(M)$ which exchanges the sheets of the covering. Then the Riemannian metric $\alpha^*\gamma$ has Riemannian volume $\alpha^*D(\text{vol}) = D(\text{vol})$ and makes the curves of $\alpha^{-1}(D(\mathcal{C})) = D(\mathcal{C})$ geodesics. Since the equations of geodesics are linear in the metric, the curves of $D(\mathcal{C})$ are also geodesic for the metric $\frac{1}{2}(\gamma + \alpha^*\gamma)$. This metric has also Riemannian volume $D(\text{vol})$. Moreover, it is invariant by α , hence descends to a Riemannian metric on M with the desired properties.

Finally, if M has a boundary, then we may glue a disk along each boundary component of M and extend vol to a measure of area on the surface without boundary thus obtained. \square

1.4.2. Partially ordered sets of graphs. — The set of graphs on a compact surface carries a natural partial order.

Definition 1.4.5. — Let M be a topological compact surface. Let \mathbb{G}_1 and \mathbb{G}_2 be two graphs on M . We say that \mathbb{G}_2 is finer than \mathbb{G}_1 and write

$$\mathbb{G}_1 \preceq \mathbb{G}_2 \quad \text{if} \quad C(\mathbb{G}_1) \subset C(\mathbb{G}_2).$$

It is straightforward that \mathbb{G}_2 is finer than \mathbb{G}_1 if and only if $\mathbb{E}_1 \subset C(\mathbb{G}_2)$. The inclusion $\mathbb{E}_1 \subset \mathbb{E}_2$ implies $\mathbb{G}_1 \preceq \mathbb{G}_2$ but the converse is false.




FIGURE 16. The graph of Figure 5 and a refinement of it.

As a poset, the set of graphs on a surface has few good properties. In particular, it is not directed, which means that it contains pairs without upper bound.

Lemma 1.4.6. — Let M be a topological compact surface. The set of graphs on M endowed with the partial order \preceq is not directed. In other words, there exist two graphs \mathbb{G}_1 and \mathbb{G}_2 on M such that no third graph \mathbb{G} satisfies $\mathbb{G}_1 \preceq \mathbb{G}$ and $\mathbb{G}_2 \preceq \mathbb{G}$.

Proof. — Let U be an open subset of M identified through a homeomorphism with the disk of \mathbb{R}^2 centred at the origin and of radius 2. Let e_1 and e_2 be the parametrised curves defined by

$$\forall t \in [0, 1], \quad e_1(t) = (t, 0) \quad \text{and} \quad e_2(t) = (t, t^2 \sin \pi/t).$$

For all $k \geq 1$, let A_k denote the open domain of \mathbb{R}^2 delimited by the restrictions of e_1 and e_2 to $[\frac{1}{k+1}, \frac{1}{k}]$. The sets $A_k, k \geq 1$ are also the bounded connected components of the complement of the union of the ranges of e_1 and e_2 in the plane.

The curves e_1 and e_2 are edges, so by Lemma 1.3.7, there exist two graphs \mathbb{G}_1 and \mathbb{G}_2 on M such that e_1 is an edge of \mathbb{G}_1 and e_2 is an edge of \mathbb{G}_2 . Assume that

there exists a graph \mathbb{G} such that $\mathbb{G}_1 \preceq \mathbb{G}$ and $\mathbb{G}_2 \preceq \mathbb{G}$. Then $\text{Sk}(\mathbb{G})$ would contain the union of the ranges of e_1 and e_2 . Since the range of an edge has an empty topological interior, none of the sets A_k would be contained in $\text{Sk}(\mathbb{G})$. Hence, \mathbb{G} would have infinitely many faces. We have observed after Proposition 1.3.8 that this is impossible. \square

In the lemma above, it is not difficult to check that one of the two graphs, for instance \mathbb{G}_1 , can even be chosen arbitrarily. This lack of upper bounds will anyway be a serious problem for us at a later stage. A better-behaved substitute for the set of graphs is the set of graphs with piecewise geodesic edges.

Definition 1.4.7. — *Let (M, \mathcal{C}) be a marked surface endowed with a Riemannian metric γ . We define $\text{Gr}_\gamma(M, \mathcal{C})$ as the set of graphs on (M, \mathcal{C}) with piecewise geodesic edges, that is*

$$\text{Gr}_\gamma(M, \mathcal{C}) = \{ \mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F}) \text{ graph on } (M, \mathcal{C}) : \mathbb{E} \subset \mathbb{A}_\gamma(M) \}.$$

The set $\text{Gr}_\gamma(M, \mathcal{C})$ can be non-empty only if the marks on M are geodesic curves. We know by Proposition 1.4.3 that it is always possible to choose a Riemannian metric on M for which this is the case.

The next result states that $\text{Gr}_\gamma(M, \mathcal{C})$ is indeed a better set of graphs than the set of all graphs.

Proposition 1.4.8. — *Let (M, \mathcal{C}) be a marked surface endowed with a Riemannian metric for which the marks are geodesic curves. Any finite family of piecewise geodesic paths on M is a subset of $\text{P}(\mathbb{G})$ for some graph \mathbb{G} on (M, \mathcal{C}) with geodesic edges. In particular, the poset $(\text{Gr}_\gamma(M, \mathcal{C}), \preceq)$ is directed.*

Proof. — We have observed after the definition of a graph (Definition 1.3.2) that there exists a graph on M with geodesic edges. Now, by induction on the number of curves in the finite family of curves that we consider, it suffices to prove that, given a piecewise geodesic path c and a graph \mathbb{G}_0 with geodesic edges, there exists a graph \mathbb{G} with geodesic edges such that $\text{P}(\mathbb{G}) \supset \text{P}(\mathbb{G}_0) \cup \{c\}$.

For this, let us subdivide c into a product of geodesic edges $e = e_1 \cdots e_m$ in such a way that each edge e_k either is contained in one edge of \mathbb{G}_0 or has its interior contained in the interior of a face. By adding finitely many vertices to \mathbb{G}_0 , which means subdividing some of its edges, we produce a new graph \mathbb{G}_1 which is such that each e_k which is contained in an edge of \mathbb{G}_0 is an edge of \mathbb{G}_1 . Each other e_k has its interior contained in a single face of \mathbb{G}_1 . By lifting the curves e_1, \dots, e_m to a split pattern of \mathbb{G}_1 , we reduce the problem to the case of a finite collection of geodesic segments contained in the interior of a disk with piecewise geodesic boundary. In this case, since the skeleton of \mathbb{G}_1 contains the boundary of the disk, it follows from

Proposition 1.3.10 that any pre-graph obtained by adding edges to \mathbb{G}_1 and whose skeleton is connected is a graph. Thus, it suffices to join one endpoint of each of the curves e_1, \dots, e_m to a point on the boundary of the disk by a geodesic segment and then to add a vertex at every point where two distinct geodesic curves meet. Hence, \mathbb{G}_1 can be refined into a graph \mathbb{G} with geodesic edges such that e_1, \dots, e_m belong to $\mathcal{P}(\mathbb{G})$.

In order to prove that $\text{Gr}_\gamma(M, \mathcal{E})$ is directed, consider two graphs \mathbb{G}_1 and \mathbb{G}_2 with piecewise geodesic edges. The property that we have just proved applied to $\mathbb{E}_1 \cup \mathbb{E}_2$ provides us with a graph which is finer than both \mathbb{G}_1 and \mathbb{G}_2 . \square

1.4.3. Approximation of graphs. — In this section, we prove that any graph can be approximated in a strong sense by a sequence of graphs with piecewise geodesic edges. We start by defining the lasso decomposition of a piecewise geodesic path, which is a variant of the more familiar operation of loop-erasure. Recall the definition of equivalence of paths and lassos (see Section 1.3.30).

Proposition 1.4.9. — *Let (M, γ) be a Riemannian compact surface. Let c be an element of $\mathcal{A}_\gamma(M)$ such that $\underline{c} \neq \bar{c}$. There exists in $\mathcal{A}_\gamma(M)$ a finite sequence of lassos l_1, \dots, l_p with meanders m_1, \dots, m_p and an injective path d with the same endpoints as c such that*

- 1) $c \simeq l_1 \cdots l_p d$,
- 2) $\ell(c) \geq \ell(m_1) + \cdots + \ell(m_p) + \ell(d)$.

If c is a loop, the same decomposition holds with the single difference that d is a simple loop. In both cases, we call d the loop-erasure of c and denote it by $\text{LE}(c)$.




FIGURE 17. A path, the two lassos involved in its lasso decomposition and its loop-erasure.

Proof. — By Lemma 1.4.8, we may assume that c is a path in a graph. By adding vertices to this graph, we may also assume that no edge of this graph is a simple loop. Let us write $c = e_{i_1} \cdots e_{i_n}$ as a product of edges. We proceed by induction on n . If $n = 1$, then c is its own loop-erasure. Now assume that $n > 1$. If c is not reduced, that is, if it contains at least one sequence ee^{-1} , then we reduce it. This can only shorten c . Now set $r = \min\{j > 1 : \exists k \in \{1, \dots, j-1\}, \bar{e}_{i_j} = \bar{e}_{i_k}\}$. This is the first time at which c hits itself. Assume that $\bar{e}_{i_r} = \bar{e}_{i_p}$ with $1 \leq p < r$. By definition of r , p is uniquely determined by this relation. Set $t = e_{i_1} \cdots e_{i_{p-1}}$, $m = e_{i_p} \cdots e_{i_r}$ and

$c' = te_{i_{r+1}} \cdots e_{i_n}$. Then $c \simeq tmt^{-1}c'$. By construction, m is a simple loop and c' is a path shorter than c with the same endpoints. Moreover, $\ell(c) = \ell(c') + \ell(m)$. The result follows. \square

Proposition 1.4.10. — *Let $(M, \text{vol}, \gamma, \mathcal{C})$ be a marked Riemannian surface. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on (M, \mathcal{C}) . Let $\varepsilon > 0$ be a real number. There exists a graph $\mathbb{G}' = (\mathbb{V}', \mathbb{E}', \mathbb{F}')$ on (M, \mathcal{C}) with piecewise geodesic edges and two bijections $S : \mathbb{E} \rightarrow \mathbb{E}'$ and $S : \mathbb{F} \rightarrow \mathbb{F}'$, denoted by the same letter, such that the following properties hold.*

- 1) $\mathbb{V}' = \mathbb{V}$.
- 2) *The bijection $S : \mathbb{E} \rightarrow \mathbb{E}'$ commutes with the inversion and preserves the endpoints: for all $e \in \mathbb{E}$, $S(e^{-1}) = S(e)^{-1}$, $\underline{S(e)} = \underline{e}$ and $\overline{S(e)} = \bar{e}$.*
- 3) *The bijection $S : \mathbb{P}(\mathbb{G}) \rightarrow \mathbb{P}(\mathbb{G}')$ induced by S is such that $\partial(S(F)) = S(\partial F)$, for all $F \in \mathbb{F}$.*
- 4) *For all $e \in \mathbb{E}$, $d_\ell(e, S(e)) < \varepsilon$ and for all $F \in \mathbb{F}$, $\text{vol}((F \cup S(F)) \setminus (F \cap S(F))) < \varepsilon$.*

Notice that it would not solve the problem to simply replace each edge by the loop-erasure of a fine dyadic approximation of it. Indeed, the new edges thus produced may behave badly near the vertices. For example, several edges may form a complicated spiral near a vertex that they share. In this case, it is neither certain that their loop-erased dyadic piecewise geodesic approximations do not intersect each other nor that they leave the vertex in the same cyclic order as the original edges.

If $x, y \in M$ and $d(x, y)$ is smaller than the injectivity radius of (M, γ) , we denote by $[x, y]$ the segment of minimizing geodesic joining x to y .

Proof. — Let R be the injectivity radius of M . Let $r_0 \in (0, R)$ be such that the balls $B(v, r_0), v \in \mathbb{V}$ are pairwise disjoint, an edge e meets a ball $B(v, r_0)$ only if v is an endpoint of e , and the cyclic order of the outgoing edges at every vertex v is the cyclic order of their last exit points from the ball $B(v, r)$ for all $r \in (0, r_0)$. The existence of such an r_0 is granted by Lemma 1.3.16. Let us choose an orientation of each ball $B(v, r_0), v \in \mathbb{V}$. Thanks to the Riemannian metric, we parametrise all edges at constant speed.

For all real $r \in (0, r_0)$ and all $e \in \mathbb{E}$, define

$$t_e(r) = \sup \{ t \in [0, 1] : d(\underline{e}, e(t)) = r \}.$$

Observe that $1 - t_{e^{-1}} = \inf \{ t \in [0, 1] : d(e(t), \bar{e}) = r \}$. Define

$$A_r(e) = [\underline{e}, e(t_e(r))] e_{|[t_e(r), 1-t_{e^{-1}}(r)] [e(1-t_{e^{-1}}(r)), \bar{e}].$$

The path $A_r(e)$ is the concatenation of three injective paths which meet only at their endpoints, so that it is injective. Thus, $A_r(e)$ is an edge with the same endpoints as e . Let e and e' be two edges such that $e' \notin \{e, e^{-1}\}$. The central portions of $A_r(e)$

and $A_r(e')$ are contained in e and e' respectively, so that they are disjoint. By the assumption made on r , they do not enter any ball $B(v, r), v \in \mathbb{V}$. Hence, $A_r(e)$ and $A_r(e')$ meet, if at all, in one of these balls and this can occur only at one of their endpoints.




FIGURE 18. The successive stages of the definition of the approximation of an edge.

Let e be an edge. The continuity and injectivity of e imply that $t_e(r) \rightarrow 0$ as $r \rightarrow 0$. Since $d_\ell(A_r(e), e) \leq 2(t_e(r) + t_{e^{-1}}(r))\ell(e)$, this implies that $A_r(e)$ tends to e as r tends to 0. Moreover, one always has the inequality $\ell(A_r(e)) \leq \ell(e)$.

Let $r \in (0, r_0)$ be fixed. Let $n_0 \geq 1$ be an integer such that for all $e \in \mathbb{E}$, $2^{-n_0}\ell(e) < r < R$. For all integer $n \geq n_0$, define

$$C_{r,n}(e) = D_n(A_r(e)),$$

where D_n is the dyadic approximation of order n (see Definition 1.2.11). The path $C_{r,n}(e)$ is piecewise geodesic, with the same endpoints as e , but it may not be injective, even for large n . On the other hand, it coincides with $A_r(e)$ near its endpoints, more precisely, on a segment of length at least $r - 2^{-n}\ell(e)$. When n tends to infinity, $C_{r,n}(e)$ converges to $A_r(e)$. We claim that for all $e, e' \in \mathbb{E}$ such that $e' \notin \{e, e^{-1}\}$, the paths $C_{r,n}(e)$ and $C_{r,n}(e')$ intersect only at some of their endpoints for n large enough. Indeed, consider for all $e \in \mathbb{E}$ the segment

$$A'_r(e) = \{m \in A_r(e) : d(m, \underline{e}) \geq \frac{1}{2}r, d(m, \bar{e}) \geq \frac{1}{2}r\}.$$

Choose $n_1 \geq n_0$ such that

$$2^{-n_1} \max\{\ell(e) : e \in \mathbb{E}\} < \min\left(\frac{1}{2}r, \frac{1}{2} \min\{d(A_r(e), A'_r(e')) : e, e' \in \mathbb{E}, e' \notin \{e, e^{-1}\}\}\right).$$

Choose $e, e' \in \mathbb{E}$ such that $e' \notin \{e, e^{-1}\}$. For all $n \geq n_1$, $C_{r,n}(e)$ and $C_{r,n}(e')$ are respectively contained in the sets $A_r(e) \cup \{m \in M : d(m, A'_r(e)) < 2^{-n}\ell(e)\}$ and $A_r(e') \cup \{m \in M : d(m, A'_r(e')) < 2^{-n}\ell(e)\}$, whose intersection is the same as the intersection of e and e' .

For each integer $p \geq 1$, choose $r \in (0, r_0)$ such that $d_\ell(e, A_r(e)) < \frac{1}{2p}$ for all $e \in \mathbb{E}$. Then choose $n \geq n_1$ such that $d_\ell(A_r(e), C_{r,n}(e)) < \frac{1}{2p}$ for all e . Set $S_p^0(e) = C_{r,n}(e)$. Then for each $e \in \mathbb{E}$, the sequence $(S_p^0(e))_{p \geq 1}$ converges to e with fixed endpoints, and satisfies $\ell(S_p^0(e)) \leq \ell(e)$. For each $e \in \mathbb{E}$ and each $p \geq 1$, set $S_p^1(e) = \text{LE}(S_p^0(e))$. We claim that a subsequence of the sequence $(S_p^1(e))_{p \geq 1}$ tends to e when p tends to infinity.

Indeed the sequence $(S_p^1(e))_{p \geq 1}$ is uniformly bounded in length by $\ell(e)$. Hence, the paths being parametrised at constant speed, it is relatively compact in the uniform topology and we can extract a sequence $(S_{p_q}^1(e))_{q \geq 1}$ which converges uniformly to a path \tilde{e} . The image of \tilde{e} is contained in the image of e and it joins \underline{e} to \bar{e} . Hence, the images of \tilde{e} and e coincide. In particular, $\ell(\tilde{e}) \geq \ell(e)$. Using the lower semi-continuity of the length with respect to the uniform convergence, we find

$$\ell(e) \leq \ell(\tilde{e}) \leq \liminf \ell(S_{p_q}^1(e)) \leq \sup \ell(S_{p_q}^1(e)) \leq \ell(e).$$

It follows from these inequalities that $\ell(S_{p_q}^1(e))$ converges to $\ell(e)$, hence $S_{p_q}^1(e)$ to e , as q tends to infinity. Let us choose a subsequence $(p_q)_{q \geq 0}$ such that the convergence holds for each edge $e \in \mathbb{E}$ and define $S_q(e) = S_{p_q}^1(e)$ for all $q \geq 0$.

For all $e \in \mathbb{E}$ and all $q \geq 0$, $S_q(e)$ is a piecewise geodesic edge with the same endpoints as e . Moreover, by construction, $S_q(e^{-1}) = S_q(e)^{-1}$. If e is geodesic, then $S_q(e) = e$. Finally, the set $\mathbb{E}_q = \{S_q(e) : e \in \mathbb{E}\}$ is the set of edges of a pre-graph on (M, \mathcal{C}) . By Corollary 1.3.12, \mathbb{E}_q is in fact the set of edges of a graph on (M, \mathcal{C}) , which we denote by \mathbb{G}_q .

By construction, the bijection S_q between \mathbb{E} and \mathbb{E}_q preserved the cyclic order at every vertex of \mathbb{V} . Let us apply Proposition 1.3.22. Since $\text{vol}(\text{Sk}(\mathbb{G})) = 0$, the graph \mathbb{G}_q satisfies all the desired properties for q large enough. \square

CHAPTER 2

MULTIPLICATIVE PROCESSES INDEXED BY PATHS

In this chapter, we begin the study of the class of objects to which Markovian holonomy fields belong: stochastic processes indexed by paths on a compact surface and with values in a compact Lie group, which satisfy a condition of multiplicativity. We discuss the canonical space of such processes, two distinct σ -fields on it, and prove the version of Kolmogorov's theorem that is best suited to our situation. We then study a kind of uniform measure on the canonical space, thus providing the first example of what will be called in the next chapter a discrete Markovian holonomy field. We conclude by constructing a set of generators of the group of reduced loops in a graph for which we are able to determine the finite-dimensional marginal of the uniform Markovian holonomy field.

From now on, the expression *compact surface* will mean *smooth compact surface*.

2.1. Multiplicative functions

Let G be a group, on which we make for the moment no assumption at all. Recall that $P(M)$ denotes the set of paths on the surface M (see Definition 1.2.2).

Definition 2.1.1. — *Let M be a compact surface. Let P be a subset of $P(M)$. A function $h : P \rightarrow G$ is said to be multiplicative if $h(c^{-1}) = h(c)^{-1}$ for all $c \in P$ such that $c, c^{-1} \in P$ and $h(c_1c_2) = h(c_2)h(c_1)$ for all $c_1, c_2 \in P$ such that $\bar{c}_1 = \underline{c}_2$ and $c_1c_2 \in P$. The set of multiplicative functions from P to G is denoted by $\mathcal{M}(P, G)$.*

For an explanation of the reversed order in this definition, see the introduction, Section 0.5.

The gauge group is the symmetry group of the physical theory from which our objects are issued.

Definition 2.1.2. — *The group G^M of all mappings from M to G is called the gauge group of M and it acts by gauge transformations on the space $\mathcal{M}(P(M), G)$, as follows.*

If $j = (j_m)_{m \in M}$ belongs to G^M and h belongs to $\mathcal{M}(P(M), G)$, then $j \cdot h$ is defined by

$$\forall c \in P(M), \quad (j \cdot h)(c) = j_{\bar{c}}^{-1} h(c) j_c.$$

More generally, given a subset P of $P(M)$, the group G^V acts on $\mathcal{M}(P, G)$ in the same way, where V is defined as the set of endpoints of the paths of P .

Example 2.1.3. — Assume that l_1, \dots, l_n are n simple loops based at the same point m of M . Write $L = \{l_1, \dots, l_n\}$. The concatenation of two loops of L is never a simple loop, hence never an element of L . If we assume moreover that the inverse of a loop of L is never in L , then any G -valued function on L is multiplicative, so that $\mathcal{M}(L, G) = G^L$. The action of the gauge group on $\mathcal{M}(L, G)$ is simply the action of $G = G^{\{m\}}$ on G^L by simultaneous conjugation of each factor. This fundamental example should be kept in mind when one reads Lemma 2.1.5. It explains the importance of this action of G in our context.

We cannot do much if we do not make a few assumptions on G . For the rest of this section, we assume that it is a compact topological group. We do not assume that it is connected, so that it can in particular be finite. Also, for the moment, we do not assume that it is a Lie group. The group G carries its normalised Haar measure and we denote simply by $\int_G f(x) dx$ the integral of a function $f : G \rightarrow \mathbb{R}$ with respect to this measure.

Let P be a subset of $P(M)$. The object of this paragraph is to discuss two natural σ -fields on $\mathcal{M}(P, G)$. The simplest one is the cylinder σ -field, denoted by \mathcal{C} , which is defined as the smallest σ -field which makes the evaluation mapping $h \mapsto h(c)$ measurable for all $c \in P$. The gauge group acts by bi-measurable transformations and it makes sense to speak of measures on $\mathcal{M}(P, G)$ which are invariant under gauge transformations.

It is also natural to consider a smaller σ -field which consists in events which are invariant under the action of the gauge group. In order to discuss this σ -field, let us first associate an abstract graph to each subset of $P(M)$.

Definition 2.1.4. — An abstract graph is a pair of sets (V, E) , whose elements are called vertices and edges, endowed with two mappings $s, t : E \rightarrow V$, called respectively source and target.

Let P be a subset of $P(M)$. The configuration graph of P is the abstract graph (V, E, s, t) defined by setting $V = \bigcup_{c \in P} \{c\} \cup \{\bar{c}\} \subset M$, $E = P$ and, for each $e \in E$, $s(e) = \underline{e}$ and $t(e) = \bar{e}$.

In what follows, we will make use of the diagonal adjoint action of G on G^n defined by $g \cdot (x_1, \dots, x_n) = (gx_1g^{-1}, \dots, gx_ng^{-1})$.

Lemma 2.1.5. — *Let M be a compact surface. Let P be a subset of $\mathcal{P}(M)$. Assume that P is stable by concatenation and inversion, and the configuration graph of P is connected. Let m be a vertex of this configuration graph. Let c_1, \dots, c_n be n elements of P . There exists n loops l_1, \dots, l_n in P all based at m such that the following property holds. For all continuous function $f : G^n \rightarrow \mathbb{R}$ such that the function $h \mapsto f(h(c_1), \dots, h(c_n))$ is gauge-invariant on $\mathcal{M}(P, G)$, there exists a continuous function $\tilde{f} : G^n \rightarrow \mathbb{R}$ invariant under the diagonal action of G such that*

$$\forall h \in \mathcal{M}(P, G), \quad f(h(c_1), \dots, h(c_n)) = \tilde{f}(h(l_1), \dots, h(l_n)).$$

Proof. — Consider the configuration graph of $\{c_1, \dots, c_n\}$. If it is not connected, let us choose one vertex in each connected component not containing m and add to the collection $\{c_1, \dots, c_n\}$ a path of P joining m to this vertex. That such a path exists follows from the assumptions made on P . The collection has become $\{c_1, \dots, c_r\}$ for some $r \geq n$. We denote the configuration graph of this enlarged collection by (V, E) .

Let $T \subset E$ be a spanning tree of (V, E) . Recall the notation $[v_1, v_2]_T$ introduced in Definition 1.3.31. For each $i \in \{1, \dots, n\}$, the path $[m, \underline{c}_i]_T c_i [\bar{c}_i, m]_T$ is a loop based at m . When written as a product of edges, it becomes a word in c_1, \dots, c_r and this word makes sense as an element of P . More precisely, it is a loop of P based at m , which we denote by l_i .

Let h be an element of $\mathcal{M}(P, G)$. Let us define a gauge transformation $j \in G^M$ by setting $j(p) = 1$ if $p \notin V$ and, for all $v \in V$, $j(v) = h([m, v]_T)$. If e is an edge which belongs to T , then it is easy to check that $(j \cdot h)(e) = 1$. Hence, for all $i \in \{1, \dots, n\}$, $(j \cdot h)(l_i) = (j \cdot h)(c_i)$ and, by the invariance property of f , $f(h(c_1), \dots, h(c_n)) = f(h(l_1), \dots, h(l_n))$.

Choose $g \in G$. Since the loops l_1, \dots, l_n are all based at m , the action of the gauge transformation j defined by $j(p) = 1$ if $p \neq m$ and $j(m) = g^{-1}$ transforms $h(l_1), \dots, h(l_n)$ into $gh(l_1)g^{-1}, \dots, gh(l_n)g^{-1}$. Let us define $\tilde{f} : G^n \rightarrow \mathbb{R}$ by

$$\tilde{f}(x_1, \dots, x_n) = \int_G f(gx_1g^{-1}, \dots, gx_ng^{-1}) dg.$$

Then $f(h(l_1), \dots, h(l_n)) = \tilde{f}(h(l_1), \dots, h(l_n))$ and \tilde{f} is invariant under the diagonal action of G . The result is proved. \square

This result motivates the following definition.

Definition 2.1.6. — *Let P be a subset of $\mathcal{P}(M)$ stable by concatenation and inversion. The invariant σ -field on $\mathcal{M}(P, G)$, denoted by \mathcal{I} , is the smallest σ -field such that for all $m \in M$, all integer $n \geq 1$, all finite collection l_1, \dots, l_n of loops based at m and all continuous function $f : G^n \rightarrow \mathbb{R}$ invariant under the diagonal action of G , the mapping $h \mapsto f(h(l_1), \dots, h(l_n))$ is measurable.*

If P can be written as the disjoint union of P_1 and P_2 which are both stable by concatenation and inversion and whose configuration graphs are disjoint, then $\mathcal{M}(P, G)$ is canonically isomorphic to $\mathcal{M}(P_1, G) \times \mathcal{M}(P_2, G)$ and the invariant σ -field on $\mathcal{M}(P, G)$ is the tensor product of the invariant σ -fields on $\mathcal{M}(P_1, G)$ and $\mathcal{M}(P_2, G)$.

Let M and M' be two surfaces. Let $\psi : M' \rightarrow M$ be a smooth mapping. Let P' be a subset of $\mathbf{P}(M')$ and $P = \psi(P')$. Then ψ induces a map from P' to P , hence a map from $\mathcal{M}(P, G)$ to $\mathcal{M}(P', G)$.

Lemma 2.1.7. — *Let M and M' be two surfaces. Let $\psi : M' \rightarrow M$ be a smooth mapping. Let P' be a subset of $\mathbf{P}(M')$ and $P = \psi(P')$. Then the induced map $\psi : \mathcal{M}(P, G) \rightarrow \mathcal{M}(P', G)$ is measurable with respect to the cylinder σ -fields, and also with respect to the invariant σ -fields.*

Proof. — Let $f : G \rightarrow \mathbb{R}$ be continuous and consider $c' \in P'$. The function $h \mapsto f(\psi(h)(c'))$ on $\mathcal{M}(P, G)$ is equal to the function $h \mapsto f(h(\psi(c')))$, which is measurable with respect to \mathcal{C} because $\psi(c')$ belongs to P . This proves the first assertion.

Now let $f : G^n \rightarrow \mathbb{R}$ be continuous and invariant under the diagonal action of G by conjugation. Let l'_1, \dots, l'_n be n loops of P' based at the same point. Then $\psi(l'_1), \dots, \psi(l'_n)$ are n loops of P based at the same point and the function $h \mapsto f(\psi(h)(l'_1), \dots, \psi(h)(l'_n)) = f(h(\psi(l'_1), \dots, \psi(l'_n)))$ on $\mathcal{M}(P, G)$ is measurable with respect to \mathcal{J} . This proves the second assertion. \square

Let us conclude this paragraph by discussing the case where P is the set of paths in a graph, in the sense of Definition 1.3.2. Let M be a surface endowed with a graph \mathbb{G} . Let \mathbb{E}^+ be an orientation of \mathbb{G} , that is, a subset of \mathbb{E} which contains exactly one element of each of the pairs $\{e, e^{-1}\}, e \in \mathbb{E}$. It is plain that a multiplicative function on $\mathbf{P}(\mathbb{G})$ is determined by its values on the edges \mathbb{E} or even just those of \mathbb{E}^+ . More precisely, the natural surjective mapping $\mathcal{M}(\mathbf{P}(\mathbb{G}), G) \rightarrow \mathcal{M}(\mathbb{E}^+, G)$ induced by the inclusion $\mathbb{E}^+ \subset \mathbf{P}(\mathbb{G})$ is one-to-one. Indeed, if c belongs to $\mathbf{P}(\mathbb{G})$, then $c = e_1^{\varepsilon_1} \cdots e_n^{\varepsilon_n}$ for some $e_1, \dots, e_n \in \mathbb{E}^+$ and $\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}$. Then, for all multiplicative function h , one has $h(c) = h(e_n)^{\varepsilon_n} \cdots h(e_1)^{\varepsilon_1}$.

Since the interior of an edge contains no vertex, the concatenation of at least two edges is never an edge. Hence, every mapping from \mathbb{E}^+ to G is multiplicative. We will often make the identifications

$$\mathcal{M}(\mathbf{P}(\mathbb{G}), G) = \mathcal{M}(\mathbb{E}, G) = \mathcal{M}(\mathbb{E}^+, G) = G^{\mathbb{E}^+}$$

without further comment. In particular, we will sometimes use a collection $(g_e)_{e \in \mathbb{E}^+}$ of elements of G to denote an element of $\mathcal{M}(\mathbf{P}(\mathbb{G}), G)$. Recall the definition of the group of reduced loops (Definition 1.3.30).

Proposition 2.1.8. — *Let M be a compact surface endowed with a graph \mathbb{G} . Let v be a vertex of \mathbb{G} . Let $\{l_1, \dots, l_r\}$ be a generating subset of the group $\text{RL}_v(\mathbb{G})$ of reduced loops in \mathbb{G} based at v . The invariant σ -field \mathcal{I} on $\mathcal{M}(\text{P}(\mathbb{G}), G)$ is generated by the functions of the form $f(h(l_1), \dots, h(l_r))$ where $f : G^r \rightarrow \mathbb{R}$ is continuous and invariant by diagonal conjugation.*

Proof. — By definition of the invariant σ -field and by Lemma 2.1.5, it suffices to prove that for all $l'_1, \dots, l'_n \in \text{L}_v(\mathbb{G})$ and all continuous $f' : G^n \rightarrow \mathbb{R}$ invariant by diagonal conjugation, the function $f'(h(l'_1), \dots, h(l'_n))$ can be put under the form $f(h(l_1), \dots, h(l_r))$ for some invariant function f . This is easily done by expressing, modulo the equivalence relation on paths, the loops l'_1, \dots, l'_n as words in the generators l_1, \dots, l_r . One then uses the multiplicativity of the elements of $\mathcal{M}(\text{P}(\mathbb{G}), G)$ and the fact that the group operation of G , seen as a map $G^2 \rightarrow G$, is equivariant with respect to the diagonal actions of G by conjugation. \square

2.2. Multiplicative families of random variables

Let M be a compact surface and P a subset of $\text{P}(M)$. A probability measure on $(\mathcal{M}(P, G), \mathcal{C})$ determines a family of G -valued random variables $(H_c)_{c \in P}$ which are just the evaluation functions on $\mathcal{M}(P, G)$, defined by $H_c(h) = h(c)$. These random variables form a multiplicative family in the sense that $H_{c^{-1}} = H_c^{-1}$ and $H_{c_1 c_2} = H_{c_2} H_{c_1}$ almost surely whenever this makes sense. If P is countable, then the converse is true since one can dismiss the negligible event on which the equalities do not hold. We prove in this section that the converse is in fact true even if P is not countable. Let us recall the definition of a projective family of probability spaces.

Definition 2.2.1. — *A projective family of probability spaces is the data of the following ingredients.*

- ▷ A partially ordered set (Λ, \preceq) .
- ▷ For each $\lambda \in \Lambda$, a probability space $(\Omega_\lambda, \mathcal{C}_\lambda, m_\lambda)$.
- ▷ For each pair $(\lambda, \mu) \in \Lambda^2$ such that $\lambda \preceq \mu$, a measurable mapping $\rho_{\lambda\mu} : \Omega_\mu \rightarrow \Omega_\lambda$.

These ingredients are assumed to satisfy the following conditions.

- 1) *The poset (Λ, \preceq) is directed: for all $\lambda, \mu \in \Lambda$, there exists $\nu \in \Lambda$ such that $\lambda \preceq \nu$ and $\mu \preceq \nu$.*
- 2) *For all $\lambda, \mu, \nu \in \Lambda$ such that $\lambda \preceq \mu \preceq \nu$, one has the equality $\rho_{\lambda\mu} \circ \rho_{\mu\nu} = \rho_{\lambda\nu}$.*
- 3) *For all $\lambda, \mu \in \Lambda$ such that $\lambda \preceq \mu$, one has $m_\mu \circ \rho_{\lambda\mu}^{-1} = m_\lambda$.*

Let us state a general result of existence and uniqueness.

Proposition 2.2.2. — *We keep the notation of Definition 2.2.1. Let Ω denote the set-theoretic projective limit of the family $(\{\Omega_\lambda\}, \{\rho_{\lambda\mu}\})$, endowed for each $\lambda \in \Lambda$ with the canonical mapping $\rho_\lambda : \Omega \rightarrow \Omega_\lambda$. Let \mathcal{C} denote the smallest σ -field on Ω which makes all the mappings ρ_λ measurable.*

Assume that for each $\lambda \in \Lambda$, $(\Omega_\lambda, \mathcal{C}_\lambda)$ is a compact metric space endowed with the Borel σ -field. Assume also that for all $\lambda, \mu \in \Lambda$ such that $\lambda \preceq \mu$, the mapping $\rho_{\lambda\mu}$ is continuous.

Then there exists a unique probability measure m on \mathcal{C} such that for all $\lambda \in \Lambda$, one has $m \circ \rho_\lambda^{-1} = m_\lambda$.

This result belongs to a wide family of theorems whose prototype is due to Kolmogorov and whose common ground is Caratheodory's extension theorem. The most common versions assert the existence and uniqueness of m under less restrictive conditions on the probability spaces but more restrictive conditions on the poset. Typically, Λ is the set of finite subsets of a finite-dimensional Euclidean space and the probability spaces are Polish spaces. The form which we have stated is in fact fairly easy to prove thanks to the stronger assumptions which we make on the probability spaces. We think wiser to give a proof than to refer to several places in the literature from which the reader would have to collect the various pieces of the argument.

Proof. — Let Π_Λ denote the Cartesian product of the sets Ω_λ , $\lambda \in \Lambda$. Recall that the set-theoretic projective limit of the family $(\{\Omega_\lambda\}, \{\rho_{\lambda\mu}\})$ is, by definition, the set

$$\Omega = \{(\omega_\lambda)_{\lambda \in \Lambda} \in \Pi_\Lambda : \forall \lambda, \mu \in \Lambda, \lambda \preceq \mu \Rightarrow \rho_{\lambda\mu}(\omega_\mu) = \omega_\lambda\}.$$

The set Π_Λ , endowed with the product topology, is a compact topological space of which Ω , as intersection of closed subsets, is a compact topological subspace. It is endowed with the continuous coordinate mappings $\rho_\lambda : \Omega \rightarrow \Omega_\lambda$, $\lambda \in \Lambda$.

For each $\lambda \in \Lambda$, let S_λ denote the closed support of the probability measure m_λ . It is a non-empty compact subset of Ω_λ . Consider $\lambda \preceq \mu$. The equality $m_\mu \circ \rho_{\lambda\mu}^{-1} = m_\lambda$ implies that $\rho_{\lambda\mu}(S_\mu) = S_\lambda$. We claim that for all $\xi \in \Lambda$, $\rho_\xi(\Omega) \supset S_\xi$. Indeed, choose $\xi \in \Lambda$ and $s_\xi \in S_\xi$. Define, for all $\mu \preceq \nu$,

$$\Omega(\mu, \nu; \xi) = \{(\omega_\lambda)_{\lambda \in \Lambda} \in \Pi_\Lambda : \rho_{\mu\nu}(\omega_\nu) = \omega_\mu, \omega_\xi = s_\xi\}.$$

Then on one hand, $\bigcap_{\mu \preceq \nu} \Omega(\mu, \nu; \xi) = \rho_\xi^{-1}(s_\xi)$. On the other hand, no finite intersection of the sets $\Omega(\mu, \nu; \xi)$ is empty. Indeed, given $\mu_1 \preceq \nu_1, \dots, \mu_k \preceq \nu_k$, consider $\zeta \in \Lambda$ such that $\mu_1, \nu_1, \dots, \mu_k, \nu_k, \xi \preceq \zeta$. Since $\rho_{\xi\zeta}(S_\zeta) = S_\xi$, there exists $s_\zeta \in \Omega_\zeta$ such that $\rho_{\xi\zeta}(s_\zeta) = s_\xi$. Then, the collection $(\rho_{\mu_1\zeta}(s_\zeta), \rho_{\nu_1\zeta}(s_\zeta), \dots, \rho_{\mu_k\zeta}(s_\zeta), \rho_{\nu_k\zeta}(s_\zeta), s_\xi)$, arbitrarily extended to an element of Π_Λ , belongs to $\Omega(\mu_1, \nu_1; \xi) \cap \dots \cap \Omega(\mu_k, \nu_k; \xi)$. Since the sets $\Omega(\mu, \nu; \xi)$ are compact, their intersection is non-empty. Hence $s_\xi \in \rho_\xi(\Omega)$.

Define a collection \mathcal{C}_Λ of subsets of Ω by setting $\mathcal{C}_\Lambda = \bigcup_{\lambda \in \Lambda} \rho_\lambda^{-1}(\mathcal{C}_\lambda)$. The collection \mathcal{C}_Λ is not a σ -field, but it is stable by complementation, finite unions and finite

intersections. Consider $A \in \mathcal{C}_\Lambda$. Assume that $A = \rho_\lambda^{-1}(A_\lambda)$ and $A = \rho_\mu^{-1}(A_\mu)$ for some $\lambda, \mu \in \Lambda$ and $A_\lambda \in \mathcal{C}_\lambda$, $A_\mu \in \mathcal{C}_\mu$. Let ν be such that $\lambda \preceq \nu$ and $\mu \preceq \nu$. Then

$$A = \rho_\nu^{-1}(\rho_{\lambda\nu}^{-1}(A_\lambda)) = \rho_\nu^{-1}(\rho_{\mu\nu}^{-1}(A_\mu)).$$

Since $\rho_\nu(\Omega) \supset S_\nu$, this equality implies $\rho_{\lambda\nu}^{-1}(A_\lambda) \cap S_\nu = \rho_{\mu\nu}^{-1}(A_\mu) \cap S_\nu$. Hence,

$$\begin{aligned} m_\lambda(A_\lambda) &= m_\nu(\rho_{\lambda\nu}^{-1}(A_\lambda)) = m_\nu(\rho_{\lambda\nu}^{-1}(A_\lambda) \cap S_\nu) \\ &= m_\nu(\rho_{\mu\nu}^{-1}(A_\mu) \cap S_\nu) = m_\nu(\rho_{\mu\nu}^{-1}(A_\mu)) = m_\mu(A_\mu). \end{aligned}$$

We have proved that $\rho_\lambda^{-1}(A_\lambda) = \rho_\mu^{-1}(A_\mu)$ implies $m_\lambda(A_\lambda) = m_\mu(A_\mu)$. Hence, for all $A \in \mathcal{C}_\Lambda$, it is legitimate to call $m(A)$ the common value of all $m_\lambda(A_\lambda)$ for $\lambda \in \Lambda$ and $A_\lambda \in \mathcal{C}_\lambda$ such that $A = \rho_\lambda^{-1}(A_\lambda)$. Thus, we have defined $m : \mathcal{C}_\Lambda \rightarrow [0, 1]$. It is not difficult to check that m is finitely additive.

We claim that m is σ -additive on \mathcal{C}_Λ . In order to prove this, it is sufficient to prove that if $(A_n)_{n \geq 1}$ is a decreasing sequence of elements of \mathcal{C}_Λ such that $\bigcap_{n \geq 1} A_n = \emptyset$, then $\lim m(A_n) = 0$. Let us choose such a sequence $(A_n)_{n \geq 1}$. Let us also choose $\varepsilon > 0$. For each n , let us write $A_n = \rho_{\lambda_n}^{-1}(A_{\lambda_n})$ for some $\lambda_n \in \Lambda$ and some $A_{\lambda_n} \in \mathcal{C}_{\lambda_n}$. For each n , the inner regularity of the measure m_{λ_n} implies the existence of a compact subset K_{λ_n} of A_{λ_n} such that $m_{\lambda_n}(A_{\lambda_n} - K_{\lambda_n}) \leq \varepsilon 2^{-n}$. Since Ω endowed with the trace of the product topology of Π_Λ is compact, the sets $K_n = \rho_{\lambda_n}^{-1}(K_{\lambda_n})$, $n \geq 0$ are compact. Since $K_n \subset A_n$ for all $n \geq 1$, the intersection $\bigcap_{n \geq 1} K_n$ is empty. Hence, there exists $N \geq 1$ such that $\bigcap_{n=1}^N K_n = \emptyset$. Now, for each $\omega \in A_N$, there exists at least one $n \in \{1, \dots, N\}$ such that $\omega \notin K_n$, so that $\omega \in A_n - K_n$. Hence, $m(A_N) \leq \sum_{n=1}^N m(A_n - K_n) \leq \varepsilon$. This proves that $m(A_n)$ tends to 0 when n tends to infinity.

Carathéodory's extension theorem asserts that m admits a unique σ -additive extension to the σ -field generated by \mathcal{C}_Λ , which is by definition the smallest σ -field on Ω such that the mappings ρ_λ are measurable. This extension of m , which we still denote by m , satisfies $m \circ \rho_\lambda^{-1} = m_\lambda$ for all $\lambda \in \Lambda$ by definition. \square

In our setting, this theorem can be applied as follows.

Proposition 2.2.3. — *Let P be a subset of $\mathbf{P}(M)$. Let \mathcal{F} be a collection of finite subsets of P whose union is P and which, when ordered by the inclusion, is directed. For all $J \in \mathcal{F}$, let m_J be a probability measure on $(\mathcal{M}(J, G), \mathcal{C})$. Assume that the probability spaces $(\mathcal{M}(J, G), \mathcal{C}, m_J)$ endowed with the restriction mappings $\rho_{JK} : \mathcal{M}(K, G) \rightarrow \mathcal{M}(J, G)$ defined for $J \subset K$ form a projective system. Then there exists a unique probability measure m on $(\mathcal{M}(P, G), \mathcal{C})$ such that for all $J \subset P$, the image of m by the restriction mapping $\rho_J : \mathcal{M}(P, G) \rightarrow \mathcal{M}(J, G)$ is m_J .*

In particular, if $(H_c)_{c \in P}$ is a collection of G -valued random variables such that

- 1) $\forall c \in P$, $c^{-1} \in P$ implies $H_{c^{-1}} = H_c^{-1}$ a.s.;

2) $\forall c_1, c_2 \in P, c_1 c_2 \in P$ implies $H_{c_1 c_2} = H_{c_2} H_{c_1}$ a.s.,
 then there exists a unique probability measure m on $(\mathcal{M}(P, G), \mathcal{C})$ such that the distribution of the canonical process under m is the same as the distribution of $(H_c)_{c \in P}$.

Proof. — Let J, K be elements of \mathcal{F} such that $J \subset K$. Then $\mathcal{M}(J, G)$ is a compact subset of G^J , actually a smooth compact submanifold. The evaluation mappings on $\mathcal{M}(J, G)$ generate both the topology and the σ -field \mathcal{C} . Hence, \mathcal{C} is the Borel σ -field. Moreover, the restriction mapping $\mathcal{M}(K, G) \rightarrow \mathcal{M}(J, G)$ is continuous.

Proposition 2.2.2 ensures the existence of a probability measure on the projective limit of the sets underlying our probability spaces, endowed with a certain σ -field. In the present case, the projective limit of the sets $\mathcal{M}(J, G)$ is easily identified with the set $\mathcal{M}(P, G)$ in such a way that the mappings $\rho_J : \mathcal{M}(P, G) \rightarrow \mathcal{M}(J, G)$ are simply the restrictions. Through this identification, the σ -field on which Proposition 2.2.2 constructs a measure is the usual cylinder σ -field. The first assertion follows.

Let $(H_c)_{c \in P}$ be a collection of random variables which satisfies the assumptions 1) and 2). For each finite subset J of P , the distribution of $(H_c)_{c \in J}$ is a Borel probability measure on G^J , which we denote by m_J and which is actually supported by $\mathcal{M}(J, G)$. By applying the first assertion to the collection of probability spaces $(\mathcal{M}(J, G), \mathcal{C}, m_J)$ where J spans the collection of finite subsets of P , we find the desired probability measure m on $\mathcal{M}(P, G)$. \square

2.3. Uniform multiplicative functions on a graph

Let M be a compact surface. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on M . In this paragraph, we discuss the uniform measure on $\mathcal{M}(\mathbb{P}(\mathbb{G}), G)$ and some of its natural disintegrations. The disintegrations that we have in mind are associated with random variables associated to marking curves or boundary components. Before we define the measures, let us set up some notation.

2.3.1. Constraints on marked surfaces. — Let us denote by $\text{Conj}(G)$ the set of conjugacy classes of G . The inversion map $x \mapsto x^{-1}$ on G descends to an involution of $\text{Conj}(G)$ which determines an action of $\mathbb{Z}/2\mathbb{Z}$ on $\text{Conj}(G)$. Recall that if (M, \mathcal{C}) is a marked surface (see Definition 1.1.8), then $\mathbb{Z}/2\mathbb{Z}$ acts on $\mathcal{C} \cup \mathcal{B}(M)$ by reversing the orientation.

Definition 2.3.1. — Let (M, \mathcal{C}) be a marked surface. A set of G -constraints on (M, \mathcal{C}) is a $\mathbb{Z}/2\mathbb{Z}$ -equivariant mapping

$$C : \mathcal{C} \cup \mathcal{B}(M) \longrightarrow \text{Conj}(G).$$

A set of G -constraints on a marked surface determines a set of G -constraints on any splitting of this surface. In the case of unary splittings, we need the following observation. If \mathcal{O} is a conjugacy class of G , then the set of squares of elements of \mathcal{O} is also a conjugacy class, which we denote by \mathcal{O}^2 .

Definition 2.3.2. — Let (M, \mathcal{C}, C) be a marked surface endowed with a set of G -constraints. Consider $l \in \mathcal{C}$. Let $f : \text{Spl}_l(M) \rightarrow M$ be the elementary gluing with joint $\{l, l^{-1}\}$. The marked surface $(\text{Spl}_l(M), \text{Spl}_l(\mathcal{C}))$ carries the set of G -constraints $\text{Spl}_l(C)$ defined by

$$\text{Spl}_l(C)(l') = C(f(l')), \quad \forall l' \in \text{Spl}_l(\mathcal{C}),$$

with the following exception: if f is a unary gluing and $f(l') = l^{\pm 1}$, then

$$\text{Spl}_l(C)(l') = C(l)^{\pm 2}.$$

Consider a marked surface (M, \mathcal{C}, C) with G -constraints. The set of elements of $\mathcal{M}(\mathcal{P}(M), G)$ which satisfy the constraints C , that is, the set of elements h such that for all $l \in \mathcal{C} \cup \mathcal{B}(M)$, $h(l) \in C(l)$, is globally invariant under the action of the gauge group.

2.3.2. Uniform measures. — Let us start with a simple definition.

Definition 2.3.3. — Let M be a compact surface. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on M . Let \mathbb{E}^+ be an orientation of \mathbb{G} . The Haar measure on $G^{\mathbb{E}^+}$, seen as a probability measure on $(\mathcal{M}(\mathcal{P}(\mathbb{G}), G), \mathcal{C})$, is called the uniform measure and denoted by $\mathbb{U}_{M, \emptyset}^{\mathbb{G}}$.

Plainly, the uniform measure does not depend on the choice of \mathbb{E}^+ . The reason for the subscript \emptyset will become apparent soon. We would like now to incorporate boundary conditions and constraints along marking curves into the uniform measure $\mathbb{U}_{M, \emptyset}^{\mathbb{G}}$.

Let $\mathcal{O} \subset G$ be a conjugacy class. Let $n \geq 1$ be an integer. The set

$$\mathcal{O}(n) = \{(x_1, \dots, x_n) \in G^n : x_1 \dots x_n \in \mathcal{O}\}$$

is a G^n -homogeneous space under the action

$$(g_1, \dots, g_n) \cdot (x_1, \dots, x_n) = (g_1 x_1 g_2^{-1}, \dots, g_n x_n g_1^{-1}).$$

Let $\delta_{\mathcal{O}(n)}$ denote the extension to G^n of the unique G^n -invariant probability measure on $\mathcal{O}(n) \subset G^n$. In particular, $\mathcal{O}(1) = \mathcal{O}$ and $\delta_{\mathcal{O}(1)}$ is the G -invariant probability measure on \mathcal{O} , which we also denote simply by $\delta_{\mathcal{O}}$. For each element x of G , we denote by \mathcal{O}_x the conjugacy class of x . Thus,

$$(20) \quad \forall x \in G, \quad \delta_{\mathcal{O}_x} = \int_G \delta_{zxz^{-1}} dz.$$

Lemma 2.3.4. — Consider $f \in C^0(G^n)$, $g \in C^0(G^{n-1})$, $u \in C^0(G)$ and a continuous map $h : G^{n-1} \rightarrow G$. Let \mathcal{O} be a conjugacy class of G . The following relations hold:

$$(21) \quad \int_{G^n} f d\delta_{\mathcal{O}(n)} = \int_{G^n} f(x_1, \dots, x_{n-1}, (x_1 \dots x_{n-1})^{-1}y) dx_1 \cdots dx_{n-1} \delta_{\mathcal{O}}(dy),$$

$$(22) \quad \int_{G^n} u(h(x_1, \dots, x_{n-1})x_1 \cdots x_n h(x_1, \dots, x_{n-1})^{-1}) \\ \times g(x_1, \dots, x_{n-1}) \delta_{\mathcal{O}(n)}(dx_1 \cdots dx_n) \\ = \int_G u d\delta_{\mathcal{O}} \int_{G^{n-1}} g(x_1, \dots, x_{n-1}) dx_1 \cdots dx_{n-1},$$

$$(23) \quad \int_{G^{n+1}} f(x_1 x_2, x_3, \dots, x_n) \delta_{\mathcal{O}(n+1)}(dg_1 \cdots dg_{n+1}) = \int_{G^n} f d\delta_{\mathcal{O}(n)},$$

$$(24) \quad \int_{G^n} f(x_2, \dots, x_n, x_1) \delta_{\mathcal{O}(n)}(dx_1 \cdots dx_n) = \int_{G^n} f d\delta_{\mathcal{O}(n)},$$

$$(25) \quad \int_G \left[\int_{G^n} f(x_1, \dots, x_n) \delta_{\mathcal{O}_y(n)}(dx_1 \dots dx_n) \right] dy \\ = \int_{G^n} f(x_1, \dots, x_n) dx_1 \cdots dx_n.$$

Proof. — The right-hand side of (21) defines a measure supported by $\mathcal{O}(n)$ and invariant under the action of G^n . Hence, it is $\delta_{\mathcal{O}(n)}$. The relations (22) and (23) follow from the fact that the mappings $(x_1, \dots, x_{n+1}) \mapsto (x_1 x_2, x_3, \dots, x_n)$ and $(x_1, \dots, x_n) \mapsto (x_2, \dots, x_n, x_1)$ are G -equivariant maps, respectively from $\mathcal{O}(n+1)$ to $\mathcal{O}(n)$ and from $\mathcal{O}(n)$ to $\mathcal{O}(n)$. The relation (25) follows from (20) and a simple change of variables. \square

Let (M, \mathcal{C}, C) be a marked surface with a set of G -constraints, endowed with a graph \mathbb{G} . Let us choose q simple loops l_1, \dots, l_q in $L(\mathbb{G})$ which represent the unoriented cycles of $\mathcal{C} \cup \mathcal{B}(M)$, that is, such that $\mathcal{C} \cup \mathcal{B}(M) = \{l_1, l_1^{-1}, \dots, l_q, l_q^{-1}\}$. Recall that for the sake of simplicity, we denote in the same way a cycle and the corresponding simple loop. We label the elements of \mathbb{E} in such a way that $l_i = e_{i,1} \dots e_{i,n_i}$ for $i \in \{1, \dots, q\}$. Let \mathbb{E}^+ be an orientation of \mathbb{G} such that $e_{i,j} \in \mathbb{E}^+$ for all $i \in \{1, \dots, q\}$ and $j \in \{1, \dots, n_i\}$. Let us label e_1, \dots, e_m the other edges of \mathbb{E}^+ .

Lemma 2.3.5. — Let (M, \mathcal{C}, C) be a marked surface with G -constraints. Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on (M, \mathcal{C}) . The uniform measure on $\mathbb{G}^{\mathbb{E}^+}$ with G -constraints C is defined as the probability measure

$$dg_1 \otimes \cdots \otimes dg_m \otimes \delta_{C(l_1)(n_1)}(dg_{1,n_1} \cdots dg_{1,1}) \otimes \cdots \otimes \delta_{C(l_q)(n_q)}(dg_{q,n_q} \cdots dg_{q,1}).$$

It is denoted by $U_{M,\mathcal{C},C}^G(dg)$. The corresponding probability measure on $\mathcal{M}(\mathbb{P}(\mathbb{G}), G)$, also denoted by $U_{M,\mathcal{C},C}^G$, depends neither on the choice of the simple loops which represent the marking of M nor on the choice of \mathbb{E}^+ .

Proof. — The invariance of the measures $\delta_{\theta(n)}$ by cyclic permutation of the variables, granted by (24), ensures that the measure $U_{M,\mathcal{C},C}^G$ does not depend on the simple loops which we have chosen to represent the cycles of the marking of M . Since the Haar measure on G is invariant by inversion, the measure induced on $\mathcal{M}(\mathbb{P}(\mathbb{G}), G)$ by $U_{M,\mathcal{C},C}^G$ does not depend on the choice of \mathbb{E}^+ . \square

Let us state two basic properties of the measure $U_{M,\mathcal{C},C}^G$.

Proposition 2.3.6. — *Recall the notation of Lemma 2.3.5.*

- 1) The event $\mathcal{N} = \{\exists l \in \mathcal{C} \cup \mathcal{B}(M), h(l) \notin C(l)\}$ satisfies $U_{M,\mathcal{C},C}^G(\mathcal{N}) = 0$.
- 2) The action of $G^\mathbb{V}$ on $\mathcal{M}(\mathbb{P}(\mathbb{G}), G)$ preserves the probability measure $U_{M,\mathcal{C},C}^G$.

Proof. — 1) By definition of $U_{M,\mathcal{C},C}^G$, the support of $U_{M,\mathcal{C},C}^G$ is contained in the closed set $\{\forall l \in \mathcal{C} \cup \mathcal{B}(M), h(l) \in C(l)\} = \mathcal{N}^c$.

2) Choose $v \in \mathbb{V}$ and $x \in G$. Set $j_w = 1$ if $w \neq v$ and $j_v = x$. If v is not located on a curve of \mathcal{C} , then the translation invariance of the Haar measure implies that j leaves the measure $U_{M,\mathcal{C},C}^G$ invariant. Assume now that v is on the curve l_1 which is represented by the cycle $e_1 \cdots e_n$. Assume that e_1 is outgoing at v and e_n is incoming. Then the action of j translates the variables associated to the edges adjacent to v other than e_1 and e_n and it replaces (e_n, \dots, e_1) by $(x^{-1}e_n, \dots, e_1x)$. This leaves the measure $\delta_{C(l_1)(n)}(dg_n \cdots dg_1)$ invariant. Finally, the action of j preserves the measure $U_{M,\mathcal{C},C}^G$. Since the group $G^\mathbb{V}$ is generated by elements which are equal to 1 at all vertices but one, the result follows. \square

Let us state precisely the fact that the measures $U_{M,\mathcal{C},C}^G$ disintegrate each other.

We use the notation $\mathcal{C} = \{l_1, l_1^{-1}, \dots, l_q, l_q^{-1}\}$. In contrast with the notation used before Lemma 2.3.5, this set does not include the boundary components of M anymore. Let us define $\mathcal{C}_0 = \emptyset$ and for each $r \in \{1, \dots, q\}$, $\mathcal{C}_r = \{l_1, l_1^{-1}, \dots, l_r, l_r^{-1}\}$. For each $r \in \{1, \dots, q\}$, any collection $(\theta_1, \dots, \theta_r)$ of r conjugacy classes of G determines the set of G -constraints on \mathcal{C}_r which maps l_i on θ_i for all $i \in \{1, \dots, r\}$. We denote this set of constraints simply by $(\theta_1, \dots, \theta_r)$.

The following lemma is a direct consequence of the definition of $U_{M,\mathcal{C},C}^G$ and (25). It shows that the measures $U_{M,\mathcal{C},(\theta_1, \dots, \theta_r, \theta_{x_{r+1}}, \dots, \theta_{x_q})}^G$, seen as functions of x_{r+1}, \dots, x_q , provide us with a regular disintegration of $U_{M,\mathcal{C}_r,(\theta_1, \dots, \theta_r)}^G$ with respect to the random variables $h(l_{r+1}), \dots, h(l_q)$.

Lemma 2.3.7. — *Let r be an integer between 0 and $q-1$. Let $\theta_1, \dots, \theta_r$ be conjugacy classes of G . Let $f : G^{q-r} \rightarrow \mathbb{R}$ be a continuous function. Then*

$$(26) \quad \int_{\mathcal{M}(\mathbb{P}(G), G)} f(h(l_{r+1}), \dots, h(l_q)) \mathbb{U}_{M, \mathcal{C}, (\theta_1, \dots, \theta_r)}^G(dh) = \int_{G^{q-r}} f(x_{r+1}, \dots, x_q) dx_{r+1} \cdots dx_q,$$

$$(27) \quad \mathbb{U}_{M, \mathcal{C}, (\theta_1, \dots, \theta_r)}^G = \int_{G^{q-r}} \mathbb{U}_{M, \mathcal{C}, (\theta_1, \dots, \theta_r, \theta_{x_{r+1}}, \dots, \theta_{x_q})}^G dx_{r+1} \cdots dx_q.$$

2.3.3. Surgery of uniform measures. — Let us investigate the behaviour of the uniform measures that we have just defined under the basic operations of surgery. So far, we have used the letter f to denote gluing maps and also test functions on G . In the proof of the next result, we need to use both. This is why we change our notation for gluings.

Proposition 2.3.8. — *Let (M, \mathcal{C}, C) be a marked surface with G -constraints endowed with a graph \mathbb{G} . Let (M', \mathcal{C}', C') be a splitting of M and let $\pi : M' \rightarrow M$ denote the gluing map. Let \mathbb{G}' be the graph on M' obtained by lifting \mathbb{G} . Then the mapping $\pi : (\mathcal{M}(\mathbb{P}(G), G), \mathcal{I}) \rightarrow (\mathcal{M}(\mathbb{P}(G'), G), \mathcal{I}')$ induced by π satisfies*

$$(28) \quad \mathbb{U}_{M', \mathcal{C}', C'}^{\mathbb{G}'} = \mathbb{U}_{M, \mathcal{C}, C}^{\mathbb{G}} \circ \pi^{-1}.$$

In this proposition, it is crucial that we consider the invariant σ -fields, not only because in general one should consider exclusively gauge-invariant quantities, but because the equality (28), although meaningful with cylinder σ -fields thanks to Lemma 2.1.7, would simply be false. For instance, consider a binary gluing along two curves b'_1 and b'_2 with joint $b = f(b'_1) = f(b'_2)$. Then the event $\{h' : h'(b'_1) = h'(b'_2)\}$ belongs to \mathcal{C}' and has measure zero. On the other hand, the pull-back by ψ of this event is the event $\{h : h(b) = h(b)\}$ of \mathcal{C} which has full measure. The point is that the first event does not belong to \mathcal{I}' .

Proof. — Let $f : \mathcal{M}(\mathbb{P}(G'), G) \rightarrow \mathbb{R}$ be a continuous function invariant under the action of $G^{V'}$. Then f induces a continuous function $\tilde{f} : \mathcal{M}(\mathbb{P}(G), G) \rightarrow \mathbb{R}$ and we need to prove that

$$(29) \quad \int_{\mathcal{M}(\mathbb{P}(G), G)} \tilde{f} d\mathbb{U}_{M, \mathcal{C}, C}^{\mathbb{G}} = \int_{\mathcal{M}(\mathbb{P}(G'), G)} f d\mathbb{U}_{M', \mathcal{C}', C'}^{\mathbb{G}'}$$

We treat separately the cases of binary and unary gluings.

1) *Binary gluing.* — Let $\{b, b^{-1}\}$ be the joint of the gluing. Let b'_1 and b'_2 be the two boundary components of M' which are identified by π , oriented in such a way that $\pi(b'_1) = \pi(b'_2) = b$.

Let \mathbb{E}_b , \mathbb{E}'_1 and \mathbb{E}'_2 denote respectively the set of edges located on b , b'_1 and b'_2 . These three sets are naturally identified by the gluing. Set $\mathbb{E}'_o = \mathbb{E}' - (\mathbb{E}'_1 \cup \mathbb{E}'_2)$ and $\mathbb{E}_o = \mathbb{E} - \mathbb{E}_b = \pi(\mathbb{E}'_o)$. The gluing identifies naturally \mathbb{E}'_o with \mathbb{E}_o . With this notation, we have the partitions

$$(30) \quad \mathbb{E} = \mathbb{E}_o \cup \mathbb{E}_b \quad \text{and} \quad \mathbb{E}' = \mathbb{E}'_o \cup \mathbb{E}'_1 \cup \mathbb{E}'_2.$$

The partition of \mathbb{E} above determines the equality $\mathcal{M}(\mathbb{E}, G) = \mathcal{M}(\mathbb{E}_o, G) \times \mathcal{M}(\mathbb{E}_b, G)$, according to which we denote by $g = (g_o, g_b)$ the generic element of $\mathcal{M}(\mathbb{E}, G)$. Similarly, we have the decomposition $\mathcal{M}(\mathbb{E}', G) = \mathcal{M}(\mathbb{E}'_o, G) \times \mathcal{M}(\mathbb{E}'_1, G) \times \mathcal{M}(\mathbb{E}'_2, G)$ and we write $g' = (g'_o, g'_1, g'_2)$ for the generic element of $\mathcal{M}(\mathbb{E}', G)$.

With this notation and these identifications, the function \tilde{f} is defined by the equality $\tilde{f}(g_o, g_b) = f(g_o, g_b, g_b)$.

Since each curve of \mathcal{C} is covered either by \mathbb{E}_o or by \mathbb{E}_b , the decomposition of $\mathcal{M}(\mathbb{E}, G)$ above determines a decomposition of the measure $U_{M, \mathcal{C}, G}^G$ as the tensor product of two measures U_o and U_b on $\mathcal{M}(\mathbb{E}_o, G)$ and $\mathcal{M}(\mathbb{E}_b, G)$ respectively, each of which is invariant under the action of the gauge group G^V . Let us assume that b is the product of n edges of \mathbb{E} . Then $\mathcal{M}(\mathbb{E}_b, G)$ can be identified with G^n and the measure U_b corresponds to $\delta_{C(b)(n)}$ under this identification.

Similarly, the measure $U_{M', \mathcal{C}', G'}^G$ splits as the tensor product of three measures U'_o , U'_1 and U'_2 , on $\mathcal{M}(\mathbb{E}'_o, G)$, $\mathcal{M}(\mathbb{E}'_1, G)$ and $\mathcal{M}(\mathbb{E}'_2, G)$ respectively. The last two spaces can be identified with G^n and the measures U'_1 and U'_2 both correspond to $\delta_{C(b)(n)}$ under this identification.

The measures U_o and U'_o correspond to each other via the identification of $\mathcal{M}(\mathbb{E}_o, G)$ and $\mathcal{M}(\mathbb{E}'_o, G)$. Hence, the equality which we need to prove is

$$\begin{aligned} & \int_{\mathcal{M}(P(G), G)} f(g_o, g_b, g_b) U_o(dg_o) U_b(dg_b) \\ &= \int_{\mathcal{M}(\mathbb{E}_o, G) \times \mathcal{M}(\mathbb{E}_b, G)^2} f(g_o, g_1, g_2) U_o(dg_o) U_b(dg_1) U_b(dg_2). \end{aligned}$$

Let \mathbb{V}'_2 denote the subset of \mathbb{V}' consisting of the vertices which lie on b'_2 . The group $G^{\mathbb{V}'_2}$ is a compact Lie group and the invariant measure of its transitive action on the subset $C(b)(n)$ of $\mathcal{M}(\mathbb{E}'_2, G)$ is the measure $\delta_{C(b)(n)} = U'_2$. Let us denote by j'_2 the generic element of $G^{\mathbb{V}'_2}$ and by dj'_2 the Haar measure on this group. Since f is invariant under the action of the gauge group, we have for all (g'_o, g'_1, g'_2) in $\mathcal{M}(P(G'), G)$ the equality $f(j'_2 \cdot g'_o, g'_1, j'_2 \cdot g'_2) = f(g'_o, g'_1, g'_2)$. Observe that some edges of \mathbb{E}'_o have some of their endpoints in \mathbb{V}'_2 , so that the term g'_o is affected by the gauge transformation j'_2 . On the contrary, the term g'_1 is not affected because b'_1 and b'_2 are disjoint. Hence,

with the identifications made earlier,

$$\begin{aligned} \int_{\mathcal{M}(\mathbb{P}(\mathbb{G}), G)} f(g_o, g_b, g_b) U_o(dg_o) U_b(dg_b) \\ = \int_{\mathcal{M}(\mathbb{P}(\mathbb{G}), G) \times G^{v_2}} f(j'_2 \cdot g_o, g_b, j'_2 \cdot g_b) U_o(dg_o) U_b(dg_b) dj'_2. \end{aligned}$$

Since the measure U_o is invariant under the action of G^V , the last term is equal to

$$\int_{\mathcal{M}(\mathbb{P}(\mathbb{G}), G) \times G^{v_2}} f(g_o, g_b, j'_2 \cdot g_b) U_o(dg_o) U_b(dg_b) dj'_2.$$

It suffices now to prove an equality about the function $\int_{\mathcal{M}(\mathbb{E}_o, G)} f(g_o, \cdot, \cdot) U_o(dg_o)$ on G^{2n} . Indeed, let us denote this function by $u : G^{2n} \rightarrow \mathbb{R}$. All we need to prove is that, for all conjugacy class θ ,

$$\begin{aligned} \int_{G^{2n}} u(g_1, \dots, g_n, j_1 g_1 j_2^{-1}, \dots, j_n g_n j_1^{-1}) \delta_{\theta(n)}(dg_1 \cdots dg_n) dj_1 \cdots dj_n \\ = \int_{G^{2n}} u d(\delta_{\theta(n)} \otimes \delta_{\theta(n)}). \end{aligned}$$

But the left-hand side is equal to

$$\int_{G^n} \int_{G^n} u(g_1, \dots, g_n, g_{n+1}, \dots, g_{2n}) \delta_{\theta_{g_1 \dots g_n}(n)}(dg_{n+1} \cdots dg_{2n}) \delta_{\theta(n)}(dg_1 \cdots dg_n)$$

by the very definition of $\delta_{\theta_{g_1 \dots g_n}(n)}$, hence it is equal to the right-hand side.

2) *Unary gluing*. — Let $\{b, b^{-1}\}$ be the joint of the gluing. Let b' be the component of the boundary of M' such that $\pi(b') = b$. Let us write $b = e_1 \cdots e_n$ and $b' = e'_{1,1} \cdots e'_{n,1} e'_{1,2} \cdots e'_{n,2}$ in such a way that $\pi(b'_{i,1}) = \pi(b'_{i,2}) = b_i$ for all $i \in \{1, \dots, n\}$.

Let \mathbb{E}_b denote the set of edges located on b . Set $\mathbb{E}'_1 = \{e'_{1,1}{}^{\pm 1}, \dots, e'_{n,1}{}^{\pm 1}\}$ and $\mathbb{E}'_2 = \{e'_{1,2}{}^{\pm 1}, \dots, e'_{n,2}{}^{\pm 1}\}$. Set $\mathbb{E}'_o = \mathbb{E}' - (\mathbb{E}'_1 \cup \mathbb{E}'_2)$ and $\mathbb{E}_o = \mathbb{E} - \mathbb{E}_b = \pi(\mathbb{E}'_o)$. We will identify freely \mathbb{E}_o with \mathbb{E}'_o .

With this notation, the equalities (30) hold, as well as the subsequent decompositions of $\mathcal{M}(\mathbb{E}, G)$ and $\mathcal{M}(\mathbb{E}', G)$. The function \tilde{f} is also still defined by the equality $\tilde{f}(g_o, g_b) = f(g_o, g_b, g_b)$.

The decomposition $U_{M, \mathcal{E}, C}^G = U_o \otimes U_b$ is valid just as in the binary case, but the decomposition of $U_{M', \mathcal{E}', C'}^G$ is now different. Indeed, this measure splits into the tensor product of U'_o and a measure U'_{12} on $\mathcal{M}(\mathbb{E}'_1, G) \times \mathcal{M}(\mathbb{E}'_2, G)$ which is $\delta_{C(b)^2(2n)}$ under the natural identification of $\mathcal{M}(\mathbb{E}'_1, G) \times \mathcal{M}(\mathbb{E}'_2, G)$ with G^{2n} . The formula which we need to prove is

$$\begin{aligned} \int_{\mathcal{M}(\mathbb{P}(\mathbb{G}), G)} f(g_o, g_b, g_b) U_o(dg_o) U_b(dg_b) \\ = \int_{\mathcal{M}(\mathbb{E}_o, G) \times \mathcal{M}(\mathbb{E}'_1 \cup \mathbb{E}'_2, G)} f(g_o, g_1, g_2) U_o(dg_o) U'_{12}(dg_1, dg_2). \end{aligned}$$

Let V'_{12} denote the set of vertices which lie on b' . By using the invariance of f under the action of the subgroup $G^{V'_{12}}$ of the gauge group and the invariance of the measure U_\circ under the same action, we find just as in the binary case that the left-hand side of the equality to prove is equal to

$$\int_{\mathcal{M}(\mathbb{E}_\circ, G) \times \mathcal{M}(\mathbb{E}_b, G) \times G^{V'_{12}}} f(g_\circ, j'_{12} \cdot g_b, j'_{12} \cdot g_b) U_\circ(dg_\circ) U_b(dg_b) dj'_{12}.$$

The notation here is misleading, since the two occurrences of $j'_{12} \cdot g_b$ do not denote the same object. Indeed, the two occurrences of g_b in the arguments of f are identified respectively with an element of $\mathcal{M}(\mathbb{E}'_1, G)$ and an element of $\mathcal{M}(\mathbb{E}'_2, G)$, on which $G^{V'_{12}}$ acts differently.

Now what we have to prove is really an equality about the function

$$\int_{\mathcal{M}(\mathbb{E}_\circ, G)} f(g_\circ, \dots) U_\circ(dg_\circ)$$

on G^{2n} . Let us call this function $u : G^{2n} \rightarrow \mathbb{R}$. We need to prove that for all conjugacy class θ in G the following equality holds:

$$\begin{aligned} \int_{G^{3n}} u(j_0^{-1} g_1 j_1, \dots, j_{n-1} g_n j_n, j_n^{-1} g_1 j_{n+1}, \dots, j_{2n-1} g_n j_0) \delta_{\theta(n)}(dg_1 \cdots dg_n) dj_0 \cdots dj_{2n} \\ = \int_{G^{2n}} u \delta_{\theta^2(2n)}. \end{aligned}$$

Recall that θ^2 is the conjugacy class constituted by the squares of the elements of θ . We claim that this equality holds for all continuous function u . Indeed, the integral

$$\int_{G^{2n}} u(j_0^{-1} g_1 j_1, \dots, j_{n-1} g_n j_n, j_n^{-1} g_1 j_{n+1}, \dots, j_{2n-1} g_n j_0) dj_0 \cdots dj_{2n}$$

is the integral of u with respect to the measure $\delta_{\theta_{(g_1 \dots g_n)^2(2n)}}$, by the very definition of this measure as the invariant measure under the natural action of G^{2n} on $\theta_{(g_1 \dots g_n)^2(2n)}$. Hence, by a simple particular case of (22), the integral that we are computing is equal to

$$\begin{aligned} \int_{G^n} \left(\int_{G^{2n}} u(x) \delta_{\theta_{(g_1 \dots g_n)^2(2n)}}(dx) \right) \delta_{\theta(n)}(dg_1 \cdots dg_n) \\ = \int_{G^{2n}} u(x) \left(\int_G \delta_{\theta_{g^2(2n)}} \delta_\theta(dg) \right) (dx). \end{aligned}$$

The measure $\delta_{\theta_g^2(2n)}$ is equal to $\delta_{O^2(2n)}$ for δ_θ -almost all g . Hence, the measure between the brackets is $\delta_{\theta^2(2n)}$. This concludes the proof. \square

Let us conclude this section by a much simpler result, which is the simplest instance of invariance under subdivision.

Proposition 2.3.9. — *Let (M, \mathcal{C}, C) be a marked surface with G -constraints. Let \mathbb{G}_1 and \mathbb{G}_2 be graphs on (M, \mathcal{C}) . Assume that $\mathbb{G}_1 \preceq \mathbb{G}_2$. Then the inclusion $\mathcal{P}(\mathbb{G}_1) \subset \mathcal{P}(\mathbb{G}_2)$ induces a measurable restriction map $r : \mathcal{M}(\mathcal{P}(\mathbb{G}_2), G) \rightarrow \mathcal{M}(\mathcal{P}(\mathbb{G}_1), G)$ which satisfies*

$$U_{M, \mathcal{C}, C}^{\mathbb{G}_2} \circ r^{-1} = U_{M, \mathcal{C}, C}^{\mathbb{G}_1}.$$

Proof. — Let us choose an orientation for \mathbb{G}_1 and \mathbb{G}_2 . The restriction map, seen as a map from $G^{\mathbb{E}_2^+}$ to $G^{\mathbb{E}_1^+}$, multiplies the components which correspond to the edges of \mathbb{E}_2 which constitute each edge of \mathbb{E}_1 and forgets about the components which correspond to edges which do not lie in the skeleton of \mathbb{E}_1 .

Since the product of independent uniform variables on G is still uniform, the only non-trivial thing to check is what happens along the marking curves or the boundary components. There, the invariance follows from (23). □

2.4. Tame generators of the group of reduced loops

Consider a surface (M, \emptyset, C) with G -constraints along its boundary (see Definition 2.3.1), endowed with a graph \mathbb{G} . Our objective in this section is to exhibit a family of lassos (see Definition 1.2.4) which generates the group of reduced loops in \mathbb{G} (see Definition 1.3.30) and to compute the distribution of the G -valued random variables associated with these lassos under the constrained uniform measure defined in the previous section. Moreover, we are going to do this in a way which is consistent with the partial order on the set of graphs. The main result of this section will not be used until Chapter 5.

Definition 2.4.1. — *Let M be a compact surface endowed with a graph \mathbb{G} . Let v be a vertex of \mathbb{G} . A lasso $l \in \text{RL}_v(\mathbb{G})$ is said to be facial if its meander represents a facial cycle of \mathbb{G} . It is said to be bounding if its meander covers a connected component of ∂M .*

We want to prove the existence of systems of generators of $\text{RL}_v(\mathbb{G})$ which consist in one bounding lasso for each connected component of ∂M , one facial lasso for each face of \mathbb{G} , and as many supplementary lassos with non-contractible meander as the genus of M (see Section 1.1.1).

Recall the notation $W(M)$ from Section 1.3.5. If w is a word in some set of letters, we denote by \overleftarrow{w} the word obtained by reversing the order of the letters of w .

Proposition 2.4.2. — *Let M be a connected compact surface. Let $g = g(M)$ denote the genus of M . Let $p = p(M)$ be the number of connected components of ∂M . Write $\mathcal{B}(M) = \{b_1, b_1^{-1}, \dots, b_p, b_p^{-1}\}$. If M is oriented, we assume that b_1, \dots, b_p bound M*

positively. Let \mathbb{G} be a graph on M . Let v be a vertex of \mathbb{G} . Let f denote the number of faces of \mathbb{G} .

1) There exists, in the group $\text{RL}_v(\mathbb{G})$, g lassos a_1, \dots, a_g ; p bounding lassos c_1, \dots, c_p whose meanders are equivalent to b_1, \dots, b_p up to permutation; f facial lassos l_1, \dots, l_f whose meanders bound the f faces of \mathbb{G} , positively if M is oriented; and there exists an word w in the letters a_1, \dots, a_g which belongs to $W(M)$ such that the group $\text{RL}_v(\mathbb{G})$ admits the presentation

$$\langle a_1, \dots, a_g, c_1, \dots, c_p, l_1, \dots, l_f \mid w(a_1, \dots, a_g)c_1 \dots c_p = l_1 \dots l_f \rangle$$

and, for all continuous function $f : G^{\mathfrak{g}+\mathfrak{p}+\mathfrak{f}} \rightarrow \mathbb{R}$ and all set C of G -constraints along the boundary components of \mathbb{G} ,

$$\begin{aligned} \int_{\mathcal{M}(\mathbb{G}, G)} f(h(a_1), \dots, h(a_g), h(c_1), \dots, h(c_p), h(l_1), \dots, h(l_f)) U_{M, \emptyset, C}^{\mathbb{G}}(dh) \\ = \int_{G^{\mathfrak{g}+\mathfrak{p}+\mathfrak{f}-1}} f(x_1, \dots, x_g, y_1, \dots, y_p, z_1, \dots, z_{f-1}, z_f) \\ dx_1 \dots dx_g \delta_C(b_1)(dy_1) \dots \delta_C(b_p)(dy_p) dz_1 \dots dz_{f-1}, \end{aligned}$$

where we have set $z_f = y_p \dots y_1 \overleftarrow{w}(x_1, \dots, x_g)(z_{f-1} \dots z_1)^{-1}$.

A collection of lassos such as $\{a_1, \dots, a_g, c_1, \dots, c_p, l_1, \dots, l_f\}$ will be called a tame system of generators associated with the word w .

2) Let \mathbb{G}_1 and \mathbb{G}_2 be two graphs on M such that $\mathbb{G}_1 \preccurlyeq \mathbb{G}_2$. Set $f_1 = f(\mathbb{G}_1)$. Let v be a vertex of \mathbb{G}_1 . Let $\{a_1, \dots, a_g, c_1, \dots, c_p, l_1, \dots, l_{f_1}\}$ be a tame system of generators of $\text{RL}_v(\mathbb{G}_1)$ associated with the word w . Assume that the faces of \mathbb{G}_1 are labelled $\mathbb{F}_1 = \{F_i^1 : i \in \{1, \dots, f_1\}\}$ in such a way that for all $i \in \{1, \dots, f_1\}$, the meander of the lasso l_i bounds F_i^1 . For all $i \in \{1, \dots, n\}$, let $\mathbb{F}_{2,i} = \{F^2 \in \mathbb{F}_2 : F^2 \subset F_i^1\}$ be the set of faces of \mathbb{G}_2 contained in F_i^1 and set $f_{2,i} = \#\mathbb{F}_{2,i}$.

Then there exists a set of facial lassos $\{l_{i,j} : i \in \{1, \dots, f_1\}, j \in \{1, \dots, f_{2,i}\}\}$ in $\text{RL}_v(\mathbb{G}_2)$ such that for all i, j , the meander of the lasso $l_{i,j}$ bounds a face of $\mathbb{F}_{2,i}$, for all i the equality $l_i = l_{i,1} \dots l_{i,f_{2,i}}$ holds and

$$\{a_1, \dots, a_g, c_1, \dots, c_p, l_{1,1}, \dots, l_{f_1, f_{2, f_1}}\}$$

is a tame system of generators of $\text{RL}_v(\mathbb{G}_2)$ associated with the word w .

Until the end of this chapter we will use the notation

$$g = g(M), \quad p = p(M) \quad \text{and} \quad f = f(\mathbb{G}).$$

By Lemma 1.3.33, we know that $\text{RL}_v(\mathbb{G})$ is free of rank $e(\mathbb{G}) - v(\mathbb{G}) + 1$. By Euler's relation for the graph \mathbb{G} , which writes $v(\mathbb{G}) - e(\mathbb{G}) + f(\mathbb{G}) = \chi(M) = 2 - g - p$, the rank of $\text{RL}_v(\mathbb{G})$ can also be written as $g + p + f - 1$. In order to find families of generators

whose cardinal decomposes naturally as $g + p + f - 1$, we introduce the dual graph of G .

Definition 2.4.3. — Let M be a compact surface endowed with a graph G . Let (M', G', ι, f) be a split pattern of (M, G) . Let \widehat{V} denote the set of connected components of M' , which is in canonical bijection with \mathbb{F} . Then, set

$$\widehat{\mathbb{E}} = \{\{e', e'^{-1}\} : e' \in \mathbb{E}'\}.$$

For each element $\hat{e} = \{e', e'^{-1}\}$ of $\widehat{\mathbb{E}}$, define the source $s(\hat{e})$ of \hat{e} as the connected component of M' which contains e' and the target $t(\hat{e})$ of \hat{e} is the connected component of M' which contains $\iota(e')$.

The abstract graph $\widehat{G} = (\widehat{V}, \widehat{\mathbb{E}}, s, t)$ is called the dual graph of G .

Properly speaking, the dual graph of G depends on the choice of the split pattern of G and it is unique only up to an obvious notion of isomorphism. We shall in fact choose a split pattern and work with the associated dual graph.

The involution ι of the split pattern induces an involution on $\widehat{\mathbb{E}}$, which we still denote by ι . This involution is similar to an orientation reversal, but one should observe that it may have fixed points. For example, if an edge $e' \in \mathbb{E}'$ is sent by f on a boundary component of M , then $\iota(e') = e'$ and the dual edge $\hat{e} = \{e', e'^{-1}\}$ also satisfies $\iota(\hat{e}) = \hat{e}$.

The orbits of ι on $\widehat{\mathbb{E}}$, which we call unoriented edges of the dual graph, correspond bijectively by f with the unoriented edges of the graph G , that is, the pairs $\{e, e^{-1}\}$ for $e \in \mathbb{E}$.

Definition 2.4.4. — Let M be a compact surface endowed with a graph G . Let (M', G', ι, f) be a split pattern of (M, G) . Recall that a spanning tree of the dual graph $(\widehat{V}, \widehat{\mathbb{E}}, s, t)$ is a subset $\widehat{T} \subset \widehat{\mathbb{E}}$ stable by the involution ι and such that any two vertices of the dual graph are joined by a unique injective path made with edges of \widehat{T} .

Let \widehat{T} be a spanning tree of the dual graph. An orientation of M' is said to be adapted to \widehat{T} if for all $\hat{e} = \{e, e^{-1}\} \in \widehat{T}$, the edges e and $\iota(e)$ are neither both positively oriented nor both negatively oriented with respect to this orientation of M' .

Assume that M' is endowed with an orientation adapted to \widehat{T} . Then we define the image by f of an edge $\hat{e} = \{e', e'^{-1}\}$ of \widehat{T} as an edge of G by setting $f(\hat{e}) = f(e')$ if e' is positively oriented as a subset of $\partial M'$ and $f(\hat{e}) = f(e'^{-1})$ otherwise.

It is not difficult to check that there are exactly two orientations of M' , among the 2^f possible, which are adapted to any given spanning tree in the dual graph. Moreover, for all $\hat{e} \in \widehat{T}$, we have the equality $f(\iota(\hat{e})) = f(\hat{e})^{-1}$ in \mathbb{E} .

It follows from Proposition 1.3.15 that $\mathbb{E} \setminus f(\widehat{T})$ is the set of edges of a graph G_0 on M with a single face. This graph has $e_0 = e(G) - (f(G) - 1)$ edges, so by Euler's




FIGURE 1. The split pattern of our favourite graph as its dual graph. The edge $f(\hat{e})$ of \mathbb{G} bounds positively the face corresponding to the face of the split pattern on the boundary of which \hat{e} sits.

relation, it still has $v_0 = v(\mathbb{G})$ vertices. Moreover, its skeleton contains ∂M and every edge of \mathbb{G} which lies on ∂M is also an edge of \mathbb{G}_0 . Let $B^+ \subset \mathbb{E}_0$ be a collection of edges of \mathbb{G}_0 comprising exactly one edge on each connected component of ∂M . If M is oriented, we assume that the edges of B bound M positively. The set of edges of \mathbb{G}_0 located on ∂M and which do not belong to $B = B^+ \cup (B^+)^{-1}$ form a cycle-free subgraph of \mathbb{G}_0 . Hence, it is possible to extend this subgraph to a spanning tree T of \mathbb{G}_0 such that $T \cap B = \emptyset$ (see Figure 2).

Recall that for all $e \in \mathbb{E}_0$, the lasso $l_{e,T}$ is defined as $l_{e,T} = [v, \underline{e}]_T e [\bar{e}, v]_T$.

Lemma 2.4.5. — *With the notation above, for each edge $e \in B$, the lasso $l_{e,T}$ is a bounding lasso whose meander covers the connected component of ∂M on which e lies and whose spoke contains no edge lying on this connected component.*

Proof. — Let b be the connected component of ∂M on which e lies. If the base point v is located on b , then $l_{e,T}$ is a simple loop which represents the cycle b . Otherwise, since any two vertices located on b can be joined by a path in T which stays in b , there exists a unique vertex w on b which is joined to the base point v by a path in T with no edge lying on b . Hence, $l_{e,T} = [v, w]_T ([w, \underline{e}]_T e [\bar{e}, w]_T) [w, v]_T$. The three

paths between the brackets form the meander of the lasso, which is a simple loop covering b . \square

Let us carry on with our construction. We have a partition $\mathbb{E}_0 = T \cup B \cup R$, where $R = \mathbb{E}_0 \setminus (T \cup B)$. It follows from Euler's relation that R contains exactly g unoriented edges. From the fact that \mathbb{E}_0 has $v(\mathbb{G})$ vertices, we deduce that T is in fact a spanning tree of \mathbb{G} . Hence, Lemma 1.3.33 implies the first assertion of the following result.

Proposition 2.4.6. — *Let M be a compact surface endowed with a graph \mathbb{G} . Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . Let \widehat{T} be a spanning tree of the dual graph of \mathbb{G} . Let B^+ be a collection of edges comprising exactly one edge on each connected component of ∂M , which we assume to be positively oriented if M is oriented. Set*

$$B = B^+ \cup (B^+)^{-1}.$$

Let T be a spanning tree of $\mathbb{E} \setminus f(\widehat{T})$ such that $T \cap B = \emptyset$. Set

$$R = \mathbb{E} \setminus (f(\widehat{T}) \cup B \cup T).$$

Choose an orientation \mathbb{E}^+ of \mathbb{G} such that $B^+ \subset \mathbb{E}^+$. Write $B^+ = \{e_1, \dots, e_p\}$. Let b_1, \dots, b_p be the boundary components of M enumerated in such a way that for all $i \in \{1, \dots, p\}$, b_i is the meander of the lasso $l_{e_i, T}$. Then R^+ , B^+ and $f(\widehat{T})^+$ contain respectively g , p and $f - 1$ edges and the following properties hold.

- 1) *The group $\text{RL}_v(\mathbb{G})$ is freely generated by the loops $\{l_{e, T} : e \in R^+ \cup B^+ \cup f(\widehat{T})^+\}$.*
- 2) *Under the probability measure $\text{U}_{M, \emptyset, C}^{\mathbb{G}}$ on $\mathcal{M}(\mathbb{P}(\mathbb{G}), G)$, the collection of random variables $\{h(l_{e, T}) : e \in R^+ \cup B^+ \cup f(\widehat{T})^+\}$ is a collection of independent variables. For all $e \in R^+ \cup f(\widehat{T})^+$, the variable $h(l_{e, T})$ is uniformly distributed on G and for all $i \in \{1, \dots, p\}$, the variable $h(l_{c_i, T})$ has the distribution $\delta_C(b_i)$.*

Proof. — Let us prove the second assertion. By definition of the measure $\text{U}_{M, \emptyset, C}^{\mathbb{G}}$, the random variables $\{h(e) : e \in R^+ \cup f(\widehat{T})^+\}$ are independent, uniformly distributed on G , and independent of $\{h(e) : e \in \mathbb{E}^+ \setminus (R^+ \cup f(\widehat{T})^+)\}$. It is a general fact that if X_1, \dots, X_n are independent uniformly distributed random variables on G , independent of $Y_1, \dots, Y_n, Z_1, \dots, Z_n$, then the n variables $Y_1 X_1 Z_1, \dots, Y_n X_n Z_n$ are independent, uniformly distributed, and independent of $Y_1, \dots, Y_n, Z_1, \dots, Z_n$. Hence, the variables $\{h(l_{e, T}) : e \in R^+ \cup f(\widehat{T})^+\}$ are independent, uniformly distributed and independent of the variables $\{h(l_{e, T}) : e \in B^+\}$, which do not involve any edge of $R \cup f(\widehat{T})$.

Now choose $i \in \{1, \dots, p\}$. Consider the loop $l_{e_i, T}$ and the corresponding bounding cycle b_i . By Lemma 2.4.5, the spoke of the lasso $l_{e_i, T}$ does not involve any edge lying on the boundary component b_i of M . However, it may involve edges located on other boundary components of M .

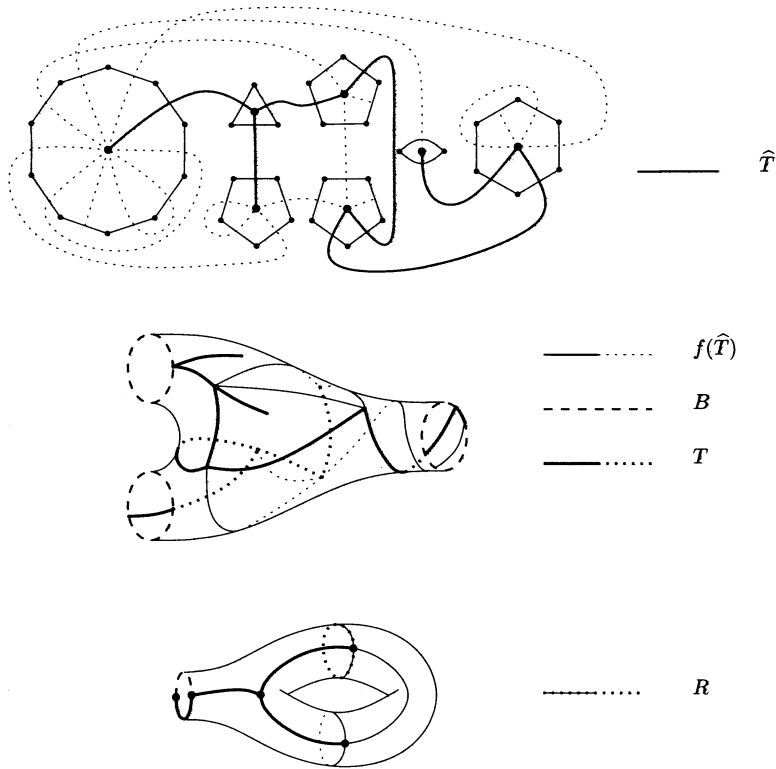


FIGURE 2. Two examples of partitions of the set of edges of a graph as $R \cup B \cup f(\widehat{T}) \cup T$. The first example is our favourite graph, and since it is a graph on a sphere, albeit with three holes, R is empty. In the much simpler second example, R contains one edge.

We claim that for every subset $I \subset \{1, \dots, p\}$, there exists $i_0 \in I$ such that for all $j \in I - \{i_0\}$, the spoke of $l_{e_j, T}$ does not involve edges located on b_{i_0} . Assume to the contrary that for some subset I there does not exist such an i_0 . Then there would exist $i_1, \dots, i_{k-1} \in I$ all distinct and $i_k = i_1$ such that for all $j \in \{1, \dots, k-1\}$, the spoke of the lasso $l_{e_{i_{j+1}}, T}$ involves an edge which lies on the boundary component b_{i_j} . This would in particular imply that for each $j \in \{1, \dots, k-1\}$ there exists a path in T from b_{i_j} to $b_{i_{j+1}}$ and, since $i_k = i_1$, that T contains a cycle.

By relabelling the boundary components of M , we may assume that for all k in $\{1, \dots, p\}$, the element k of the subset $\{k, \dots, p\} \subset \{1, \dots, p\}$ has the property described above. Since under $U_{M, \emptyset, C}^G$, the random variable $h(b_k)$ has the distribution $\delta_{C(b_k)}$ and is independent of the variables $h(e)$ for $e \in \mathbb{E}$ not located

on b_k , the variable $h(l_{e_k, T})$ itself has the distribution $\delta_{C(b_k)}$ and is independent of $h(l_{e_{k+1}, T}), \dots, h(l_{e_p, T})$. This implies easily the result. \square

With the proof of Proposition 2.4.2 in mind, the next step is to express the loops $l_{e, T}$ for $e \in f(\widehat{T})$ in function of facial lassos. The exact way in which these lassos decompose into products of facial lassos depends on, and in fact encodes completely, the geometry of the tree \widehat{T} .

Let M' be endowed with an orientation adapted with \widehat{T} . Let $\hat{v} = M'_F$ be a vertex of $\widehat{\mathbb{G}}$ which corresponds to a face F . The set of edges of $\widehat{\mathbb{G}}$ whose source is \hat{v} is in one-to-one correspondence with the set of edges of \mathbb{G}' located on the boundary of M'_F and which bound it positively. This set carries a natural cyclic order, which is the order in which the edges are traversed by a cycle bounding M'_F . By restriction, the set of edges of \widehat{T} which share \hat{v} as their source is endowed with a cyclic order.

Let us root \widehat{T} by not only choosing a reference vertex but also by choosing among the edges issued from this vertex which one is the first. The simplest way to do this is to choose a vertex of \mathbb{G}' . This determines a root vertex for \widehat{T} , namely the connected component of M' which carries this vertex, and this also breaks the cyclic symmetry of the edges issued from this connected component, which are now totally ordered.

The object that we are now contemplating is a tree (that is, an abstract graph without simple cycle) endowed with a distinguished vertex, a total order on the edges issued from this distinguished vertex and a cyclic order on the set of edges issued from any other vertex. Such an object is called a rooted planar tree and it has a canonical representation as a set of words of integers, according to a formalism due to J. Neveu. Let us simply describe how the vertices are labelled by words of integers, that is, finite sequences of non-negative integers.

The root vertex is labelled by the empty word \emptyset . Let $k(\emptyset)$ be the number of children of \emptyset , that is, the number of vertices to which it is joined by an edge. These vertices are labelled by words of length 1, namely $1, 2, \dots, k(\emptyset)$, according to the total order on these edges.

Then, consider a vertex labelled by a word $u = u_1 \dots u_n$. The integer n is called the *height* of u and it is denoted by $h(u)$. Let $\pi(u) = u_1 \dots u_{n-1}$ denote the predecessor of u and let $k(u)$ denote the number of vertices other than $\pi(u)$ to which u is joined by an edge. Then the edge from u to $\pi(u)$, denoted by $\langle u, \pi(u) \rangle$, breaks the cyclic symmetry among the edges issued from u and determines a total order on the $k(u)$ other edges issued from u . The $k(u)$ targets of these edges are labelled $u_1 \dots u_n 1, \dots, u_1 \dots u_n k(u)$ in this order.

Thus, the choice of a rooting of \widehat{T} and an adapted orientation of M' determines a labelling of \mathbb{F} by words of integers (see Figure 3). As we shall see now, it determines

also a specific facial lasso for each face of \mathbb{G} and one of the loops $l_{e,T}$ for each face distinct from the root face.

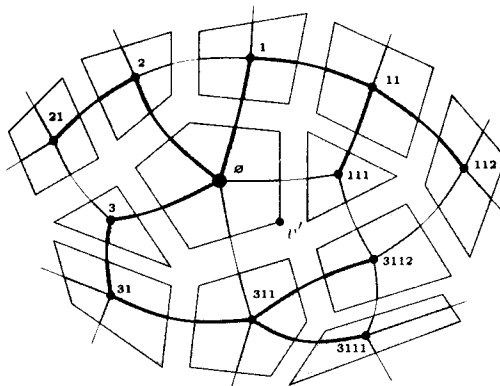


FIGURE 3. The labelling of the faces of a graph by words of integers. The vertex v' determines the rooting of \widehat{T} .

Definition 2.4.7. — Let M be a compact surface endowed with a graph \mathbb{G} . Let $(M', \mathbb{G}', \iota, f)$ be a split pattern of (M, \mathbb{G}) . Let \widehat{T} be a spanning tree of the dual graph of \mathbb{G} . Endow M' with an orientation adapted to \widehat{T} . Choose a root for \widehat{T} . Label the faces of \mathbb{G} with words of integers accordingly.

Consider a face F of \mathbb{G} labelled by the word u . If $u = \emptyset$, let $\partial_{\widehat{T}}F$ be the unique representative of $\partial M'_F$ which starts at the root vertex chosen in \mathbb{V}' and bounds positively $\partial M'_F$. Otherwise, consider the edge $\langle u, \pi(u) \rangle$ of \widehat{T} , which we identify with an edge of \mathbb{G}' which bounds M'_F positively. Let $\partial_{\widehat{T}}F$ be the unique simple loop representing $\partial M'_F$ which starts with this edge. Write $\partial_{\widehat{T}}F$ as a product of edges $e_1 \cdots e_r$. We define the facial lasso $l_u \in \text{RL}_v(\mathbb{G})$ (also denoted by l_F) by

$$l_u = l_{e_1, T} \cdots l_{e_r, T}.$$

Consider a face labelled by $u \neq \emptyset$. The edge $\langle u, \pi(u) \rangle$ of \widehat{T} determines the edge $f(\langle u, \pi(u) \rangle)$ of \mathbb{G} and we define

$$l_{u, \pi(u)} = l_{f(\langle u, \pi(u) \rangle), T}.$$

We also define $l_{\emptyset, \pi(\emptyset)}$ as the constant loop at v .

Let u and u' be two vertices of \widehat{T} . We say that u is a prefix of u' if there exists a word of integers u'' such that $u' = uu''$, the concatenation of u and u'' . Genealogically, this can be phrased by saying that u is an ancestor of u' . This is also equivalent to the fact that the unique injective path in \widehat{T} from v to u' visits u .

Lemma 2.4.8. — Recall the notation of Definition 2.4.7. Let F be a face of \mathbb{G} labelled by u . There exist $k(u)+1$ elements $t_{u,\pi(u)}, t_{u,u1}, \dots, t_{u,uk(u)}$ of the subgroup of $\text{RL}_v(\mathbb{G})$ generated by $\{l_{e,T} : e \in R \cup B\}$ such that

$$(31) \quad l_u = l_{u,\pi(u)} t_{u,\pi(u)}^{-1} l_{u1,u}^{-1} t_{u,u1}^{-1} \cdots l_{uk(u),u}^{-1} t_{u,uk(u)}^{-1}.$$

Moreover, there exists an element t of the subgroup of $\text{RL}_v(\mathbb{G})$ generated by $\{l_u : u \text{ prefix of } v\}$ and $\{l_{e,T} : e \in R \cup B\}$ such that

$$(32) \quad l_u = l_{u,\pi(u)} t.$$

In the course of the proof of this result, we use an argument of downwards induction along the branches of \widehat{T} . We say that a vertex u of \widehat{T} is a *leaf* if $k(u) = 0$. The induction argument says that if a property is satisfied by all leaves of \widehat{T} and if it is satisfied by a vertex as soon as it is satisfied by all the vertices of which it is a predecessor, then it is satisfied by all vertices.

Proof. — Let us write $\partial_{\widehat{T}} F = e_1 \cdots e_n$. By definition, e_1 is the edge of \mathbb{G} which corresponds to the edge $\langle u, \pi(u) \rangle$ of \widehat{T} . Hence, $l_{e_1,T} = l_{u,\pi(u)}$. Then, the list (e_2, \dots, e_n) consists of the images by f of the edges of \widehat{T} which join u to $u1, \dots, uk(u)$, in this order, and in the midst of which occur some edges of T and some edges of $R \cup B$. The claimed expression for l_u follows. In the case where F is the root of \widehat{T} , the edge e_1 does not play a special role, it is either the image by f of the edge which joins \emptyset to 1 in \widehat{T} , or an edge of $T \cup R \cup B$.

The second expression of l_u reduces to the first if u is a leaf. Let us now consider a vertex u and assume that (32) holds for each vertex of which u is the predecessor, that is, for $u1, \dots, uk(u)$. For every vertex v , the fact that uj is a prefix of v for some $j \in \{1, \dots, k(u)\}$ implies that u is a prefix of v . Hence, by applying (31) to u and then (32) to $u1, \dots, uk(u)$, we find that (32) holds for u .

By downwards induction along the branches of \widehat{T} , the second expression of l_u holds for all $u \in \widehat{T}$. \square

Corollary 2.4.9. — Recall the notation of Proposition 2.4.6 and Definition 2.4.7. The group $\text{RL}_v(\mathbb{G})$ is freely generated by $\{l_{e,T} : e \in R^+ \cup B^+\}$ and $\{l_u : u \neq \emptyset\}$. Moreover, under $\text{U}_{M,\emptyset,C}^{\mathbb{G}}$, the random variables $\{h(l_u) : u \neq \emptyset\}$ are mutually independent, uniformly distributed and independent of $\{h(l_{e,T}) : e \in R^+ \cup B^+\}$.

Proof. — Let us call the height of \widehat{T} and denote by $h(\widehat{T})$ the maximal height of a vertex of \widehat{T} , which is necessarily a leaf. We claim that for all integer $n \geq 1$, $\text{RL}_v(\mathbb{G})$ is freely generated by the set $X(n)$ defined by

$$X(n) = \{l_{e,T} : e \in R^+ \cup B^+\} \cup \{l_{u,\pi(u)} : h(u) < n\} \cup \{l_u : h(u) \geq n\}.$$

For $n = 1$, it is the statement which we wish to prove. For $n > h(\widehat{T})$, it is the content of Proposition 2.4.6. The fact that the claim is true for all n is easily proved by descending induction. Indeed, for all $n \geq 1$, the set $X(n)$ is deduced from $X(n+1)$ by replacing, for all labels u of \widehat{T} such that $h(u) = n$, the loop $l_{u,\pi(u)}$ by the facial lasso l_u . If $X(n+1)$ is a basis of $\text{RL}_v(\mathbb{G})$, then it follows immediately from (32) that $X(n)$ is another basis of $\text{RL}_v(\mathbb{G})$.

Since the sets $\{v : u \text{ prefix of } v\}$ are disjoint for the distinct labels u of height n , the assertion on the distribution of the variables $\{h(l) : l \in X(n)\}$ follows easily from (32) and the distribution of $\{h(l) : l \in X(n+1)\}$. \square

At this point, we have proved most of the first assertion of Proposition 2.4.2. We have exhibited a family of generators of the group $\text{RL}_v(\mathbb{G})$ which generate it freely. We know that if we add one element to this family, we get a presentation of $\text{RL}_v(\mathbb{G})$ with one relation. The next result helps us to get a relation of the form that we expect.

Lemma 2.4.8 defines $t_{u,v}$ for each ordered pair of adjacent vertices (u, v) in \widehat{T} , and also an extra element $t_{\emptyset,\pi(\emptyset)}$. Now for each pair of vertices (u, u') of \widehat{T} , not necessarily adjacent, let us define $t_{u,u'} = t_{v_1,v_2} \cdots t_{v_{m-1},v_m}$, where $u = v_1, \dots, v_m = u'$ is the unique injective path in \widehat{T} from u to u' .

Lemma 2.4.10. — *Endow the set of vertices of \widehat{T} with the lexicographic order associated to the reversed order on \mathbb{N} . Enumerate its elements accordingly:*

$$\mathbb{F} = \{u_1 \leq \cdots \leq u_f\}.$$

Then

$$(33) \quad l_{u_1} t_{u_1,u_2} l_{u_2} t_{u_2,u_3} \cdots t_{u_{f-1},u_f} l_{u_f} t_{u_f,\emptyset} t_{\emptyset,\pi(\emptyset)} = 1.$$

Moreover, the loop $(t_{u_1,u_2} t_{u_2,u_3} \cdots t_{u_{f-1},u_f} t_{u_f,\emptyset} t_{\emptyset,\pi(\emptyset)})^{-1}$ is the boundary of the unique face of \mathbb{G}_0 . It is equal to a word in the lassos $\{l_{e,T} : e \in R^+ \cup B^+\}$ where each lasso of the set $\{l_{e,T} : e \in R^+\}$ appears exactly twice, possibly sometimes with exponent -1 , and each lasso of the set $\{l_{e,T} : e \in B^+\}$ appears exactly once, possibly with exponent -1 . If M is oriented, then each lasso of the set $\{l_{e,T} : e \in B^+\}$ appears with exponent 1.

Proof. — Equation (31) can be rewritten as

$$l_{u,\pi(u)} = l_u t_{u,uk(u)} l_{uk(u),u} \cdots t_{u,u_1} l_{u_1,u} t_{u,\pi(u)}.$$

Let us apply this relation to $u = \emptyset$. We find

$$(34) \quad 1 = l_{\emptyset} t_{\emptyset,k(\emptyset)} l_{k(\emptyset),\emptyset} \cdots t_{\emptyset,1} l_{1,\emptyset} t_{\emptyset,\pi(\emptyset)}.$$




FIGURE 4. The lexicographic order on \mathbb{F} associated to the reversed order on \mathbb{N} is the same order in which the faces are discovered for the first time by an explorer who wanders around the dual tree along the track depicted above.

Let us define, for all integer $n \geq 1$ and all vertex u of \widehat{T} such that $h(u) \leq n$ an element $\tilde{l}_u^{(n)}$ of $\text{RL}_v(\mathbb{G})$ by setting

$$\tilde{l}_u^{(n)} = \begin{cases} l_u & \text{if } h(u) < n \\ l_{u, \pi(u)} & \text{if } h(u) = n. \end{cases}$$

Also, for each n , let $u_1^n \leq \dots \leq u_{r_n}^n$ be the vertices of \widehat{T} of height at most n listed in the lexicographic order corresponding to the reversed order on \mathbb{N} . Then (34) can be written as

$$1 = \tilde{l}_{u_1^1}^{(1)} t_{u_1^1, u_2^1} \tilde{l}_{u_2^1}^{(1)} \cdots t_{u_{r_1-1}^1, u_{r_1}^1} \tilde{l}_{u_{r_1}^1}^{(1)} t_{\emptyset, \pi(\emptyset)}.$$

By applying (31) iteratively to the terms of the form $l_{u, \pi(u)}$ in this equality, one finds that the equality

$$1 = \tilde{l}_{u_1^n}^{(n)} t_{u_1^n, u_2^n} \tilde{l}_{u_2^n}^{(n)} \cdots t_{u_{r_n-1}^n, u_{r_n}^n} \tilde{l}_{u_{r_n}^n}^{(n)} t_{\emptyset, \pi(\emptyset)}$$

holds for all $n \geq 1$. Here, $t_{u, u'}$ has the meaning explained before the statement of Lemma 2.4.10. For n larger than the height of \widehat{T} , this formula is exactly what we wanted to prove.

Finally, $(t_{u_1, u_2} t_{u_2, u_3} \cdots t_{u_{f-1}, u_f} t_{u_f, \emptyset} t_{\emptyset, \pi(\emptyset)})^{-1}$ is the product of the loops $l_{e, T}$ where e goes around the unique face of the pattern of \mathbb{G}_0 obtained from M' by sewing the edges of \widehat{T} . The result follows. \square

We are now ready to prove the main result of this section.

Proof of Proposition 2.4.2. — Let us consider the relation given by Lemma 2.4.10. Let us define, for all $i \in \{1, \dots, f\}$,

$$l_i = (t_{u_i, u_{i+1}} \cdots t_{u_{f-1}, u_f} t_{u_f, \emptyset} t_{\emptyset, \pi(\emptyset)})^{-1} l_{u_i} (t_{u_i, u_{i+1}} \cdots t_{u_{f-1}, u_f} t_{u_f, \emptyset} t_{\emptyset, \pi(\emptyset)}).$$

Then relation (33) becomes

$$(35) \quad (t_{u_1, u_2} \cdots t_{u_{f-1}, u_f} t_{u_f, \emptyset} t_{\emptyset, \pi(\emptyset)})^{-1} = l_1 \cdots l_f.$$

By Corollary 2.4.9 and the general result of independence invoked at the beginning of the proof of Proposition 2.4.6, $\text{RL}_v(\mathbb{G})$ is freely generated by

$$\{l_{e, T} : e \in R^+ \cup B^+\} \cup \{l_1, \dots, l_{f-1}\}$$

and, under $\text{U}_{M, \emptyset, C}^{\mathbb{G}}$, the variables $h(l_1), \dots, h(l_{f-1})$ are independent, uniformly distributed and independent of $\{h(l_{e, T}) : e \in R^+ \cup B^+\}$.

Let us write $B^+ = \{b_1, \dots, b_p\}$. By the last part of Lemma 2.4.10, there exists $p + 1$ elements t_0, \dots, t_p of the subgroup generated by $\{l_{e, T} : e \in R^+\}$, p signs $\varepsilon_1, \dots, \varepsilon_p \in \{-1, 1\}$, and a permutation $\sigma \in \mathfrak{S}_p$ such that

$$(t_{u_1, u_2} \cdots t_{u_{f-1}, u_f} t_{u_f, \emptyset} t_{\emptyset, \pi(\emptyset)})^{-1} = t_0 l_{b_{\sigma(1)}, T}^{\varepsilon_1} t_1 \cdots t_{p-1} l_{b_{\sigma(p)}, T}^{\varepsilon_p} t_p,$$

with $\varepsilon_1 = \dots = \varepsilon_p = 1$ if M is oriented.

By Lemma 2.4.10 again, the loop $t_0 \cdots t_p$ can be written as a word in $\{l_e : e \in R^+\}$ where each loop appears exactly twice, possibly sometimes with exponent -1 . Recall that $\#R^+ = g$. Let us name a_1, \dots, a_g the loops $\{l_e : e \in R^+\}$ and let w denote the element of the free group of rank g such that $w(a_1, \dots, a_g) = t_0 \cdots t_p$. Let us also define, for all $k \in \{1, \dots, p\}$, $c_i = (t_i \cdots t_p)^{-1} l_{b_{\sigma(i)}, T}^{\varepsilon_i} (t_i \cdots t_p)$. Then relation (35) becomes

$$(36) \quad w(a_1, \dots, a_g) c_1 \cdots c_p = l_1 \cdots l_f.$$

The first assertion result now from Proposition 2.4.6, Corollary 2.4.9 and the definition of l_1, \dots, l_f and c_1, \dots, c_p .

2) Let us prove the second assertion. By adding vertices to \mathbb{G}_1 , we do not change the group $\text{RL}_v(\mathbb{G}_1)$ nor the distribution of the associated random variables, according to the property of invariance by subdivision (Lemma 2.3.9). Hence, we may add to \mathbb{G}_1 the vertices of \mathbb{G}_2 located on $\text{Sk}(\mathbb{G}_1)$ and, without loss of generality, assume that $\mathbb{E}_1 \subset \mathbb{E}_2$.

For the sake of simplicity, let us treat the case where only one face of \mathbb{G}_1 contains the interior of an edge of $\mathbb{E}_2 \setminus \mathbb{E}_1$. Once the result is proved under this restrictive assumption, the general result follows by iteration. So, let us denote by F_1 the only face which contains the interior of some edges of \mathbb{G}_2 .

Let $(M', \mathbb{G}'_1, \iota, f)$ be a split pattern of (M, \mathbb{G}_1) . Let M'_{F_1} be the closure of the preimage by f of F_1 . Let \mathbb{G}'_2 be the restriction to M'_{F_1} of the lift to M' of \mathbb{G}_2 (this lift can be rigorously defined by successive applications of the second assertion of Proposition 1.3.5). Let v_1 be the finishing point of the spoke of the lasso l_1 . Let us choose a vertex v'_1 of \mathbb{G}'_2 which is sent by f on v_1 .

The first assertion of the proposition that we are proving applied to the graph \mathbb{G}'_2 on the disk M'_{F_1} , whose reduced genus is 0, provides us with f_1 facial lassos $l'_1, \dots, l'_{f_2,1}$ based at v'_1 which bound the faces of \mathbb{G}'_2 and such that $l'_1 \cdots l'_{f_2,1}$ bounds M'_2 . By projecting these lassos on M by f and conjugating them by the spoke of l_1 , we get lassos based at v which we denote by $l_{1,1}, \dots, l_{1,f_2,1}$.

Let \mathcal{G}_1 denote the tame set of generators of $\text{RL}_v(\mathbb{G}_1)$ that we are given. Let us prove that the set \mathcal{G}_2 of loops obtained by replacing l_1 by $l_{1,1}, \dots, l_{1,f_2,1}$ in \mathcal{G}_1 is a set which generates $\text{RL}_v(\mathbb{G}_2)$. Let c be a loop based at v in \mathbb{G}_2 . Let us split c into a concatenation of paths which are either paths in \mathbb{G}_1 or concatenation of edges of $\mathbb{E}_2 \setminus \mathbb{E}_1$. We get an expression of the form $c = c_1 d_1 \cdots c_n d_n$, where the paths c_1, \dots, c_n are in \mathbb{G}_1 and the paths d_1, \dots, d_n are concatenations of edges of $\mathbb{E}_2 \setminus \mathbb{E}_1$. Choose $k \in \{1, \dots, n\}$ and consider the path d_k . It can be lifted in a unique way to a path in \mathbb{G}'_2 , which we denote by d'_k . Even if d_k is a loop, d'_k needs not be a loop. However, there exist two paths a'_k and b'_k in \mathbb{G}'_2 which stay on the boundary of M'_{F_1} and such that $a'_k d'_k b'_k$ is a loop based at v'_1 . Let us write $a_k = f(a'_k)$ and $b_k = f(b'_k)$, and denote the spoke of the lasso l_1 by s_1 . Finally, let f_k be a path in \mathbb{G}_1 from \bar{d}_k to v . We have the equality in $\text{RL}_v(\mathbb{G}_2)$

$$c = (c_1 a_1^{-1} s_1^{-1}) [s_1 a_1 d_1 b_1 s_1^{-1}] (s_1 b_1^{-1} f_1) \cdots (f_{n-1}^{-1} c_n a_n^{-1} s_1^{-1}) [s_1 a_n d_n b_n s_1^{-1}] (s_1 b_n^{-1}).$$

The loops between brackets are loops of $\text{RL}_v(\mathbb{G}_1)$ and the loops between square brackets are the image by f of loops in \mathbb{G}'_2 , conjugated by the spoke of l_1 , hence, equal to words in the loops $l_{1,1}, \dots, l_{1,f_2,1}$. Hence, the loops of \mathcal{G}_2 generate $\text{RL}_v(\mathbb{G}_2)$. They satisfy the equation

$$w(a_1, \dots, a_g) c_1 \cdots c_p = l_{1,1} \cdots l_{1,f_1} l_2 \cdots l_{f_1}.$$

In particular, the set $\mathcal{G}_2 \setminus \{l_{1,1}\}$ for instance has cardinal $\mathbf{g} + \mathbf{p} + \mathbf{f}(\mathbb{G}_2) - 1$ and generates $\text{RL}_v(\mathbb{G}_2)$ which is free of rank $\mathbf{g} + \mathbf{p} + \mathbf{f}(\mathbb{G}_2) - 1$. Hence, it is a free basis of this group (see Proposition 2.7 in Chapter 1 of [29]).

There remains to determine the distribution of $\{h(l) : l \in \mathcal{G}_2\}$ under $\mathbf{U}_{M, \emptyset, C}^{\mathbb{G}_2}$. It follows from the way in which the loops $l_{1,1}, \dots, l_{1,f_2,1}$ were constructed that the random variables $h(l_{1,1}), \dots, h(l_{1,f_2,1-1})$ are independent, uniformly distributed, and independent of $\sigma(h(l) : l \in \mathcal{G}_1)$. Moreover, $h(l_1)$ is independent of $\sigma(h(l) : l \in \mathcal{G}_1 \setminus \{l_1\})$. Hence, the three σ -fields

$$\sigma(h(l_{1,1}), \dots, h(l_{1,f_2,1-1})), \quad \sigma(h(l_1)), \quad \sigma(h(l) : l \in \mathcal{G}_1 \setminus \{l_1\})$$

are independent. Since $h(l_{1,1}), \dots, h(l_{1,f_2,1-1})$ and $h(l_1)$ are uniformly distributed, and $l_{1,f_2,1} = (l_{1,1} \cdots l_{1,f_2,1-1})^{-1} l_1$, it follows that $h(l_{1,1}), \dots, h(l_{1,f_2,1})$ are uniformly distributed and independent of $\sigma(h(l) : l \in \mathcal{G}_2 \setminus \{l_{1,1}, \dots, l_{1,f_2,1}\})$. \square

CHAPTER 3

MARKOVIAN HOLONOMY FIELDS

In this chapter, which is the core of this work, we define Markovian holonomy fields and their discrete analogues. We prove in full generality that the partition functions of a discrete Markovian holonomy field do not depend on the graph in which they are computed. We then prove the first main result of this work, which asserts that any discrete Markovian holonomy field which satisfies some regularity conditions can be extended in a unique way to a Markovian holonomy field.

3.1. Definition

Recall the definitions of a marked surface (Definition 1.1.8) and of a set of G -constraints (Definition 2.3.1).

Definition 3.1.1. — *A measured marked surface with G -constraints is a quadruple $(M, \text{vol}, \mathcal{C}, C)$ where (M, \mathcal{C}) is a marked surface, vol is a smooth non-vanishing density on M and C is a set of G -constraints on (M, \mathcal{C}) .*

Two measured marked surfaces (resp. oriented measured marked surfaces) with G -constraints $(M, \text{vol}, \mathcal{C}, C)$ and $(M', \text{vol}', \mathcal{C}', C')$ are isomorphic if there exists a diffeomorphism (resp. an orientation preserving diffeomorphism) $\psi : M \rightarrow M'$ such that $\psi_\text{vol} = \text{vol}'$, ψ sends each curve of \mathcal{C} to a curve of \mathcal{C}' and, for all $l \in \mathcal{C}$,*

$$C'(\psi(l)) = C(l).$$

From now on, we will make the assumption that G is a compact Lie group, not necessarily connected. If $\dim G \geq 1$, we endow G with a bi-invariant Riemannian metric which we normalise in such a way that the total Riemannian volume of G is 1. In this way, the Riemannian density coincides with the normalised Haar measure on G . If G is finite, we endow it with the distance $d_G(x, y) = \delta_{x,y}$ and with the uniform probability measure.

The set $\text{Conj}(G)$ of conjugacy classes of G inherits the quotient topology from G and the corresponding Borel σ -field. The space of G -constraints on a marked surface (M, \mathcal{C}) , which we denote by $\text{Const}_G(M, \mathcal{C})$, is a subset of $\text{Conj}(G)^{\mathcal{C} \cup \mathcal{B}(M)}$ and thus carries the trace topology and σ -field.

Let us introduce a last piece of notation. Let $(M, \text{vol}, \mathcal{C}, C)$ be a measured marked surface with G -constraints. Let l be a curve which belongs to $\mathcal{C} \cup \mathcal{B}(M)$. Let x be an element of G . We define a new set of G -constraints $C_{l \rightarrow x}$ as follows:

$$C_{l \rightarrow x}(l) = \vartheta_x, \quad C_{l \rightarrow x}(l^{-1}) = \vartheta_{x^{-1}} \quad \text{and} \quad C_{l \rightarrow x} = C \text{ on } \mathcal{C} \cup \mathcal{B}(M) - \{l, l^{-1}\}.$$

The main definition of this work is the following.

Definition 3.1.2. — *A G -valued two-dimensional Markovian holonomy field is the data, for each quadruplet $(M, \text{vol}, \mathcal{C}, C)$ consisting of a marked surface endowed with a density and a set of G -constraints, of a finite measure $\text{HF}_{M, \text{vol}, \mathcal{C}, C}$ on $(\mathcal{M}(\text{P}(M), G), \mathcal{J})$ such that the following properties are satisfied.*

A₁) For all $(M, \text{vol}, \mathcal{C}, C)$, $\text{HF}_{M, \text{vol}, \mathcal{C}, C}(\exists l \in \mathcal{C} \cup \mathcal{B}(M), h(l) \notin C(l)) = 0$.

A₂) For all $(M, \text{vol}, \mathcal{C})$ and all event Γ in the invariant σ -field of $\mathcal{M}(\text{P}(M), G)$, the function $C \mapsto \text{HF}_{M, \text{vol}, \mathcal{C}, C}(\Gamma)$ is a measurable function on the set $\text{Const}_G(M, \mathcal{C})$.

A₃) For all $(M, \text{vol}, \mathcal{C}, C)$ and all $l \in \mathcal{C}$,

$$\text{HF}_{M, \text{vol}, \mathcal{C} - \{l, l^{-1}\}, C|_{\mathcal{B}(M) \cup \mathcal{C} \setminus \{l, l^{-1}\}}} = \int_G \text{HF}_{M, \text{vol}, \mathcal{C}, C_{l \rightarrow x}} dx.$$

A₄) Let $\psi : (M, \text{vol}, \mathcal{C}, C) \rightarrow (M', \text{vol}', \mathcal{C}', C')$ be a homeomorphism such that $\text{vol} \circ \psi^{-1} = \text{vol}'$, $\psi(\mathcal{C}) = \mathcal{C}'$ and $C = C' \circ \psi$. Let l_1, \dots, l_n be loops based at the same point on M . Assume that their images l'_1, \dots, l'_n by ψ are also rectifiable loops. Then, for all continuous function $f : G^n \rightarrow G$ invariant under the diagonal action of G by conjugation,

$$\begin{aligned} \int_{\mathcal{M}(\text{P}(M), G)} f(h(l_1), \dots, h(l_n)) \text{HF}_{M, \text{vol}, \mathcal{C}, C}(dh) \\ = \int_{\mathcal{M}(\text{P}(M'), G)} f(h'(l'_1), \dots, h'(l'_n)) \text{HF}_{M', \text{vol}', \mathcal{C}', C'}(dh'). \end{aligned}$$

In particular, if ψ is a diffeomorphism, then the mapping from $\mathcal{M}(\text{P}(M'), G)$ to $\mathcal{M}(\text{P}(M), G)$ induced by ψ sends the measure $\text{HF}_{M', \text{vol}', \mathcal{C}', C'}$ to $\text{HF}_{M, \text{vol}, \mathcal{C}, C}$.

A₅) For all $(M_1, \text{vol}_1, \mathcal{C}_1, C_1)$ and $(M_2, \text{vol}_2, \mathcal{C}_2, C_2)$, one has the identity

$$\text{HF}_{M_1 \cup M_2, \text{vol}_1 \cup \text{vol}_2, \mathcal{C}_1 \cup \mathcal{C}_2, C_1 \cup C_2} = \text{HF}_{M_1, \text{vol}_1, \mathcal{C}_1, C_1} \otimes \text{HF}_{M_2, \text{vol}_2, \mathcal{C}_2, C_2},$$

where $M_1 \cup M_2$ denotes the disjoint union of M_1 and M_2 , and $\text{vol}_1 \cup \text{vol}_2$, $\mathcal{C}_1 \cup \mathcal{C}_2$ and $C_1 \cup C_2$ denote respectively the natural density, set of marks, and set of G -constraints on $M_1 \cup M_2$.

A₆) For all $(M, \text{vol}, \mathcal{C}, C)$, all $l \in \mathcal{C}$ and all gluing $\psi : \text{Spl}_l(M) \rightarrow M$ along l , one has

$$\text{HF}_{\text{Spl}_l(M), \text{Spl}_l(\text{vol}), \text{Spl}_l(\mathcal{C}), \text{Spl}_l(C)} = \text{HF}_{M, \text{vol}, \mathcal{C}, C} \circ \psi^{-1}.$$

A₇) For all $(M, \text{vol}, \emptyset, C)$ and for all $l \in \mathcal{B}(M)$,

$$\int_G \text{HF}_{M, \text{vol}, \emptyset, C_{l \rightarrow x}}(1) dx = 1.$$

A G -valued two-dimensional oriented Markovian holonomy field is the data, for each quadruplet $(M, \text{vol}, \mathcal{C}, C)$ consisting of an oriented marked surface endowed with a density and a set of G -constraints, of a finite measure $\text{HF}_{M, \text{vol}, \mathcal{C}, C}$ on the measurable space $(\mathcal{M}(\text{P}(M), G), \mathcal{J})$ such that the seven properties above are satisfied.

It is important to notice that the measures HF are *not* probability measures in general. They are finite measures, whose masses carry a lot of information about the holonomy field. It is actually possible that they characterise it completely, but this is a question which has yet to be answered.

Let us discuss briefly the significance of these axioms. The axioms A₁, A₂ and A₃ express the fact that the measure $\text{HF}_{M, \text{vol}, \mathcal{C}, C}$, seen as a function of the G -constraints, provides a regular disintegration of $\text{HF}_{M, \text{vol}, \emptyset, C_{|\mathcal{B}(M)}}$ with respect to the value of the holonomy field along the curves of \mathcal{C} . The simple expression of A₃ is permitted by the fact that we consider finite measures rather than probability measures. The axiom A₄ expresses the invariance of the field under area-preserving diffeomorphisms. The axioms A₅ and A₆ express the Markov property of the field. Finally, A₇ is a normalization axiom. Without it, if HF was a given Markovian holonomy field, then for any real α , the measures $e^{\alpha \text{vol}(M)} \text{HF}_{M, \text{vol}, \mathcal{C}, C}$ would constitute another Markovian holonomy field.

Our purpose is not to study Markovian holonomy fields in full generality. In the rest of this paper, we are going to make strong regularity assumptions and investigate the corresponding fields.

Recall that d_G denotes a bi-invariant distance on G . Let c and c' be two paths with the same endpoints. Then, although h_c and $h_{c'}$ are not measurable with respect with the invariant σ -field \mathcal{J} , unless c and c' are loops, the function $h \mapsto d_G(h(c), h(c'))$ is \mathcal{J} -measurable, because it can also be written as $h \mapsto d_G(1, h(c)^{-1}h(c')) = d_G(1, h(c'c^{-1}))$ and $c'c^{-1}$ is a loop.

Definition 3.1.3. — Let HF be a G -valued two-dimensional Markovian holonomy field.

1) We say that HF is stochastically continuous if, for all $(M, \text{vol}, \mathcal{C}, C)$ and for all sequence $(c_n)_{n \geq 0}$ of elements of $\text{P}(M)$ which converges with fixed endpoints to $c \in \text{P}(M)$,

$$\int_{\mathcal{M}(\text{P}(M), G)} d_G(h(c_n), h(c)) \text{HF}_{M, \text{vol}, \mathcal{C}, C}(dh) \xrightarrow{n \rightarrow \infty} 0.$$

2) We say that HF is Fellerian if, for all $(M, \text{vol}, \mathcal{C})$, the function

$$(t, C) \mapsto \text{HF}_{M, \text{tvol}, \mathcal{C}, C}(1)$$

defined on $\mathbb{R}_+^* \times \text{Const}_G(M, \mathcal{C})$ is continuous.

3) We say that HF is regular if it is both stochastically continuous and Fellerian.

In the definition of stochastic continuity, we use L^1 convergence of G -valued random variables. Since G is compact, this is equivalent to convergence in probability and to convergence in L^p for any $p \in [1, +\infty)$.

From now on, the expression (regular) Markovian holonomy field will stand for two-dimensional G -valued (regular) Markovian holonomy field.

3.2. Discrete Markovian holonomy fields

It is not easy to construct a Markovian holonomy field. Indeed, one has to construct a stochastic process indexed by loops. To do this, one must naturally specify the finite-dimensional marginals of this process. Thus, to each finite collection of loops, one has to associate a probability measure on some power of G . But unlike points on a time interval, loops on a surface may form a very complicated picture. In fact, in most cases, it is impossible to determine a probability measure from a finite set of loops. A way around this problem is to start by describing a restriction of the process to a class of loops which are nice enough, like piecewise geodesic loops, and then to extend the process by continuity. In fact, the very first step is to build a process indexed by the set of loops in a graph for every graph on a surface. This is what we call a discrete holonomy field. Recall the partial order on the set of graphs (Definition 1.4.5).

Definition 3.2.1. — A G -valued two-dimensional discrete Markovian holonomy field is the data, for each measured marked surface $(M, \text{vol}, \mathcal{C}, C)$ with G -constraints, and each graph \mathbb{G} on (M, \mathcal{C}) , of a finite measure $\text{DF}_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}$ on $(\mathcal{M}(\text{P}(\mathbb{G}), G), \mathcal{I})$ such that the following properties hold.

D₁) For all $(M, \text{vol}, \mathcal{C}, C)$ and all \mathbb{G} ,

$$\text{DF}_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}(\exists l \in \mathcal{C} \cup \mathcal{B}(M), h(l) \notin C(l)) = 0.$$

D₂) For all $(M, \text{vol}, \mathcal{C})$, all \mathbb{G} and all event Γ in the invariant σ -field of $\mathcal{M}(\text{P}(\mathbb{G}), G)$, $C \mapsto \text{DF}_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}(\Gamma)$ is a measurable function on $\text{Const}_G(M, \mathcal{C})$.

D₃) For all $(M, \text{vol}, \mathcal{C}, C)$, all \mathbb{G} and all $l \in \mathcal{C}$,

$$\text{DF}_{M, \text{vol}, \mathcal{C} - \{l, l^{-1}\}, C_{|\mathcal{B}(M) \cup \mathcal{C} - \{l, l^{-1}\}}}^{\mathbb{G}} = \int_G \text{DF}_{M, \text{vol}, \mathcal{C}, C_{l \rightarrow z}}^{\mathbb{G}} dz.$$

D₄) Consider $(M, \text{vol}, \mathcal{C}, C)$ and $(M', \text{vol}', \mathcal{C}', C')$, endowed respectively with \mathbb{G} and \mathbb{G}' . Let $\psi : M \rightarrow M'$ be a homeomorphism. Assume that $\text{vol} \circ \psi^{-1}$ and vol' give

the same area to each face of \mathbb{G}' , that $\psi(\mathcal{C}) = \mathcal{C}'$, $C' \circ \psi = C$ and $\psi(\mathbb{G}) = \mathbb{G}'$. Then the mapping from $\mathcal{M}(\mathbb{P}(\mathbb{G}'), G)$ to $\mathcal{M}(\mathbb{P}(\mathbb{G}), G)$ induced by ψ satisfies

$$\mathrm{DF}_{M', \mathrm{vol}', \mathcal{C}', C'}^{\mathbb{G}'} \circ \psi^{-1} = \mathrm{DF}_{M, \mathrm{vol}, \mathcal{C}, C}^{\mathbb{G}}.$$

D₅) For all $(M_1, \mathrm{vol}_1, \mathcal{C}_1, C_1)$ and $(M_2, \mathrm{vol}_2, \mathcal{C}_2, C_2)$, endowed respectively with two graphs \mathbb{G}_1 and \mathbb{G}_2 one has the identity

$$\mathrm{DF}_{M_1 \cup M_2, \mathrm{vol}_1 \cup \mathrm{vol}_2, \mathcal{C}_1 \cup \mathcal{C}_2, C_1 \cup C_2}^{\mathbb{G}_1 \cup \mathbb{G}_2} = \mathrm{DF}_{M_1, \mathrm{vol}_1, \mathcal{C}_1, C_1}^{\mathbb{G}_1} \otimes \mathrm{DF}_{M_2, \mathrm{vol}_2, \mathcal{C}_2, C_2}^{\mathbb{G}_2}.$$

D₆) For all $(M, \mathrm{vol}, \mathcal{C}, C)$, all \mathbb{G} , all $l \in \mathcal{C}$ and all gluing $\psi : \mathrm{Spl}_l(M) \rightarrow M$ along l , one has

$$\mathrm{DF}_{\mathrm{Spl}_l(M), \mathrm{Spl}_l(\mathrm{vol}), \mathrm{Spl}_l(\mathcal{C}), \mathrm{Spl}_l(C)}^{\mathrm{Spl}_l(\mathbb{G})} = \mathrm{DF}_{M, \mathrm{vol}, \mathcal{C}, C}^{\mathbb{G}} \circ \psi^{-1}.$$

D₇) For all $(M, \mathrm{vol}, \emptyset, C)$, all \mathbb{G} and all $l \in \mathcal{B}(M)$,

$$\int_G \mathrm{DF}_{M, \mathrm{vol}, \emptyset, C_{l \rightarrow x}}^{\mathbb{G}}(1) dx = 1.$$

D_I) For all $(M, \mathrm{vol}, \mathcal{C}, C)$, all \mathbb{G}_1 and \mathbb{G}_2 graphs on (M, \mathcal{C}) such that $\mathbb{G}_1 \preceq \mathbb{G}_2$, the restriction map $r : \mathcal{M}(\mathbb{P}(\mathbb{G}_2), G) \rightarrow \mathcal{M}(\mathbb{P}(\mathbb{G}_1), G)$ satisfies

$$\mathrm{DF}_{M, \mathrm{vol}, \mathcal{C}, C}^{\mathbb{G}_2} \circ r^{-1} = \mathrm{DF}_{M, \mathrm{vol}, \mathcal{C}, C}^{\mathbb{G}_1}.$$

A G -valued two-dimensional discrete oriented Markovian holonomy field is the data, for each oriented measured marked surface $(M, \mathrm{vol}, \mathcal{C}, C)$ with G -constraints and all graph \mathbb{G} on (M, \mathcal{C}) , of a finite measure $\mathrm{DF}_{M, \mathrm{vol}, \mathcal{C}, C}^{\mathbb{G}}$ on $(\mathcal{M}(\mathbb{P}(\mathbb{G}), G), \mathcal{I})$ such that the properties above are satisfied.

The axiom D_I is specific to discrete holonomy fields. It is an axiom of consistency and is usually called the property of *invariance under subdivision*.

Lemma 3.2.2. — *Any (oriented) Markovian holonomy field determines a discrete (oriented) Markovian holonomy field.*

Proof. — For all surface M endowed with a graph \mathbb{G} , the set $\mathbb{P}(\mathbb{G})$ is a subset of $\mathbb{P}(M)$. Hence, a Markovian holonomy field determines by restriction a discrete Markovian holonomy field, except perhaps for the axiom D_I. In fact, this axiom is satisfied by the restriction of a Markovian holonomy field because the finite dimensional marginals of a stochastic process constitute a consistent system of probability measures. \square

Our main goal in this chapter is to prove a result in the other direction and to extend, when it is possible, a discrete holonomy field to a continuous one. For the moment, let us discuss briefly discrete holonomy fields in themselves. It turns out that we have already constructed a fundamental example of a discrete Markovian holonomy field.

Proposition 3.2.3. — *The family of measures $U_{M,\mathcal{C},C}^G$ is a discrete Markovian holonomy field. We call it the uniform G -valued discrete holonomy field.*

Proof. — The measure $U_{M,\mathcal{C},C}^G$ is a probability measure on the cylinder σ -field of $\mathcal{M}(P(G), G)$. By restriction, it defines a measure on the invariant σ -field. The axiom D_1 is satisfied by Proposition 2.3.6. By Proposition 2.1.8 and the definition of the measures $\delta_{\theta(n)}$, it is possible to write $U_{M,\mathcal{C},C}^G(\Gamma)$ as an expression which is explicitly measurable with respect to the G -constraints. Hence, the axiom D_2 is satisfied. Axiom D_3 is satisfied by (25) (Lemma 2.3.4). That axioms D_4 , D_5 and D_7 hold is straightforward. Axiom D_6 is satisfied thanks to Proposition 2.3.8. The invariance property D_I holds by Proposition 2.3.9. \square

This discrete Markovian holonomy field is very special in that it consists only in probability measures.

Definition 3.2.4. — *Let DF be a collection of finite measures as in Definition 3.2.1, which does not necessarily satisfy any of the axioms listed in this definition. To each $(M, \mathcal{C}, \text{vol}, C)$ and each graph \mathbb{G} , we associate the number*

$$Z_{M,\text{vol},\mathcal{C},C}^{\mathbb{G}} = DF_{M,\text{vol},\mathcal{C},C}^{\mathbb{G}}(1),$$

which is called the partition function.

We have said earlier that these numbers carry a lot of information about the field. A crucial property is that they do not depend on the graph \mathbb{G} .

Proposition 3.2.5. — *Let DF be a collection of finite measures as in Definition 3.2.1, which satisfies axioms D_4 , D_5 , D_6 and D_I . Consider $(M, \text{vol}, \mathcal{C}, C)$ and two graphs \mathbb{G}_1 and \mathbb{G}_2 on (M, \mathcal{C}) . Then*

$$(37) \quad Z_{M,\text{vol},\mathcal{C},C}^{\mathbb{G}_1} = Z_{M,\text{vol},\mathcal{C},C}^{\mathbb{G}_2}.$$

The common value of all $Z_{M,\text{vol},\mathcal{C},C}^{\mathbb{G}}$, where \mathbb{G} spans the set of graphs on (M, \mathcal{C}) , is denoted by $Z_{M,\text{vol},\mathcal{C},C}$.

If there exists a graph \mathbb{G}_3 such that $\mathbb{G}_1 \preceq \mathbb{G}_3$ and $\mathbb{G}_2 \preceq \mathbb{G}_3$, then (37) is a straightforward consequence of the axiom D_I . Unfortunately, by Lemma 1.4.6, the set of graphs on M is not directed and such a graph \mathbb{G}_3 does not always exist.

Proof. — By Proposition 1.1.12, it is possible to split M along each of the curves of \mathcal{C} . Thus, there exists a measured surface $(M', \text{vol}', \emptyset, C')$ with G -constraints and no marks and a gluing $f : M' \rightarrow M$ whose joint is \mathcal{C} . We can lift \mathbb{G}_1 to a graph \mathbb{G}'_1 on M' (see the comment before Lemma 1.3.6) and the axiom D_6 enforces the equality $Z_{M',\text{vol}',\emptyset,C'}^{\mathbb{G}'_1} = Z_{M,\text{vol},\mathcal{C},C}^{\mathbb{G}_1}$. Since we can do the same for the graph \mathbb{G}_2 , it suffices to prove the result when $\mathcal{C} = \emptyset$. Then, the axiom D_5 implies that the partition function

associated to \mathbb{G}_1 (resp. \mathbb{G}_2) is the product of the partition functions associated to the connected components of M endowed with the corresponding restrictions of \mathbb{G}_1 (resp. \mathbb{G}_2).

Hence, it suffices to prove the result when M is connected and $\mathcal{C} = \emptyset$. In this case, the axiom D_I implies that we can remove or add edges to \mathbb{G}_1 , in the sense of Propositions 1.3.15 and 1.3.26, without altering the partition function.

By Proposition 1.3.36, by such transformations we can go from \mathbb{G}_1 to a graph which is sent to \mathbb{G}_2 by a homeomorphism of M . By D_4 , this implies the result. \square

Let DF be a discrete holonomy field. In order to produce a Markovian holonomy field from DF , the first natural step is to apply Proposition 2.2.2 to put together the measures which DF associates to a directed subset of the set of graphs on a surface. For this, we need to consider Riemannian metrics.

Let $(M, \text{vol}, \gamma, \mathcal{C})$ be a Riemannian marked surface (see Definition 1.4.4). Axiom D_I and Proposition 2.2.3 imply that the collection of the measures $DF_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}$, where \mathbb{G} spans the set $\text{Gr}_{\gamma}(M, \mathcal{C})$ of graphs with piecewise geodesic edges, determines a finite measure on $\mathcal{M}(\mathcal{A}_{\gamma}(M), G)$, where $\mathcal{A}_{\gamma}(M)$ is the set of piecewise geodesic paths on M .

Definition 3.2.6. — *Let DF be a discrete holonomy field. Let $(M, \text{vol}, \gamma, \mathcal{C}, C)$ be a Riemannian measured marked surface with G -constraints. The finite measure on $\mathcal{M}(\mathcal{A}_{\gamma}(M), G)$ obtained by taking the projective limit of the measures $DF_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}$ where \mathbb{G} spans $\text{Gr}_{\gamma}(M, \mathcal{C})$ is denoted by $DF_{M, \text{vol}, \mathcal{C}, C}^{\gamma}$.*

At this point, there are two things to do. Firstly, one needs to extend the measure $DF_{M, \text{vol}, \mathcal{C}, C}^{\gamma}$ to $\mathcal{M}(\text{P}(M), G)$ for all $(M, \text{vol}, \mathcal{C}, C)$ and secondly, one needs to prove that the result of this procedure would have been the same with another Riemannian metric. The first step requires some regularity from the holonomy field. We denote by 1 the unit element of G .

Definition 3.2.7. — *Let DF be a discrete holonomy field. We say that DF is stochastically $\frac{1}{2}$ -Hölder continuous if the following property holds.*

Let $(M, \text{vol}, \gamma, \mathcal{C}, C)$ be a Riemannian marked surface with G -constraints. Then there exists a constant $K > 0$ such that for all graph \mathbb{G} on M and all simple loop $l \in \text{P}(\mathbb{G})$ with $l(l) \leq K^{-1}$ bounding a disk D ,

$$\int_{\mathcal{M}(\text{P}(\mathbb{G}), G)} d_G(1, h(l)) DF_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}(dh) \leq K \sqrt{\text{vol}(D)}.$$

Our unique example so far of a discrete holonomy field, the uniform holonomy field, does unfortunately not satisfy this property. Indeed, it assigns uniform random variables to arbitrary small simple loops. We will construct other more regular discrete holonomy fields later.

The second step of the program outlined above requires another regularity condition, which is more global but less quantitative.

Definition 3.2.8. — *Let DF be a discrete holonomy field.*

1) *We say that DF is continuously area-dependent if the following property holds.*

Let $(M, \text{vol}, \mathcal{C}, C)$ be a marked surface with G -constraints. Let $(\text{vol}_n)_{n \geq 1}$ be a sequence of surface measures on M . Assume that for all $F \in \mathbb{F}$, $\text{vol}_n(F)$ converges to $\text{vol}(F)$. Then we have the weak convergence

$$DF_{M, \text{vol}_n, \mathcal{C}, C}^G \xrightarrow[n \rightarrow \infty]{} DF_{M, \text{vol}, \mathcal{C}, C}^G.$$

2) *We say that that DF is Fellerian if for all measured marked surface $(M, \text{vol}, \mathcal{C})$, the mapping which to a set of G -constraints C associates the number $Z_{M, \text{vol}, \mathcal{C}, C}$ is continuous on $\text{Const}_G(M, \mathcal{C})$.*

3) *We say that DF is regular if it is stochastically $\frac{1}{2}$ -Hölder continuous, continuously area-dependent and Fellerian.*

It is tempting to conjecture that a stochastically $\frac{1}{2}$ -Hölder continuous discrete Markovian holonomy field is continuously area-dependent. At least, the two properties are not equivalent, as our unique example so far shows. For the moment, let us state our main result concerning the construction of Markovian holonomy fields.

Theorem 3.2.9. — *Every regular discrete Markovian holonomy field is the restriction of a unique regular Markovian holonomy field.*

The proof of this theorem occupies the rest of this chapter.

3.3. An abstract extension theorem

The core of the proof of Theorem 3.2.9 is the next result, which we formulate in a way which is mostly independent of the context of Markovian holonomy fields. Recall the notation $A_\gamma(M)$ for the set of piecewise geodesic paths on M (Definition 1.2.10).

Theorem 3.3.1. — *Let (M, γ) be a compact Riemannian surface. Let vol denote the Riemannian volume of γ . Let (Γ, d) be a complete metric group on which translations and inversion are isometries. Let $H \in \mathcal{M}(A_\gamma(M), \Gamma)$ be a multiplicative function. Assume that there exists $K > 0$ such that for all simple loop $l \in A_\gamma(M)$ bounding a disk D and such that $\ell(l) \leq K^{-1}$, the inequality $d(1, H(l)) \leq K\sqrt{\text{vol}(D)}$ holds.*

Then H admits a unique extension to an element of $\mathcal{M}(P(M), G)$, also denoted by H , such that if a sequence $(c_n)_{n \geq 0}$ of paths converges with fixed endpoints to a path c , then $H(c_n) \xrightarrow[n \rightarrow \infty]{} H(c)$.

Let us state right now the application of this theorem to the extension of holonomy fields.

Corollary 3.3.2. — *Let DF be a stochastically $\frac{1}{2}$ -Hölder continuous discrete Markovian holonomy field. Let $(M, \text{vol}, \gamma, \mathcal{E}, C)$ be a Riemannian marked surface with G -constraints. There exists a unique probability measure $DF_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)}$ on $(\mathcal{M}(P(M), G), \mathcal{E})$ whose image by the restriction mapping*

$$\mathcal{M}(P(M), G) \longrightarrow \mathcal{M}(A_\gamma(M), G)$$

is the measure $DF_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)}$ and such that for all sequence $(c_n)_{n \geq 0}$ of paths which converges with fixed endpoints to a path c , one has

$$\int_{\mathcal{M}(P(M), G)} d_G(h(c_n), h(c)) DF_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)}(dh) \xrightarrow{n \rightarrow \infty} 0.$$

Proof. — The canonical process $(H_c)_{c \in A_\gamma(M)}$ on $(\mathcal{M}(A_\gamma(M), G), \mathcal{E}, DF_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)})$ can be seen as a mapping from $A_\gamma(M)$ to the set $\Gamma = L^1(\mathcal{M}(A_\gamma(M), G), \mathcal{E}, DF_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)}; G)$. In general, the set $L^1(\Omega, \mathcal{U}, \mathbb{P}; G)$ of G -valued random variables on any probability space $(\Omega, \mathcal{U}, \mathbb{P})$ is a group for the multiplication of random variables. The metric $d(H_1, H_2) = \mathbb{E}[d_G(H_1, H_2)]$ makes it a complete metric space. The two structures are compatible in that the product and inverse mappings are continuous. Moreover, the translations and the inversion map are isometries.

The assumption of stochastic Hölder continuity ensures that the regularity condition is satisfied by the family $(H_c)_{c \in A_\gamma(M)}$. Hence, we can apply Theorem 3.3.1. It produces a family of random variables $(H_c)_{c \in P(M)}$ which is multiplicative and continuous with respect to the convergence of paths with fixed endpoints. Proposition 2.2.3 applied to this family asserts that there exists a probability measure on $(\mathcal{M}(P(M), G), \mathcal{E})$ under which the canonical process has the distribution of $(H_c)_{c \in P(M)}$.

The uniqueness of the measure $DF_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)}$ follows from Proposition 1.2.12 which asserts that $A_\gamma(M)$ is dense in $P(M)$ for the convergence with fixed endpoints. \square

The rest of this section is devoted to the proof of Theorem 3.3.1. A basic tool for this proof is the lasso decomposition of a piecewise geodesic path (see Proposition 1.4.9).

Proposition 3.3.3. — *Under the assumptions of Theorem 3.3.1, there exists a constant $K > 0$ such that for every loop $l \in A_\gamma(M)$ with $\ell(l) \leq K^{-1}$,*

$$d(1, H(l)) \leq K\ell(l).$$

Proof. — Assume first that l is a simple loop. If $\ell(l)$ is small enough, then l bounds a disk, which we denote by D . Moreover, a local isoperimetric inequality holds on M : if $\ell(l)$ is small enough, say $\ell(l) < L$, then $\sqrt{\text{vol}(D)} \leq K_1\ell(l)$ for some constant

$K > 0$. Then, if K denotes the constant given by the assumptions of Theorem 3.3.1, $d(1, H(l)) \leq K\sqrt{\text{vol}(\overline{D})} \leq KK_1\ell(l)$. The result is proved in this case.

Let us now treat the general case. Let us apply Proposition 1.4.9 to find the lasso decomposition $l = l_1 \cdots l_p d$ of l . By the multiplicativity of H , which is part of the assumptions of Theorem 3.3.1, $H(l) = H(d)H(l_p) \cdots H(l_1)$. Since the distance d on Γ is invariant by left translations and inversion, we have, for all $x, y \in \Gamma$,

$$d(1, xy) = d(x^{-1}, y) \leq d(1, x^{-1}) + d(1, y) = d(1, x) + d(1, y).$$

Hence, $d(1, H(l)) \leq d(1, H(d)) + \sum_{i=1}^p d(1, H(l_i)) \leq K\ell(d) + K \sum_{i=1}^p \ell(l_i)$. By Proposition 1.4.9, $\ell(d) + \sum_{i=1}^p \ell(l_i) \leq \ell(l)$ and the result follows. \square

This result tells us that if l is a simple piecewise geodesic loop close to the constant loop, then $H(l)$ is close to the image by H of the constant loop. Our next generalises this statement to the case of a piecewise geodesic path which is close to a geodesic segment.

Proposition 3.3.4. — *Under the assumptions of Theorem 3.3.1, there exists a constant $K > 0$ such that the following property holds. Let $s \in A_\gamma(M)$ be a segment of minimizing geodesic. Let c be a piecewise geodesic path with the same endpoints as s . Assume that $\ell(c) \leq K^{-1}$. Then*

$$d(H(c), H(s)) \leq K\ell(c)^{\frac{3}{4}} |\ell(c) - \ell(s)|^{\frac{1}{4}}.$$

The assumptions imply that s is shorter than c . If we take $\ell(s) = 0$, then c is a loop and we recover the estimate of Proposition 3.3.3.

In order to prove this proposition, we use a decomposition result similar to the lasso decomposition and the following isoperimetric inequality. Needless to say, the constant $\pi/3\sqrt{2}$ which appears in this Euclidean case is not optimal.

Lemma 3.3.5. — *Let $R, r \geq 0$ be real numbers.*

Let J be a rectifiable Jordan curve of length $2R + r$ in the Euclidean plane such that a piece of this curve is a segment of length R . Then the area A of the bounded component of $\mathbb{R}^2 \setminus J$ satisfies the inequality

$$A \leq \frac{\pi}{3\sqrt{2}} r^{\frac{1}{2}} (R + r)^{\frac{3}{2}}.$$

Let M be a compact Riemannian surface. There exists a constant $K > 0$ such that the following property holds. Let J be a rectifiable Jordan curve of length $2R + r < K^{-1}$ such that a piece of this curve is a segment of minimizing geodesic of length R . Then the area A of the smallest disk bounded by J satisfies the inequality

$$A \leq Kr^{\frac{1}{2}} (R + r)^{\frac{3}{2}}.$$




FIGURE 1. The generic case and two optimal cases.

Proof. — The Riemannian case can easily be deduced from the Euclidean case by working in normal coordinates at one end point of the geodesic part of J . The compactness of M ensures that the resulting distortions of lengths and areas are bounded.

Let us consider the Euclidean case. Under the constraints on J , A is maximal when J is the concatenation of a segment of length R and an arc of circle of length $R + r$. In this case, let ρ be the radius of this circle and $\alpha \in [0, 2\pi)$ the angle under which the arc of circle is seen from the centre of the circle. Then $R + r = \rho\alpha$ and $R = 2\rho \sin \frac{1}{2}\alpha$. Now one has the relations

$$(38) \quad A = \frac{\rho^2}{2}(\alpha - \sin \alpha) \leq \frac{\rho^2 \alpha^3}{12} = \frac{\alpha}{12}(R + r)^2,$$

$$(39) \quad \frac{\sin \frac{1}{2}\alpha}{\frac{1}{2}\alpha} = \frac{R}{R + r}.$$

One checks easily that for all $x \in [0, \pi]$, $\sin x/x \leq 1 - x^2/\pi^2$. Hence, (39) implies $\alpha \leq 2\pi\sqrt{r/R}$, and since $\alpha \in [0, 2\pi)$, we have $\alpha \leq 2\pi(\sqrt{r/R} \wedge 1)$. Combining this with (38), we find

$$A \leq \frac{\pi}{6}(R + r)^{\frac{3}{2}}\sqrt{r}(\sqrt{1 + r/R} \wedge \sqrt{1 + R/r}) \leq \frac{\pi}{3\sqrt{2}}r^{\frac{1}{2}}(R + r)^{\frac{3}{2}}.$$

This is the expected result. □

The generalization of the lasso decomposition that we need is the following.

Proposition 3.3.6. — *Let M be a Riemannian compact surface. Let s be a geodesic segment on M and c an injective piecewise geodesic path with the same endpoints as s . It is possible to decompose c and s as concatenations $c = c_1 \cdots c_p$ and $s = s_1 \cdots s_p$ in such a way that, for each $k = 1, \dots, p$, $c_k = s_k$ or $c_k s_k^{-1}$ is a simple loop. In particular, if we set $l_k = (c_1 \cdots c_{k-1})c_k s_k^{-1}(c_1 \cdots c_{k-1})^{-1}$, then l_k is a lasso and $cs^{-1} \simeq l_1 \cdots l_p$.*

Proof. — Consider the set $c([0, 1]) \cap s([0, 1])$. It is a reunion of isolated points and closed subintervals of $s([0, 1])$. Let V be the reunion of these isolated points and the end points of these intervals. The set V contains the two end points of s .




FIGURE 2. The pieces c_6 and c_9 of c (thicker on the picture) intersect s , but not s_6 nor s_9 respectively (dotted on the picture).

Set $n = \#V - 1$. Then V dissects both c and s into n edges: $c = e_1 \cdots e_n$ and $s = f_1 \cdots f_n$.

For each $i \in \{1, \dots, n\}$, define $j(i)$ by the relation $\bar{e}_i = \overline{f_{j(i)}}$. Since c is injective, the mapping j from $\{1, \dots, n\}$ to itself is one-to-one. Hence, it is a permutation. We look for records in the sequence $j(1), j(2), \dots, j(n)$: define

$$I = \{i : 1 \leq i \leq n, j(i) = \max(j(1), \dots, j(i))\}$$

and write $I = \{i_1 \leq \dots \leq i_p\}$. Set $J = j(I) = \{j(i_1) \leq \dots \leq j(i_p)\}$. Observe that $i_p = n = j(i_p)$. Set $i_0 = j(i_0) = 0$. For each $k = 1, \dots, p$, set $c_k = e_{i_{k-1}+1} \cdots e_{i_k}$ and $s_k = f_{j(i_{k-1})+1} \cdots f_{j(i_k)}$. By construction, $c = c_1 \cdots c_p$ and $s = s_1 \cdots s_p$.

Choose $k \in \{1, \dots, p\}$ and consider $c_k = e_{i_{k-1}+1} \cdots e_{i_k}$. Assume first that $i_k = i_{k-1} + 1$. Then $c_k = e_{i_k}$. Either this edge is contained in s , in which case $c_k = e_{i_k} = f_{j(i_k)} = s_k$, or it meets s only at its endpoints, which are also those of s_k . In this case, $c_k s_k^{-1}$ is a simple loop. Assume now that $i_k \geq i_{k-1} + 2$. We claim that any point of c_k other than one of its end points which is located on s is in fact located on one of the edges $f_1, \dots, f_{j(i_{k-1})}$. Indeed, if this was not the case, there would exist $l \in \{i_{k-1} + 1, \dots, i_k - 1\} \neq \emptyset$ such that \bar{e}_l is located on s between \bar{e}_{i_k} and \bar{s} . But then we would have $j(l) > j(i_k)$ which contradicts the definition of i_k . Hence, in this case, $c_k s_k^{-1}$ is a simple loop. \square

We are now ready to prove Proposition 3.3.4.

Proof of Proposition 3.3.4. — Assume that c and s are shorter than the bound K^{-1} of Proposition 3.3.3. Let $\text{LE}(c)$ be the loop-erasure of c . By the properties of the lasso decomposition of c (Proposition 1.4.9) and Proposition 3.3.3,

$$d(H(c), H(\text{LE}(c))) \leq K(\ell(c) - d(c(0), c(1))) = K|\ell(c) - \ell(s)|.$$

Now we are reduced to consider $c' = \text{LE}(c)$ which is an injective path. Let $c' = c'_1 \cdots c'_p$ and $s = s_1 \cdots s_p$ be the decomposition given by Proposition 3.3.6. We have

$$d(H(c'), H(s)) \leq \sum_{i=1}^p d(1, H(c'_i s_i^{-1})) \leq K \sum_{i=1}^p \sqrt{A_i},$$

where A_i is the area enclosed by $c'_i s_i^{-1}$. By Lemma 3.3.5 and since $\ell(s) \leq \ell(c') \leq \ell(c)$,

$$A_i \leq K \ell(c_i)^{\frac{3}{2}} |\ell(c_i) - \ell(s_i)|^{\frac{1}{2}}.$$

By Hölder inequality, it follows that

$$d(H(c'), H(s)) \leq K \ell(c)^{\frac{3}{4}} |\ell(c) - \ell(s)|^{\frac{1}{4}}.$$

Hence,

$$\begin{aligned} d(H(c), H(s)) &\leq K (|\ell(c) - \ell(s)| + \ell(c)^{\frac{3}{4}} |\ell(c) - \ell(s)|^{\frac{1}{4}}) \\ &\leq 2K \ell(c)^{\frac{3}{4}} (\ell(c) - \ell(s))^{\frac{1}{4}}. \end{aligned}$$

□

The estimate given by Proposition 3.3.4 will allow us to associate an element of Γ to every element of $P(M)$. Recall the definition of the dyadic piecewise geodesic approximation of a path (Definition 1.2.11). For a given path c , $D_n(c)$ is in general only defined for n larger than a certain integer $n_0(c)$. Nevertheless, by Proposition 1.2.12, the sequence $(D_n(c))_{n \geq n_0(c)}$ converges to c in 1-variation.

Proposition 3.3.7. — *Let $c \in P(M)$. Under the assumptions of Theorem 3.3.1, the sequence $(H(D_n(c)))_{n \geq n_0(c)}$ is a Cauchy sequence in Γ .*

Proof. — Let K be the constant given by Proposition 3.3.4. Let $n_1(c)$ be an integer such that $2^{-n_1(c)} \ell(c) < K^{-1}$. Choose $m \geq n \geq \max(n_0(c), n_1(c))$. Write

$$D_n(c) = \sigma_0 \cdots \sigma_{2^n - 1}$$

according to the notation of Definition 1.2.11. Write also $D_m(c) = \eta_0 \cdots \eta_{2^m - 1}$, where for each $k \in \{0, \dots, 2^n - 1\}$, $\eta_k = D_{m-n}(c_{|[k2^{-n}, (k+1)2^{-n}]})$. By Proposition 3.3.4, for all $k \in \{0, \dots, 2^n - 1\}$,

$$d(H(\eta_k), H(\sigma_k)) \leq K \ell(\eta_k)^{\frac{3}{4}} |\ell(\eta_k) - \ell(\sigma_k)|^{\frac{1}{4}}.$$

Hence, by Hölder inequality, and since $\ell(D_n(c)) \leq \ell(D_m(c)) \leq \ell(c)$,

$$d(H(D_n(c)), H(D_m(c))) \leq K \ell(c)^{\frac{3}{4}} |\ell(c) - \ell(D_n(c))|^{\frac{1}{4}}.$$

The result follows now from the fact that $\ell(D_n(c))$ converges to $\ell(c)$. □

By Proposition 3.3.7 and the assumption that (Γ, d) is complete, it is now legitimate to set the following definition.

Definition 3.3.8. — *For each $c \in P(M) \setminus A_\gamma(M)$, we define*

$$H(c) = \lim_{n \rightarrow \infty} H(D_n(c)).$$

Proposition 3.3.9. — *The mapping $H : P(M) \rightarrow \Gamma$ thus defined is continuous at fixed endpoints.*

Notice that we have not proved yet that H is continuous even on $A_\gamma(M)$.

Proof. — Take $c \in P(M)$ and consider a sequence $(c_n)_{n \geq 0}$ in $A_\gamma(M)$ converging to c with fixed endpoints. We claim that $H(c_n)$ converges to $H(c)$. An elementary argument shows that this implies the continuity of H with fixed endpoints on $P(M)$.

Choose $\varepsilon > 0$. Choose an integer $m \geq 0$ such that $d(H(D_m(c)), H(c)) < \frac{1}{2}\varepsilon$. Such an integer exists by definition of $H(c)$. Now for each $n \geq 0$ and each $k \in \{0, \dots, 2^m\}$, let $\eta_{n,k}$ be the geodesic segment joining $c_n(k2^{-m})$ to $c(k2^{-m})$. If $k \neq 2^m$, let us also denote by $c_{n,k}$ the portion of c_n parametrised by the interval $[k2^{-m}, (k+1)2^{-m}]$. Observe that $\eta_{n,0}$ and $\eta_{n,2^m}$ are constant paths. The equivalence of paths

$$c_n = c_{n,0} \cdots c_{n,2^m} \simeq (\eta_{n,0}^{-1} c_{n,0} \eta_{n,1}) \cdots (\eta_{n,k}^{-1} c_{n,k} \eta_{n,k+1}) \cdots (\eta_{n,2^m-1} c_{n,2^m-1} \eta_{n,2^m})$$

implies the following inequality:

$$d(H(c_n), H(D_m(c))) \leq \sum_{k=0}^{2^m-1} d(H(\eta_{n,k}^{-1} c_{n,k} \eta_{n,k}), H(\sigma_k)),$$

where $D_m(c) = \sigma_0 \dots \sigma_{2^m-1}$ is the decomposition given by the definition of $D_m(c)$. The path $\eta_{n,k}^{-1} c_{n,k} \eta_{n,k}$ is piecewise geodesic and shares the same endpoints as the segment σ_k . Hence we can apply Proposition 3.3.4 to find

$$d(H(\eta_{n,k}^{-1} c_{n,k} \eta_{n,k}), H(\sigma_k)) \leq K(\ell(c_{n,k}) + 2d_\infty(c_n, c))^{\frac{3}{4}} (\ell(c_{n,k}) - \ell(\sigma_k) + 2d_\infty(c_n, c))^{\frac{1}{4}}.$$

By Hölder inequality again,

$$d(H(c_n), H(D_m(c))) \leq K(\ell(c_n) + 2^{m+1}d_\infty(c_n, c))^{\frac{3}{4}} (\ell(c_n) - \ell(c) + 2^{m+1}d_\infty(c_n, c))^{\frac{1}{4}}.$$

Since $\ell(c_n)$ converges to $\ell(c)$ and $d_\infty(c_n, c)$ tends to 0, the right hand side tends to zero as n tends to infinity. For n large enough, it is smaller than $\frac{1}{2}\varepsilon$. For such n , $d(H(c_n), H(c)) < \varepsilon$. \square

Remark 3.3.10. — The factors 2^{m+1} in the last expression are very unpleasant, because they give the feeling that H is not uniformly continuous on $P(M)$. In fact, the last proof reveals that, on $P(M)$ endowed with the distance d_ℓ , H is uniformly continuous on subsets where $\ell \circ D_n$ converges uniformly to ℓ . It is likely that a much better result can be achieved by considering the stronger distance d_1 on $P(M)$. Ideally, one could expect H to be $\frac{1}{4}$ -Hölder continuous on $(P(M), d_1)$. I have not been able to prove or disprove this statement.

3.4. Extension of discrete holonomy fields

Let $(M, \text{vol}, \gamma, \mathcal{E}, C)$ be a Riemannian marked surface with G -constraints. Starting from a discrete Markovian holonomy field DF satisfying a condition of regularity, we have constructed a measure $\text{DF}_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)}$ on $(\mathcal{M}(P(M), G), \mathcal{E})$ (see Corollary 3.3.2).

The construction of this measure involves a Riemannian metric and we must now prove that the result is independent of this choice. We start by identifying the distribution of $(H_c)_{c \in P(G)}$ under $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$ for an arbitrary graph G .

Proposition 3.4.1. — *Let DF be a stochastically $\frac{1}{2}$ -Hölder continuous and continuously area-dependent discrete Markovian holonomy field. Let $(M, \text{vol}, \gamma, \mathcal{C}, C)$ be a Riemannian marked surface with G -constraints. Let $G = (V, E, F)$ be a graph on (M, \mathcal{C}) . Recall the notation $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$ from Corollary 3.3.2.*

- 1) *The distribution of $(H_e)_{e \in E}$ under $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$ is $DF_{M, \text{vol}, \mathcal{C}, C}^G$.*
- 2) *The measure $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$ does not depend on the Riemannian metric γ . We denote it by $DF_{M, \text{vol}, \mathcal{C}, C}$.*
- 3) *The distribution of $(H_e)_{e \in E}$ under $DF_{M, \text{vol}, \mathcal{C}, C}$ is $DF_{M, \text{vol}, \mathcal{C}, C}^G$.*

Let us emphasise that the graph G considered in this statement is not assumed to have piecewise geodesic edges.

Proof. — 1) For each $n \geq 1$, let G_n be the graph produced by Proposition 1.4.10 with $\varepsilon = n^{-1}$. Let $S_n : \mathbb{E} \rightarrow \mathbb{E}_n$ denote the corresponding bijection. By the stochastic continuity of the process $(H_c)_{c \in P(M)}$ under $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$, which follows from Theorem 3.3.1, the distribution of $(H_e)_{e \in E}$ is the limit of the distributions of the families $(H_{S_n(e)})_{e \in \mathbb{E}}$ as n tends to infinity.

Since S_n preserves the cyclic order at each vertex of G , there exists for each n a homeomorphism of M which preserves \mathcal{C} and which sends G on G_n and which induces the bijection S_n . Let ψ_n be such a homeomorphism. Since for each face F of G , the boundary of $\psi_n(F)$ is $\psi_n(\partial F) = S_n(\partial F) = \partial S_n(F)$, the face $\psi_n(F)$ is $S_n(F)$.

Moreover, the distribution of the family $(H_{S_n(e)})_{e \in \mathbb{E}}$ under $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$ is the distribution of the same family under $DF_{M, \text{vol}, \mathcal{C}, C}^{G_n}$, hence the distribution of $(H_e)_{e \in E}$ under $DF_{M, \text{vol} \circ \psi_n, \mathcal{C}, C}^G$, where $\text{vol} \circ \psi_n$ denotes any measure of area on M which gives to each face F of G the area $\text{vol}(\psi_n(F))$.

By the fourth assertion of Proposition 1.4.10, $\text{vol}(\psi_n(F))$ tends to $\text{vol}(F)$ as n tends to infinity for all $F \in F$. Hence, the assumption that DF is continuously area-dependent implies that the distribution of the family $(H_{S_n(e)})_{e \in \mathbb{E}}$ under $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$ converges weakly to $DF_{M, \text{vol}, \mathcal{C}, C}^G$ as n tends to infinity.

2) Let γ and γ' be two Riemannian metrics on $(M, \text{vol}, \mathcal{C})$. By Lemma 1.4.8 and the property that we have just proved, the distributions of $(H_c)_{c \in A_\gamma(M)}$ under $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$ and $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma')}$ agree. Since $A_\gamma(M)$ is dense in $P(M)$ for the convergence with fixed endpoints (see Proposition 1.2.12), the continuity property granted by Theorem 3.3.1 implies that $DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)} = DF_{M, \text{vol}, \mathcal{C}, C}^{(\gamma')}$.

3) This property follows immediately from the first two. □

We can now finish the proof of the main theorem of this section.

Proof of Theorem 3.2.9. — Let DF be a regular discrete Markovian holonomy field. By applying Corollary 3.3.2 and Proposition 3.4.1, we get for all measured marked surface with G -constraints $(M, \text{vol}, \mathcal{C}, C)$ a finite measure $\text{DF}_{M, \text{vol}, \mathcal{C}, C}$ on $(\mathcal{M}(\text{P}(M), G), \mathcal{C})$ which by restriction produces a measure on the invariant σ -field. The total mass of the measure $\text{DF}_{M, \text{vol}, \mathcal{C}, C}$ is the common value of the masses of the measures $\text{DF}_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}$ for all graph \mathbb{G} on (M, \mathcal{C}) , which we have denoted by $Z_{M, \text{vol}, \mathcal{C}, C}$. In particular, this mass is finite.

Now, we check that the seven axioms of Definition 3.1.2 are satisfied. We choose a measured marked surface with G -constraints $(M, \text{vol}, \mathcal{C}, C)$. We endow (M, \mathcal{C}, C) with a Riemannian metric γ and use without further comment the fact, granted by Proposition 3.4.1, that $\text{DF}_{M, \text{vol}, \mathcal{C}, C} = \text{DF}_{M, \text{vol}, \mathcal{C}, C}^{(\gamma)}$.

A₁) Let \mathcal{N} denote the event $\{\exists l \in \mathcal{C} \cup \mathcal{B}(M), h(l) \notin C(l)\}$. Let \mathbb{G} be a graph on (M, \mathcal{C}) . By Proposition 3.4.1 and the axiom D₁ for DF ,

$$\text{DF}_{M, \text{vol}, \mathcal{C}, C}(\mathcal{N}) = \text{DF}_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}}(\mathcal{N}) = 0.$$

A₂) The set of bounded measurable functions from $(\mathcal{M}(\text{P}(M), G), \mathcal{J})$ to \mathbb{R} whose integral against $\text{DF}_{M, \text{vol}, \mathcal{C}, C}$ depends measurably of the G -constraints C is a vector space which contains the constant functions and is stable by uniformly bounded monotone limit. Thus, by a monotone class argument, and by definition of the invariant σ -field \mathcal{J} , in order to show that this space contains all bounded functions measurable with respect to \mathcal{J} , it suffices to show that it contains all functions of the form $h \mapsto f(h(l_1), \dots, h(l_n))$ where l_1, \dots, l_n are loops based at the same point and $f : G^n \rightarrow \mathbb{R}$ is continuous and invariant under the diagonal action of G by conjugation.

Let us choose l_1, \dots, l_n and f as above. For each $i \in \{1, \dots, n\}$ and all $m \geq 1$ large enough, let $D_m(l_i)$ denote the dyadic piecewise geodesic approximation of l_i of order m . Let us define a function F and a sequence of functions F_m on the set $\text{Const}_G(M, \mathcal{C})$ of G -constraints on (M, \mathcal{C}) as follows:

$$F(C) = \int_{\mathcal{M}(\text{P}(M), G)} f(h(l_1), \dots, h(l_n)) \text{DF}_{M, \text{vol}, \mathcal{C}, C}(dh)$$

and, for all $m \geq 1$,

$$F_m(C) = \int_{\mathcal{M}(\text{P}(M), G)} f(h(D_m(l_1)), \dots, h(D_m(l_n))) \text{DF}_{M, \text{vol}, \mathcal{C}, C}(dh).$$

Our goal is to prove that the function F is measurable. For all $C \in \text{Const}_G(M, \mathcal{C})$, the fact that $D_m(l_i)$ converges to l_i with fixed endpoints implies, according to the conclusion of Theorem 3.3.1, that $H_{D_m(l_i)}$ converges in probability to H_{l_i} , so that $F_m(C)$ tends to $F(C)$ as m tends to infinity. Hence, it suffices to prove that F_m is measurable for m large enough. This follows immediately from the axiom D₂.

A₃) Just as in the last point, it suffices to check the equality when it integrates a function of the form $h \mapsto f(h(l_1), \dots, h(l_n))$, where l_1, \dots, l_n are piecewise geodesic and f is invariant by diagonal conjugation. Hence, it suffices to prove the equality for $\text{DF}_{M, \text{vol}, \mathcal{E}, C}^{\mathbb{G}}$ for any graph \mathbb{G} with piecewise geodesic edges. In this last case, the property follows from the axiom D₃ satisfied by DF.

A₄) The metric $\gamma' = (\psi^{-1})^* \gamma$ is a Riemannian metric on $(M', \text{vol}', \mathcal{E}')$ in the sense of Proposition 1.4.3. We need to prove that the image measure of $\text{DF}_{M, \text{vol}, \mathcal{E}, C}^{(\gamma)}$ by the mapping induced by ψ is $\text{DF}_{M', \text{vol}', \mathcal{E}', C'}^{(\gamma')}$. Again, we may restrict ourselves to functions of the form $h \mapsto f(h(l_1), \dots, h(l_n))$, where l_1, \dots, l_n are piecewise geodesic, hence to the discrete measures associated to graphs with piecewise geodesic edges. Let \mathbb{G} be a graph on (M, \mathcal{E}) with piecewise geodesic edges. Then \mathbb{G}' , the graph constituted by the images by ψ of the edges of \mathbb{G} , is a graph on (M', \mathcal{E}') with piecewise geodesic edges and, by the axiom D₄ for DF, the distribution of $(H_e)_{e \in \mathbb{E}}$ under $\text{DF}_{M, \text{vol}, \mathcal{E}, C}^{\mathbb{G}}$ is the same as the distribution of $(H_{e'})_{e' \in \mathbb{E}'}$ under $\text{DF}_{M', \text{vol}', \mathcal{E}', C'}^{\mathbb{G}'}$. The property follows.

A₅, A₆ and A₇ follow respectively from the axioms D₅, D₆ and D₇ satisfied by DF and the same approximation argument as in the previous points.

The fact that the new Markovian holonomy field DF is stochastically continuous is a part of the conclusion of Theorem 3.3.1. Finally, the Feller property follows from the Feller property of DF. \square

CHAPTER 4

LÉVY PROCESSES AND MARKOVIAN HOLONOMY FIELDS

In this chapter, we apply the extension theorem proved in the previous chapter in order to construct a whole family of Markovian holonomy fields. Before that, we study the partition functions of an arbitrary regular Markovian holonomy field and prove that they are completely determined by a Lévy process on the group G with some nice properties, essentially a continuous density. We then construct a Markovian holonomy field for each such Lévy process. The case of the Brownian motion on a connected Lie group yields the Yang-Mills measure.

4.1. The partition functions of a Markovian holonomy field

In this section, we establish some fundamental properties satisfied by the masses of the finite measures which constitute a Markovian holonomy field.

To start with, we describe the isomorphism classes of connected surfaces with G -constraints on the boundary. If M is oriented, we denote by $\mathcal{B}^+(M)$ the subset of $\mathcal{B}(M)$ which consists of the curves which have the orientation induced by that of M .

Proposition 4.1.1. — *Let $(M, \text{vol}, \varnothing, C)$ and $(M', \text{vol}', \varnothing, C')$ be two connected measured marked surfaces with G -constraints. If M and M' are orientable, we assume that they are oriented. They are isomorphic if and only if the following conditions hold simultaneously.*

- 1) M and M' are homeomorphic.
- 2) $\text{vol}(M) = \text{vol}'(M')$.
- 3) If M and M' are oriented, there exists a bijection $\psi : \mathcal{B}^+(M) \rightarrow \mathcal{B}^+(M')$ such that $C = C' \circ \psi$ on $\mathcal{B}^+(M)$.
- 3') If M and M' are non-orientable, there exists a $\mathbb{Z}/2\mathbb{Z}$ -equivariant bijection $\psi : \mathcal{B}(M) \rightarrow \mathcal{B}(M')$ such that $C = C' \circ \psi$.

We use this result to associate to every Markovian holonomy field a family of functions of one or several variables in $\text{Conj}(G)$.

Let HF be a Markovian holonomy field. Let g and p be two non-negative integers, with g even. Let t be a positive real number. Recall the notation of Section 1.1. Let M be a surface homeomorphic to $\Sigma_{p,g}^+$, endowed with a density vol of total area t . Let b_1, \dots, b_p denote the positively oriented connected components of ∂M . Let x_1, \dots, x_p be elements of G . Let C be the unique set of G -constraints on (M, \emptyset) such that, for all $i \in \{1, \dots, p\}$, $C(b_i) = \mathcal{O}_{x_i}$. By Proposition 4.1.1 and the axiom A_4 , the number $\text{HF}_{(M, \text{vol}, \emptyset, C)}(1)$ depends on M , vol and C only through p, g, t , the unordered list $\mathcal{O}_{x_1}, \dots, \mathcal{O}_{x_p}$ and the fact that M is orientable. Hence, it is legitimate to set

$$Z_{p,g,t}^+(x_1, \dots, x_p) = \text{HF}_{(M, \text{vol}, \emptyset, C)}(1).$$

This defines a symmetric function $Z_{p,g,t}^+$ of p conjugacy classes of G . If the Markovian field is Fellerian, then this function is continuous with respect to (t, x_1, \dots, x_p) . If $p = 0$, $Z_{0,g,t}^+$ is just a number, namely the total mass of the measure $\text{HF}_{(M, \text{vol}, \emptyset, \emptyset)}$ where M is the connected sum of $\frac{1}{2}g$ tori endowed with a density of total area t .

Similarly, let g and p be two integers, respectively positive and non-negative. Let M be a surface homeomorphic to $\Sigma_{p,g}^-$, endowed with a density vol of total area t . Let b_1, \dots, b_p denote the disjoint connected components of ∂M endowed with an arbitrary orientation. Let x_1, \dots, x_p be elements of G . Let C be the unique set of G -constraints on (M, \emptyset) such that, for all $i \in \{1, \dots, p\}$, $C(b_i) = \mathcal{O}_{x_i}$. When HF is not oriented, we define

$$Z_{p,g,t}^-(x_1, \dots, x_p) = \text{HF}_{(M, \text{vol}, \emptyset, C)}(1).$$

Again, if $p = 0$, $Z_{0,g,t}^-$ is just the total mass of the measure $\text{HF}_{(M, \text{vol}, \emptyset, \emptyset)}$ where M is the connected sum of g projective planes endowed with a density of total area t .

Definition 4.1.2. — *Let HF be a Markovian holonomy field. The functions*

$$Z_{p,g,t}^\varepsilon : G^p \longrightarrow \mathbb{R}_+^*,$$

where (ε, p, g, t) spans $(\{+\} \times \mathbb{N} \times 2\mathbb{N} \times \mathbb{R}_+^*) \cup (\{-\} \times \mathbb{N} \times \mathbb{N}^* \times \mathbb{R}_+^*)$, are called the partition functions of the field HF .

In the rest of this section, we fix a Markovian holonomy field HF and study its partition functions. They are infinitely many but the Markov property of the field implies that they satisfy an infinite set of relations and that they are in fact completely determined by a small number of them.

Let us introduce several operations on functions of conjugacy classes of G . Firstly, we identify functions of a conjugacy class of G and functions on G which are constant on the conjugacy classes. Thus, we may speak of continuous or integrable functions of a conjugacy class. This point of view is consistent with our previous definition of a

topology and σ -field on $\text{Conj}(G)$ (see Section 3.1). Of particular interest is the space of square-integrable functions of one conjugacy class of G , which we identify with the space $L^2(G)^G$ of conjugation-invariant square-integrable functions on G . If $p \geq 1$ is an integer, we identify the elements of the p -th symmetric tensor power $\text{Sym}^p(L^2(G)^G)$ with symmetric functions of p conjugacy classes. We use the shorthand notation $S^p(G)$ for $\text{Sym}^p(L^2(G)^G)$.

Definition 4.1.3. — For all integers $p, q, r \geq 1$ and $s \geq 2$, the three linear mappings

$$\begin{aligned} v : S^r(G) &\longrightarrow S^{r-1}(G), & \beta_1 : S^s(G) &\longrightarrow S^{s-2}(G), \\ \beta_2 : S^p(G) \otimes S^q(G) &\longrightarrow S^{p-1}(G) \otimes S^{q-1}(G) \end{aligned}$$

are defined by

$$(40) \quad \forall f \in S^r(G), \quad (vf)(x_1, \dots, x_{r-1}) = \int_G f(x_1, \dots, x_{r-1}, x^2) dx,$$

$$(41) \quad \forall f \in S^s(G), \quad (\beta_1 f)(x_1, \dots, x_{s-2}) = \int_G f(x_1, \dots, x_{s-2}, x, x^{-1}) dx,$$

$$(42) \quad \forall f \in S^p(G), \forall f' \in S^q(G), \\ (\beta_2(f \otimes f'))(x_1, \dots, x_{p-1}, y_1, \dots, y_{q-1}) \\ = \int_G f(x_1, \dots, x_{p-1}, z) f'(y_1, \dots, y_{q-1}, z^{-1}) dz.$$

We will now use these linear mappings to formulate the relations between the partition functions of the holonomy field. The operation v expresses the transformation of the partition function under a unary gluing, and the contractions β_1 and β_2 correspond to binary gluings, respectively of two boundary components which lie on the same connected component of the surface and two boundary components which lie on two distinct connected components.

Recall that we use the notation $\varepsilon \wedge \varepsilon'$ for $\varepsilon, \varepsilon' \in \{-, +\}$ with the meaning

$$\varepsilon \wedge \varepsilon' = \begin{cases} + & \text{if } \varepsilon = \varepsilon' = +, \\ - & \text{in all other cases.} \end{cases}$$

Proposition 4.1.4. — Let $(Z_{p,g,t}^\pm)_{p,g \geq 0, t > 0}$ be the partition functions of a Markovian holonomy field HF.

1) For all $(\varepsilon, g) \in (\{+\} \times 2\mathbb{N}) \cup (\{-\} \times \mathbb{N}^*)$, all integers $p \geq q \geq 1$ and all real $t > 0$, the function $Z_{p,g,t}^\varepsilon$ is square-integrable with respect to any q of its variables for any value of the $p - q$ other.

2) The following equality holds:

$$(43) \quad v(Z_{p,g,t}^\varepsilon) = Z_{p-1,g+1,t}^-.$$

Moreover, if $p \geq 2$, then

$$(44) \quad \beta_1(Z_{p,g,t}^\varepsilon) = Z_{p-2,g+2,t}^\varepsilon.$$

Finally, for all $(\varepsilon', g') \in (\{+\} \times 2\mathbb{N}) \cup (\{-\} \times \mathbb{N}^*)$, all $p' \geq 1$ and all real $t' > 0$,

$$(45) \quad \beta_2(Z_{p,g,t}^\varepsilon \otimes Z_{p',g',t'}^{\varepsilon'}) = Z_{p+p'-2,g+g',t+t'}^{\varepsilon \wedge \varepsilon'}.$$

Proof. — Let us start by proving the second assertion. Choose ε, p, g, t as above. It follows from the axiom A_2 that $Z_{p,g,t}^\varepsilon : G^p \rightarrow \mathbb{R}_+^*$ is a measurable function. Since it takes non-negative values, the integral which defines $v(Z_{p,g,t}^\varepsilon)$ is well defined, possibly infinite. Let us prove that the identity (43) holds.

Let M be a surface homeomorphic to $\Sigma_{p,g}^\varepsilon$, endowed with a surface measure vol of total area t . Consider $b \in \mathcal{B}(M)$ and let $f : M \rightarrow M_1$ be a unary gluing along b , with joint $\{l, l^{-1}\} = \{f(b), f(b^{-1})\}$. Thus, $M = \text{Spl}_l(M_1)$. The surface M_1 is not orientable and it has $p-1$ boundary components. According to the observation made after Definition 1.1.10, it is homeomorphic to the connected sum of a projective plane and the surface obtained by gluing a disk along one boundary component of M . Thus, it is homeomorphic to $\Sigma_{p-1,g+1}^-$. Finally, vol induces on M_1 a surface measure vol_1 with total area t .

Let C be a set of G -constraints on $(M, \{l, l^{-1}\})$. Applying axioms A_3 and then A_6 gives us the relation

$$\text{HF}_{M_1, \text{vol}_1, \emptyset, C}(1) = \int_G \text{HF}_{M_1, \text{vol}_1, \{l, l^{-1}\}, C_{l \rightarrow x}}(1) dx = \int_G \text{HF}_{M, \text{vol}, \emptyset, \text{Spl}_l(C_{l \rightarrow x})}(1) dx.$$

Recall from Definition 2.3.2 that, since we are considering a unary gluing, $\text{Spl}_l(C_{l \rightarrow x})$ puts the constraint \emptyset_{x^2} on $b \in \mathcal{B}(M)$. Translating the last relation in terms of the partition functions gives (43).

The proofs of the relations (44) and (45) are very similar. For the relation (44), one considers a binary gluing in which one identifies two boundary components of a connected surface. If the surface is orientable, the two boundary components are identified by an orientation-reversing diffeomorphism. The result of this gluing can be described as the connected sum of a torus and the surface obtained by gluing two disks along two boundary components of the original surface. Thus, it has two boundary components less, and reduced genus increased by 2. For the relation (45), one considers a binary gluing in which one identifies two boundary components located on two distinct connected components of a surface. Proposition 1.1.4 settles the issue of orientation and reduced genus of the resulting surface.

The assertion of square-integrability follows from (44) and (45) as we shall see now. Firstly, observe that Theorem 1.1.7 and the axiom A_4 imply that the value of a partition function is not affected by simultaneously replacing each argument by its

inverse. Now, for all $x_{q+1}, \dots, x_p \in G$,

$$\begin{aligned} & \int_{G^q} Z_{p,g,t}^\varepsilon(x_1, \dots, x_p)^2 dx_1 \cdots dx_q \\ &= \int_{G^q} Z_{p,g,t}^\varepsilon(x_1, \dots, x_p) Z_{p,g,t}^\varepsilon(x_1^{-1}, \dots, x_p^{-1}) dx_1 \cdots dx_q \\ &= (\beta_1^{q-1} \circ \beta_2)(Z_{p,g,t}^\varepsilon \otimes Z_{p,g,t}^\varepsilon)(x_{q+1}, x_{q+1}^{-1}, \dots, x_p, x_p^{-1}) \\ &= Z_{2(p-q), 2(g+q-1), 2t}^\varepsilon(x_{q+1}, x_{q+1}^{-1}, \dots, x_p, x_p^{-1}). \end{aligned}$$

The last number is indeed finite and the assertion is proved. □

Corollary 4.1.5. — *The partition functions of a Markovian holonomy field are completely determined by the functions $(Z_{1,0,t}^+)_{t>0}$ and $(Z_{3,0,t}^+)_{t>0}$.*

Proof. — For all non-negative p and g , such that g is even and $p+g > 0$, the following identity holds:

$$(46) \quad Z_{p,g,t}^+ = (\beta_1^{\frac{1}{2}g} \circ \beta_2^{p+g-1})(Z_{1,0,t/(p+g)}^+ \otimes Z_{3,0,t/(p+g)}^+ \otimes^{(p+g-1)}).$$

Then, the equalities

$$(47) \quad \forall p, g \geq 0, \quad Z_{p,2g+1,t}^- = v(Z_{p+1,2g,t}^+) \text{ and } Z_{p,2g+2,t}^- = v^2(Z_{p+2,2g,t}^+)$$

determine all non-orientable partition functions. Finally, the equality

$$Z_{0,0,t}^+ = \beta_2(Z_{1,0,\frac{1}{2}t}^+ \otimes Z_{1,0,\frac{1}{2}t}^+)$$

determines the only remaining partition function. □

An important consequence of Proposition 4.1.4 is that for all ε, p, g, t , the function $Z_{p,g,t}^\varepsilon$ is the density with respect to the Haar measure of a measure on G^p . According to the axiom A_7 , this measure is a probability measure. In the next proposition, we will study the behaviour of these probability measures as t tends to 0, under the assumption that the holonomy field is stochastically continuous. Recall that we denote by 1 the unit element of G and that we use the notation δ_{θ_x} for the unique G -invariant probability measure on θ_x , the conjugacy class of $x \in G$.

Proposition 4.1.6. — *Let $(Z_{p,g,t}^\pm)_{p,g,t}$ be the partition functions of a stochastically continuous Markovian holonomy field. Then, as t tends to 0, one has the following weak convergences of measures on G .*

1) $Z_{1,0,t}^+(x) dx \xrightarrow[t \rightarrow 0]{} \delta_1.$

2) for all $x, y \in G$,

$$Z_{3,0,t}^+(x, y, z^{-1}) dz \xrightarrow[t \rightarrow 0]{} \delta_{\theta_x} * \delta_{\theta_y} = \int_{G^2} \delta_{vxv^{-1}wyw^{-1}} dv dw.$$

Proof. — 1) Let M be the disk of radius 1 centred at the origin in \mathbb{R}^2 endowed with the Lebesgue measure, which we denote by vol . We denote by ∂M the positively oriented boundary of M . For each $r \in [0, 1]$, let s_r be the path which goes straight from the origin to the point $(r, 0)$ and let c_r be the loop based at $(r, 0)$ which goes once counterclockwise around the circle of radius r centred at the origin. For all $y \in G$, we denote by $(h_c)_{c \in \mathcal{P}(M)}$ the canonical process on $\mathcal{M}(\mathcal{P}(M), G)$, we consider the measure $\text{HF}_{M, \text{vol}, \emptyset, \partial M \mapsto y}$ on $\mathcal{M}(\mathcal{P}(M), G)$ and we denote by \mathbb{E}_y the corresponding expectation. Let f be a continuous invariant function on G . We compute $\mathbb{E}_y[f(h_{s_r c_r s_r^{-1}})]$. By the multiplicativity property of h , it is equal to $\mathbb{E}_y[f(h_{s_r^{-1} h_{c_r} h_{s_r})] = \mathbb{E}_y[f(h_{c_r})]$. By the axiom A_3 , we can disintegrate this expectation with respect to the value of $h(c_r)$. We find

$$\begin{aligned} \mathbb{E}_y[f(h_{c_r})] &= \int_{G \times \mathcal{M}(\mathcal{P}(M), G)} f(h(c_r)) \text{HF}_{M, \text{vol}, \{c_r, c_r^{-1}\}, (\partial M \mapsto y, c_r \mapsto x)}(dh) dx \\ &= \int_G f(x) \text{HF}_{M, \text{vol}, \{c_r, c_r^{-1}\}, (\partial M \mapsto y, c_r \mapsto x)}(1) dx. \end{aligned}$$

We use now the axioms A_6 and A_5 to split M along c_r and we find

$$\mathbb{E}_y[f(h_{c_r})] = \int_G Z_{1,0,\pi r^2}^+(x) f(x) Z_{2,0,\pi(1-r^2)}^+(x^{-1}, y) dx.$$

By integrating over y and using the axiom A_7 , we find

$$(48) \quad \int_G \mathbb{E}_y[f(h_{c_r})] dy = \int_G Z_{1,0,\pi r^2}^+(x) f(x).$$

Our goal is now to prove that the left-hand side tends to $f(1)$ as r tends to 0. For this, we use the stochastic continuity of the holonomy field. Indeed, as r tends to 0, the loop $s_r c_r s_r^{-1}$ converges with fixed endpoints to the constant loop c_0 at the origin. Hence, for all $y \in G$,

$$\mathbb{E}_y[f(h_{c_r})] \longrightarrow \mathbb{E}_y[f(h_{c_0})]$$

as r tends to 0. By the multiplicativity of h and the fact that $c_0 = c_0 c_0^{-1}$, the mapping $h_{c_0} : \mathcal{P}(M) \rightarrow G$ is identically equal to 1. Hence,

$$\mathbb{E}_y[f(h_{c_r})] \xrightarrow{r \rightarrow 0} f(1) Z_{1,0,\pi}^+(y).$$

In order to integrate this convergence with respect to y , we use the fact that

$$|\mathbb{E}_y[f(h_{c_r})]| \leq \|f\|_\infty Z_{1,0,\pi}^+(y)$$

and the right-hand side is continuous, hence integrable, with respect to y . Hence, the dominated convergence theorem applies and we deduce that the left-hand side of (48) tends to $f(1)$ as r tends to 0.

2) Let M be the closed disk of radius 4 centred at the origin in \mathbb{R}^2 from which one has removed the two open disks of radius 1 centred respectively at the points $\alpha = (2, 0)$ and $\beta = (-2, 0)$. We endow M with some density denoted by vol . Let a

(resp. b) be the loop which starts at $(1, 0)$ (resp. $(-1, 0)$) and goes once around the circle of radius 1 centred at α (resp. β), counterclockwise. Let d be the path which goes straight from $(1, 0)$ to $(-1, 0)$. Choose $r \in (0, 1)$. Consider the union of the two closed disks of radius $1+r$ centred at α and β and the rectangle $[-1, 1] \times [-r, r]$. Let c_r be the loop which starts at $(2 - \sqrt{1+2r}, r)$ and bounds this domain with positive orientation. Let s_r be the path which goes straight from $(1, 0)$ to $(2 - \sqrt{1+2r}, r)$. As r tends to 0, the loop $s_r c_r s_r^{-1}$ converges with fixed endpoints to the loop $dbd^{-1}a$. However, in order to apply our axioms, we need to replace c_r by a loop based at the same point and whose image is a smooth submanifold of M . We do this in such a way that the convergence of $s_r c_r s_r^{-1}$ to $dbd^{-1}a$ is preserved.

We consider the measure $\text{HF}_{M, \text{vol}, \emptyset, C_{x,y,z}}$ on $\mathcal{M}(\text{P}(M), G)$, where $C_{x,y,z}$ is characterised by the fact that $C(a) = x$, $C(b) = y$ and C maps the circle of radius 4 centred at the origin to z . We denote the corresponding expectation by $\mathbb{E}_{x,y,z}$. Let f be a continuous invariant function on G . For all $r \in (0, 1)$, we have

$$(49) \quad \mathbb{E}_{x,y,z} [f(h_{s_r c_r s_r^{-1}})] = \int_G Z_{3,0,t_r}^+(x, y, v) f(v) Z_{2,0,T-t_r}^+(v^{-1}, z) dv,$$

where t_r is the area of the domain delimited by a , b and c_r and T is the total area of M . By the same arguments as in the previous proofs, the left-hand side converges as r tends to 0 to $\mathbb{E}_{x,y,z} [f(h_{dbd^{-1}a})]$ and the convergence is dominated with respect to z . Since the two endpoints of d are distinct and the measure $\text{HF}_{M, \text{vol}, \emptyset, C_{x,y,z}}$ is invariant under gauge transformations, the distribution of h_d is both left and right-invariant on G . Thus, h_d has the uniform distribution on G . Hence, $\mathbb{E}_{x,y,z} [f(h_{dbd^{-1}a})] = \int_G f(xwyw^{-1})dw$. By integrating (49) with respect to z , we find

$$\int_G Z_{3,0,t_r}^+(x, y, z) f(z) dz \xrightarrow{r \rightarrow 0} \int_G f(xwyw^{-1}) dw.$$

□

We will now use this result to prove that the partition functions a stochastically continuous Markovian holonomy field are completely determined by the functions $(Z_{1,0,t}^+)_{t>0}$. Let us introduce two probability measures on G .

Definition 4.1.7. — *Let η and κ be the two invariant probability measures on G defined respectively by the fact that for all continuous function f on G ,*

$$(50) \quad \int_G f d\eta = \int_{G^2} f(aba^{-1}b^{-1}) da db \quad \text{and} \quad \int_G f d\kappa = \int_G f(a^2) da.$$

The letters η and κ correspond to the words *handle* and *cross-cap*. We start by proving some important properties of these measures.

Let $\text{Irr}(G)$ denote the set of isomorphism classes of irreducible representations of G over \mathbb{C} . Given $\alpha \in \text{Irr}(G)$ with character χ_α and a measure μ on G , the

Fourier coefficient $\hat{\mu}(\alpha)$ is defined by $\hat{\mu}(\alpha) = \int_G \overline{\chi_\alpha} d\mu$. Recall that an irreducible representation is said to be complex if its character is not real valued, otherwise real (resp. quaternionic) if it preserves a non-degenerate symmetric (resp. skew-symmetric) complex bilinear form.

If μ is a measure on G , we denote by μ^\vee the measure defined by

$$\int_G f d\mu^\vee = \int_G f(g^{-1})\mu(dg).$$

By an invariant measure we mean a measure which is invariant by conjugation.

Lemma 4.1.8. — *Let μ, ν, ξ be three invariant probability measures on G . If $\hat{\mu}(\alpha) = 0$ for every complex representation α , then $\mu * \nu = \mu * \nu^\vee$. In particular, $\kappa * \xi * \nu = \kappa * \xi * \nu^\vee$. Moreover, $\kappa * \eta = \kappa^{*3}$.*

Proof. — The Fourier coefficients of η and κ can be computed easily by using the elementary theory of characters: for all $\alpha \in \text{Irr}(G)$,

$$\hat{\eta}(\alpha) = \frac{1}{\dim \alpha} \quad \text{and} \quad \hat{\kappa}(\alpha) = \begin{cases} 1 & \text{if } \alpha \text{ is real,} \\ 0 & \text{if } \alpha \text{ is complex,} \\ -1 & \text{if } \alpha \text{ is quaternionic.} \end{cases}$$

The Fourier coefficients of the convolution product of two measures is given by the relation $\widehat{\mu * \mu'}(\alpha) = \hat{\mu}(\alpha)\hat{\mu}'(\alpha)/\dim \alpha$. Moreover, the Fourier coefficients of ν^\vee are the complex conjugate of those of ν . Hence, on real and quaternionic representations, whose character is real, the Fourier coefficients of ν are real and agree with those of ν^\vee . Hence, $\mu * \nu$ and $\mu * \nu^\vee$ have the same Fourier coefficients. Since they are both invariant, they are equal. The equality $\kappa * \xi * \nu = \kappa * \xi * \nu^\vee$ follows immediately. The last assertion is proved by computing the Fourier coefficients of both sides. \square

Remark 4.1.9. — This equality $\kappa * \eta = \kappa^{*3}$ should be compared to the fact that the connected sum of a projective plane and a torus is homeomorphic to the connected sum of three projective planes (or the connected sum of a projective plane and a Klein bottle).

In the next proposition, we use the notation $\mu(f)$ for the integral of a function f against a measure μ and we denote by $*$ the convolution of probability measures.

Proposition 4.1.10. — *The partition functions of a regular Markovian holonomy field are completely determined by the functions $Z_{1,0,t}^+$ for $t > 0$. One has the following explicit formulas.*

For all $p \geq 0$, all $g \geq 0$ even, all $t > 0$ and all $x_1, \dots, x_p \in G$, one has

$$(51) \quad Z_{p,g,t}^+(x_1, \dots, x_p) = \eta^{*\frac{1}{2}g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_p}}(Z_{1,0,t}^+).$$

Moreover, for all $p \geq 0$, all $g > 0$, all $t > 0$ and all $x_1, \dots, x_p \in G$, one has

$$(52) \quad Z_{p,g,t}^-(x_1, \dots, x_p) = \kappa^{*g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_p}}(Z_{1,0,t}^+).$$

Proof. — Let us start by proving (51) when $g = 0$ and $p > 0$, by induction on p . For $p = 1$, it is a consequence of the fact that $Z_{1,0,t}^+$ is invariant by conjugation. Assume that $p > 1$ and the result has been proved for $Z_{p-1,0,t}^+$. Then, for all $x_1, \dots, x_p \in G$ and all $s \in (0, t)$, (45) yields

$$Z_{p,0,t}^+ = \int_G Z_{p-1,0,t-s}^+(x_1, \dots, x_{p-2}, y) Z_{3,0,s}^+(x_{p-1}, x_p, y^{-1}) dy.$$

Since the Markovian holonomy field that we consider is Fellerian, the function

$$(s, x_1, \dots, x_{p-2}, y) \mapsto Z_{p-1,0,t-s}^+(x_1, \dots, x_{p-2}, y)$$

is continuous on the compact set $[0, \frac{1}{2}t] \times G^{p-1}$. Hence, when s tends to 0, it converges uniformly as a function on G^{p-1} towards $Z_{p-1,0,t}^+$. Thus, using the convergence proved in Proposition 4.1.6, we find

$$Z_{p,0,t}^+ = \int_G Z_{p-1,0,t}^+(x_1, \dots, x_{p-2}, y) (\delta_{O_{x_{p-1}}} * \delta_{O_{x_p}})(dy).$$

Using the induction hypothesis, we find

$$\begin{aligned} Z_{p,0,t}^+ &= \int_G \left(\int_{G^{p-1}} Z_{1,0,t}^+(w_1 \cdots w_{p-2} z) \prod_{i=1}^{p-2} \delta_{\theta_{x_i}}(dw_i) \delta_{O_y}(dz) \right) (\delta_{O_{x_{p-1}}} * \delta_{O_{x_p}})(dy) \\ &= \int_{G^p} Z_{1,0,t}^+(w_1 \cdots w_p) \prod_{i=1}^p \delta_{\theta_{x_i}}(dw_i) \end{aligned}$$

because $\delta_{O_{x_{p-1}}} * \delta_{O_{x_p}}$ is already an invariant measure on G . This is the expected result.

Let us now treat the case where $p + g > 0$. We have

$$\begin{aligned} Z_{p,g,t}^+(x_1, \dots, x_p) &= \beta_1^{\frac{1}{2}g} (Z_{p+g,0,t}^+)(x_1, \dots, x_p) \\ &= \int_{G^{\frac{1}{2}g}} Z_{p+g,0,t}^+(x_1, \dots, x_p, y_1, y_1^{-1}, \dots, y_{\frac{1}{2}g}, y_{\frac{1}{2}g}^{-1}) dy_1 \cdots dy_{\frac{1}{2}g} \\ &= \int_{G^{p+\frac{3}{2}g}} Z_{1,0,t}^+(w_1 \cdots w_p z_1 z_1' \cdots z_{\frac{1}{2}g} z_{\frac{1}{2}g}') \\ &\quad \times \prod_{i=1}^p \delta_{\theta_{x_i}}(dw_i) \prod_{i=1}^{\frac{1}{2}g} \delta_{\theta_{y_i}}(dz_i) \delta_{\theta_{y_i^{-1}}}(dz_i') \prod_{i=1}^{\frac{1}{2}g} dy_i. \end{aligned}$$

The result follows now from the equality $\int_G \delta_{\theta_v} * \delta_{\theta_{v^{-1}}} dy = \eta$ which one checks easily using the elementary properties of the Haar measure.

In order to prove (51), we still have to prove that $Z_{0,0,t}^+ = Z_{1,0,t}^+(1)$. This follows from the equality $Z_{0,0,t}^+ = \int_G Z_{1,0,t-s}^+(y) Z_{1,0,s}^+(y^{-1}) dy$, Proposition 4.1.6 and the argument of uniform convergence that we have already used above.

By (4.1) and the first part of this proof, we have for all $p, g \geq 0$, all $t > 0$ and all $x_1, \dots, x_p \in G$,

$$\begin{aligned} Z_{p,2g+1,t}^-(x_1, \dots, x_p) &= \eta^{*g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_p}} * \left(\int_G \delta_{\theta_{y_2}} dy \right) (Z_{1,0,t}^+) \\ &= \kappa * \eta^{*g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_p}} (Z_{1,0,t}^+), \end{aligned}$$

and

$$\begin{aligned} Z_{p,2g+2,t}^-(x_1, \dots, x_p) &= \eta^{*g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_p}} * \left(\int_{G^2} \delta_{\theta_{y_1}} * \delta_{\theta_{y_2}} dy_1 dy_2 \right) (Z_{1,0,t}^+) \\ &= \kappa^{*2} * \eta^{*g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_p}} (Z_{1,0,t}^+). \end{aligned}$$

The result is now a consequence of the last assertion of Proposition 4.1.8. \square

It is now easy to complete the result obtained in Proposition 4.1.6. We leave the proof of the following corollary to the reader.

Corollary 4.1.11. — *Let $(Z_{p,g,t}^\pm)_{p,g,t}$ be the partition functions of a regular Markovian holonomy field. Then, as t tends to 0, one has the following weak convergences of measures on G .*

1) For all $p \geq 0$, all $g \geq 0$ even and all $x_1, \dots, x_{p-1} \in G$,

$$Z_{p,g,t}^+(x_1, \dots, x_{p-1}, x^{-1}) dx \xrightarrow[t \rightarrow 0]{} \eta^{*\frac{1}{2}g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_{p-1}}}.$$

2) For all $p \geq 0$, all $g > 0$ and all $x_1, \dots, x_{p-1} \in G$,

$$Z_{p,g,t}^-(x_1, \dots, x_{p-1}, x^{-1}) dx \xrightarrow[t \rightarrow 0]{} \kappa^{*g} * \delta_{\theta_{x_1}} * \dots * \delta_{\theta_{x_{p-1}}}.$$

Remark 4.1.12. — In the proofs so far, we have used the axiom A_7 only for cylinders. Let us call A'_7 the axiom A_7 restricted to cylinders. By using Propositions 4.1.6 and 4.1.10, one could easily prove now that A_7 could be replaced by A'_7 without affecting the notion of regular Markovian holonomy field.

4.2. The Lévy process associated to a Markovian holonomy field

In the previous section, we have reduced the description of the partition functions of a regular Markovian holonomy field to the description of the one-parameter family of functions $Z_{1,0,t}^+ : G \rightarrow [0, +\infty)$, $t > 0$. This allows us to state a classification result.

Proposition 4.2.1. — *Let HF be a regular Markovian holonomy field. Then the probability measures $(Z_{1,0,t}^+(x) dx)_{t>0}$ on G are the one-dimensional distributions of a unique conjugation-invariant Lévy process issued from the unit element. If the Markovian field is not oriented, then the distribution of this Lévy process is invariant by*

inversion. Moreover, this Lévy process determines completely the partition functions of HF.

It is conceivable that a regular Markovian holonomy field is completely determined by its associated Lévy process, but we have warned the reader in the introduction that we are not yet able to settle this question.

Proof. — It suffices to prove that the probability measures $\nu_t = Z_{1,0,t}^+(x)dx$ form a convolution semigroup. Let us fix $s, t > 0$. We have $Z_{2,0,s}^+(x, y) = \int_G Z_{1,0,t}^+(xwyw^{-1})$ by (51). Now,

$$\begin{aligned} \nu_t * \nu_s(dy) &= \int_G Z_{1,0,t}^+(x)Z_{1,0,s}^+(x^{-1}y)dx dy \\ &= \int_{G^2} Z_{1,0,t}^+(wxw^{-1})Z_{1,0,s}^+(w^{-1}x^{-1}wy)dw dx dy \\ &= \int_G Z_{1,0,t}^+(x)Z_{2,0,s}^+(x^{-1}, y)dx dy \\ &= \beta_2(Z_{1,0,t}^+ \otimes Z_{2,0,s}^+)(y) = Z_{1,0,s+t}^+(y)dy = \nu_{t+s}(dy). \end{aligned}$$

Proposition 4.1.6 ensures that ν_t tends to the Dirac mass at the unit element as t tends to 0. Moreover, the conjugation invariance of the partition functions implies that the measure ν_t is invariant for all $t \geq 0$.

If the Markovian field is not orientable, then for all t , it follows from the axiom A_4 applied to an orientation-reversing diffeomorphism of a disk of area t that $Z_{1,0,t}^+(x) = Z_{1,0,t}^+(x^{-1})$. It follows that the 1-dimensional distributions of the Lévy process are invariant by inversion, hence the distribution of the process itself.

The fact that the measures $(\nu_t)_{t \geq 0}$ determine the partition functions is the main content of Proposition 4.1.10. □

Let us recall some classical facts about Lévy processes in compact Lie groups and use them to prove that the function $Z_{1,0,t}^+$ is positive on the connected component of the identity of G for all $t > 0$. Our constant reference in this section is the M. Liao [27] book.

Let X be an arbitrary Lévy process on G with a conjugation-invariant distribution. Let us describe briefly the generator of X . Let \mathfrak{g} be the Lie algebra of G . Let $A : G \rightarrow \mathfrak{g}$ be a smooth mapping such that $A(1) = 0$, $d_1A = \text{id}_{\mathfrak{g}}$ and, for all $x, y \in G$, $A(xy x^{-1}) = \text{Ad}(x)A(y)$. For example, let $r > 0$ be such that the exponential mapping is a diffeomorphism from the ball $B(0, r)$ in \mathfrak{g} to the ball $B(1, r)$ in G . Let \log denote the inverse mapping. Let $\varphi : [0, +\infty) \rightarrow [0, 1]$ be a smooth function with compact support contained in $[0, \frac{1}{2}r]$ and equal to 1 in a neighbourhood of 0. Then $A(x) = \log(x)\varphi(d_G(1, x))$ satisfies the required properties.

In what follows, we identify the elements of \mathfrak{g} with left-invariant vector fields on G .

Proposition 4.2.2. — *Let X be a Lévy process on G whose distribution is invariant by conjugation. Let $\{A_1, \dots, A_d\}$ denote a basis of \mathfrak{g} . Let \mathfrak{z} denote the centre of \mathfrak{g} . There exists a symmetric non-negative definite matrix $(a_{jk})_{j,k \in \{1, \dots, d\}}$, an element $A_0 \in \mathfrak{z}$, and a Borel measure Π on G which satisfies $\Pi(\{1\}) = 0$ and $\int_G d_G(1, x)^2 \Pi(dx) < +\infty$, such that the generator L of X is the following: for all $f \in C^2(G)$, all $g \in G$,*

$$Lf(g) = \frac{1}{2} \sum_{j,k=1}^d a_{jk} (A_j A_k f)(g) + (A_0 f)(g) + \int_G [f(gh) - f(g) - (A(h)f)(g)] \Pi(dh).$$

The differential operator

$$L_D = \frac{1}{2} \sum_{j,k=1}^d a_{jk} A_j A_k$$

and the measure Π are both invariant by conjugation. They are called respectively the diffusive part of the generator of X and the Lévy measure of X . Both are independent of the choice of the mapping $A : G \rightarrow \mathfrak{g}$.

Proof. — The unique point in which this presentation differs from that of [27] is the fact that $A_0 \in \mathfrak{z}$. The mapping $A : G \rightarrow \mathfrak{g}$ has been chosen to be equivariant under the adjoint action of G . This makes the third term of the generator invariant by conjugation. Since L_D is also invariant, the second term must be invariant as well. This implies that A_0 belongs to the invariant subspace of \mathfrak{g} under the adjoint action, that is, \mathfrak{z} . \square

Let us now assume that, for all $t > 0$, the distribution of X_t has a density with respect to the Haar measure on G , which we denote by Q_t . The function Q_t is a central function and, if X is invariant by inversion, it satisfies the relation $Q_t(x) = Q_t(x^{-1})$ for all $t > 0$ and all $x \in G$.

Let $\text{Irr}(G)$ denote the set of isomorphism classes of irreducible representations of G . For each $\alpha \in \text{Irr}(G)$, let $\chi_\alpha : G \rightarrow \mathbb{C}$ denote the character of α . Also, set

$$(53) \quad \lambda_\alpha = -\frac{(L_D \chi_\alpha)(1)}{\chi_\alpha(1)}, \quad \delta_\alpha = -\frac{(A_0 \chi_\alpha)(1)}{\chi_\alpha(1)}, \quad \pi_\alpha = \int_G \left(1 - \frac{\chi_\alpha(x)}{\chi_\alpha(1)}\right) \Pi(dx).$$

The results of [27, Chapter 4] show that Q_t is square-integrable for all $t > 0$ if and only if, for all $t > 0$,

$$(54) \quad \sum_{\alpha \in \text{Irr}(G)} e^{-(\lambda_\alpha + \delta_\alpha + \pi_\alpha)t} \chi_\alpha(1)^2 < +\infty.$$

It is also proved in [27] that Q_t is square-integrable for all $t > 0$ if and only if $(t, x) \mapsto Q_t(x)$ is continuous on $(0, +\infty) \times G$. Let us assume that these equivalent

properties are satisfied. In this case, the following expansion is uniformly absolutely convergent on $[\eta, +\infty) \times G$ for all $\eta > 0$:

$$(55) \quad Q_t(x) = \sum_{\alpha \in \text{Irr}(G)} e^{-(\lambda_\alpha + \delta_\alpha + \pi_\alpha)t} \chi_\alpha(1) \overline{\chi_\alpha(x)}.$$

In the following result, we use the compactness of G to prove that, in this situation, Q_t is positive for all $t > 0$.

Proposition 4.2.3. — *Let G be a compact connected Lie group. Let $(X_t)_{t \geq 0}$ be a Lévy process on G issued from 1 and invariant in law by conjugation. Assume that, for all $t > 0$, the distribution of X_t has a square-integrable density Q_t with respect to the Haar measure on G . Then $(t, x) \mapsto Q_t(x)$ is a continuous function on $(0, +\infty) \times G$ and*

$$\forall t > 0, \forall x \in G, \quad Q_t(x) > 0.$$

Proof. — The continuity property follows from the results presented above. We focus on the assertion of positivity.

We claim that there exists $t_0 > 0$ such that $Q_t(x) > 0$ for all $x \in G$ and all $t \geq t_0$. Indeed, since $\int_G Q_1(x) dx = 1$, there exists an open subset U of G on which Q_1 is positive. Hence, for all $n \geq 1$, Q_n is positive on U^n . Since G is a compact topological group, there exists $n_0 > 1$ such that $U^{n_0} = G$. Then $t_0 = n_0$ satisfies the expected property.

We write the generator of X as the sum of the diffusive part, a drift and the jump part:

$$L = L_D + A_0 + L_J.$$

Since A_0 belongs to \mathfrak{z} and L_D is invariant by conjugation, these three operators commute to each other. Let X^D and X^J be independent Lévy processes on G with respective generators L_D and L_J . Then we have the identity in distribution

$$(56) \quad \forall t > 0, \quad X_t \stackrel{(d)}{=} \exp(tA_0) X_t^D X_t^J.$$

The term $\exp(tA_0)$ modifies the subset of G where Q_t is positive by a simple translation. Hence, we may and will assume that $A_0 = 0$.

The topological support of the distribution of X_t^D does not depend on t . We denote it by D . It is the closure of the exponential of a Lie subalgebra of \mathfrak{g} which depends on L_D . Since L_D is invariant by conjugation, D is a closed normal subgroup of G .

The topological support of the distribution of X_t^J does not depend on t either and we denote it by J . It is the closure of the submonoid of G generated the topological support of Π . Since G is compact, the closure of the submonoid generated by any element of G contains the inverse of this element. Hence, J is also the closed subgroup

generated by the support of Π . Since Π is invariant, J is also a closed normal subgroup of G . In particular, $DJ = JD$ is a closed subgroup of G .

For each $t > 0$, set

$$S_t = \{x \in G : Q_t(x) > 0\}.$$

We claim that $\overline{S_t} = DJ\overline{S_t}$. Indeed, consider $x \in S_t$, $d \in D$ and $j \in J$. Let U, V, W be three open neighbourhoods of x, d, j respectively. We claim that $\int_{UVW} Q_t(y) dy > 0$. Since $(t, x) \mapsto Q_t(x)$ is continuous, there exists $\varepsilon > 0$ such that $Q_{t-\varepsilon}(x) > 0$. Now, $\mathbb{P}(X_{\frac{1}{2}\varepsilon}^D \in V) > 0$ and $P(X_{\frac{1}{2}\varepsilon}^J \in W) > 0$. Hence, by (56), $\mathbb{P}(X_\varepsilon \in VW) > 0$ and

$$\int_{UVW} Q_t(y) dy \geq \int_U Q_{t-\varepsilon}(x) dx \int_{VW} Q_\varepsilon(y) dy > 0.$$

Since this holds for any choice of U, V, W , the integral of Q_t over any neighbourhood of $x dj$ is positive. Hence, $x dj \in \overline{S_t}$. The claimed equality follows.

Now, it follows from (56), after the simplification $A_0 = 0$, that

$$\mathbb{P}(\forall t \geq 0, X_t \in DJ) = 1.$$

If the inclusion $DJ \subset G$ was a strict one, we would find a contradiction with the fact that Q_t is eventually everywhere positive on G . Hence, $DJ = G$.

Putting our results together, we find $\overline{S_t} = G$ for all $t > 0$. Now choose $t > 0$, $x \in G$ and consider the mapping $y \mapsto Q_{\frac{1}{2}t}(y)Q_{\frac{1}{2}t}(y^{-1}x)$. It vanishes on $(G - S_{\frac{1}{2}t}) \cup x(G - S_{\frac{1}{2}t})^{-1}$ which is the union of two closed sets with empty interior. This set has thus empty interior, so the mapping which we consider is continuous, non-negative and not identically zero. By integrating it with respect to y , we find $Q_t(x) > 0$. □

If G is not connected, let G^0 denote the connected component of 1. It is a normal subgroup of G and the quotient group G/G^0 is finite. The measure Π induces a measure on the group G/G^0 which is finite excepted possibly on the unit element. This measure, restricted to the complement of the unit element, is the jump measure of the projection of X on this finite group.

Corollary 4.2.4. — *Let G be a compact Lie group. Let $(X_t)_{t \geq 0}$ be a Lévy process on G which satisfies the assumptions of Proposition 4.2.3. Let G^0 denote the connected component of the unit element of G . Let H be the subgroup of G generated by G^0 and the support of Π . Then for all $t \geq 0$, $X_t \in H$ almost surely and for all $t > 0$, all $x \in H$, $Q_t(x) > 0$.*

Proof. — Let Π^0 denote the restriction of Π to G^0 . Since G^0 is a normal subgroup of G , both Π^0 and $\Pi - \Pi^0$ are Lévy measures on G invariant by conjugation. Moreover, $\Pi - \Pi^0$ is a finite measure. Let X^0 be the Lévy process whose generator is that of X in which Π is replaced by Π^0 . It is a Lévy process in G^0 . Let X^J be the pure jump

process with jump measure $\Pi - \Pi^0$. Then the generators of X^0 and X^J commute, so that we have in distribution, for all $t \geq 0$, $X_t \stackrel{d}{=} X_t^0 X_t^J$.

For all $\alpha \in \widehat{G}$, we have $\sup_{x \in G} |\chi_\alpha(x)| = \chi_\alpha(1)$. Hence, changing the measure Π by adding or subtracting to it a finite measure of mass m changes each coefficient π_α by at most $2m$. The condition (54) is not affected by such a change, so neither is the existence of a square-integrable density. This proves that the process X^0 satisfies the assumptions, hence the conclusions, of Proposition 4.2.3.

The set of connected components of G visited by the process X^J is the set of the connected components of the elements of the submonoid of G generated by the support of $\Pi - \Pi^0$. Since G is compact, this submonoid is also the subgroup generated by the same set. The conclusion follows easily. \square

Corollary 4.2.5. — *Let HF be a regular Markovian holonomy field. There exists a subgroup H of G which contains the connected component of the unit element and such that*

$$(57) \quad \forall t > 0, \forall p, g \geq 0, \forall x_1, \dots, x_p \in H, \quad Z_{p,g,t}^+(x_1, \dots, x_p) > 0.$$

Proof. — By Propositions 4.2.1 and 4.1.4, the Lévy process associated with HF satisfies the assumptions of Corollary 4.2.4. Hence, (57) holds for $Z_{1,0,t}^+$. The general case follows by Proposition 4.1.10. \square

From now on, we will always assume that $H = G$.

Definition 4.2.6. — *Let $(X_t)_{t \geq 0}$ be a Lévy process on G . We say that X is admissible if it is issued from 1, invariant in law by conjugation, and if for all $t > 0$ the distribution of X_t admits a continuous density Q_t with respect to the Haar measure on G , such that the function $(t, x) \mapsto Q_t(x)$ is continuous and positive on $\mathbb{R}_+^* \times G$.*

Let us discuss briefly the existence of a square-integrable density for the distribution of X . If G is a finite group, this condition is always satisfied. An admissible Lévy process in this case is simply a continuous-time random walk on G whose jump distribution is invariant by conjugation and has a support which generates G . In the case of the symmetric group, where every element is conjugated to its inverse, this invariance property implies that the jump distribution, hence the distribution of X , is also invariant by inversion.

If G is connected and $\dim G \geq 1$, an assumption under which the condition (54) is always satisfied is the ellipticity of X . In general, the hypoellipticity is sufficient to ensure the existence of a density, but a conjugation-invariant hypoelliptic process is necessarily elliptic. Indeed, if G is Abelian, hypoellipticity is equivalent to ellipticity and if G is simple, the invariance of X implies that the diffusive part of the generator

of X must be a non-negative multiple of the Laplace operator, hence elliptic or zero. The general case is a combination of these two.

In the case where the process is not elliptic, the distribution of X may or may not have a density, depending on the jumps of X . The discussion of ellipticity and hypo-ellipticity above implies that if X is not elliptic and has no jumps, then X has no square-integrable density. The remark made in the course of the proof of Corollary 4.2.4 implies that this is still true if the Lévy measure of X is finite.

Let us conclude this section by giving an example of an admissible pure jump processes. Let us work on $SU(2)$. Choose a real s and consider the measure

$$\Pi(dx) = d(1, x)^s dx.$$

Since the dimension of $SU(2)$ is 3, the integral $\int_{SU(2)} d(1, x)^2 \Pi(dx)$ converges if and only if $s > -5$ and Π is a finite measure for $s > -3$. The irreps of $SU(2)$ are labelled by their dimension, which can be any positive integer. Accordingly, the Fourier coefficient π_n , which is given, thanks to Weyl's integration formula, by

$$\pi_n = \frac{2}{\pi} \int_0^\pi \left(1 - \frac{\sin(n\theta)}{n \sin \theta}\right) \sin^2(\theta) \theta^s d\theta,$$

is non-negative and grows faster than a constant times n^{-s-3} . In particular, if $s < -3$, the series $\sum_{n \geq 1} e^{-\pi_n t}$ converges and the condition (54) is satisfied. Finally, for all $s \in (-5, -3)$, the pure jump process on $SU(2)$ with Lévy measure $\Pi(dx) = d(1, x)^s dx$ is admissible.

4.3. A Markovian holonomy field for each Lévy process

In this section, we prove the following theorem, which is the second main result of the present work. Recall Definitions 3.1.3 and 4.2.6, and Proposition 4.2.1.

Theorem 4.3.1. — *Every admissible Lévy process is the Lévy process associated to a regular oriented Markovian holonomy field.*

Every admissible Lévy process whose distribution is invariant by inversion is the Lévy process associated to a regular Markovian holonomy field.

Whether this regular Markovian holonomy field is unique is a natural question which we hope to be able to answer in a future work.

In order to prove this theorem, we use the results of the previous chapter. We start by constructing a discrete Markovian holonomy field, prove that it is regular and extend it to a Markovian holonomy field.

4.3.1. A discrete Markovian holonomy field. — Let X be an admissible Lévy process. Let Q_t denote the density of the distribution of X_t . Let $(M, \text{vol}, \mathcal{C}, C)$ be a connected measured marked surface with G -constraints. Let \mathbb{G} be a graph on (M, \mathcal{C}) . For each face F of \mathbb{G} , recall that ∂F is a cycle, oriented or non-oriented depending on the orientability of M . Assume first that M is orientable. For each $h \in \mathcal{M}(\mathcal{P}(\mathbb{G}), G)$, different choices of the origin of ∂F lead to different elements $h(\partial F)$ of G , but these elements belong to the same conjugacy class of G . Hence, for all $t > 0$, the assumption that the distribution of X_t is invariant by conjugation makes the positive real number $Q_t(h(\partial F))$ well-defined.

If M is non-orientable, then $h(\partial F)$ is defined only up to conjugation and inversion. In this case, we make the further assumption that the distribution of X is invariant by inversion. Then, for all $t > 0$, the non-negative real number $Q_t(h(\partial F))$ is also well-defined.

Definition 4.3.2. — Let X be an admissible Lévy process. Let $(M, \text{vol}, \mathcal{C}, C)$ be a measured marked surface with G -constraints. Let \mathbb{G} be a graph on (M, \mathcal{C}) . We define the following measure on $(\mathcal{M}(\mathcal{P}(\mathbb{G}), G), \mathcal{C})$:

$$\text{DF}_{M, \text{vol}, \mathcal{C}}^{X, \mathbb{G}}(dh) = \prod_{F \in \mathbb{F}} Q_{\text{vol}(F)}(h(\partial F)) \text{U}_{M, \mathcal{C}, C}^{\mathbb{G}}(dh).$$

We denote the collection of these measures by DF^X .

Proposition 4.3.3. — Let X be an admissible Lévy process. The collection of measures DF^X satisfies the axioms D_1 to D_6 of a discrete Markovian holonomy field.

Proof. — For each quadruple $(M, \text{vol}, \mathcal{C}, C)$, $\text{DF}_{M, \text{vol}, \mathcal{C}}^{X, \mathbb{G}}$ is a measure on the invariant σ -field of $\mathcal{M}(\mathcal{P}(\mathbb{G}), G)$. It has a bounded density with respect to the probability measure $\text{U}_{M, \mathcal{C}, C}^{\mathbb{G}}$, so that it is a finite measure. Let us prove that the axioms D_1 to D_6 are satisfied.

The fact that the discrete Markovian holonomy field U satisfies D_1 and D_3 implies immediately that DF^X also satisfies them. The argument used for U in the proof of Proposition 3.2.3 shows that DF^X satisfies D_2 . Let ψ be a homeomorphism as in the statement of the axiom D_4 . The measure $\text{DF}_{M, \text{vol}, \mathcal{C}}^{X, \mathbb{G}}$ depends only on the combinatorial structure of the graph \mathbb{G} , on the cycles which represent the curves of \mathcal{C} , on the set of G -constraints and finally on the areas of the faces of \mathbb{G} . These characteristics are all preserved by the homeomorphism ψ . The axiom D_5 is obviously satisfied. Let us finally check that DF^X satisfies D_6 .

Let us denote by M' the surface $\text{Spl}_l(M)$ and by \mathbb{G}' the graph $\text{Spl}_l(\mathbb{G})$. Let us also denote by $D' : \mathcal{M}(\mathcal{P}(\mathbb{G}'), G) \rightarrow \mathbb{R}$ the density of the measure $\text{DF}_{M', \text{vol}', \mathcal{C}'}^{X, \mathbb{G}'}$ with

respect to $U_{M', \mathcal{C}', C'}^{\mathbb{G}'}$. Then $D' \circ \psi : \mathcal{M}(P(\mathbb{G}), G) \rightarrow \mathbb{R}$ is the density of $DF_{M, \text{vol}, \mathcal{C}, C}^{X, \mathbb{G}}$ with respect to $U_{M, \mathcal{C}, C}^{\mathbb{G}}$. Hence, the property follows from Proposition 2.3.8. \square

We prove now that the collection of measures that we consider satisfies the property of invariance under subdivision.

Proposition 4.3.4. — *Let X be an admissible Lévy process. The collection of measures DF^X satisfies the axiom D_I of a discrete Markovian holonomy field.*

Proof. — Consider $(M, \text{vol}, \mathcal{C}, C)$ endowed with two graphs \mathbb{G}_1 and \mathbb{G}_2 such that $\mathbb{G}_1 \preceq \mathbb{G}_2$. Let $r : \mathcal{M}(P(\mathbb{G}_2), G) \rightarrow \mathcal{M}(P(\mathbb{G}_1), G)$ denote the restriction map. Let us first make the assumption that $\mathbb{E}_1 \subset \mathbb{E}_2$ and choose orientations \mathbb{E}_1^+ and \mathbb{E}_2^+ of \mathbb{G}_1 and \mathbb{G}_2 such that $\mathbb{E}_1^+ \subset \mathbb{E}_2^+$. The restriction map r can be thought of as a map from $G^{\mathbb{E}_2^+}$ to $G^{\mathbb{E}_1^+}$. Let us write $\mathbb{E}_2^+ = \mathbb{E}_1^+ \cup (\mathbb{E}_2 \setminus \mathbb{E}_1)^+$ and decompose the generic element of $G^{\mathbb{E}_2^+}$ as $g = (g_1, g_2)$ accordingly. With this notation, $r(g_1, g_2) = g_1$. Let $f : G^{\mathbb{E}_1^+} \rightarrow \mathbb{R}$ be a continuous function. We need to prove that

$$\int_{G^{\mathbb{E}_2^+}} f(g_1) DF_{M, \text{vol}, \mathcal{C}, C}^{X, \mathbb{G}_2}(dg_1, dg_2) = \int_{G^{\mathbb{E}_1^+}} f(g_1) DF_{M, \text{vol}, \mathcal{C}, C}^{X, \mathbb{G}_1}(dg_1).$$

We are going to perform the integration on the left-hand side with respect to dg_2 . For this, let us observe that the curves of \mathcal{C} belong to $P(\mathbb{G}_1)$, so that the measure $U_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}_2}(dg_1, dg_2)$ on $G^{\mathbb{E}_2^+}$ can be written as $U_{M, \text{vol}, \mathcal{C}, C}^{\mathbb{G}_1}(dg_1) dg_2$. This is in fact an instance of Proposition 2.3.9. Hence, if we put together the faces of \mathbb{G}_2 according to the face of \mathbb{G}_1 in which they are contained, we find the following expression for our integral:

$$(58) \quad \int_{G^{\mathbb{E}_1^+}} f(g_1) \prod_{F_1 \in \mathbb{F}_1} \left[\int \prod_{\substack{F_2 \in \mathbb{F}_2 \\ F_2 \subset F_1}} Q_{\text{vol}(F_2)}(g(\partial F_2)) \prod_{\substack{e \in \mathbb{E}_2^+ \setminus \mathbb{E}_1^+ \\ e \subset F_1}} dg_e \right] U_{M, \mathcal{C}, C}^{\mathbb{G}_1}(dg_1).$$

The integral between the brackets is over $G^{\{e \in \mathbb{E}_2^+ \setminus \mathbb{E}_1^+ : e \subset F_1\}}$. It suffices to prove that this integral is equal to $Q_{\text{vol}(F_1)}(h_{\partial F_1}(g))$.

We proceed by induction on the number of faces of \mathbb{G}_2 contained in F_1 . Let us assume first that this number is 1 and denote by F_2 the unique face of \mathbb{G}_2 contained in F_1 . In order to treat this case, we proceed by induction on the number of edges of \mathbb{G}_2 whose interior is contained in F_1 . If this number is zero, then $F_1 = F_2$ and the expression between the brackets is exactly $Q_{\text{vol}(F_1)}(g(\partial F_1))$. Now let us assume that there is at least one edge of \mathbb{G}_2 whose interior is contained in F_1 . Let us consider a split pattern M' of (M, \mathbb{G}_1) and let M'_{F_1} denote the connected component corresponding to F_1 . Let \mathbb{G}'_2 be the graph on M' induced by \mathbb{G}_2 . The restriction of \mathbb{G}'_2 to M'_{F_1} is a graph with a single face on a disk. By Euler's formula, this graph has the same number of edges and vertices. By assumption, there is at least one vertex of degree at least 3 on the boundary of M'_{F_1} . Hence, \mathbb{G}'_2 has at least one vertex of degree 1,

which must be in the interior of M'_{F_1} and hence is sent to a vertex of \mathbb{G}_2 of degree 1 contained in F_1 .




FIGURE 1. A graph with one face on a disk is a collection of trees grafted on the vertices located on the boundary of the disk. The leaves of these trees are vertices of degree 1.

Let e be an edge of \mathbb{G}_2 adjacent to a vertex v of degree 1. The cycle ∂F_2 contains either the sequence ee^{-1} or the sequence $e^{-1}e$. This sequence can be removed without affecting the value of the expression between the brackets. The cycle ∂F_2 with the sequence removed is the boundary of the face $F_2 \cup e((0, 1)) \cup \{v\}$ of the graph whose set of edges is $\mathbb{E}_2 \setminus \{e, e^{-1}\}$. This graph has one edge less inside F_1 than \mathbb{G}_2 and the result follows by induction.

Let $n \geq 2$ be an integer and let us assume that the result has been proved when F_1 contains at most $n - 1$ faces of \mathbb{G}_2 . Consider the case where F_1 contains n faces of \mathbb{G}_2 .

Let F_2, F'_2 be two distinct adjacent faces of \mathbb{G}_2 which are both contained in F_1 . The boundaries of F_2 and F'_2 are respectively of the form $e_1 \cdots e_k e''$ and $(e'')^{-1} e'_1 \cdots e'_l$, where $\{e_1, \dots, e_k, e'_1, \dots, e'_l\} \subset \mathbb{E}_2^+$ and $e'' \in \mathbb{E}_2^+ \setminus \mathbb{E}_1^+$. When we integrate with respect to the component of g corresponding to e'' between the brackets in (58), we find

$$\int_G Q_{\text{vol}(F_2)}(g_{e''} g_{e_k} \cdots g_{e_1}) Q_{\text{vol}(F'_2)}(g_{e'_1} \cdots g_{e'_l} (g_{e''})^{-1}) dg_{e''},$$

which, by the Markov property of X , is equal to $Q_{\text{vol}(F \cup F')}(g'_1 \cdots g'_l g_k \cdots g_1)$. We are thus reduced to the graph obtained from \mathbb{G}_2 by merging F_2 and F'_2 along the edge e'' . By Proposition 1.3.15, the result of this operation is indeed a graph. The induction hypothesis applied to this new graph yields the desired result.

Finally, let us treat the case where $\mathbb{E}_1 \not\subset \mathbb{E}_2$. In this case, there are vertices of \mathbb{G}_2 located on the edges of \mathbb{G}_1 which are not vertices of \mathbb{G}_1 . Adding these vertices to \mathbb{G}_1 and splitting its edges accordingly produces a new graph \mathbb{G} such that $\mathbb{G}_1 \preceq \mathbb{G} \preceq \mathbb{G}_2$ and $\mathbb{E} \subset \mathbb{E}_2$. It remains to prove that the restriction $r : \mathcal{M}(P(\mathbb{G}), G) \rightarrow \mathcal{M}(P(\mathbb{G}_1), G)$ sends the measure $DF_{M, \text{vol}, \mathcal{G}, C}^{X, \mathbb{G}}$ to $DF_{M, \text{vol}, \mathcal{G}, C}^{X, \mathbb{G}_1}$. This follows directly from (23). \square

In order to prove that DF^X is a discrete Markovian holonomy field, there remains to prove that it satisfies the axiom D_7 . Combining Propositions 3.2.5, 4.3.3 and 4.3.4, we find the following result.

Proposition 4.3.5. — *Let X be an admissible Lévy process. Let $Z_{M,\text{vol},\mathcal{C},C}^{X,G}$ denote the partition functions associated with the collection of measures DF^X . Let $(M, \text{vol}, \mathcal{C}, C)$ be a measured marked surface with G -constraints. Then $Z_{M,\text{vol},\mathcal{C},C}^{X,G}$ does not depend on the graph \mathbb{G} . We denote it by $Z_{M,\text{vol},\mathcal{C},C}^X$.*

We are now going to compute this partition function. For this, we start by associating a probability measures on G to each connected surface (M, \emptyset, C) with G -constraints along the boundary. Recall the definition of the measures η and κ from Definition 4.1.7 and the properties that they satisfy proved in Proposition 4.1.8.

Definition 4.3.6. — *Let (M, \emptyset, C) be a connected surface with G -constraints along the boundary.*

If M is oriented, write $\mathcal{B}^+(M) = \{b_1, \dots, b_p\}$ and set

$$m_{M,\emptyset,C} = \eta^{*\frac{1}{2}g(M)} * \delta_{C(b_1)} * \dots * \delta_{C(b_p)}.$$

If M is non-orientable, write $\mathcal{B}(M) = \{b_1^{\pm 1}, \dots, b_p^{\pm 1}\}$ and set

$$m_{M,\emptyset,C} = \kappa^{*g(M)} * \delta_{C(b_1)} * \dots * \delta_{C(b_p)}.$$

Remark 4.3.7. — 1) The second definition is meaningful thanks to Lemma 4.1.8. Indeed, the orientation chosen on the boundary components of M does not affect the definition of $m_{M,\emptyset,C}$.

2) Assume that (M, \emptyset, C) is an oriented surface with G -constraints. Write M^\vee for the same surface with the opposite orientation. Then $m_{M^\vee,\emptyset,C} = m_{M,\emptyset,C}^\vee$.

Lemma 4.3.8. — *Let M be a connected compact surface.*

1) *The mapping which to a set C of G -constraints on $\mathcal{B}(M)$ associates the probability measure $m_{M,\emptyset,C}$ on G is continuous from $\text{Const}_G(M, \emptyset)$ to the space of probability measures on G endowed with the topology of weak convergence.*

2) *For all $b \in \mathcal{B}(M)$, the measure $\int_G m_{M,\mathcal{C},C_{b \rightarrow x}} dx$ is the Haar measure on G .*

Proof. — 1) This property follows from the continuity of the mapping $x \mapsto \delta_{\theta_x}$ and the continuity of the convolution of measures.

2) This follows from the fact that $\int_G \delta_{\theta_x} dx$ is the Haar measure on G . □

Proposition 4.3.9. — *Let X be an admissible Lévy process. Consider the collection of measures DF^X . Let $(M, \text{vol}, \mathcal{C}, C)$ be a connected measured marked surface. Let $(M', \text{vol}', \emptyset, C')$ be a split tubular pattern of (M, \mathcal{C}, C) endowed with the induced G -constraints. Let M'_1, \dots, M'_s be the connected components of M' and for each i*

in $\{1, \dots, s\}$, let (M'_i, \emptyset, C'_i) be the associated connected surface with constraints. If M is oriented, then M' carries the induced orientation. If M is non-orientable, then let us assign an arbitrary orientation to each orientable connected component of M' . Then the following equality holds:

$$(59) \quad Z_{M, \text{vol}, \mathcal{C}, C}^X = \prod_{i=1}^s \int_G Q_{\text{vol}'(M'_i)} \, d\mathbf{m}_{M'_i, \emptyset, C'_i}.$$

Proof. — When M is non-orientable, there is a choice made in assigning an orientation to each orientable connected component of M' . However, in this case, the distribution of X , hence the function Q_t for all $t > 0$, is invariant by inversion. Hence, by the second part of Remark 4.3.7, the right-hand side of (59) is well defined.

The proof of this equality is of the same vein as that of Proposition 3.2.5. By D_6 , $Z_{M, \text{vol}, \mathcal{C}, C}^X = Z_{M', \text{vol}', \emptyset, C'}^X$. Then, by D_5 , $Z_{M', \text{vol}', \emptyset, C'}^X = \prod_{i=1}^s Z_{M'_i, \text{vol}', \emptyset, C'_i}^X$. The problem is thus reduced to the case of a connected surface with G -constraints along the boundary.

Let us assume that M is connected and $\mathcal{C} = \emptyset$. In order to compute the partition function in this case, we choose a graph on M with a single face and, by cutting and pasting, transform it so that the boundary of its unique face has a canonical form. Then, we find, if M for instance is non-orientable of reduced genus $g \geq 1$ with p boundary components,

$$\begin{aligned} Z_{M, \text{vol}, \emptyset, C}^X &= \int_{G^{g+2p}} Q_{\text{vol}(M)}(a_1^2 \cdots a_g^2 u_1 c_1 u_1^{-1} \cdots u_p c_p u_p^{-1}) \\ &\quad da_1 \cdots da_g du_1 \cdots du_p \delta_C(b_1)(dc_1) \cdots \delta_C(b_p)(dc_p) \\ &= \int_{G^{g+p}} Q_{\text{vol}(M)}(a_1^2 \cdots a_g^2 c_1 \cdots c_p) da_1 \cdots da_g \delta_C(b_1)(dc_1) \cdots \delta_C(b_p)(dc_p) \\ &= \int_G Q_{\text{vol}(M)} \, d\mathbf{m}_{M, \emptyset, C}. \end{aligned}$$

The other cases are similar. □

We can summarise our results.

Proposition 4.3.10. — *Let X be an admissible Lévy process. The collection of measures DF^X defined in Definition 4.3.2 is a discrete Markovian holonomy field.*

Proof. — For all $(M, \text{vol}, \mathcal{C}, C)$, the measure $DF_{M, \text{vol}, \mathcal{C}, C}^{X, G}$ is a measure on the cylinder σ -field of $\mathcal{M}(P(G), G)$, hence it determines by restriction a measure on the invariant σ -field. By Proposition 4.3.9, it is a finite measure.

The collections of these measures satisfies the axioms D_1 to D_6 by Proposition 4.3.3, D_I by Proposition 4.3.4 and D_7 by the combination of Proposition 4.3.9, the second assertion of Lemma 4.3.8 and the fact that for all $t > 0$, $\int_G Q_t(x) dx = 1$. □

4.3.2. A Markovian holonomy field. — Our next goal is to prove that the discrete Markovian holonomy field DF^X is regular in the sense of Definition 3.2.8.

Proposition 4.3.11. — *Let X be an admissible Lévy process. The discrete Markovian holonomy field DF^X is continuously area-dependent and Fellerian.*

Proof. — For all segment $[s, t] \subset \mathbb{R}_+^*$, the mapping $(t, x) \mapsto Q_t(G)$ is uniformly continuous on $[s, t] \times G$. Hence, for all face $F \in \mathbb{F}$, $Q_{\text{vol}_n(F)}$ converges uniformly to $Q_{\text{vol}(F)}$ as n tends to infinity. The fact that DF^X is continuously area-dependent follows immediately.

The fact that DF^X is Fellerian, that is, that the partition function $Z_{M, \text{vol}, \mathcal{C}}^X$ depends continuously on $C \in \text{Const}_G(M, \mathcal{C})$, follows at once from Proposition 4.3.9 and the first assertion of Corollary 4.3.8. \square

In order to prove that DF^X is stochastically $\frac{1}{2}$ -Hölder continuous, we need to establish the corresponding property for the Lévy process X .

Proposition 4.3.12. — *Let $(X_t)_{t \geq 0}$ be a Lévy process on the compact Lie group G issued from 1. Then there exists a constant K such that*

$$\forall t \geq 0, \quad \mathbb{E}[d_G(1, X_t)] \leq K\sqrt{t}.$$

This property follows from Lemma 3.5 in the book of M. Liao [27], but we still offer a short proof.

Proof. — We use the Itô formula for Lévy process on Lie groups, which has been proved by Applebaum and Kunita [3]. We borrow the statement from [27, Section 1.4]. In fact we use the following weak statement. Let L be the generator of X . Let f be a smooth function on G . Then f belongs to the domain of L and

$$(60) \quad f(X_t) - f(X_0) - \int_0^t Lf(X_s) ds, \quad t \geq 0$$

is a square-integrable martingale. This is the equation (1.18) of [27]. We apply this to a function f which is close to the function $d_G(1, \cdot)^2$.

Let $\{A_1, \dots, A_d\}$ be a basis of the Lie algebra of G , which we identify with the space of left-invariant vector fields on G . Let a_1, \dots, a_d be smooth functions on G such that for all $i, j \in \{1, \dots, d\}$, $a_i(1) = 0$ and $A_i a_j(1) = \delta_{ij}$. Set $\delta = \sum_{i=1}^d a_i^2$. It follows readily from the definition of δ and the fact that G is compact that there exists a constant K_1 such that for all $x \in G$, $d_G(1, x)^2 \leq K_1 \delta(x)$. Now the fact that (60) is a martingale implies that $\mathbb{E}[\delta(X_t)] \leq \|L\delta\|_\infty t$. Hence, by Jensen's inequality, $\mathbb{E}[d_G(1, X_t)]^2 \leq K_1 \|L\delta\|_\infty t$ for all $t \geq 0$. \square

In order to deduce the stochastic Hölder continuity of DF^X from this property, we need to be able to compare the values of some integrals with and without G -constraints. Actually, we introduce random holonomy fields with free boundary conditions. Recall the definition of the uniform measure $U_{M,\rho}^G$ (Definition 2.3.3).

Definition 4.3.13. — *Let X be an admissible Lévy process. Let (M, vol) be a measured surface endowed with a graph G . We define the measure $DF_{M,\text{vol},\emptyset,\emptyset}^{X,G}$ by setting*

$$DF_{M,\text{vol},\emptyset,\emptyset}^{X,G}(dh) = \prod_{F \in \mathbb{F}} Q_{\text{vol}(F)}(h(\partial F)) U_{M,\emptyset}^G(dh).$$

In the following proofs, we use the fact that the functions $t \mapsto \sup\{Q_t(x) : x \in G\}$ and $t \mapsto \inf\{Q_t(x) : x \in G\}$ are respectively non-increasing and non-decreasing. This follows from $(Q_t)_{t>0}$ being a convolution semigroup of positive continuous functions.

Lemma 4.3.14. — *Let $(M, \text{vol}, \mathcal{C}, C)$ be a measured marked surface with G -constraints. Let G be a graph on (M, \mathcal{C}) . Consider $\mathbb{E}_1 \subset \mathbb{E}$ and $\mathbb{F}_1 \subset \mathbb{F}$. Assume that $\mathbb{E}_1 = \mathbb{E}_1^{-1}$. Assume that for each $l \in \mathcal{C} \cup \mathcal{B}(M)$, at least one edge of \mathbb{E}_1 is located on l , and each face adjacent to an edge of \mathbb{E}_1 belongs to \mathbb{F}_1 . Set $r = \#\mathbb{F}_1$,*

$$A = \min \{ \text{vol}(F) : F \in \mathbb{F}_1 \} \quad \text{and} \quad K = \sup \{ Q_A(x)/Q_A(y) : x, y \in G \}.$$

Let $f : \mathcal{M}(\mathbb{E}, G) \rightarrow [0, +\infty)$ be a non-negative continuous function. Assume that f factorises through the restriction map $\mathcal{M}(\mathbb{E}, G) \rightarrow \mathcal{M}(\mathbb{E} \setminus \mathbb{E}_1, G)$. Then

$$K^{-r} \int_{\mathcal{M}(\mathbb{E}, G)} f dDF_{M,\text{vol},\emptyset,\emptyset}^{X,G} \leq \int_{\mathcal{M}(\mathbb{E}, G)} f dDF_{M,\text{vol},\mathcal{C},C}^{X,G} \leq K^r \int_{\mathcal{M}(\mathbb{E}, G)} f dDF_{M,\text{vol},\emptyset,\emptyset}^{X,G}.$$

Proof. — Increasing the number of edges in \mathbb{E}_1 can only increase \mathbb{F}_1 , hence make A smaller and K larger. So, without loss of generality, we may assume that \mathbb{E}_1 contains exactly one non-oriented edge on each curve $l \in \mathcal{C} \cup \mathcal{B}(M)$ and \mathbb{F}_1 is exactly the set of faces adjacent to these non-oriented edges.

Let us choose an orientation \mathbb{E}^+ of G and identify $\mathcal{M}(\mathbb{E}, G)$ with $G^{\mathbb{E}^+}$. Let us enumerate \mathbb{F}_1 as $\{F_1, \dots, F_r\}$ and denote the generic element of $G^{\mathbb{E}^+}$ as $g = (g_1, g_2)$ according to the partition $\mathbb{E} = \mathbb{E}_1 \cup (\mathbb{E} \setminus \mathbb{E}_1)$. The assumption on f expresses that $f(g_1, g_2)$ depends only on g_2 .

By (21), the integration against $U_{M,\mathcal{C},C}^G$ can be decomposed into the integration with respect to the Haar measure on $\mathcal{M}(\mathbb{E} \setminus \mathbb{E}_1, G)$ and then with respect to the explicitly known conditional distribution of g_1 given g_2 , which we denote by $U_{M,\mathcal{C},C}^G(dg_1|g_2)$.

$$\begin{aligned} & \int_{\mathcal{M}(\mathbb{E},G)} f dF_{M,\text{vol},\mathcal{C},C}^{X,G} \\ &= \int_{\mathcal{M}(\mathbb{E},G)} f(g) \prod_{F \in \mathbb{F}} Q_{\text{vol}(F)}(g(\partial F)) U_{M,\mathcal{C},C}^G(dg) \\ &= \int_{G(\mathbb{E} \setminus \mathbb{E}_1)^+} f(g_2) \prod_{F \in \mathbb{F} \setminus \mathbb{F}_1} Q_{\text{vol}(F)}(g_2(\partial F)) \left[\int_{G\mathbb{E}_1^+} \prod_{F \in \mathbb{F}_1} Q_{\text{vol}(F)}(g(\partial F)) U_{M,\mathcal{C},C}^G(dg_1|g_2) \right] dg_2. \end{aligned}$$

Changing the probability measure with respect to which the integral between the brackets is taken can at most multiply the integral by $\max u / \min u$ and at most divide it by the same number, where u denotes the integrand. In the present situation, the definition of K implies that $\max u / \min u \leq K^r$.

Hence, focusing for example on the upper bound, we have

$$\begin{aligned} \int_{\mathcal{M}(\mathbb{E},G)} f dF_{M,\text{vol},\mathcal{C},C}^{X,G} &\leq K^r \int_{G\mathbb{E}^+} f(g_2) \prod_{F \in \mathbb{F}} Q_{\text{vol}(F)}(g(\partial F)) dg_1 dg_2 \\ &= K^r \int_{\mathcal{M}(\mathbb{E},G)} f dDF_{M,\text{vol},\emptyset,\emptyset}^{X,G}. \end{aligned}$$

The derivation of the lower bound is similar. □

Proposition 4.3.15. — *Let X be an admissible Lévy process. The discrete Markovian field DF^X is stochastically $\frac{1}{2}$ -Hölder continuous.*

Proof. — Let $(M, \text{vol}, \gamma, \mathcal{C}, C)$ be a measured marked surface with G -constraints. Write $\mathcal{C} \cup \mathcal{B}(M) = \{l_1, l_1^{-1}, \dots, l_q, l_q^{-1}\}$. Let M_1, \dots, M_s denote the connected components of $M \setminus \mathcal{C}$. Set $A = \frac{1}{2} \min\{\text{vol}(M_i) : i \in \{1, \dots, s\}\}$. For each $i \in \{1, \dots, q\}$, let us write l_i as the product of three edges: $l_i = e_{i,1}e_{i,2}e_{i,3}$. Let $L > 0$ be such that any Riemannian ball of radius smaller than L intersects at most one curve of $\mathcal{C} \cup \mathcal{B}(M)$ and at most two of the edges $\{e_{i,j} : i \in \{1, \dots, q\}, j \in \{1, 2, 3\}\}$, and has a Riemannian area smaller than A .

Let l be a piecewise geodesic loop such that $\ell(l) < L$ and l bounds a disk which we denote by D . If l bounds two disks, we choose the one included in the ball of radius L centred at the basepoint of l . By assumption on L , there is at most one $i \in \{1, \dots, q\}$ such that l meets one of the edges $e_{i,1}, e_{i,2}, e_{i,3}$ and it does not meet the three of them. We may assume that l meets none of the edges $e_{i,j}$ except possibly $e_{1,2}$ and $e_{1,3}$.

Let \mathbb{G} be a graph on (M, \mathcal{C}) such that the edges $e_{i,j}$ are edges of \mathbb{G} . By repeated applications of Proposition 1.3.15, we may assume that \mathbb{G} has exactly one face in each connected component of $M \setminus (\mathcal{C} \cup \{l\})$. The number of these components depends on

the relative position of l and the curve l_1 . Nevertheless, \mathbb{G} has the following property: each face adjacent to one of the edges $e_{1,1}, e_{2,1}, \dots, e_{p,1}$ or their inverses has an area greater or equal to A .

Let us define $\mathbb{E}_1 = \{e_{1,1}, e_{1,1}^{-1}, \dots, e_{q,1}, e_{q,1}^{-1}\}$ and \mathbb{F}_1 as the subset of \mathbb{F} consisting of all faces adjacent to an edge of \mathbb{E}_1 . We have $\# F_1 \leq 2q$. Set

$$K_{X,A} = \sup \{Q_A(x)/Q_A(y) : x, y \in G\}.$$

Then, by Lemma 4.3.14,

$$(61) \quad \int_{\mathcal{M}(\mathbb{P}(\mathbb{G}), G)} d_G(1, h(l)) DF_{M, \text{vol}, \emptyset, C}^{X, \mathbb{G}}(dh) \leq K_{X,A}^{2q} \int_{\mathcal{M}(\mathbb{P}(\mathbb{G}), G)} d_G(1, h(l)) DF_{M, \text{vol}, \emptyset, \emptyset}^{X, \mathbb{G}}(dh).$$

By the axiom D_I , we can remove edges from \mathbb{G} so that it becomes a graph \mathbb{G}_1 with only two faces, D and another one, denoted by F , of area $\text{vol}(M \setminus D)$, without altering the value of the integral above. Hence, by Proposition 4.3.12,

l.h.s. of (61)

$$\begin{aligned} &\leq K_{X,A}^{2q} \int_{\mathcal{M}(\mathbb{P}(\mathbb{G}_1), G)} d_G(1, h(l)) Q_{\text{vol}(D)}(h(l)) Q_{\text{vol}(M-D)}(h(\partial F)) U_{M, \emptyset}^{\mathbb{G}}(dh) \\ &\leq K_{X,A}^{2q+1} \int_G d_G(1, x) Q_{\text{vol}(D)}(x) dx \\ &= K_{X,A}^{2q+1} \mathbb{E}[d_G(1, X_{\text{vol}(D)})] \leq K K_{X,A}^{2q+1} \sqrt{\text{vol}(D)}. \end{aligned}$$

This is the expected result. □

We can conclude this chapter by proving Theorem 4.3.1.

Proof of Theorem 4.3.1. — Let X be an admissible Lévy process. Let DF^X be defined by Definition 4.3.2. By Proposition 4.3.10, it is an oriented discrete Markovian holonomy field, and a discrete Markovian holonomy field if X is invariant in law by inversion. By Propositions 4.3.11 and 4.3.15, it is regular. By Theorem 3.2.9, DF^X is the restriction of a regular Markovian holonomy field, which we denote by HF^X . By Proposition 4.3.9, the Lévy process associated with HF^X is indeed X . □

CHAPTER 5

RANDOM RAMIFIED COVERINGS

In this chapter, we investigate the Markovian holonomy field that we have associated to a Lévy process in the case where G is a finite group. In this case, the structure of the Lévy process is particularly simple. It is a continuous time random walk with a jump distribution invariant by conjugation and, depending on the issue of orientability, by inversion.

It turns out that in this case, the canonical process associated to the Markovian holonomy field is the process of monodromy in a random ramified covering picked under a probability measure which depends in a simple way on the Lévy process. This is consistent with the usual heuristic interpretation of the Yang-Mills measure as a probability measure on the space of connections on a principal bundle. Indeed, ramified coverings can be naturally interpreted as discrete models for principal bundles, endowed with a connection which is flat everywhere but at the ramification points, where its curvature is concentrated.

5.1. Ramified G -bundles

Let us choose once for all a finite group G . Let $(M, \text{vol}, \emptyset, \mathcal{C})$ be a measured surface with G -constraints on the boundary. For the sake of simplicity, we treat the case $\mathcal{C} = \emptyset$.

Let $Y \subset M \setminus \partial M$ be a finite subset. A *principal G -bundle* over $M - Y$ is a smooth covering $\pi : P \rightarrow M \setminus Y$ of $M \setminus Y$ by a surface P on which G acts freely on the right, by smooth automorphisms of covering and transitively on each fibre. The surface P is not compact unless $Y = \emptyset$ and in general it is not connected. Two G -bundles $\pi : P \rightarrow M \setminus Y$ and $\pi' : P' \rightarrow M \setminus Y$ are isomorphic if there exists a G -equivariant diffeomorphism $h : P \rightarrow P'$ such that $\pi' \circ h = \pi$.

A *ramified covering* over M with ramification locus Y is a continuous mapping $\pi : P \rightarrow M$ from a surface P such that the restriction of π to $\pi^{-1}(M \setminus Y)$ is a covering and, for all $y \in Y$ and all $p \in \pi^{-1}(y)$, there exists a neighbourhood U of p

and an integer $n \geq 1$ such that the mapping $\pi|_U : (U, p) \rightarrow (\pi(U), y)$ is topologically conjugated to the mapping $z \mapsto z^n : (\mathbb{C}, 0) \rightarrow (\mathbb{C}, 0)$. This means that there exists homeomorphisms $h_1 : U \rightarrow \mathbb{C}$ and $h_2 : \pi(U) \rightarrow \mathbb{C}$ such that $h_2 \circ \pi|_U \circ h_1^{-1}$ is the mapping $z \mapsto z^n$. The integer n is called the order of ramification of p . We assume that for all $y \in Y$, there exists $p \in \pi^{-1}(y)$ whose order of ramification is at least 2. Two ramified coverings $\pi : P \rightarrow M$ and $\pi' : P' \rightarrow M$ are isomorphic if there exists a homeomorphism $h : P \rightarrow P'$ such that $\pi' \circ h = \pi$.

From the classical fact that the only connected coverings of finite degree of \mathbb{C}^* are, up to isomorphism, the mappings $z \mapsto z^n : \mathbb{C}^* \rightarrow \mathbb{C}^*$ for $n \geq 1$, it follows that a principal G -bundle $\pi : P \rightarrow M \setminus Y$ can always be extended to a ramified covering of M by a suitable compactification of P , and that any two such extensions give rise to isomorphic ramified coverings. Moreover, it is possible to endow the total space of the ramified covering with a differentiable structure in such a way that the covering map is smooth.

Definition 5.1.1. — A ramified principal G -bundle over M with ramification locus Y is a smooth ramified covering $\pi : R \rightarrow M$ of M with ramification locus Y , together with an action of G on $\pi^{-1}(M \setminus Y)$ which endows the restriction of π to $\pi^{-1}(M \setminus Y)$ with the structure of a principal G -bundle.

Two ramified principal G -bundles $\pi : R \rightarrow M$ and $\pi : R' \rightarrow M$ with ramification locus Y are isomorphic if their restrictions to $M \setminus Y$ are isomorphic as principal G -bundles.

Remark 5.1.2. — By the discussion before the definition, two isomorphic ramified G -bundles are also isomorphic as ramified coverings. However, an isomorphism of ramified coverings between two ramified G -bundles is not necessarily an isomorphism of ramified G -bundles. Consider for example, for $n \geq 3$, the trivial $\mathbb{Z}/n\mathbb{Z}$ -bundle $R = M \times \mathbb{Z}/n\mathbb{Z}$. The group \mathfrak{S}_n acts on R by permuting the sheets and this is an action by automorphisms of covering. Nevertheless, only a cyclic permutation of the sheets is an isomorphism of $\mathbb{Z}/n\mathbb{Z}$ -bundle. In general, the group of covering automorphisms of a principal G -bundle is bigger than G . This is related to the fact that the total space of the covering is not always connected, so that the group of covering automorphisms does not always act freely.

Remark 5.1.3. — In the case where G is the symmetric group \mathfrak{S}_n , ramified G -bundles are the same thing as ramified coverings of degree n . Indeed, let $\pi : P \rightarrow M$ be a principal \mathfrak{S}_n -bundle, ramified over Y . The group \mathfrak{S}_n acts naturally on $\{1, \dots, n\}$ and the associated bundle $P \times_{\mathfrak{S}_n} \{1, \dots, n\}$, which is the quotient of $P \times \{1, \dots, n\}$ by the relation

$$(p, k) \simeq (p', k') \iff \exists \sigma \in \mathfrak{S}_n, (p', k') = (p\sigma, \sigma^{-1}(k)),$$

is a ramified covering of M of degree n with ramification locus Y .

Conversely, let $\pi : R \rightarrow M$ be a ramified covering of degree n with ramification locus Y . Let m be a point of $M \setminus Y$. By a *labelling* of R at m we mean a bijection between $\pi^{-1}(m)$ and the set $\{1, \dots, n\}$. For each $m \in M \setminus Y$, let $\text{Lab}_m(R)$ denote the set of labellings of R at m . The group \mathfrak{S}_n acts on the right, transitively and freely on $\text{Lab}_m(R)$. Then $\text{Lab}(R) = \cup_{m \in M} \text{Lab}_m(R)$ endowed with the natural topology and projection on M , is a principal \mathfrak{S}_n -bundle over M ramified over Y . At a point $y \in Y$, the ramification type of R is a partition of n and the monodromy of $\text{Lab}(R)$ around y is the corresponding conjugacy class of \mathfrak{S}_n .

It is easy to check that, if P is a \mathfrak{S}_n -bundle, then $\text{Lab}(P \times_{\mathfrak{S}_n} \{1, \dots, n\})$ is canonically isomorphic to P .

Let $\mathcal{R}(M)$ (resp. $\mathcal{R}(M, Y)$) denote the set of isomorphism classes of ramified principal G -bundles over M (resp. with ramification locus Y).

Definition 5.1.4. — *A based ramified G -bundle is a pair (R, p) where $\pi : R \rightarrow M$ is a ramified bundle and $p \in R$ is a point such that $\pi(p)$ does not belong to the ramification locus of R . The pair (R, p) is said to be based at $\pi(p)$.*

Two based ramified G -bundles (R, p) and (R', p') are isomorphic if there exists an isomorphism $f : R \rightarrow R'$ of ramified G -bundles such that $f(p) = p'$.

The importance of this notion comes from the fact that the automorphism group of a based ramified bundle over a connected surface is trivial.

Let m be a point of M . We denote by $\mathcal{R}_m(M)$ (resp. $\mathcal{R}_m(M, Y)$) the set of isomorphism classes of based ramified G -bundles based at m (resp. with ramification locus Y).

5.2. Monodromy of ramified G -bundles

From now on we will assume that the surface M is connected. Consider R in $\mathcal{R}(M, Y)$. Choose $m \in M \setminus Y$. Choose $p \in \pi^{-1}(m)$. For each loop $l \in \mathbb{L}_m(M)$ which does not meet Y , the lift of l starting at p finishes at pg for a unique $g \in G$, called the monodromy of R along l with respect to p . This monodromy depends on l only through its homotopy class in $M \setminus Y$. Hence, the choice of p determines a group homomorphism $\text{mon}_p : \pi_1(M \setminus Y, m) \rightarrow G$, which characterises the based ramified G -bundle (R, p) up to isomorphism. Another choice of p would lead to another homomorphism, which differs from mon_p by composition by an inner automorphism of G . The class of the homomorphism mon_p modulo inner automorphisms of G characterises the ramified G -bundle R up to isomorphism (see [36]).

The words *monodromy* and *holonomy* are synonymous in this work, but we use the first in the context of ramified bundles and the second in the framework of Markovian holonomy fields.

An automorphism of the ramified G -bundle R is completely determined by the point to which it sends p . This point is of the form pg for a unique $g \in G$, hence the choice of p allows us also to identify the group $\text{Aut}(R)$ of automorphisms of R with a subgroup of G which we denote by $\text{Aut}_p(R)$. Let $\text{Mon}_p(R)$ be the image of the homomorphism $\text{mon}_p : \pi_1(M \setminus Y, m) \rightarrow G$. Then $\text{Aut}_p(R)$ is the centraliser of $\text{Mon}_p(R)$. Again, changing p to ph for some $h \in G$ would conjugate $\text{Mon}_p(R)$ and $\text{Aut}_p(R)$ by h .

In order to study $\mathcal{R}(M, Y)$, it is convenient to choose a system of generators of $\pi_1(M \setminus Y, m)$. To do this, let us first assume that Y is not empty and set $k = \# Y$. Let us choose on M a graph \mathbb{G} such that m is a vertex of \mathbb{G} and each face of \mathbb{G} contains exactly one point of Y . Throughout this chapter, we use the notation $g = g(M)$ for the genus of M , $p = p(M)$ for the number of boundary components of M , and $f = f(\mathbb{G})$ for the number of faces of \mathbb{G} . In the present situation, $f = k$. By Lemma 1.3.33, the group $\text{RL}_v(\mathbb{G})$ of reduced loops in \mathbb{G} (see Definition 1.3.30) is naturally isomorphic to $\pi_1(M \setminus Y, m)$. Thanks to Proposition 2.4.2, let us choose a tame system $\mathcal{S} = \{a_1, \dots, a_g, c_1, \dots, c_p, l_1, \dots, l_k\}$ of generators of $\text{RL}_v(\mathbb{G})$, associated with a certain word w in the free group of rank g . This means that the loops a_1, \dots, a_g generate the fundamental group of the surface obtained by gluing a disk along each boundary component of M , the loops c_1, \dots, c_p are lassos whose meanders cover the boundary components of M , the loops l_1, \dots, l_k are lassos whose meanders cover the boundary of the faces of \mathbb{G} , and all these loops generate $\text{RL}_v(\mathbb{G})$ with the single relation $w(a_1, \dots, a_g)c_1 \cdots c_p = l_1 \cdots l_k$. For all ramification point $y \in Y$, we denote by $\mathcal{O}(R, y)$ the conjugacy class of the monodromy along the facial lasso whose meander goes around y . This is also the conjugacy class of the monodromy along any small loop which circles once around y , positively if M is oriented. In particular, it does not depend on the choice of \mathbb{G} .

Recall that the surface M is endowed with a set C of G -constraints along its boundary. Thus, to each oriented connected component b of ∂M , the G -constraints C associate a conjugacy class $C(b)$ of G . Let us write $\mathcal{B}(M) = \{b_1, b_1^{-1}, \dots, b_p, b_p^{-1}\}$ and, for all $i \in \{1, \dots, p\}$, $\mathcal{O}_i = C(b_i)$. We define the sets $\mathcal{R}(M, C)$, (resp. $\mathcal{R}(M, Y, C)$), $\mathcal{R}_m(M, C)$, $\mathcal{R}_m(M, Y, C)$) as the sets of isomorphism classes of ramified G -bundles (resp. with ramification locus Y , based at m , based at m with ramification locus Y) such that the monodromy along b_i belongs to \mathcal{O}_i for all $i \in \{1, \dots, p\}$.

Let us use the information gathered so far to build concrete models for the various spaces of isomorphism classes of ramified G -bundles. Let us define

$$\mathcal{H}(M, k, C, w) = \{ (a_1, \dots, a_g, c_1, \dots, c_p, d_1, \dots, d_k) \in G^g \times \mathcal{O}_1 \times \dots \times \mathcal{O}_p \times (G \setminus \{1\})^k : w(a_1, \dots, a_g)c_1 \dots c_p = d_1 \dots d_k \}.$$

We denote by $a_i, c_i, d_i : \mathcal{H}(M, k, C, w) \rightarrow G$ the obvious coordinate mappings.

If Y is empty, then we choose \mathbb{G} with a single face. In this case the appropriate concrete model is the following space:

$$\mathcal{H}(M, 0, C, w) = \{ (a_1, \dots, a_g, c_1, \dots, c_p) \in G^g \times \mathcal{O}_1 \times \dots \times \mathcal{O}_p : w(a_1, \dots, a_g)c_1 \dots c_p = 1 \}.$$

The group G acts on $\mathcal{H}(M, k, C, w)$ by simultaneous conjugation on each factor. Let us consider the diagram

$$(62) \quad \begin{array}{ccc} \mathcal{R}_m(M, Y, C) & \xrightarrow{\sim} & \mathcal{H}(M, k, C, w) \\ \downarrow & & \downarrow \\ \mathcal{R}(M, Y, C) & \xrightarrow{\sim} & \mathcal{H}(M, k, C, w)/G \end{array}$$

The vertical arrow on the left is the map which forgets the base point. The vertical arrow on the right is the quotient map. The top horizontal arrow is given by the monodromy with respect to the base point along the elements of \mathcal{G} . The bottom horizontal arrow is also given by this monodromy, but since no base point is specified, it is defined up to global conjugation.

This diagram is commutative and, according to the discussion at the beginning of this section, its horizontal arrows are bijections.

The preimage of an element $R \in \mathcal{R}(M, Y, C)$ by the vertical arrow consists in $\# G / \# \text{Aut}(R)$ elements. It follows that, for all function $f : \mathcal{R}(M, Y, C) \rightarrow \mathbb{C}$, which can alternatively be seen as an invariant function on $\mathcal{H}(M, k, C, w)$, we have the counting formula

$$(63) \quad \begin{aligned} \sum_{R \in \mathcal{R}(M, Y, C)} \frac{1}{\# \text{Aut}(R)} f(R) &= \frac{1}{\# G} \sum_{(R, p) \in \mathcal{R}_m(M, Y, C)} f(R) \\ &= \frac{1}{\# G} \sum_{h \in \mathcal{H}(M, k, C, w)} f(h). \end{aligned}$$

5.3. Measured spaces of ramified G -bundles

Let us start by putting a topology on the sets of ramified G -bundles. For each based ramified G -bundle (R, p) based at $\pi(p) = m$ with ramification locus Y , and

each open subset U of $M \setminus \{m\}$ containing Y , we define

$$\mathcal{V}((R, p), U) = \{(R', p') \in \mathcal{R}_m(M, C) : (R, p)|_{M \setminus U} \simeq (R', p')|_{M \setminus U} \text{ as } G\text{-bundles}\}.$$

The sets $\mathcal{V}((R, p), U)$ form a basis of a topology on $\mathcal{R}_m(M, C)$ and from now on we consider this space endowed with that topology. Similarly, we endow $\mathcal{R}(M, C)$ with the topology generated by the sets

$$\mathcal{V}(R, U) = \{R' \in \mathcal{R}(M, C) : R|_{M \setminus U} \simeq R'|_{M \setminus U}\},$$

where U contains the ramification locus of R . These topologies make the projection $\mathcal{R}_m(M, C) \rightarrow \mathcal{R}(M, C)$ continuous. However, observe that the number of ramification points is not a continuous function with respect to these topologies, it is only lower semi-continuous. In fact, these topologies are the roughest which make the monodromy along any loop on M a continuous functions on its definition set.

Let $F(M)$ denote the set of finite subsets of M . For each $k \geq 0$, let $\Delta_k \subset M^k$ denote the subset of M^k on which at least two components are equal. We endow $F(M)$ with the topology which makes the bijection $F(M) \simeq \bigsqcup_{k \geq 0} (M^k \setminus \Delta_k) / \mathfrak{S}_k$ a homeomorphism. Once again, the natural mapping $\text{Ram} : \mathcal{R}(M, C) \rightarrow F(M)$ which associates to a covering its ramification locus is not continuous.

It is now time to introduce the Lévy process. Let X be a continuous time random walk on G , with jump measure invariant by conjugation, and also invariant by inversion if M is non-orientable. Its Lévy measure Π is a finite invariant measure supported by $G \setminus \{1\}$. We assume that the support of Π generates G .

We denote by Π_1 the probability measure $\Pi/\Pi(G)$ on G . We define now the weight of a ramified G -bundle with respect to Π . Recall that if R is ramified at y , then $\mathcal{O}(R, y)$ denotes the conjugacy class of the monodromy of a small circle around y , positively oriented if M is oriented.

Definition 5.3.1. — Consider $R \in \mathcal{R}(M)$. Let Y denote the ramification locus of R . The weight of R is the non-negative real number

$$\Pi_1(R) = \prod_{y \in Y} \frac{\Pi_1(\mathcal{O}(R, y))}{\#\mathcal{O}(R, y)}.$$

If R is represented by an element h of $\mathcal{H}(M, k, C, w)$, then $\Pi_1(R) = \prod_{i=1}^k \Pi_1(\{d_i(h)\})$.

The notion of weight of a ramified G -bundle allows us to define positive measures on the spaces of bundles. The choice of the normalization will be justified by later results.

Definition 5.3.2. — The Borel measure $\text{BB}_{M,m,Y,C}^X$ on $\mathcal{R}_m(M, Y, C)$ is defined by

$$\text{BB}_{M,m,Y,C}^X = \frac{\#G^{1-g}}{\#\mathcal{O}_1 \cdots \#\mathcal{O}_p} \sum_{(R,p) \in \mathcal{R}_m(M,Y,C)} \Pi_1(R) \delta_{(R,p)}.$$

By the left vertical arrow of (62), this measure is projected on the Borel measure $\text{RB}_{M,Y,C}^X$ on $\mathcal{R}(M, Y, C)$ defined by

$$\text{RB}_{M,Y,C}^X = \frac{\# G^{2-g}}{\# \mathcal{O}_1 \cdots \# \mathcal{O}_p} \sum_{R \in \mathcal{R}(M,Y,C)} \frac{\Pi_1(R)}{\# \text{Aut}(R)} \delta_R.$$

Thanks to the counting formula (63), we can roughly bound above the total mass of $\text{BB}_{M,m,Y,C}^X$ by

$$(64) \quad \text{BB}_{M,m,Y,C}^X(1) = \frac{\# G^{1-g}}{\# \mathcal{O}_1 \cdots \# \mathcal{O}_p} \sum_{h \in \mathcal{H}(M,k,C,w)} \Pi_1(h) \leq \# G.$$

Our next objective is to put measures on $\mathcal{R}_m(M, C)$ and $\mathcal{R}(M, C)$, the sets of ramified bundles in which the ramification loci are not fixed. We have endowed both spaces with topologies. Thus, they carry a Borel σ -field. Let $\mathcal{M}_+(\mathcal{R}(M, C))$ and $\mathcal{M}_+(\mathcal{R}_m(M, C))$ denote the spaces of positive Borel measures on $\mathcal{R}(M, C)$ and $\mathcal{R}_m(M, C)$ respectively, endowed with the topology of weak convergence.

Proposition 5.3.3. — *The mapping from $F(M)$ to $\mathcal{M}_+(\mathcal{R}(M, C))$ which sends Y to $\text{RB}_{M,Y,C}^X$ is continuous. Similarly, the mapping from $F(M)$ to $\mathcal{M}_+(\mathcal{R}_m(M, C))$ which sends Y to $\text{BB}_{M,m,Y,C}^X$ is continuous on its definition set.*

Proof. — We prove only the first statement. The second one is very similar.

By definition of the topology on $F(M)$, it suffices to prove that the mapping from $M^k \setminus \Delta_k$ to $\mathcal{M}_+(\mathcal{R}(M, C))$ which sends $Y = (y_1, \dots, y_k)$ to $\text{RB}_{M,Y,C}^X$ is continuous for all $k \geq 0$. Consider $k \geq 0$, $Y = \{y_1, \dots, y_k\}$ and a bounded continuous function $f : \mathcal{R}(M, C) \rightarrow \mathbb{R}$. Choose $\varepsilon > 0$. For simplicity, assume that M is endowed with a Riemannian metric.

Since $\mathcal{R}(M, Y, C)$ is a finite set, the continuity of f implies the existence of $r > 0$ such that the balls $B(y_i, r)$ are contained in $M \setminus \partial M$, pairwise disjoint and such that the neighbourhood $U = B(y_1, r) \times \cdots \times B(y_k, r)$ of Y in $M^k \setminus \Delta_k$ satisfies

$$\forall R \in \mathcal{R}(M, Y, C), \forall R' \in \mathcal{V}(R, U), \quad |f(R') - f(R)| < \frac{\varepsilon}{\# \mathcal{R}(M, Y, C)}.$$

Let $Y' = \{y'_1, \dots, y'_k\}$ be an element of U . Let ϕ be a diffeomorphism of M such that $\phi|_{M \setminus U} = \text{id}_{M \setminus U}$ and $\phi(y_i) = y'_i$ for all $i \in \{1, \dots, k\}$. For each bundle $\pi : R \rightarrow M$ belonging to $\mathcal{R}(M, Y, C)$, the bundle $\phi(R) = (\phi \circ \pi : R \rightarrow M)$ belongs to $\mathcal{R}(M, Y', C)$. Replacing ϕ by its inverse in the definition of $\phi : \mathcal{R}(M, Y, C) \rightarrow \mathcal{R}(M, Y', C)$ yields the inverse mapping, hence ϕ is a bijection. Moreover, for each $i \in \{1, \dots, k\}$, $\mathcal{O}(\phi(R), y'_i) = \mathcal{O}(R, y_i)$, so that $\Pi_1(\phi(R)) = \Pi_1(R)$. Also, the conjugation by ϕ determines an isomorphism between $\text{Aut}(R)$ and $\text{Aut}(\phi(R))$. Finally, R

and $\phi(R)$ are isomorphic outside U . Altogether,

$$|\text{RB}_{M,Y',C}^X(f) - \text{RB}_{M,Y,C}^X(f)| \leq \sum_{R \in \mathcal{R}(M,Y,C)} \frac{\Pi_1(R)}{\#\text{Aut}(R)} |f(\phi(R)) - f(R)| < \varepsilon.$$

Since k, Y, f and ε were arbitrary, the result follows. □

We choose for the ramification locus a very simple probability distribution which incorporates the measure vol on M . Let Ξ be the distribution of a Poisson point process of intensity $\Pi(G)\text{vol}$ on M . It is a Borel probability measure on $F(M)$. Moreover, for all $m \in M, \Xi(\{Y : m \in Y\}) = 0$. According to Proposition 5.3.3, the following definition is legitimate.

Definition 5.3.4. — *The Borel measures $\text{RB}_{M,\text{vol},C}^X$ on $\mathcal{R}(M,C)$ and $\text{BB}_{M,m,\text{vol},C}^X$ on $\mathcal{R}_m(M,C)$ are defined by*

$$\begin{aligned} \text{RB}_{M,\text{vol},C}^X &= \int_{F(M)} \text{RB}_{M,Y,C}^X \Xi(dY) = \int_{F(M)} \left(\sum_{R \in \mathcal{R}(M,Y,C)} \frac{\Pi_1(R)}{\#\text{Aut}(R)} \delta_R \right) \Xi(dY), \\ \text{BB}_{M,m,\text{vol},C}^X &= \int_{F(M)} \text{BB}_{M,m,C}^X \Xi(dY) \\ &= \int_{F(M)} \left(\frac{1}{\#G} \sum_{(R,p) \in \mathcal{R}_m(M,Y,C)} \Pi_1(R) \delta_{(R,p)} \right) \Xi(dY). \end{aligned}$$

Since $\Xi(\{Y : m \in Y\}) = 0$, the subset of $\mathcal{R}(M,C)$ which consists in bundles ramified over m is negligible for the measure $\text{RB}_{M,\text{vol},C}^X$. Hence, the measure $\text{BB}_{M,m,\text{vol},C}^X$ projects on $\text{RB}_{M,\text{vol},C}^X$ by the left vertical arrow of the diagram (62). In particular, these measures have the same total mass. Thanks to (64), this total mass is finite. Hence, $\text{RB}_{M,\text{vol},C}^X$ and $\text{BB}_{M,m,\text{vol},C}^X$ are finite measures. We will denote by $\text{NRB}_{M,\text{vol},C}^X$ and $\text{NBB}_{M,m,\text{vol},C}^X$ the corresponding probability measures.

Although this is not absolutely necessary, let us compute $\text{BB}_{M,m,\text{vol},C}^X(1)$. Recall that (55) gives an expression of the density of the 1-dimensional marginals of the Lévy process X with respect to the uniform measure on G : setting, for all $\alpha \in \text{Irr}(G)$,

$$\widehat{\Pi}(\alpha) = \sum_{x \in G} \overline{\chi_\alpha(x)} \Pi(\{x\}),$$

we have

$$(65) \quad \forall t > 0, \forall x \in G, \quad Q_t(x) = e^{-t\Pi(G)} \sum_{\alpha \in \text{Irr}(G)} e^{t\widehat{\Pi}(\alpha)/\chi_\alpha(1)} \chi_\alpha(1) \chi_\alpha(x).$$

In the present context, this equality can be checked by an elementary computation, using the following formula, which we will need again later and which is proved by using the standard properties of characters.

Lemma 5.3.5. — For all $k \geq 1$ and all $x \in G$, the following equality holds:

$$\sum_{x_1, \dots, x_k \in G} \Pi(\{x_1\}) \cdots \Pi(\{x_k\}) \mathbb{1}_{x_1 \cdots x_k = x} = \frac{1}{\#G} \sum_{\alpha \in \text{Irr}(G)} \left(\frac{\widehat{\Pi}(\alpha)}{\chi_\alpha(1)} \right)^k \chi_\alpha(1) \chi_\alpha(x).$$

We can now compute the mass of $\text{BB}_{M,m,\text{vol},C}^X(1)$.

Proposition 5.3.6. — The total mass of the measure $\text{BB}_{M,m,\text{vol},C}^X$ is equal to

$$(66) \quad \frac{1}{\#G^g \prod_{i=1}^p \# \theta_i} \sum_{\substack{a_1, \dots, a_g \in G \\ c_1 \in \theta_1, \dots, c_p \in \theta_p}} Q_{\text{vol}(M)}(w(a_1 \cdots a_g) c_1 \cdots c_p),$$

which, with the notation of Definition 4.3.6, is equal to

$$\int_G Q_{\text{vol}(M)} dm_{M,\emptyset,C}.$$

Proof. — Choose $Y \in F(M)$. Set $k = \#Y$. Choose a graph \mathbb{G} , a vertex v of \mathbb{G} and a tame system of generators of $\text{RL}_v(\mathbb{G})$ associated with some word w , as we did in Section 5.2. This determines a bijection $\mathcal{R}_m(M, Y, C) \simeq \mathcal{H}(M, k, C, w)$. By the counting formula (63),

$$\begin{aligned} \text{BB}_{M,m,Y,C}^X(1) &= \frac{\#G^{1-g}}{\prod_{i=1}^p \# \theta_i} \sum_{h \in \mathcal{H}(M,k,C,w)} \prod_{i=1}^k \Pi_1(\{d_i(h)\}) \\ &= \frac{\#G^{1-g}}{\prod_{i=1}^p \# \theta_i} \sum_{\substack{a_1, \dots, a_g \in G \\ c_1 \in \theta_1, \dots, c_p \in \theta_p}} \sum_{d_1, \dots, d_k \in G} \prod_{i=1}^k \frac{\Pi(\{d_i(h)\})}{\Pi(G)} \mathbb{1}_{d_1 \cdots d_k = w(a_1 \cdots a_g) c_1 \cdots c_p}. \end{aligned}$$

By Lemma 5.3.5, this is equal to

$$\frac{1}{\#G^g \prod_{i=1}^p \# \theta_i} \sum_{\substack{a_1, \dots, a_g \in G \\ c_1 \in \theta_1, \dots, c_p \in \theta_p \\ \alpha \in \text{Irr}(G)}} \frac{1}{\Pi(G)^k} \left(\frac{\widehat{\Pi}(\alpha)}{\chi_\alpha(1)} \right)^k \chi_\alpha(1) \chi_\alpha(w(a_1 \cdots a_g) c_1 \cdots c_p).$$

Integrating this expression with respect to Y under the probability measure Ξ amounts to replacing k by a Poisson random variable with parameter $\Pi(G)\text{vol}(M)$ and taking the expectation. Using (65), we find that this expectation is equal to (66). \square

5.4. The monodromy process as a Markovian holonomy field

Let m be a point of M . Let $l \in \mathbb{L}_m(M)$ be a loop based at m . Since l is rectifiable, its range is negligible for the measure vol . Hence, the ramification locus of a ramified G -bundle based at m distributed according to the probability measure $\text{NBB}_{M,m,\text{vol},C}^X$ is almost surely disjoint from the range of l . The mapping $P_l : \mathcal{R}_m(M, C) \rightarrow G$ which sends a pair (R, p) to the monodromy of R along l with respect to p is defined on

the subset where the ramification locus is disjoint from l and thus is a well-defined random variable under the probability measure $\text{NBB}_{M,m,\text{vol},C}^X$.

Let $l_1, l_2 \in \mathcal{L}_m(M)$ be two loops. Let (R, p) be an element of $\mathcal{R}_m(M, C)$. Let g_1 and g_2 be the monodromies of l_1 and l_2 respectively. Let us compute the monodromy of $l_1 l_2$. The point p is sent to pg_1 by the parallel transport along l_1 . Then, on one hand the parallel transport along l_2 sends p to pg_2 and on the other hand the the parallel transport commutes to the action of G on the right on $\pi^{-1}(m)$. Thus, the parallel transport along l_2 sends pg_1 to $pg_2 g_1$. It appears that monodromies are multiplied in the reversed order of concatenation. Coming back to the probabilistic setting, this implies that

$$\forall l_1, l_2 \in \mathcal{L}_m(M), \quad P_{l_1 l_2} = P_{l_2} P_{l_1} \text{ almost surely.}$$

It is even easier to check that for all $l \in \mathcal{L}_m(M)$, $P_{l^{-1}} = P_l^{-1}$ almost surely.

Thanks to Proposition 2.2.3, these two relations ensure that the collection of random variables $(P_l)_{l \in \mathcal{L}_m(M)}$ defined on the probability space $(\mathcal{R}_m(M, C), \text{NBB}_{M,m,\text{vol},C}^X)$ determines a probability measure on the space $(\mathcal{M}(\mathcal{L}_m(M), G), \mathcal{E})$ (see Section 2.1). We denote this probability measure by $\text{NMF}_{M,m,\text{vol},C}^X$. By restriction, this probability measure is also defined on the invariant σ -field and, by Lemma 2.1.5, determines a probability measure on the measurable space $(\mathcal{M}(\mathcal{P}(M), G), \mathcal{J})$, which we denote by $\text{NMF}_{M,(m),\text{vol},C}^X$. Finally, we define a finite measure on $(\mathcal{M}(\mathcal{P}(M), G), \mathcal{J})$ by

$$\text{MF}_{M,(m),\text{vol},C}^X = \text{BB}_{M,\text{vol},C}^X(1) \text{NMF}_{M,(m),\text{vol},C}^X.$$

Lemma 5.4.1. — *The measure $\text{MF}_{M,(m),\text{vol},C}^X$ on $(\mathcal{M}(\mathcal{P}(M), G), \mathcal{J})$ does not depend on the point m . We denote it by $\text{MF}_{M,\text{vol},C}^X$.*

Proof. — By Definition 2.1.6, it suffices to show that if l_1, \dots, l_n are loops on M based at the same point m_0 , and $f : G^n \rightarrow \mathbb{C}$ is a function invariant under the action of G by diagonal conjugation, then the distribution of $f(h(l_1), \dots, h(l_n))$ under $\text{NMF}_{M,(m),\text{vol},C}^X$ does not depend on m . By definition, this distribution is that of $f(P_{cl_1c^{-1}}, \dots, P_{cl_n c^{-1}})$ under $\text{BB}_{M,m,\text{vol},C}^X$, where c is an arbitrary path from m to m_0 . Let m and m' be two points. Let us chose a path c from m to m_0 and a path c' from m' to m . It suffices to prove that the distributions of $f(P_{cl_1c^{-1}}, \dots, P_{cl_n c^{-1}})$ under $\text{BB}_{M,m,\text{vol},C}^X$ and $f(P_{c'cl_1c^{-1}c'^{-1}}, \dots, P_{c'cl_n c^{-1}c'^{-1}})$ under $\text{BB}_{M,m',\text{vol},C}^X$ coincide.

Let Y be a finite subset of M which does not meet c' . Let (R, p') be an element of $\mathcal{R}_{m'}(M, Y, C)$. Then, for each $i \in \{1, \dots, n\}$, the monodromy of R along $c'cl_i c^{-1}c'^{-1}$ relatively to p' is equal to the monodromy of R along $cl_i c^{-1}$ relatively to the image of p' by parallel transport along c' , which we denote by p . Thus, it suffices to prove that the mapping $\mathcal{R}_{m'}(M, Y, C) \rightarrow \mathcal{R}_m(M, Y, C)$ which sends (R, p') to (R, p) , where p is the image of p' by parallel transport along c' , sends the measure $\text{BB}_{M,m',Y,C}$ to the

measure $\text{BB}_{M,m,Y,C}$. This follows from the definition of these measures and the fact that the mapping which we consider is a bijection which preserves the weight. \square

The main result of this chapter is the following.

Theorem 5.4.2. — *The finite measures $\text{HF}_{M,\text{vol},\emptyset,C}^X$ and $\text{MF}_{M,\text{vol},C}^X$ on the measurable space $(\mathcal{M}(\mathbb{P}(M), G), \mathcal{J})$ are equal.*

This theorem expresses, at least when the surfaces carry only G -constraints along their boundary, the fact that the holonomy process associated with the Markovian holonomy field HF^X is the monodromy process associated to a random ramified G -bundle taken under the appropriate distribution.

The proof of this theorem consists in two main steps. In the first step, we prove that the monodromy process is stochastically continuous. Then, we prove that the holonomy process and the monodromy process coincide in distribution on the set of piecewise geodesic loops for some Riemannian metric on M .

Proposition 5.4.3. — *The measure $\text{MF}_{M,\text{vol},C}^X$ is stochastically continuous, in the sense that it satisfies the first property of Definition 3.1.3.*

Proof. — We need to prove that for all $m \in M$, all $l \in \mathbb{L}_m(M)$ and all sequence $(l_n)_{n \geq 0}$ of loops based at m converging to l , the sequence $(P_{l_n})_{n \geq 0}$ converges in measure to P_l . Let us endow M with a Riemannian metric and choose m , l and $(l_n)_{n \geq 0}$ as above. We assume that all loops are parametrised at constant speed, so that the sequence of parametrised paths $(l_n)_{n \geq 0}$ converges uniformly to l .

Choose $\varepsilon > 0$. For each $r > 0$, let $\mathcal{N}_r(l)$ denote the r -neighbourhood of the image of l . Since the distribution of the ramification locus $\text{Ram}(R)$ of R under the finite measure $\text{BB}_{M,m,\text{vol},C}^X$ is absolutely continuous with respect to Ξ ,

$$\lim_{r \rightarrow 0} \text{BB}_{M,m,\text{vol},C}^X(\{(R, p) : \text{Ram}(R) \cap \mathcal{N}_r(l) \neq \emptyset\}) = 0.$$

Choose $r > 0$ such that this probability is smaller than ε . Assume also that r is smaller than the convexity radius of our Riemannian metric on M . Finally, let n_0 be such that $n \geq n_0$ implies $d_\infty(l_n, l) < r$. Then, if $n \geq n_0$, l_n and l are homotopic inside $\mathcal{N}_r(l)$ which, with probability greater than $1 - \varepsilon$, does not contain any ramification point. Hence,

$$\forall n \geq n_0, \quad \text{BB}_{M,m,\text{vol},C}^X(\{(R, p) : P_{l_n}(R) \neq P_l(R)\}) < \varepsilon.$$

Since ε is arbitrary, this proves that P_{l_n} converges to P_l in measure. \square

Theorem 5.4.2 asserts the equality of two finite measures. We consider the two stochastic processes $(H_l)_{l \in \mathbb{L}_m(M)}$ and $(P_l)_{l \in \mathbb{L}_m(M)}$ which are both the canonical process on $\mathcal{M}(\mathbb{L}_m(M), G)$, the first considered under the measure $\text{HF}_{M,\text{vol},\emptyset,C}^X$ and the

second under the measure $\mathbf{MF}_{M,(m),\text{vol},C}^X$. Although these measures are not in general probability measures, we use the language of stochastic processes for H and P .

By Proposition 5.4.3 and Theorem 3.3.1, it suffices to endow M with a Riemannian metric and to show that the restrictions of P and H to piecewise geodesic loops agree in distribution. For this, as we have already observed several times, it suffices to show that they agree in distribution when restricted to the set of loops in a graph with piecewise geodesic edges, or in fact any graph.

Proposition 5.4.4. — *Let $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ be a graph on M such that $m \in \mathbb{V}$. The families of random variables $(P_l)_{l \in \mathbb{L}_m(\mathbb{G})}$ and $(H_l)_{l \in \mathbb{L}_m(\mathbb{G})}$ have the same distribution.*

Proof. — It suffices to prove that the equality holds when the processes are restricted to a family of loops which generate the group $\text{RL}_m(M)$. Consider a tame family of generators $\mathcal{G} = \{a_1, \dots, a_g, c_1, \dots, c_p, l_1, \dots, l_f\}$ of $\text{RL}_m(\mathbb{G})$ associated with a word w . The loop l_f is a function of all other loops, so that it suffices to compute the distribution of $H = (H_{a_1}, \dots, H_{a_g}, H_{c_1}, \dots, H_{c_p}, H_{l_1}, \dots, H_{l_{f-1}})$. Let us choose $h = (g_{a_1}, \dots, g_{a_g}, g_{c_1}, \dots, g_{c_p}, g_{l_1}, \dots, g_{l_{f-1}})$ in $G^g \times \vartheta_1 \times \dots \times \vartheta_p \times G^{f-1}$. By Proposition 2.4.2,

$$\mathbf{HF}_{M,\text{vol},\varnothing,C}^X(H = h) = \frac{\# G^{1-g-f}}{\prod_{i=1}^p \# \vartheta_i} \prod_{i=1}^f Q_{\text{vol}(F_i)}(g_{l_i}),$$

where we have set $g_{l_f} = w(g_{a_1}, \dots, g_{a_g})g_{c_1} \dots g_{c_p}(g_{l_1} \dots g_{l_{f-1}})^{-1}$.

Now let us compute the corresponding quantity for the monodromy field. Let Y be a finite subset of M which does not meet \mathbb{G} . Let us refine \mathbb{G} inside each face which meets Y in order to get a new graph \mathbb{G}' , finer than \mathbb{G} , such that each face of \mathbb{G}' either does not meet Y and is equal to a face of \mathbb{G} , or contains exactly one point of Y .

By applying the second assertion of Proposition 2.4.2, we can construct a tame family of generators \mathcal{G}' of the group of reduced loops of \mathbb{G}' which is finer than \mathcal{G} in the sense that for each face F , the facial lasso of \mathcal{G} corresponding to F is the product in a certain order of the facial lassos of \mathcal{G}' corresponding to the faces of \mathbb{G}' contained in F .

For each $i \in \{1, \dots, f\}$, let us write $Y_i = Y \cap F_i = \{y_{i,1}, \dots, y_{i,k_i}\}$. Let

$$\{l_{i,j} : i \in \{1, \dots, f\}, j \in \{1, \dots, k_i\}\}$$

be the set of facial lassos of \mathcal{G}' indexed accordingly. We may assume that for all $i \in \{1, \dots, f\}$, $l_i = l_{i,1} \dots l_{i,k_i}$.

The family \mathcal{G}' determines a bijection between $\mathcal{R}_m(M, Y, C)$ and $\mathcal{H}(M, k, C, w)$. This allows us to compute the distribution of the random variable

$$M = (M_{a_1}, \dots, M_{a_g}, M_{c_1}, \dots, M_{c_p}, M_{l_1}, \dots, M_{l_{f-1}})$$

under $\mathbb{B}\mathbb{B}_{M,m,Y,C}^X$. We find

$$(67) \quad \mathbb{B}\mathbb{B}_{M,m,Y,C}^X(M = h) = \frac{\# G^{1-g}}{\prod_{i=1}^p \# \theta_i} \prod_{i=1}^f \sum_{g_{i,1}, \dots, g_{i,k_i} \in G} \Pi_1(\{g_{i,1}\}) \cdots \Pi_1(\{g_{i,k_i}\}) \mathbb{1}_{g_{i,k_i} \cdots g_{i,1} = g_i}.$$

Using Lemma 5.3.5, we find that the quantity (67) is equal to

$$\frac{\# G^{1-g-f}}{\prod_{i=1}^p \# \theta_i} \prod_{i=1}^f \sum_{\alpha_i \in \text{Irr}(G)} \left(\frac{\widehat{\Pi}(\alpha_i)}{\Pi(G)\chi_{\alpha_i}(1)} \right)^{k_i} \chi_{\alpha_i}(1)\chi_{\alpha_i}(g_i).$$

By integrating this expression with respect to Y under the measure Ξ , we find

$$\begin{aligned} \mathbb{B}\mathbb{B}_{M,m,\text{vol},C}^X(M = h) &= \frac{\# G^{1-g-f}}{\prod_{i=1}^p \# \theta_i} \prod_{i=1}^f e^{-\text{vol}(F_i)\Pi(G)} \sum_{\alpha \in \text{Irr}(G)} e^{\text{vol}(F_i) \frac{\widehat{\Pi}(\alpha)}{\chi_{\alpha}(1)}} \chi_{\alpha}(1)\chi_{\alpha}(g_i) \\ &= \frac{\# G^{1-g-f}}{\prod_{i=1}^p \# \theta_i} \prod_{i=1}^f Q_{\text{vol}(F_i)}(g_i). \end{aligned}$$

This proves that $\mathbb{B}\mathbb{B}_{M,m,\text{vol},C}^X(M = h) = \text{HF}_{M,\text{vol},\emptyset,C}^X(H = h)$. □

INDEX

- $A_\gamma(M)$, 33
- abstract graph, 80
- admissible
 - word, 68
- area, 69
- $\mathcal{B}(M)$, 26
- Banchoff-Pohl inequality, 11
- binary, *see also* gluing
- boundary of a face, *see also* graph, 48–59
- Brouwer property of the sphere, 61
- Brownian motion
 - indexed by loops, 9
 - on G , 15, 21
- $\text{Conj}(G)$, 108
- $\text{Const}(G)$, 108
- C , 86
- \mathcal{C} , 80
- \mathcal{C} , 27
- cellular embedding of a graph, 42
- charge of a particle, 13
- cobordism, 5
- concatenation, 4, 14, 30
- conjugation
 - diagonal action by, 80
 - Lévy process invariant by, 136
 - measures invariant by, 132
- connection
 - Levi-Civita, 35
 - on a principal bundle, 12, 14, 21
- constraints, *see also* surface
- curve, 30
 - equivalence, 63
- curve accessible, 59
- cycle, 31
 - facial, 49
- $\text{DF}_{M, \text{vol}, \mathcal{C}, C}^{X, G}$, 141
- $D_n(c)$, 33
- ∂F , 49
- $\delta_{\theta(n)}$, 87
- directed set of graphs, *see also* graph
- \mathbb{E} , 41
- \mathbb{E}^+ , 66
- $e(G)$, 66
- η, κ , 131
- edge, *see also* graph
- \mathbb{F} , 41
- $f(G)$, 66
- face, *see also* graph
- framing, *see also* graph
- G , 80
- \mathbb{G} , 41
- $\text{Gr}_\gamma(M, \mathcal{C})$, 73
- \mathcal{G} , 154
- $\widehat{\mathbb{G}}$, 96
- gauge group, 80
- G -constraints, *see also* surface
- genus, *see also* surface
- gluing, 28
- graph, 41
 - abstract, 80
 - approximation by piecewise geodesic graphs, 75
 - dual, 96
 - edge, 41
 - adjunction, 62
 - cyclic order, 52
 - face
 - boundary (combinatorial description), 55
 - boundary (definition), 49
 - topology, 46
 - with the same bounding cycle, 56
 - framing, 55
 - partial order on the set of, 72

- spanning tree, 96
- vertex
 - orientation, 54
- group of reduced loops, 65
 - tame generators, 94
- $\mathcal{H}(M, k, C, w)$, 155
- $\text{HF}_{M, \text{vol}, \mathcal{C}, C}$, 108
- \mathcal{J} , 82
- isoperimetric inequality, 116
- Klein bottle, 132
- Kolmogorov's extension theorem, 84
- Kolmogorov-Chapman equation, 128
- Lévy process, 134–140
 - admissible, 139
 - generator, 136
 - pure jump admissible, 140
 - with positive density, 137
- $\text{LE}(c)$, 74
- $\text{L}(M)$, 31
- lasso, 31
 - bounding, 94
 - decomposition of a path, 74
 - facial, 94
 - meander, 31
 - spoke, 31
- loop, 31
 - reduced, 65
 - simple, 31
- loop erasure, 74
- (M, \mathcal{C}) , 27
- $\text{MF}_{M, \text{vol}, C}^X$, 160
- $\mathcal{M}(\mathbb{P}(M), G)$, 79
- Markovian holonomy field
 - definition, 108
 - discrete, 110, 141–149
 - continuously area-dependent, 114
 - Fellerian, 114
 - regular, 114
 - stochastically $\frac{1}{2}$ -Hölder continuous, 113
 - Fellerian, 109
 - partition function, 126–134, 139
 - partition functions, 112
 - regular, 109
 - stochastically continuous, 109
- meander, *see also* lasso
- multiplicative function, 79
 - invariant σ -field, 82
- uniform, 87
- uniform with constraints, 89–94
- \mathcal{O}_x , 86
- $\mathcal{O}(R, y)$, 154
- $\text{PP}(M), \mathbb{P}(M)$, 31
- Π , 136
- $\Pi(R)$, 156
- path, 31
 - convergence, 35
 - convergence with fixed endpoints, 33
 - dyadic approximation, 33
 - equivalence, 63
 - piecewise geodesic, 33
- pattern, 42
 - split, 45
 - tubular, 27
- planar tree, 101
- Poisson point process, 158
 - indexed by loops, 9
- projective family of probability spaces, 83
- Q_t , 136
- ramified principal bundle, 152
 - automorphisms, 154
 - counting, 155
 - monodromy, 153
 - random, 158
 - topology, 156
 - weight, 156
- $\text{BB}_{M, m, Y, C}^X, \text{RB}_{M, Y, C}^X$, 157
- $\text{RL}_m(\mathbb{G})$, 65
- $\mathcal{R}(M), \mathcal{R}(M, Y)$, 153
- $\mathcal{R}_m(M), \mathcal{R}_m(M, Y)$, 153
- Riemannian metric, 70
- $\text{Sk}(\mathbb{G})$, 41
- $\text{Spl}_i(M)$, 29
- skeleton, *see also* graph
- spanning tree, 96
- split genus, *see also* surface
- splitting, 28
- spoke, *see also* lasso
- surface
 - G -constraints, 87
 - genus, 24
 - classification, 25, 126
 - marked, 27
 - split genus, 30
 - surgery, 28, 87
- $\text{UG}_{M, \mathcal{C}, C}^G$, 89

unary, *see also* gluing
 \mathbb{V} , 41
 $v(G)$, 66
vertex, *see also* graph

Ξ , 158

$F(M)$, 156

$Z_{p,g,t}^+$, $Z_{p,g,t}^-$, 126

BIBLIOGRAPHY

- [1] ALBEVERIO (SERGIO), HØEGH-KROHN (RAPHAEL) & HOLDEN (HELGE) – *Stochastic Lie group-valued measures and their relations to stochastic curve integrals, gauge fields and Markov cosurfaces*, in *Stochastic processes—mathematics and physics (Bielefeld, 1984)*, Lecture Notes in Math., vol. 1158, Springer, 1986, pp. 1–24.
- [2] ———, *Stochastic multiplicative measures, generalized Markov semigroups, and group-valued stochastic processes and fields*, J. Funct. Anal., t. **78** (1988), pp. 154–184.
- [3] APPLEBAUM (DAVID) & KUNITA (HIROSHI) – *Lévy flows on manifolds and Lévy processes on Lie groups*, J. Math. Kyoto Univ., t. **33** (1993), pp. 1103–1123.
- [4] ATIYAH (MICHAEL) – *Topological quantum field theories.*, Publ. Math., Inst. Hautes Étud. Sci., t. **68** (1988), pp. 175–186.
- [5] BAEZ (JOHN) & MUNIAIN (JAVIER P.) – *Gauge fields, knots and gravity*, Series on Knots and Everything, vol. 4, World Scientific Publishing Co. Inc., 1994.
- [6] BANCHOFF (THOMAS F.) & POHL (WILLIAM F.) – *A generalization of the isoperimetric inequality*, J. Differential Geometry, t. **6** (1971/72), pp. 175–192.
- [7] BLEECKER (DAVID) – *Gauge theory and variational principles*, Global Analysis Pure and Applied Series A, vol. 1, Addison-Wesley Publishing Co., Reading, Mass., 1981.
- [8] D’ADDA (ALESSANDRO) & PROVERO (PAOLO) – *Two-dimensional gauge theories of the symmetric group S_n in the large- n limit*, Comm. Math. Phys., t. **245** (2004), pp. 1–25.
- [9] DAVID (GUY) – *Singular sets of minimizers for the Mumford-Shah functional*, Progress in Math., vol. 233, Birkhäuser, 2005.
- [10] DRIVER (BRUCE K.) – *YM_2 : continuum expectations, lattice convergence, and lassos*, Comm. Math. Phys., t. **123** (1989), pp. 575–616.

- [11] ———, *Two-dimensional Euclidean quantized Yang-Mills fields*, in *Probability models in mathematical physics (Colorado Springs, CO, 1990)*, World Sci. Publ., Teaneck, NJ, 1991, pp. 21–36.
- [12] DUQUESNE (THOMAS) – *The coding of compact real trees by real valued functions.*, Preprint (2006).
- [13] FINE (DANA S.) – *Quantum Yang-Mills on the two-sphere*, *Comm. Math. Phys.*, t. **134** (1990), pp. 273–292.
- [14] ———, *Quantum Yang-Mills on a Riemann surface*, *Comm. Math. Phys.*, t. **140** (1991), pp. 321–338.
- [15] GAMBINI (RODOLFO) & PULLIN (JORGE) – *Loops, knots, gauge theories and quantum gravity*, Cambridge Monographs on Mathematical Physics, Cambridge Univ. Press, 1996.
- [16] GROSS (DAVID J.) & MATYTSIN (ANDREI) – *Some properties of large- N two-dimensional Yang-Mills theory*, *Nuclear Phys. B*, t. **437** (1995), pp. 541–584.
- [17] GROSS (DAVID J.) & TAYLOR (WASHINGTON IV) – *Two-dimensional QCD is a string theory*, *Nuclear Phys. B*, t. **400** (1993), pp. 181–208.
- [18] GROSS (LEONARD) – *A Poincaré lemma for connection forms*, *J. Funct. Anal.*, t. **63** (1985), pp. 1–46.
- [19] ———, *The Maxwell equations for Yang-Mills theory*, in *Mathematical quantum field theory and related topics (Montreal, PQ, 1987)*, CMS Conf. Proc., vol. 9, Amer. Math. Soc., 1988, pp. 193–203.
- [20] GROSS (LEONARD), KING (CHRISTOPHER) & SENGUPTA (AMBAR N.) – *Two-dimensional Yang-Mills theory via stochastic differential equations*, *Ann. Physics*, t. **194** (1989), pp. 65–112.
- [21] HAMBLY (BEN) & LYONS (TERRY J.) – *Uniqueness for the signature of a path of bounded variation and the reduced path group*, Preprint (2006).
- [22] KOBAYASHI (SHOSHICHI) & NOMIZU (KATSUMI) – *Foundations of differential geometry. Vol. I*, Wiley Classics Library, John Wiley & Sons Inc., 1996, Reprint of the 1963 original, A Wiley-Interscience Publication.
- [23] LANDO (SERGEI K.) & ZVONKIN (ALEXANDER K.) – *Graphs on surfaces and their applications*, *Encyclopaedia of Math. Sciences*, vol. 141, Springer, 2004.

- [24] LÉVY (THIERRY) – *Yang-Mills measure on compact surfaces*, Mem. Amer. Math. Soc., t. **166** (2003).
- [25] ———, *Discrete and continuous Yang-Mills measure for non-trivial bundles over compact surfaces*, Probab. Theory Related Fields, t. **136** (2006), pp. 171–202.
- [26] ———, *Schur-Weyl duality and the heat kernel measure on the unitary group.*, Adv. Math., t. **218** (2008), pp. 537–575.
- [27] LIAO (MING) – *Lévy processes in Lie groups*, Cambridge Tracts in Mathematics, vol. 162, Cambridge Univ. Press, 2004.
- [28] LUUKKAINEN (JOUNI) & VÄISÄLÄ (JUSSI) – *Elements of Lipschitz topology*, Ann. Acad. Sci. Fenn. Ser. A I Math., t. **3** (1977), pp. 85–122.
- [29] LYNDON (ROGER C.) & SCHUPP (PAUL E.) – *Combinatorial group theory*, Classics in Mathematics, Springer, 2001, Reprint of the 1977 edition.
- [30] MASSEY (WILLIAM S.) – *Algebraic topology: an introduction*, Springer, 1977, Reprint of the 1967 edition, Graduate Texts in Mathematics, Vol. 56.
- [31] MIGDAL (ALEXANDER A.) – *Recursion equations in gauge field theories*, Sov. Phys. JETP, t. **42** (1975), pp. 413–418.
- [32] MOHAR (BOJAN) & THOMASSEN (CARSTEN) – *Graphs on surfaces*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, 2001.
- [33] MOISE (EDWIN E.) – *Geometric topology in dimensions 2 and 3*, Springer, 1977, Graduate Texts in Mathematics, Vol. 47.
- [34] SENGUPTA (AMBAR N.) – *The Yang-Mills measure for S^2* , J. Funct. Anal., t. **108** (1992), pp. 231–273.
- [35] ———, *Gauge theory on compact surfaces*, Mem. Amer. Math. Soc., t. **126** (1997).
- [36] STEENROD (NORMAN) – *The topology of fibre bundles*, Princeton Landmarks in Mathematics, Princeton Univ. Press, 1999, Reprint of the 1957 edition, Princeton Paperbacks.
- [37] VOGT (ANDREW) – *The isoperimetric inequality for curves with self-intersections*, Canad. Math. Bull., t. **24** (1981), pp. 161–167.
- [38] WILDER (RAYMOND L.) – *Topology of Manifolds*, American Mathematical Society Colloquium Publications, vol. 32, Amer. Math. Soc., 1949.

- [39] WITTEN (EDWARD) – *On quantum gauge theories in two dimensions*, *Comm. Math. Phys.*, t. **141** (1991), pp. 153–209.
- [40] ———, *Two-dimensional gauge theories revisited*, *J. Geom. Phys.*, t. **9** (1992), pp. 303–368.
- [41] YANG (CHEN N.) & MILLS (ROBERT L.) – *Conservation of isotopic spin and isotopic gauge invariance*, *Physical Rev.*, t. **96** (1954), pp. 191–195.
- [42] YOUNG (LAURENCE C.) – *An inequality of the Hölder type, connected with Stieltjes integration*, *Acta Math.*, t. **67** (1936), pp. 251–282.