



HAL
open science

Large deviations for the Yang-Mills measure on a compact surface

Thierry Lévy, James Norris

► **To cite this version:**

Thierry Lévy, James Norris. Large deviations for the Yang-Mills measure on a compact surface. 2005.
hal-00002089v3

HAL Id: hal-00002089

<https://hal.science/hal-00002089v3>

Preprint submitted on 2 May 2005

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Large deviations for the Yang-Mills measure on a compact surface

Thierry Lévy

CNRS, Département de Mathématiques et Applications

École Normale Supérieure

45, rue d'Ulm, F-75230 Paris Cedex 05

levy@dma.ens.fr

James R. Norris

Statistical Laboratory

Center for Mathematical Sciences

Wilberforce Road, Cambridge, CB3 0WB, UK

j.r.norris@statslab.cam.ac.uk

May 2, 2005

Abstract

We prove the first mathematical result relating the Yang-Mills measure on a compact surface and the Yang-Mills energy. We show that, at the small volume limit, the scaled Yang-Mills measures satisfy a large deviation principle with a rate function which is expressed in a simple and natural way in terms of the Yang-Mills energy.

Introduction

The Yang-Mills measure is the distribution of a stochastic process indexed by paths on a smooth manifold and with values in a Lie group. This process is at a heuristical level the random holonomy process of a random connection on a principal bundle over the manifold, the connection being distributed according to the Gibbs measure of the Yang-Mills energy. This is usually expressed by the following inspiring but meaningless formula, where S is the Yang-Mills energy, T a positive real parameter, Z_T a normalization constant and P_T the Yang-Mills measure:

$$dP_T(\omega) = \frac{1}{Z_T} e^{-\frac{1}{2T}S(\omega)} d\omega.$$

In the case of a compact two-dimensional base manifold and compact structure group, two different constructions of the Yang-Mills measure are known. The first one proceeds by an infinite-dimensional approach [3, 15] and the second one by a finite-dimensional approach [9]. Neither of these constructions involve the Yang-Mills energy more than at an informal level, as a guide for the intuition. Both begin from the specifications given by physicists of certain characterizing properties of the distribution of the stochastic process, namely its finite-dimensional distributions, which physicists have of course derived from the Yang-Mills energy, but by some non-rigorous arguments. These finite-dimensional distributions seem to have been first described by A. Migdal in [12] and they are also discussed in [23]. The papers [12] and [23] have thus so far

played the role of a touchstone for the constructions of the Yang-Mills measure. Nevertheless, a rigorous link between the measure and the energy was still lacking.

That the constructions of the Yang-Mills measure do not incorporate explicitly the Yang-Mills energy is not more surprising for example than the fact that one can and usually does construct Brownian motion without referring to the Sobolev H^1 norm. However, essential links between the Wiener measure and the H^1 norm are attested for instance by Cameron-Martin's quasi-invariance theorem and Schilder's large deviation principle.

The point of this paper is to relate at a mathematical level the Yang-Mills measure in two dimensions and the Yang-Mills energy by a large deviation principle. We consider a base space which is a compact surface M with or without boundary and a structure group which is any compact connected Lie group G . We choose a principal G -bundle P over M and consider the space $H^1\mathcal{A}(P)$ of H^1 connections on P , which is the most natural space of connections with finite Yang-Mills energy. If M has a boundary, we consider only those connections which satisfy certain boundary conditions. Then, up to gauge transformations, we embed the space $H^1\mathcal{A}(P)$ of H^1 connections into the canonical space of the random holonomy process and define a natural non-negative functional I^{YM} on this canonical space by extending the Yang-Mills energy by $+\infty$ outside the range of the embedding. The Yang-Mills measure, denoted by P_T , depends on a positive parameter T which is the total area of M . The main result of this paper says that, as T tends to 0 and for every measurable subset A of the canonical space of the process, one has

$$-\inf_{x \in A^\circ} I^{\text{YM}}(x) \leq \liminf_{T \rightarrow 0} T \log P_T(A) \leq \limsup_{T \rightarrow 0} T \log P_T(A) \leq -\inf_{x \in \bar{A}} I^{\text{YM}}(x),$$

where A° and \bar{A} denote respectively the interior and the closure of A with respect to the product topology on the canonical space.

The paper is divided into four sections. In the first one, we give a precise statement of the two main results, one for the case where M has a boundary and one for the case where it is closed. For this, we recall how the Yang-Mills measure is constructed in both cases, what Sobolev connections are and then explain how the rate functions for the large deviation principles are defined.

In the second section, we collect the technical properties of H^1 connections which we will need at various stages of the proof. In particular, we study the holonomy that they determine, the way they are transformed by gauge transformations, and check that they satisfy a classical energy inequality.

In the third section, we prove the large deviation principles. The starting point is the classical short-time estimate of the heat kernel on a compact Riemannian manifold, that we apply to the group G . Then, we build a large deviation principle for the random holonomy process by following step by step its construction described in [9, 10].

However, proving rather abstractly the existence of a large deviation principle is not enough and we must, at each step, identify the rate function in terms of the Yang-Mills energy. Apart from standard results from the theory of large deviations, this relies mainly on three technical tools. The first one is the energy inequality mentioned above and of which we give a complete proof in the appendix. The second one is a beautiful compactness theorem of K. Uhlenbeck that we recall at the end of the second section. The third one is, as far as we know, original, and asserts the existence of a connection of minimal energy with prescribed holonomy along the edges of a graph on M . The proof that such a minimizer exists occupies the fourth and last section of the paper.

1 The large deviation principles

1.1 The Yang-Mills measures

1.1.1 The space of paths

Let M be an oriented compact connected surface, possibly with boundary. Let σ be a volume 2-form on M consistent with the orientation. We will often identify σ with its density which is a Borel measure on M . For technical purposes, let us assume that M is endowed with a Riemannian metric whose Riemannian volume is the density of σ .

By an *edge* on M we mean a segment of a smooth oriented 1-dimensional submanifold. If e is an edge, we call *inverse* of e and denote by e^{-1} the edge obtained by reversing the orientation of e . We also denote respectively by \underline{e} and \bar{e} the starting and finishing point of e . Let e_1, \dots, e_n be n edges. If, for each i between 1 and $n-1$, one has $\bar{e}_i = \underline{e}_{i+1}$, then one can form the concatenation $e_1 \dots e_n$. If moreover f_1, \dots, f_m are also edges which can be concatenated, we declare $e_1 \dots e_n$ equivalent to $f_1 \dots f_m$ if and only if there exists a continuous mapping $c : [0, 1] \rightarrow M$ and two finite sequences $0 = t_0 < t_1 < \dots < t_n = 1$ and $0 = s_0 < s_1 < \dots < s_m = 1$ of real numbers such that, for each $i = 1 \dots n$, the restriction of c to the interval $[t_{i-1}, t_i]$ is a smooth embedding of image e_i and, for each $j = 1 \dots m$, the restriction of c to the interval $[s_{j-1}, s_j]$ is a smooth embedding of image f_j . By a *path* we mean an equivalence class of finite concatenations of edges. We denote the set of paths by PM . Loops, starting and finishing points of paths, their concatenation, are defined in the obvious way. We say that a loop $e_1 \dots e_n$ is *simple* if the vertices $\underline{e}_1, \dots, \underline{e}_n$ are pairwise distinct. If c is a path, we denote respectively by \underline{c} and \bar{c} its starting and finishing point.

Let l_1 and l_2 be two loops. We say that l_1 and l_2 are cyclically equivalent if there exist two paths c and d in PM such that $l_1 = cd$ and $l_2 = dc$. We call *cycle* an equivalence class of loops for this relation. Informally, a cycle is a loop on which one has forgotten the starting point. We say that a cycle is *simple* if its representatives are simple loops.

For any two paths c and c' , denote by $\ell(c)$ and $\ell(c')$ their respective lengths and set $d_\infty(c, c') = \inf \sup_{t \in [0, 1]} d(c(t), c'(t))$, where d is the Riemannian distance on M and the infimum is taken over all continuous parametrizations of c and c' by the interval $[0, 1]$. The function $(c, c') \mapsto d_\ell(c, c') = d_\infty(c, c') + |\ell(c) - \ell(c')|$ is a distance on PM (see [9], Lemmas 2.22, 2.23 and Remark 2.24) which depends on the Riemannian metric on M . However, since M is compact, the topology induced by d_ℓ on PM is independent of the metric. In this paper, we simply say that a sequence of paths *converges* to indicate that it converges in the topology induced by d_ℓ . We will sometimes use a stronger notion of convergence, saying that a sequence $(c_n)_{n \geq 0}$ *converges to c with fixed endpoints* if c_n converges to c and c_n and c share the same starting points and the same finishing points.

1.1.2 The measurable space

Let G be a connected compact Lie group endowed with a bi-invariant metric γ of total volume 1.

Consider a subset $J \subset PM$. We say that a function $f : J \rightarrow G$ is *multiplicative* if the following properties hold. Firstly, for all c_1, c_2 belonging to J such that $\bar{c}_1 = \underline{c}_2$ and $c_1 c_2 \in J$, one has $f(c_1 c_2) = f(c_2) f(c_1)$. Secondly, if both c and c^{-1} belong to J , then $f(c^{-1}) = f(c)^{-1}$. The set of multiplicative functions from J to G is denoted by $\mathcal{M}(J, G)$.

For each path $c \in PM$, denote by $H_c : \mathcal{M}(PM, G) \rightarrow G$ the evaluation at c defined by $H_c(f) = f(c)$. Let \mathcal{C} be the cylinder σ -field on $\mathcal{M}(PM, G)$, that is, the smallest σ -field with respect to which all the mappings H_c , $c \in PM$ are measurable. The Yang-Mills measures¹ are probability measures on the measurable space $(\mathcal{M}(PM, G), \mathcal{C})$.

Consider again a subset $J \subset PM$. Let U be a subset of M such that the endpoints of every path of J belong to U . The group $\mathcal{F}(U, G)$ of all G -valued functions on U acts on $\mathcal{M}(J, G)$ according to the following rule:

$$(j \cdot f)(c) = j(\bar{c})^{-1} f(c) j(\underline{c}),$$

which is inspired by the way a gauge transformation affects the holonomy of a connection. In particular, the group $\mathcal{F}(M, G)$ acts on the measurable space $(\mathcal{M}(PM, G), \mathcal{C})$. This action is measurable and all the probability measures which we shall consider are invariant under this action. This fact has important implications which we shall discuss later.

1.1.3 Graphs

In order to characterize the different instances of the Yang-Mills measure that we are considering here, we need to describe some of their finite dimensional marginals. This involves putting graphs on M and associating to each of them a probability measure on a finite product of copies of G . More details can be found in [9], Sections 1.1 and 1.4, and in [10], from which what follows is inspired.

Definition 1.1 *A graph is a triple $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$, where*

1. \mathbb{E} is a finite collection of edges stable by inversion and such that two distinct edges are either inverse of each other or intersect, if at all, only at some of their endpoints. Moreover, if two edges share the same starting point, then their angle at this point is different from 0 modulo 2π .
2. \mathbb{V} is the set of endpoints of the elements of \mathbb{E} .
3. \mathbb{F} is the set of the closures of the connected components of $M \setminus \bigcup_{e \in \mathbb{E}} e$.
4. Each open connected component of $M \setminus \bigcup_{e \in \mathbb{E}} e$ is diffeomorphic to the open unit disk of \mathbb{R}^2 .
5. The boundary of M is contained in $\bigcup_{e \in \mathbb{E}} e$.

The elements of $\mathbb{V}, \mathbb{E}, \mathbb{F}$ are respectively called vertices, edges and faces of \mathbb{G} . We call *open faces* the connected components of $M \setminus \bigcup_{e \in \mathbb{E}} e$. Beware that an open face could be strictly contained in the interior of its closure. In fact, we make a further assumption that makes this impossible.

We denote by \mathbb{E}^* the set of paths that can be represented by a concatenation of elements of \mathbb{E} . For example, each face of a graph has a boundary which is a cycle in \mathbb{E}^* . We say that a graph is *simple* if the boundary of each one of its faces is a simple cycle. In this paper, we shall always assume that the graphs are simple.

¹There is indeed a whole family of Yang-Mills measures, indexed by the fundamental group of G if M is closed, or by one conjugacy class of G for each connected component of the boundary of M if it is not empty, and in all cases by a positive real parameter which we call temperature and which is really a scaling factor for the volume form σ on M .

We will use the mapping $L : \mathbb{E} \longrightarrow \mathbb{F} \cup \{\emptyset\}$ defined by the fact that, for each edge e , $L(e)$ is the unique face of \mathbb{G} which e bounds with positive orientation, in other words the face located on the left of e . If $e \subset \partial M$ and M is on the right of e , we set $L(e) = \emptyset$.

Finally, let an *unoriented edge* of \mathbb{G} be a pair $\{e, e^{-1}\}$ where $e \in \mathbb{E}$. We call *orientation* of \mathbb{G} a subset \mathbb{E}^+ of \mathbb{E} which contains exactly one element of each unoriented edge.

Lemma 1.2 *Let \mathbb{G} be a graph on M . There exists an orientation \mathbb{E}^+ of \mathbb{G} such that, depending on whether M has a boundary or not, one of the two following properties holds:*

1. *If M is closed, then for each face F , there exists $e \in \mathbb{E}^+$ such that $F = L(e^{-1})$.*
2. *If M has a boundary, then for each $e \in \mathbb{E}^+$ such that $e \subset \partial M$, one has $L(e) \neq \emptyset$.*

Proof – In the case with boundary, the assertion is obvious. In the closed case, it is proved in [10] by using a spanning tree in the dual graph to the fat graph induced by \mathbb{G} . \square

We will always assume that the graphs that we consider are oriented in a way which satisfies the relevant one of these two properties.

1.1.4 The discrete measures

Choose once for all in this section a simple graph \mathbb{G} and a positive real number T . Choose an orientation $\mathbb{E}^+ = \{e_1, \dots, e_r\}$ of \mathbb{G} . We want to describe the discrete Yang-Mills measure at temperature T associated to \mathbb{G} . If M is closed, it is a Borel probability measure on $G^{\mathbb{E}^+}$ absolutely continuous with respect to the Haar measure. If M has a boundary, it is supported by a subset of $G^{\mathbb{E}^+}$ which depends on some boundary conditions one has to specify.

Choose $c \in \mathbb{E}^*$. It can be written $c = e_{i_1}^{\varepsilon_1} \dots e_{i_n}^{\varepsilon_n}$ for some $i_1, \dots, i_n \in 1, \dots, r$ and $\varepsilon_1, \dots, \varepsilon_n = \pm 1$. Define $h_c : G^{\mathbb{E}^+} \longrightarrow G$, the discrete holonomy along c , by setting $h_c(g_1, \dots, g_r) = g_{i_n}^{\varepsilon_n} \dots g_{i_1}^{\varepsilon_1}$. If F is a face of \mathbb{G} , the mapping $h_{\partial F}$ is ill-defined because the cycle ∂F lacks a base point. However, let Ad denote the action of G on itself by conjugation: $\text{Ad}(x)y = xyx^{-1}$. Then, if c is a loop which represents ∂F , the composition of h_c with the canonical projection $G \longrightarrow G/\text{Ad}$ does not depend on the choice of c . We denote this composed mapping by $h_{\partial F} : G^{\mathbb{E}^+} \longrightarrow G/\text{Ad}$.

Finally, let p be the fundamental solution of the heat equation on G , that is, the unique smooth function $p : (0, \infty) \times G \longrightarrow (0, \infty)$, $(t, g) \mapsto p_t(g)$, which is a solution of the equation $(\frac{1}{2}\Delta - \partial_t)p = 0$ and satisfies, for every continuous function f on G , the initial condition $\int_G f(g)p_t(g) dg \xrightarrow{t \rightarrow 0} f(1)$, where 1 is the unit of G . For all $t > 0$, the function p_t is invariant by conjugation on G , so that if F is a face of \mathbb{G} , the function $p_t \circ h_{\partial F} : G^{\mathbb{E}^+} \longrightarrow (0, +\infty)$ is well defined.

We can now define the discrete Yang-Mills measure under the assumptions that M is closed and that G is simply connected.

Definition 1.3 (Closed surface, simply connected group) *Assume that G is simply connected. The discrete Yang-Mills measure associated to \mathbb{G} at temperature T is the Borel probability measure $P_T^{\mathbb{G}}$ on $G^{\mathbb{E}^+}$ defined by*

$$dP_T^{\mathbb{G}}(g) = \frac{1}{Z_T^{\mathbb{G}}} \prod_{F \in \mathbb{F}} p_{T\sigma(F)}(h_{\partial F}(g)) dg,$$

where dg is the Haar measure on $G^{\mathbb{E}^+}$.

Let us drop the assumption that G is simply connected. In order to define the measure, we need a few more definitions.

Let $\pi : \tilde{G} \rightarrow G$ be a universal covering of G and set $\Pi = \pi^{-1}(1)$. Recall that Π is a discrete central subgroup of \tilde{G} canonically isomorphic to the fundamental group of G . If M is closed, then principal G -bundles over M are classified up to bundle isomorphism² by Π . On the other hand, if M has a non-empty boundary, then all principal G -bundles over M are trivial. Let us assume for the moment that M is closed. If P is a principal G -bundle over M , we denote by $\mathfrak{o}(P)$ the element of Π which represents the isomorphism class of P .

If c is a path in \mathbb{E}^* and if $F \in \mathbb{F}$ is a face, then the definitions of the mappings $h_c : G^{\mathbb{E}^+} \rightarrow G$ and $h_{\partial F} : G^{\mathbb{E}^+} \rightarrow G/\text{Ad}$ still make sense when G is replaced by \tilde{G} . When we use the new mappings thus defined, we put a superscript to their names to indicate in which group the mappings take their values, writing for example $h_{\partial F}^{\tilde{G}}$.

Recall that G is endowed with a Riemannian metric γ . Let us endow \tilde{G} with the Riemannian metric $\pi^*\gamma$ and the corresponding Riemannian volume³. Let \tilde{p} be the fundamental solution of the heat equation on \tilde{G} . It is a smooth positive function on $(0, +\infty) \times \tilde{G}$, and for each $t > 0$, \tilde{p}_t is invariant by conjugation. The functions p and \tilde{p} are related by the equality $p_t(x) = \sum_{\pi(\tilde{x})=x} \tilde{p}_t(\tilde{x})$ which holds for all $t > 0$ and $x \in G$.

For each $z \in \Pi$, define a subset $\Pi_z^{\mathbb{F}}$ of $\tilde{G}^{\mathbb{F}}$ by $\Pi_z^{\mathbb{F}} = \{z_{\mathbb{F}} = (z_F)_{F \in \mathbb{F}} \in \Pi^{\mathbb{F}} \mid \prod_{F \in \mathbb{F}} z_F = z\}$. A proof of the following result can be found in [10], Proposition 2.4 and Lemma 2.7.

Proposition 1.4 *Choose $z \in \Pi$. Let g be an element of $G^{\mathbb{E}^+}$. Let $\tilde{g} \in \tilde{G}^{\mathbb{E}^+}$ be a lift of g . Then the number*

$$\sum_{z_{\mathbb{F}} \in \Pi_z^{\mathbb{F}}} \prod_{F \in \mathbb{F}} \tilde{p}_{T\sigma(F)}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F)$$

is finite and does not depend on the choice of \tilde{g} . We denote it by $D_{T,z}^{\mathbb{G}}(g)$. Moreover, the function $D_{T,z}^{\mathbb{G}}$ is bounded on $G^{\mathbb{E}^+}$.

This proposition provides us with a positive function on $G^{\mathbb{E}^+}$ which is the density of the discrete Yang-Mills measure.

Definition 1.5 (Closed surface, the general case) *Choose $T > 0$ and $z \in \Pi$. The discrete Yang-Mills measure at temperature T associated to \mathbb{G} and to the isomorphism class of G -bundles corresponding to z is the Borel probability measure $P_{T,z}^{\mathbb{G}}$ on $G^{\mathbb{E}^+}$ defined by*

$$dP_{T,z}^{\mathbb{G}}(g) = \frac{1}{Z_{T,z}^{\mathbb{G}}} D_{T,z}^{\mathbb{G}}(g) dg.$$

Finally, the case where M has a boundary requires also a few specific definitions. Whether or not G is simply connected does not matter any more since any G -bundle over M is trivial.

²This classical fact is proved in [17], §§29–34. In the case where $G = S^1$, a more accessible proof can be found in [13], Section 6.2. Appendix A of [8] is also worth reading. Finally, a good general reference on principal bundles and smooth connections is [7].

³This metric and this invariant measure are not normalized: the total volume of \tilde{G} is equal to the cardinality of Π , which is finite if and only if G is semi-simple.

On the other hand, it is natural to put constraints on the measure and to insist that the discrete holonomy along each boundary component belong to some fixed conjugacy class in G .

Let X be a conjugacy class in G . For each $n \geq 1$, we are interested in the subset of G^n defined by the equation $g_n \dots g_1 \in X$. This subset is an orbit of the action of G^n on itself defined by $(x_1, \dots, x_n) \cdot (g_1, \dots, g_n) = (x_1 g_1 x_n^{-1}, x_2 g_2 x_1^{-1}, \dots, x_n g_n x_{n-1}^{-1})$. As such, it carries a natural measure which is the image of the Haar measure on G^n . This measure, which we denote by ν_X^n or simply ν_X , can alternatively be characterized by the fact that, for every continuous function f on G^n ,

$$\nu_X^n(f) = \lim_{t \rightarrow 0} \int_{G^{n+1}} f(g_1, \dots, g_n) p_t(g_n \dots g_1 y x^{-1} y^{-1}) dy dg_1 \dots dg_n.$$

The main properties of ν_X^n are the fact that the relation $g_n \dots g_1 \in X$ holds ν_X^n -almost surely and its invariance under circular permutation of the factors in G^n .

Assume now that ∂M has p connected components N_1, \dots, N_p . For each $i = 1, \dots, p$, set $\mathbb{E}_{N_i}^+ = \{e \in \mathbb{E}^+ | e \subset N_i\}$. Set also $\mathbb{E}_{int}^+ = \{e \in \mathbb{E}^+ | e \not\subset \partial M\}$, so that \mathbb{E}^+ is the disjoint union of the subsets we have just defined.

Let N be a component of ∂M . It is also the image of a cycle in \mathbb{E}^* . Let $e_{i_1} \dots e_{i_n}$ be a loop representing this cycle. Then, if X is a conjugacy class in G , the measure $\nu_X^n(g_{i_1}, \dots, g_{i_n})$ is well defined on $G^{\mathbb{E}_N^+}$ and does not depend on the choice of the loop representing ∂N . We denote it by ν_X^N .

Now if we choose p conjugacy classes X_1, \dots, X_p in G , then we can define the measure $\nu_{X_1}^{N_1} \otimes \dots \otimes \nu_{X_p}^{N_p}$ on $G^{\mathbb{E}_{N_1}^+} \times \dots \times G^{\mathbb{E}_{N_p}^+}$. Finally, let us denote by dg_{int} the Haar measure on the remaining factors, namely on $G^{\mathbb{E}_{int}^+}$.

Definition 1.6 (Surface with boundary) Choose p conjugacy classes X_1, \dots, X_p in G . The discrete Yang-Mills measure at temperature T associated to \mathbb{G} with boundary conditions X_1, \dots, X_p along N_1, \dots, N_p respectively is the Borel probability measure $P_{T; X_1, \dots, X_p}^{\mathbb{G}}$ on $G^{\mathbb{E}^+}$ defined by

$$dP_{T; X_1, \dots, X_p}^{\mathbb{G}}(g) = \frac{1}{Z_{T; X_1, \dots, X_p}^{\mathbb{G}}} \prod_{F \in \mathbb{F}} p_{T\sigma(F)}(h_{\partial F}(g)) d\nu_{X_1}^{N_1} \dots d\nu_{X_p}^{N_p} dg_{int}.$$

Whether or not M has a boundary, $G^{\mathbb{E}^+} = \mathcal{M}(\mathbb{E}^+, G)$, so that the group $\mathcal{F}(M, G)$ acts on $G^{\mathbb{E}^+}$. The proof of the following proposition can be found in [9] and [10].

Proposition 1.7 1. The probability measures of Definitions 1.3, 1.5 and 1.6 are all invariant under the action of the group $\mathcal{F}(M, G)$.

2. When M is closed and G is simply connected, then the measures $P_T^{\mathbb{G}}$ and $P_{T,1}^{\mathbb{G}}$ are identical for all $T > 0$.

3. None of the normalization constants which appear in the definitions above depend on \mathbb{G} . They depend only on the genus of M , denoted here by g , its total area $\sigma(M)$ and, when M has a boundary, on X_1, \dots, X_p . Their values are listed below. We use the notation $[a, b] = aba^{-1}b^{-1}$ for $a, b \in G$. Recall also that, when $a, b \in G$ and $\tilde{a}, \tilde{b} \in \tilde{G}$ satisfy $\pi(\tilde{a}) = a$ and $\pi(\tilde{b}) = b$, then $[\tilde{a}, \tilde{b}]$ depends only on a and b . We denote it by $[\tilde{a}, \tilde{b}]$. Finally, we choose $x_1 \in X_1, \dots, x_p \in X_p$.

$$Z_T = \int_{G^{2g}} p_{T\sigma(M)}([a_1, b_1] \dots [a_g, b_g]) da_1 db_1 \dots da_g db_g.$$

$$\begin{aligned}
Z_{T,z} &= \int_{G^{2g}} \tilde{p}_{T\sigma(M)}(\widetilde{[a_1, b_1]} \dots \widetilde{[a_g, b_g]} z) da_1 db_1 \dots da_g db_g. \\
Z_{T;X_1, \dots, X_p} &= \int_{G^{2g+p}} p_{T\sigma(M)}([a_1, b_1] \dots [a_g, b_g] c_1 x_1 c_1^{-1} \dots c_p x_p c_p^{-1}) da_1 db_1 \dots da_g db_g dc_1 \dots dc_p.
\end{aligned}$$

1.1.5 The Yang-Mills measures

The discrete measures we have just defined provide us with many, but not all, finite dimensional marginals of the Yang-Mills measure. For example, they do not allow us to write down the joint distribution of H_c and $H_{c'}$ if c and c' are two paths which intersect so often that $M \setminus (c \cup c')$ has infinitely many connected components. A continuity requirement fills this gap.

We say that a sequence of G -valued random variables $(H_n)_{n \geq 0}$ defined under a probability P converges in probability to H if, for all $\varepsilon > 0$, $P(d(H_n, H) > \varepsilon) \xrightarrow{n \rightarrow \infty} 0$, where d denotes the Riemannian distance on G . We can now state the theorem of existence and uniqueness of the Yang-Mills measure ([9], Theorem 2.58).

Theorem 1.8 (Closed surface) *Choose $T > 0$ and $z \in \Pi$. There exists a unique probability measure $P_{T,z}$ on $(\mathcal{M}(PM, G), \mathcal{C})$ such that the two following properties hold:*

1. *For every graph $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ on M , with orientation $\mathbb{E}^+ = \{e_1, \dots, e_r\}$, the law of $(H_{e_1}, \dots, H_{e_r})$ under $P_{T,z}$ is equal to $P_{T,z}^{\mathbb{G}}$.*
2. *Whenever c belongs to PM and $(c_n)_{n \geq 0}$ is a sequence of PM converging to c with fixed endpoints, $(H_{c_n})_{n \geq 0}$ converges in probability to H_c .*

Moreover, the measure $P_{T,z}$ is invariant under the action of $\mathcal{F}(M, G)$.

Theorem 1.9 (Surface with boundary) *Choose $T > 0$. Let X_1, \dots, X_p be p conjugacy classes in G . There exists a unique probability measure $P_{T;X_1, \dots, X_p}$ on $(\mathcal{M}(PM, G), \mathcal{C})$ such that the two following properties hold:*

1. *For every graph $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ on M , with orientation $\mathbb{E}^+ = \{e_1, \dots, e_r\}$, the law of $(H_{e_1}, \dots, H_{e_r})$ under $P_{T;X_1, \dots, X_p}$ is equal to $P_{T;X_1, \dots, X_p}^{\mathbb{G}}$.*
2. *Whenever c belongs to PM and $(c_n)_{n \geq 0}$ is a sequence of PM converging to c with fixed endpoints, $(H_{c_n})_{n \geq 0}$ converges in probability to H_c .*

Moreover, the measure $P_{T;X_1, \dots, X_p}$ is invariant under the action of $\mathcal{F}(M, G)$.

1.2 Sobolev connections and the Yang-Mills energy

In this paper, we shall often be dealing with connections on principal G -bundles which are not smooth nor even continuous, but belong to some larger Sobolev spaces. We spend this paragraph introducing them carefully. Let us fix a principal G -bundle P .

We call *local section* of P a pair (U, s) , where U is an open subset of M and s is a smooth section of P over U . We often denote such a pair simply by s and use the notation $\text{Dom}(s) = U$ when the domain has to be made explicit.

Let \mathfrak{g} denote the Lie algebra of G . The Riemannian metric γ on G determines a scalar product on \mathfrak{g} invariant by the adjoint action of G . Let $\|\cdot\|$ denote the corresponding Euclidean norm on \mathfrak{g} .

Definition 1.10 *Choose $k \geq 1$ and $p \in [1, +\infty]$. A $W^{k,p}$ connection on P is the data, for each local section s of P , of a $W^{k,p}$ \mathfrak{g} -valued 1-form ω_s on $\text{Dom}(s)$. These locally defined forms must*

satisfy the following compatibility condition. Let s and s' be any two local sections of P . Let ψ be the unique smooth G -valued function on $\text{Dom}(s) \cap \text{Dom}(s')$ such that, on this domain, $s' = s\psi$. Then $\omega_{s'} = \text{Ad}(\psi^{-1})\omega_s + \psi^{-1}d\psi$.

Remark 1.11 A $W^{k,p}$ connection on P is completely specified by the data of the locally defined 1-forms associated to a family of local sections of P whose domains cover M .

We denote the set of $W^{k,p}$ connections on P by $W^{k,p}\mathcal{A}(P)$. We shall in particular consider $W^{s,2}$ connections, which we also call H^s connections. The norms on the spaces $W^{k,p}\mathcal{A}(P)$ depend on the Riemannian metric on M , but since M is compact, different metrics give rise to equivalent norms. Hence, the spaces themselves and their topologies are intrinsically attached to the differentiable structure of M .

Example 1.12 On the trivial bundle $P = M \times G$, one can choose a smooth global section and identify $W^{k,p}$ connections with $W^{k,p}$ \mathfrak{g} -valued 1-forms on M .

In the present context, the most important space of connections is $H^1\mathcal{A}(P)$. It will turn out that this space plays to some extent the role of a Cameron-Martin space for the Yang-Mills measure.

Let ω be a H^1 connection on some bundle P . Choose a local section s of P over the domain of a coordinate chart of M , with coordinates (x, y) . Then one can write $\omega_s = \omega_{s,1}dx + \omega_{s,2}dy$, where $\omega_{s,1}$ and $\omega_{s,2}$ are H^1 \mathfrak{g} -valued functions. Since M is compact and two-dimensional, Sobolev embeddings imply that $\omega_{s,1}$ and $\omega_{s,2}$ are also L^p functions for every finite p . In particular, they are L^4 and the formula

$$\Omega_s = (\partial_x\omega_{s,2} - \partial_y\omega_{s,1} + [\omega_{s,1}, \omega_{s,2}])dx \wedge dy \quad (1)$$

defines locally an L^2 \mathfrak{g} -valued 2-form⁴. As in the smooth case, if s' is another local section and ψ is the smooth G -valued function such that $s' = s\psi$ on $\text{Dom}(s) \cap \text{Dom}(s')$, then $\Omega_{s'} = \text{Ad}(\psi^{-1})\Omega_s$.

Let $\text{Ad}(P)$ denote the vector bundle associated with P by the adjoint action of G on \mathfrak{g} ⁵. The relation between Ω_s and $\Omega_{s'}$ stated above shows that the locally defined 2-forms Ω_s build up into an L^2 $\text{Ad}(P)$ -valued 2-form on M , which is the curvature of ω and is denoted by Ω .

Let us consider the unique L^2 section $*\Omega$ of $\text{Ad}(P)$ such that $\Omega = *\Omega\sigma$. Then the Yang-Mills energy of ω is defined by the formula

$$S(\omega) = \int_M \|*\Omega\|^2 d\sigma.$$

We have just defined a functional $S : H^1\mathcal{A}(P) \rightarrow [0, +\infty)$. If one multiplies the volume form σ by a positive real number T , then $\|*\Omega\|^2$ is multiplied by $\frac{1}{T^2}$ and the energy S is multiplied by $\frac{1}{T}$. Thus, the large deviation principle is concerned with the asymptotic behaviour of the Yang-Mills measure as the area of M is scaled by a factor which tends to zero.

⁴This argument remains valid if the dimension of M is 3 or 4.

⁵Take $m \in M$. Let P_m denote the fibre of P over m . Then the fibre of $\text{Ad}(P)$ over m is the vector space of all mappings $\varphi : P_m \rightarrow \mathfrak{g}$ such that, for all $p \in P_m$ and all $g \in G$, $\varphi(pg) = \text{Ad}(g^{-1})\varphi(p)$. Thus this fibre is isomorphic to \mathfrak{g} , though not canonically. Still, if φ belongs to the fibre, then $\|\varphi(p)\|$ does not depend on $p \in P_m$ and can safely be denoted by $\|\varphi\|$. In other words, the scalar product on \mathfrak{g} induces a metric on $\text{Ad}(P)$.

Example 1.13 *Let us assume that $P = M \times G$. Then the curvature of an H^1 connection ω can be identified with an L^2 \mathfrak{g} -valued 2-form Ω on M , defined in local coordinates by (1). Then Ω can be written in a unique way as $\Omega = *\Omega\sigma$, where $*\Omega$ is a square-integrable \mathfrak{g} -valued function on M . The energy $S(\omega)$ is then nothing but the square of the L^2 norm of this function.*

1.3 The rate functions

In order to define the rate functions of the large deviation principles satisfied by the Yang-Mills measure, we need to understand some properties of the holonomy induced by H^1 connections. This is a bit technical and explaining it completely now would distract us from our goal which is to state the main theorems. Thus, we present here only the main ideas and postpone the technical aspects until Section 2.

Assume that P is a trivial G -bundle over M . Identify H^1 connections on P with H^1 \mathfrak{g} -valued 1-forms on M . Let ω be such a connection. The crucial property of ω as far as holonomy is concerned is the fact that it admits a trace along any smooth 1-dimensional submanifold of M , which is an L^2 \mathfrak{g} -valued function along this submanifold. Thus, if a path $c \in PM$ is parametrized by the interval $[0, 1]$, it makes sense to solve the differential equation $\dot{a}_t a_t^{-1} = -\omega(\dot{c}(t))$, $a_0 = 1$, where the unknown function is $a : [0, 1] \rightarrow G$. If we set $f(c) = a_1$ and do this for each element of PM , we get an element f of the space $\mathcal{M}(PM, G)$. This element is called the holonomy of ω .

If M has a boundary and if N is a component of ∂M , then the holonomy of ω along N is well defined as a conjugacy class in G , namely that of $f(c)$ if c is a simple loop which represents the cycle N . If ∂M has p components N_1, \dots, N_p and if X_1, \dots, X_p are p conjugacy classes in G , then we denote by $H^1\mathcal{A}_{X_1, \dots, X_p}(P)$ the set of H^1 connections on P whose holonomy along N_i is equal to X_i for each $i = 1, \dots, p$.

When P is not trivial, an H^1 connection induces only a class in $\mathcal{M}(PM, G)$ modulo the action of $\mathcal{F}(M, G)$. This class still contains a lot of information about ω . For example, if two connections, not even necessarily defined on the same bundle, induce the same class in $\mathcal{M}(PM, G)$, then they have the same energy.

Let us say that an H^1 connection ω on some G -bundle over M and an element f of $\mathcal{M}(PM, G)$ agree if they determine the same class, that is if f belongs to the class induced by ω . Let us denote this by $\omega \sim f$. The claim made a few lines above implies that, given f in $\mathcal{M}(PM, G)$, if there exists an H^1 connection ω which agrees with f , then $S(\omega)$ depends only on f , not on ω .

Recall that, if M is closed and P is a principal G -bundle over M , then $\mathfrak{o}(P)$ denotes the element of Π which corresponds to the isomorphism class of P among all principal G -bundles.

Definition 1.14 (Rate functions) 1. *Assume that M is closed. Let z be an element of Π . Let P be a principal G -bundle over M such that $\mathfrak{o}(P) = z$. Define the function $I_z^{\text{YM}} : \mathcal{M}(PM, G) \rightarrow [0, +\infty]$ by*

$$I_z^{\text{YM}}(f) = \begin{cases} \frac{1}{2}S(\omega) & \text{if } \exists \omega \in H^1\mathcal{A}(P), f \sim \omega, \\ +\infty & \text{otherwise.} \end{cases}$$

2. *Assume that M has a boundary with p components N_1, \dots, N_p . Let P be a G -bundle over M . Define for every choice of p conjugacy classes X_1, \dots, X_p in G the function $I_{X_1, \dots, X_p}^{\text{YM}} : \mathcal{M}(PM, G) \rightarrow [0, +\infty]$ by*

$$I_{X_1, \dots, X_p}^{\text{YM}}(f) = \begin{cases} \frac{1}{2}S(\omega) & \text{if } \exists \omega \in H^1\mathcal{A}_{X_1, \dots, X_p}(P), f \sim \omega, \\ +\infty & \text{otherwise.} \end{cases}$$

1.4 Statement of the main results

Theorem 1.15 (Surface with boundary) *Let N_1, \dots, N_p be the connected components of ∂M . Let X_1, \dots, X_p be p conjugacy classes in G . The family $(P_{T;X_1, \dots, X_p})_{T>0}$ of probability measures on $(\mathcal{M}(PM, G), \mathcal{C})$ satisfies, as T tends to 0, a large deviation principle of speed T with rate function $I_{X_1, \dots, X_p}^{\text{YM}}$.*

Theorem 1.16 (Closed surface) *Let z be an element of Π . The family $(P_{T,z})_{T>0}$ of probability measures on $(\mathcal{M}(PM, G), \mathcal{C})$ satisfies, as T tends to 0, a large deviation principle of speed T with rate function I_z^{YM} .*

A good general reference on large deviation principles is [2]. For the sake of clarity, let us explain the meaning of the second theorem. The space $\mathcal{M}(PM, G)$ is endowed with the trace of the product topology on G^{PM} . That the function I_z^{YM} is a rate function⁶ means that it is lower semi-continuous on $\mathcal{M}(PM, G)$. Then, the large deviation principle asserts that, for each measurable subset $A \in \mathcal{C}$ of $\mathcal{M}(PM, G)$, with closure \bar{A} and interior A° , the following inequalities hold:

$$-\inf_{f \in A^\circ} I_{YM}^z(f) \leq \liminf_{T \rightarrow 0} T \log P_{T,z}(A) \leq \limsup_{T \rightarrow 0} T \log P_{T,z}(A) \leq -\inf_{f \in \bar{A}} I_{YM}^z(f).$$

2 Holonomy of H^1 connections

2.1 Local results

Let P be a principal G -bundle over M . Let (U, s) be a local section of P . An H^1 connection ω determines by definition a \mathfrak{g} -valued 1-form ω_s on U . In this first section, we will focus on what can be done with this single locally defined 1-form. We use the notation $\Omega^1(U)$ and $H^1\Omega^1(U)$ respectively for the spaces of smooth and H^1 real-valued 1-forms on U . We put a subscript \mathfrak{g} to indicate \mathfrak{g} -valued forms or functions.

The covering $\pi : \tilde{G} \rightarrow G$ induces an isomorphism of Lie algebras through which we identify \mathfrak{g} with the Lie algebra of \tilde{G} . Most of the results we are about to prove hold in \tilde{G} as well as in G . Hence we decide, until the end of Section 2.1, that G is any Lie group with Lie algebra \mathfrak{g} . Such a group is the direct product of a compact Lie group and a group isomorphic to \mathbb{R}^m for some $m \geq 0$.

2.1.1 Holonomy

Let e be an edge contained in U . Let us choose a smooth parametrization $e : [0, 1] \rightarrow M$ of e . If ω_0 belongs to $\Omega_{\mathfrak{g}}^1(U)$, then $e^*\omega_0$ is a smooth 1-form on $[0, 1]$ with values in \mathfrak{g} . We identify this 1-form with the \mathfrak{g} -valued function $t \mapsto e^*\omega_0(\partial_t) = \omega_0(\dot{e}_t)$.

If now $(\omega_n)_{n \geq 0}$ is a sequence of smooth 1-forms which converges in $H^1\Omega_{\mathfrak{g}}^1(U)$ to an H^1 1-form ω , then a classical result ([1], Theorem 5.22) asserts that the sequence $(e^*\omega_n)_{n \geq 0}$ of \mathfrak{g} -valued functions on $[0, 1]$ converges in $H_{\mathfrak{g}}^{1/2}([0, 1])$ to a limit which depends only on ω . This limit is denoted by $e^*\omega$ and it is called the *trace* of ω along the parametrized edge e . Moreover, the mapping $e^* : H^1\Omega_{\mathfrak{g}}^1(U) \rightarrow H_{\mathfrak{g}}^{1/2}([0, 1])$ is continuous. By composing this mapping with the

⁶Since $\mathcal{M}(PM, G)$ is a compact space, it does not tell much to say that it is also a *good* rate function.

compact embedding $H_{\mathfrak{g}}^{1/2} \hookrightarrow L_{\mathfrak{g}}^2$, we get a compact linear mapping $e^* : H^1\Omega_{\mathfrak{g}}^1(U) \longrightarrow L_{\mathfrak{g}}^2([0, 1])$. We define the holonomy of ω along e by rolling the curve $e^*\omega$ onto G . For this, consider the space of G -valued functions on $[0, 1]$

$$H_{\bullet}^1([0, 1]; G) = \{a \in H^1([0, 1]; G) | a_0 = 1\}.$$

Although this space is not a vector space, let us call H^1 norm of one of its elements a the number $\|a\|_{H_{\bullet}^1}^2 = \int_0^1 \|\dot{a}_t a_t^{-1}\|^2 dt$. The following result is classical and explains what is meant by rolling a curve in \mathfrak{g} onto G .

Proposition 2.1 *The rolling map $R : L_{\mathfrak{g}}^2([0, 1]) \longrightarrow H_{\bullet}^1([0, 1]; G)$ which assigns to each $\alpha \in L_{\mathfrak{g}}^2([0, 1])$ the unique element $a \in H_{\bullet}^1([0, 1]; G)$ such that*

$$\dot{a}_t a_t^{-1} = -\alpha_t \text{ a.e. } t \in [0, 1] \quad (2)$$

is a norm-preserving homeomorphism.

Let ω and e be given as above. According to this proposition, we can consider the H^1 G -valued function $R(e^*\omega)$. It is straightforward to check that if we change the parametrization of e , replacing e by $e \circ \varphi$ for some diffeomorphism φ of $[0, 1]$, then $R(e^*\omega)$ is replaced by $R((e \circ \varphi)^*\omega) = R(e^*\omega) \circ \varphi$. In particular, the element $R(e^*\omega)(1)$ of G is independent of the parametrization of e . We denote this element by $\langle \omega, e \rangle$ or $\langle \omega, e \rangle_G$ if there is any ambiguity.

Choose $t \in (0, 1)$. Consider the two edges $e_1 = e|_{[0, t]}$ and $e_2 = e|_{[t, 1]}$. Then it follows from (2) that $\langle \omega, e \rangle = \langle \omega, e_2 \rangle \langle \omega, e_1 \rangle$. This relation implies the following: if the concatenations of edges $e_1 \dots e_n$ and $f_1 \dots f_n$ are equivalent, then the products $\langle \omega, e_n \rangle \dots \langle \omega, e_1 \rangle$ and $\langle \omega, f_n \rangle \dots \langle \omega, f_1 \rangle$ are equal. Thus, if $c \in PM$ is a path contained in U , then $\langle \omega, c \rangle$ is well defined. We call it the holonomy of ω along c and denote it sometimes by $\langle \omega, c \rangle_G$. We use the notation PU for the set of paths contained in U .

Proposition 2.2 *Let ω be an element of $H^1\Omega_{\mathfrak{g}}^1(U)$.*

1. *The mapping $\langle \omega, \cdot \rangle : PU \longrightarrow G$ belongs to $\mathcal{M}(PU, G)$.*
2. *Let $(\omega_n)_{n \geq 1}$ be a sequence of elements of the Hilbert space $H^1\Omega_{\mathfrak{g}}^1(U)$ which converges weakly to ω . Then, for each $c \in PU$, $\langle \omega_n, c \rangle$ converges to $\langle \omega, c \rangle$.*

Proof – The first statement is straightforward. The second one is a consequence of the compactness of the mapping c^* . Indeed, such a compact mapping sends a weakly convergent sequence to a strongly convergent one. Hence, $R(c^*\omega_n)$ converges in the H^1 topology to $R(c^*\omega)$, in particular uniformly, and the result holds. \square

Lemma 2.3 *Let ω and ω' two elements of $H^1\Omega_{\mathfrak{g}}^1(U)$. Then the functions $\langle \omega, \cdot \rangle$ and $\langle \omega', \cdot \rangle$ are equal on PU if and only if $\omega = \omega'$.*

Proof – Write locally $\omega = \alpha dx + \beta dy$ and $\omega' = \alpha' dx + \beta' dy$, where $\alpha, \alpha', \beta, \beta'$ belong to $H_{\mathfrak{g}}^1(U)$. Since ω and ω' have the same holonomy along every vertical segment in U , the forms β and β' have the same trace along every vertical segment. Thus, by Fubini's theorem, the integral over any rectangle of the difference $\beta - \beta'$ is equal to 0. Hence, $\beta = \beta'$ on U . The same argument with horizontal segments shows that $\alpha = \alpha'$, hence $\omega = \omega'$. \square

2.1.2 Energy inequality and the continuity of the holonomy

We keep considering an element ω of $H^1\Omega_{\mathfrak{g}}^1(U)$. Let Ω be the element of $L^2\Omega_{\mathfrak{g}}^2(U)$ defined in local coordinates by (1). Let $*\Omega$ denote the unique element of $L_{\mathfrak{g}}^2(U)$ such that $\Omega = *\Omega\sigma$. If $V \subset U$ is an open subdomain of U , let us define the Yang-Mills energy of ω on V by $S_V(\omega) = \int_V \|*\Omega\|^2 d\sigma$. For each $x \in G$, we denote by $\rho(x)$ the Riemannian distance in G between x and the unit element. The following result was proved by A. Sengupta in [16].

Proposition 2.4 (Energy inequality) *Let l be a simple loop in U which bounds a domain V diffeomorphic to a disk. Assume that ω is smooth. Then one has the inequality*

$$\rho(\langle\omega, l\rangle_G)^2 \leq \sigma(V)S_V(\omega). \quad (3)$$

For the sake of completeness, we give a proof of this inequality in the appendix (Corollary A-2).

Proposition 2.5 *Proposition 2.4 is still true under the weaker assumption that ω belongs to $H^1\Omega_{\mathfrak{g}}^1(U)$.*

Proof – Let us approximate ω in H^1 norm by a sequence $(\omega_n)_{n \geq 0}$ of smooth 1-forms. On one hand, the functional $S_V : H^1\Omega_{\mathfrak{g}}^1(U) \rightarrow [0, +\infty)$ is continuous because the embedding $H_{\mathfrak{g}}^1(U) \hookrightarrow L_{\mathfrak{g}}^4(U)$ is continuous. On the other hand, by Proposition 2.2, $\langle\omega_n, l\rangle_G$ converges to $\langle\omega, l\rangle_G$. Hence, inequality (3) passes to the limit. \square

This energy inequality is essential. For example, it suffices to imply the following continuity result.

Proposition 2.6 *Consider $\omega \in H^1\Omega_{\mathfrak{g}}^1(U)$. The multiplicative mapping $\langle\omega, \cdot\rangle : PU \rightarrow G$ is continuous with fixed endpoints. This means that, whenever c_n converges to c with fixed endpoints, $\langle\omega, c_n\rangle$ converges to $\langle\omega, c\rangle$.*

We prefer to state and prove the following slightly more general result.

Proposition 2.7 *Let U be an open domain contained in the interior of M . Let f be an element of $\mathcal{M}(PU, G)$. Assume that there exists a constant K such that, for every sub-domain V of U diffeomorphic to an open disk and bounded by a simple loop ∂V , one has the inequality*

$$\rho(f(\partial V))^2 \leq K\sigma(V). \quad (4)$$

Then f is continuous with fixed endpoints.

Proof – The arguments for this proof are spread in the sections 2.4 to 2.6 of [9], but this result was not stated there. We give here a full sketch of proof and refer the reader to [9] for the details.

Let $f \in \mathcal{M}(PM, G)$ satisfy (4). From now on in this proof, all references are to be found in [9]. The first step is to use an isoperimetric inequality on M (Proposition 2.15), which holds locally, to deduce the existence of a new constant, still denoted by K , such that, for any simple loop l of sufficiently small length $\ell(l)$, $\rho(f(l)) \leq K\ell(l)$.

Since f is multiplicative, this inequality can be extended to short loops with finite self-intersection, defined in Definition 2.11. It holds in particular for piecewise geodesic loops. Hence, for every short enough piecewise geodesic loop ζ , one has $\rho(f(\zeta)) \leq K\ell(\zeta)$.

Now, let c be an edge. It is possible to find a sequence of piecewise geodesic paths $(\alpha_n)_{n \geq 1}$ converging to c with fixed endpoints and such that, for all n , $\alpha_n^{-1}c$ is a simple loop bounding a domain diffeomorphic to a disk, whose area tends to zero as n tends to infinity. Hence, $f(\alpha_n)$ tends to $f(c)$. This is explained in Section 2.5.3.

Then, the arguments of the proofs of Lemma 2.36 and Proposition 2.35 show that, whenever (ζ_n) is a sequence of piecewise geodesic paths converging to c with fixed endpoints, $f(\zeta_n)$ tends to $f(c)$. This is the main step of the proof. It involves cutting each ζ_n in three parts, two short loops based at the endpoints of c and one path with the same endpoints as c and staying in a tubular neighbourhood of c . The two short loops do not contribute asymptotically to $f(\zeta_n)$ and the path in the tubular neighbourhood is shown to have a value under f close to that of c , by comparing it with an appropriate term in the first approximating sequence (α_n) .

After extending the result to the case of a piecewise embedded path c (Section 2.6.4), one concludes that f is continuous with fixed endpoints (Proposition 2.42). \square

2.1.3 Gauge transformations

Let $k \geq 1$ be an integer and $p \geq 2$ be a real number or $+\infty$. In what follows, we consider $W^{k,p}$ connections. Since M is compact, they are in particular H^1 , so that the results of the preceding sections apply.

The group $W^{k+1,p}(U; G)$ of $W^{k+1,p}$ G -valued functions on U acts on $W^{k,p}\Omega_{\mathfrak{g}}^1(U)$ as follows. If j belongs to $W^{k+1,p}(U; G)$ and ω to $W^{k,p}\Omega_{\mathfrak{g}}^1(U)$, then

$$j \cdot \omega = \text{Ad}(j^{-1})\omega + j^{-1}dj.$$

Since M is 2-dimensional, the Sobolev embedding $W^{k+1,p}(U) \hookrightarrow C^{k-1}(U)$ holds and $W^{k+1,p}$ functions can be evaluated at any point. If c belongs to PU , then the relation

$$\langle j \cdot \omega, c \rangle = j(\bar{c})^{-1} \langle \omega, c \rangle j(\underline{c}) \tag{5}$$

is classical if all objects are smooth and easy to check in our setting. It fits with the action of $\mathcal{F}(M, G)$ on $\mathcal{M}(PU, G)$.

We need to establish a result which allows us to determine when two elements of $H^1\Omega_{\mathfrak{g}}^1(U)$ differ by the action of an element of $H^2(U; G)$ and how regular this element is.

Lemma 2.8 *Let $V \subset U$ be a connected subset of U , not necessarily open, such that any two points of V can be joined by a path contained in V . Let m be a point of V .*

Let ω and ω' be two elements of $H^1\Omega_{\mathfrak{g}}^1(U)$. Assume that there exists an element x of G such that, for each loop l based at m and contained in V , one has $\langle \omega', l \rangle = \text{Ad}(x^{-1})\langle \omega, l \rangle$.

Then there exists $j \in \mathcal{F}(V, G)$ such that $j(m) = x$ and the multiplicative functions $j \cdot \langle \omega, \cdot \rangle$ and $\langle \omega', \cdot \rangle$ coincide on the set PV of paths contained in V .

Proof – The mapping $j : V \rightarrow G$ must satisfy the following condition: for each path starting at m and finishing at some point n , $j(n) = \langle \omega, c \rangle j(m) \langle \omega', c \rangle^{-1}$.

Set $j(m) = x$. Then the assumption on ω and ω' shows that, for each $n \in V$, the value of $j(n)$ determined by the equation above does not depend on the choice of the path c joining m to n . This defines an element of $\mathcal{F}(V, G)$ which satisfies the required property. \square

Proposition 2.9 *Assume that U is connected and that its boundary is locally the graph of a Lipschitz function. Endow M with an auxiliary Riemannian metric and let $P_g U$ be the set of paths contained in U which are piecewise geodesic for this metric. Let $k \geq 1$ be an integer and $p \in [2, +\infty]$ be a real number or ∞ . Let ω and ω' be two elements of $W^{k,p}\Omega_{\mathfrak{g}}^1(U)$. Assume that some element $j \in \mathcal{F}(U, G)$ satisfies, for each path $c \in P_g U$, the relation*

$$\langle \omega', c \rangle = j(\bar{c})^{-1} \langle \omega, c \rangle j(\underline{c}). \quad (6)$$

Then $j \in W^{k+1,p}(U; G)$ and $j \cdot \omega = \omega'$.

Since for every smooth open domain V whose closure is contained in U , one has $C^\infty(\bar{V}) = \bigcap_{k \geq 1} W^{k,2}(V)$, the proposition implies in particular that, if ω and ω' are smooth on U , then j is smooth on U .

Proof – We prove this result by induction on k . Proving that j belongs to $W^{1,p}$ is a non-trivial step of the proof since we assume no regularity at all a priori on j . However, (6) allows us to prove that j is continuous and indeed $W^{1,q}$ for all $q < \infty$.

We will use the fact that every point of U has a neighbourhood diffeomorphic to $(-1, 1)^2$ in such a way that all horizontal and all vertical segments are geodesic. To construct such a neighbourhood, choose two geodesic segments γ_h and γ_v which cross orthogonally at m . Let T_h and T_v be two tubular neighbourhoods of γ_h and γ_v respectively (see [5] for a definition of tubular neighbourhoods and Fermi or normal coordinates). Let (x_h, r_h) and (x_v, r_v) respectively be normal coordinates on these tubes, such that r_h is the distance to γ_h and r_v the distance to γ_v . The mapping $T_h \cap T_v \rightarrow \mathbb{R}^2$ which sends a point n to $(x_h(n), x_v(n))$ is smooth and its differential at m is invertible. On a small neighbourhood of m , this mapping is a coordinate chart such that all segments parallel to the axes are geodesic.

For more convenience in dealing with Sobolev spaces, we will also assume that G is a subgroup of a vector space of matrices. This puts no further restriction on G , since it is by our assumptions the direct product of a compact group and a group isomorphic to $(\mathbb{R}^m, +)$. In this way, \mathfrak{g} is a subspace of the same vector space of matrices. So, we consider that all forms and functions on U are matrix-valued.

• Let us prove that j is continuous. Let m be a point of U . Let us restrict ourselves to a neighbourhood of m as above.

For each $x \in (-1, 1)$, let h_x denote the horizontal segment joining m to $(x, 0)$. For each $(x, y) \in (-1, 1)^2$, let $v_{x,y}$ denote the vertical segment joining $(x, 0)$ to (x, y) . Finally, for each $(x, y) \in (-1, 1)^2$, set $c_{x,y} = h_x v_{x,y}$. This is a piecewise geodesic path and this allows us to write

$$j(x, y) = \langle \omega, c_{x,y} \rangle j(0, 0) \langle \omega', c_{x,y} \rangle^{-1}.$$

From this equality we deduce that

$$\begin{aligned} d_G(j(0, 0), j(x, y)) &\leq \rho(\langle \omega, c_{x,y} \rangle) + \rho(\langle \omega', c_{x,y} \rangle) \\ &\leq \rho(\langle \omega, h_x \rangle) + \rho(\langle \omega, v_{x,y} \rangle) + \rho(\langle \omega', h_x \rangle) + \rho(\langle \omega', v_{x,y} \rangle). \end{aligned}$$

Assume that $0 \leq x \leq \frac{1}{2}$. Then $x \mapsto \langle \omega, h_x \rangle = R(h_{\frac{1}{2}}^* \omega)(x)$ is an absolutely continuous function of x . Now let $\varphi : (-1, 1)^2 \rightarrow [0, 1]$ be a smooth cutoff function such that $\varphi(x, y) = 1$ if $\max(|x|, |y|) \leq \frac{1}{2}$ and $\varphi(x, y) = 0$ if $\max(|x|, |y|) \geq \frac{3}{4}$. Then, if $0 \leq x, y \leq \frac{1}{2}$, then $\langle \omega, v_{x,y} \rangle = \langle \varphi \omega, v_{x,y} \rangle$. Let us write $\omega = \omega_1 dx + \omega_2 dy$. Then, by an elementary estimate ([1], Lemma 5.7), there exists a constant C independent of y such that

$$\int_0^y \|\omega_2(x, t)\|^2 dt \leq C \int_{(-1,1) \times [0,y]} \|\omega_2(s, t)\|^2 + \|\partial_x \omega_2(s, t)\|^2 ds dt.$$

Since ω belongs to H^1 , the right hand side tends to 0 as y tends to 0. Thus, the L^2 norm of $v_{x,y}^* \omega$ tends to 0 as y tends to 0, uniformly in x . This implies that $\rho(\langle \omega, v_{x,y} \rangle)$ tends to 0 as (x, y) tends to $(0, 0)$. The same arguments applied to ω' finish the proof that j is continuous at m .

• Let us prove that j belongs to $W^{1,q}(U)$ for all $q < \infty$. We begin by proving that j admits a weak derivative and for this we restrict again to a neighbourhood diffeomorphic to $(-1, 1)^2$ of some point m . For each $(x, y) \in (-1, 1)^2$, we have

$$j(x, y) = \langle \omega, v_{x,y} \rangle j(x, 0) \langle \omega', v_{x,y} \rangle^{-1}.$$

Hence, for any fixed x , the map $y \mapsto j(x, y)$, as a product of two H^1 functions of one variable, belongs to H^1 . Let us compute its derivative:

$$\begin{aligned} \partial_y j(x, y) &= -\omega(\partial_y) \langle \omega, v_{x,y} \rangle j(x, 0) \langle \omega', v_{x,y} \rangle^{-1} + \langle \omega, v_{x,y} \rangle j(x, 0) \langle \omega', v_{x,y} \rangle^{-1} \omega'(\partial_y) \\ &= j(x, y) \omega'(\partial_y) - \omega(\partial_y) j(x, y). \end{aligned}$$

This shows that j is absolutely continuous along every vertical segment and admits there an almost everywhere derivative which is the trace of the function $j\omega'(\partial_y) - \omega(\partial_y)j$. Hence, j admits a weak partial derivative with respect to y , namely $j\omega' - \omega j$ evaluated on the vector field ∂_y . A similar statement holds in the direction of x and j admits the weak differential $dj = j\omega' - \omega j$.

Since j is continuous, this weak differential belongs to $L_{loc}^q \Omega^1(U)$ for all $q < \infty$. Hence, $j \in W_{loc}^{1,q}(U)$ for all $q < \infty$. If G is compact, j is in fact bounded, so that $dj \in L^q \Omega^1(U)$ and $j \in W^{1,q}(U)$ for all $q < \infty$. If G is not compact, it is the direct product of a compact group by \mathbb{R}^m for some integer $m \geq 1$. The \mathfrak{g} -valued 1-forms and the action of G -valued functions on them split between the compact part of G and the part isomorphic to \mathbb{R}^m . It is thus enough to check that $j \in W^{1,q}(U)$ when $G = \mathbb{R}^m$.

In this case, one has $dj = \omega' - \omega$, which belongs to $H^1(U)$. Set $f = *d*(\omega' - \omega) \in L^2(U)$. Let ν be the outer normal vector field along the boundary of U and let $g = (\omega' - \omega)(\nu) \in H^{1/2}(\partial U)$. Then, up to an additive constant, j is the unique solution to the inhomogeneous Neumann problem

$$\begin{cases} \Delta j = f & \text{on } U \\ \partial_\nu j = g & \text{on } \partial U. \end{cases}$$

Hence, according to standard results on elliptic boundary value problems, j belongs to $H^2(U)$. This implies that j admits a continuous extension on the closure of U . In particular, it is bounded on U . Finally, for all $q < \infty$, $dj \in L^q \Omega^1(U)$, so that $j \in W^{1,q}(U)$.

• Let us prove that j belongs to $W^{k+1,p}(U)$. In a first step, let $r \geq 2$ be such that ω and ω' belong to $W^{1,r} \Omega^1(U)$. By using Leibnitz's rule, it is easy to check that the product of a function

of $W^{1,r}(U)$ by a function of $C^0(\overline{U}) \cap W^{1,2r}(U)$ belongs to $W^{1,r}(U)$. Hence, $dj \in W^{1,r}\Omega^1(U)$ and $j \in W^{2,r}(U)$. If $k = 1$, the proof is finished.

In general, we use a simple iteration argument. Set $W^{l,\infty-}(U) = \bigcap_{2 \leq p < \infty} W^{l,p}(U)$. We use the fact that $W^{k,p}(U)$ is stable by multiplication as soon as $kp > 2$. In particular, if ω and ω' belong to $W^{l,\infty-}$ for some $l \geq 1$, and if j belongs to $W^{l,\infty-}$, then dj belongs to $W^{l,\infty-}\Omega^1(U)$ and j belongs actually to $W^{l+1,\infty-}(U)$. Now we use the Sobolev embedding $W^{l,q}(U) \hookrightarrow W^{l-1,\infty-}$, which is valid for all $q \geq 2$ and all $l \geq 1$. By this embedding, ω and ω' belong to $W^{l,\infty-}(U)$ for $l = 0, \dots, k-1$. Assume that $k > 1$. Since ω and ω' belong to $W^{1,\infty-}(U)$, we have proved above that j belongs to $W^{1,\infty-}(U)$ and we conclude by iteration that $j \in W^{k,\infty-}(U)$. In particular, j and hence dj belong to $W^{k,p}(U)$. Finally, $j \in W^{k+1,p}(U)$.

The fact that $j \cdot \omega = \omega'$ is obvious. \square

2.2 Global results

Let P be a principal G -bundle over M . Let ω be an element of $H^1\mathcal{A}(P)$. In this second section, we explain how the holonomy of the locally defined 1-forms ω_s indexed by local sections of P fit together into a global object.

2.2.1 Holonomy

For each point m of M , we denote by P_m the fibre of P over m . Let m and n be two points of M . Recall that a mapping $\tau : P_m \rightarrow P_n$ is said to be G -equivariant if, for all $p \in P_m$ and all $g \in G$, one has $\tau(pg) = \tau(p)g$.

Definition 2.10 *A holonomy or parallel transport on P is a collection of mappings $(\tau_c, c \in PM)$ indexed by PM with the following properties:*

1. For each $c \in PM$, τ_c is a G -equivariant mapping from $P_{\underline{c}}$ to $P_{\overline{c}}$.
2. The collection $(\tau_c, c \in PM)$ is multiplicative, that is, for each path c , $\tau_{c^{-1}} = \tau_c^{-1}$ and, if c_1 and c_2 are two paths such that $\overline{c_1} = \underline{c_2}$, then $\tau_{c_1 c_2} = \tau_{c_2} \circ \tau_{c_1}$.

The set of holonomies on P is denoted by $\mathcal{I}(P)$.

Remark 2.11 *Let m and n be two points of M . Since G acts transitively on P_m , a G -equivariant mapping $\tau : P_m \rightarrow P_n$ is determined by the image of any single point of P_m . If a point p is chosen in P_m and a point q in P_n , there exists a unique element g of G such that $\tau(p) = qg$. Conversely, for each $g \in G$ there exists a unique G -equivariant mapping $\tau : P_m \rightarrow P_n$ such that $\tau(p) = qg$. This one-to-one correspondence defines a natural topology on the set of equivariant mappings from P_m to P_n .*

Proposition 2.12 *Every element of $H^1\mathcal{A}(P)$ determines a holonomy on P . Moreover, if two elements of $H^1\mathcal{A}(P)$ determine the same holonomy, then they are equal.*

Proof – Let ω be an H^1 connection on P . Let c be a path on M . Assume that there exists a local section (U, s) of P such that c is contained in the domain of U . Then define τ_c^s to be the unique G -equivariant mapping from $P_{\underline{c}}$ to $P_{\overline{c}}$ such that $\tau_c^s(s(\underline{c})) = s(\overline{c})\langle\omega_s, c\rangle$. The collection of mappings $(\tau_c^s, c \in PU)$ is clearly multiplicative in the sense of Definition 2.10. In particular, if $c = c_1 c_2$, then $\tau_c^s = \tau_{c_2}^s \tau_{c_1}^s$.

If s' is another local section of P on a domain which contains c , then one checks easily by using Definition 1.10 and (5) that the mappings τ_c^s and $\tau_c^{s'}$ are the same. Let us denote them both by τ_c .

Now pick any path c in PM . Write c as a concatenation of shorter paths $c = c_1 \dots c_n$ in such a way that each shorter path is contained in the domain of a local section of P . Then the mapping $\tau_{c_n} \dots \tau_{c_1}$ does not depend on the decomposition of c . Indeed, if $c = c'_1 \dots c'_m$ is another decomposition of c , then there exists a third decomposition $c = c''_1 \dots c''_r$ which is finer than the two other ones, in the sense that each c_i and each c'_j can be written as a concatenation of some c''_k 's. Then the multiplicativity stated above inside each domain of a local section implies that the two first decompositions give rise to the same mapping $P_{\underline{c}} \longrightarrow P_{\bar{c}}$ as the third decomposition. It is thus legitimate to call this mapping τ_c .

Two H^1 connections ω and ω' induce the same holonomy if and only if, for each local section (U, s) of P , the multiplicative functions $\langle \omega_s, \cdot \rangle$ and $\langle \omega'_s, \cdot \rangle$ are equal on PU . According to Lemma 2.3, this is equivalent to the fact that $\omega_s = \omega'_s$ on the interior of M . If M has a boundary, this shows that $\omega = \omega'$ almost everywhere, hence $\omega = \omega'$. \square

The proof of the following lemma is straightforward.

Lemma 2.13 *Let ω be an element of $H^1\mathcal{A}(P)$. Let l and l' be two loops on M , based respectively at m and m' . Assume that l and l' are cyclically equivalent, that is, that they differ only by the choice of their base points. Choose $p \in P_m$ and $p' \in P_{m'}$. Let g and g' be the elements of G such that $\tau_l(p) = pg$ and $\tau_{l'}(p') = p'g'$. Then g and g' are conjugate.*

If c is a cycle on M , for example a component of ∂M or the boundary of a face of a graph on M , then this lemma allows us to define $\langle \omega, c \rangle$ as a *conjugacy class* of G . In particular, if ∂M has p components N_1, \dots, N_p and if X_1, \dots, X_p are p conjugacy classes in G , then we set

$$H^1\mathcal{A}_{X_1, \dots, X_p}(P) = \{\omega \in H^1\mathcal{A}(P) : \langle \omega, N_1 \rangle = X_1, \dots, \langle \omega, N_p \rangle = X_p\}.$$

The space $H^1\mathcal{A}(P)$ is an affine space with direction $H^1\Omega^1(M) \otimes \text{Ad}(P)$. We endow $H^1\mathcal{A}(P)$ with the corresponding topology. Concretely, this topology is generated by the subsets $\{\omega \mid \|\omega_s - \eta_s\|_{H^1} < \varepsilon\}$, where ε runs over the positive reals, s over the local sections of P and η_s over $H^1\Omega^1_{\mathfrak{g}}(\text{Dom}(s))$. This topology can be metrized by choosing a finite covering of M . The weak topology on $H^1\mathcal{A}(P)$ is defined similarly.

Proposition 2.14 *Let $(\omega_n)_{n \geq 0}$ be a sequence of H^1 connections on P which converges weakly to a connection ω_∞ . For each n with $0 \leq n \leq \infty$, let $(\tau_c^n, c \in PM)$ be the holonomy induced by ω_n . Then, for each path $c \in PM$, the mappings τ_c^n converge to τ_c^∞ .*

Proof – If c is contained in the domain of a local section of P , then the result is a direct consequence of Proposition 2.2. If c is not contained in the domain of a local section, then we decompose it as a concatenation of shorter paths to which the local argument can be applied. The result follows by multiplicativity of the holonomy induced by a connection. \square

2.2.2 Gauge transformations

Definition 2.15 A gauge transformation on P is a collection of mappings $(\gamma_m, m \in M)$ indexed by the points of M such that, for each $m \in M$, γ_m is a G -equivariant mapping of P_m onto itself. The set of gauge transformations on P is denoted by $\mathcal{J}(P)$.

Let $j = (\gamma_m, m \in M)$ be a gauge transformation. Let (U, s) be a local section of P . Define j_s as the unique G -valued function on U such that, for each $m \in U$, $\gamma_m(s(m)) = s(m)j_s(m)$. It is easily checked that, if s' is another local section of P and if $s' = s\psi$ on $\text{Dom}(s) \cap \text{Dom}(s')$, then $j_{s'} = \psi^{-1}j_s\psi$.

In fact, a gauge transformation is completely determined by the family of locally defined G -valued mappings j_s , where s runs over local sections of P . It can be defined as such a family which satisfies the compatibility condition stated above. If in addition one puts regularity conditions on the mappings j_s , this allows us to define Sobolev gauge transformations. In particular, we shall consider the space $H^2\mathcal{J}(P)$ of H^2 gauge transformations.

The set $\mathcal{J}(P)$ and the space $H^2\mathcal{J}(P)$ are groups under pointwise composition. These groups act respectively on $\mathcal{T}(P)$ and $H^1\mathcal{A}(P)$ as follows⁷.

Let $T = (\tau_c, c \in PM)$ be a holonomy and $j = (\gamma_m, m \in M)$ be a gauge transformation. We define a new holonomy $j \cdot T = (j \cdot \tau_c, c \in PM)$ by setting, for each path c , $(j \cdot \tau)_c = \gamma_{\bar{c}}^{-1} \circ \tau_c \circ \gamma_{\underline{c}}$.

On the other hand, let ω be an H^1 connection and j an H^2 gauge transformation. Then we define a new element $j \cdot \omega$ of $H^1\mathcal{A}(P)$ by setting, for each local section s of P , $(j \cdot \omega)_s = j_s \cdot \omega_s = \text{Ad}(j_s^{-1})\omega_s + j_s^{-1}dj_s$.

According to (6), the mapping $H^1\mathcal{A}(P) \longrightarrow \mathcal{T}(P)$ which sends a connection to its holonomy induces a mapping between quotient spaces:

$$H^1\mathcal{A}(P)/H^2\mathcal{J}(P) \longrightarrow \mathcal{T}(P)/\mathcal{J}(P). \quad (7)$$

Let us choose a reference point $p(m)$ in the fibre P_m for each $m \in M$. Then, according to the Remark 2.11, a holonomy on P determines an element of $\mathcal{F}(PM, G)$. It follows from the multiplicativity of a holonomy that this function is actually multiplicative. Now it is easily checked that changing the reference point in each fibre or changing the holonomy by the action of a gauge transformation modifies the multiplicative function by the action of an element of $\mathcal{F}(M, G)$. Hence, there is a second mapping:

$$\mathcal{T}(P)/\mathcal{J}(P) \longrightarrow \mathcal{M}(PM, G)/\mathcal{F}(M, G). \quad (8)$$

Combining (7) with (8), we get a mapping which we denote by \mathcal{H}_P :

$$\mathcal{H}_P : H^1\mathcal{A}(P)/H^2\mathcal{J}(P) \longrightarrow \mathcal{M}(PM, G)/\mathcal{F}(M, G).$$

Recall that a function $f \in \mathcal{M}(PM, G)$ is said to be continuous with fixed endpoints if it is sequentially continuous along sequences of paths converging with fixed endpoints. Observe that this property is not affected by the action of $\mathcal{F}(M, G)$. Our main result in this section is the following.

Proposition 2.16 1. *Let P be a principal G -bundle over M . Then the mapping \mathcal{H}_P is a one-to-one mapping whose range contains only functions continuous with fixed endpoints.*

⁷These are actions on the right but we denote them on the left. This will not cause any ambiguity.

2. Let Q be another G -bundle over M . If P and Q are isomorphic, and if $\varphi : P \rightarrow Q$ is a bundle isomorphism, then φ induces an energy-preserving map $\varphi^* : H^1\mathcal{A}(Q)/H^2\mathcal{J}(Q) \rightarrow H^1\mathcal{A}(P)/H^2\mathcal{J}(P)$ which satisfies the relation $\mathcal{H}_P \circ \varphi^* = \mathcal{H}_Q$.

If P and Q are not isomorphic, then the ranges of \mathcal{H}_P and \mathcal{H}_Q are disjoint.

Remark 2.17 It could be argued that it is very unpleasant to choose a reference point in each fibre of P because, unless P is trivial, this cannot be done in a continuous way. It is possible to avoid this by choosing a point m on M and restricting oneself to the space L_mM of loops based at m . The action of $\mathcal{F}(M, G)$ on $\mathcal{M}(L_mM, G)$ reduces to the diagonal action of G by conjugation. Then, by choosing only a reference point in P_m , one is able to construct a mapping $\tilde{\mathcal{H}}_P : H^1\mathcal{A}(P)/H^2\mathcal{J}(P) \rightarrow \mathcal{M}(L_mM, G)/G$. This point of view is exactly equivalent to ours since the spaces $\mathcal{M}(PM, G)/\mathcal{F}(M, G)$ and $\mathcal{M}(L_mM, G)/G$ equipped with the traces of the product topologies are canonically homeomorphic. Moreover, the canonical homeomorphism sends functions continuous with fixed endpoints to continuous functions. Nevertheless, we choose to consider PM instead of L_mM because it is easier to work with the Yang-Mills measure if it is defined on $\mathcal{M}(PM, G)$.

Proof of Proposition 2.16 – 1. Let ω be an H^1 connection. For each $m \in M$, let $p(m)$ be a reference point in P_m . Let f be the element of $\mathcal{M}(PM, G)$ determined by ω and the set-theoretic section $m \mapsto p(m)$. Let us prove that f is continuous with fixed endpoints.

Let $(c_n)_{n \geq 0}$ be a sequence of paths converging with fixed endpoints to c . Let us assume that c is contained in the domain of a local section s of P . Then, for n large enough, c_n is also contained in the domain of this local section. Let us also assume that $s(\underline{c}) = p(\underline{c})$ and $s(\bar{c}) = p(\bar{c})$. Since G is connected, this causes no loss of generality. Now, for each n , we have $f(c_n) = \langle \omega_s, c_n \rangle$ and the similar equality for c . By Proposition 2.6, this implies that $f(c_n)$ tends to $f(c)$ as n tends to infinity.

If c is not contained in the domain of a local section, let us decompose it as $c = c^1 \dots c^r$ in such a way that, for each $k = 1, \dots, r$, c^k is contained in the domain U_k of some local section s_k . We assume that, for each k , the section s_k coincides with p at the endpoints of c^k . Let $R > 0$ be such that every geodesic ball on M of radius smaller than R is geodesically convex and such that, for each k , the R -neighbourhood of c^k , denoted by $c_R^k = \{m \in M \mid d_M(m, c^k) < R\}$ is contained in U_k . Then, for n large enough, c_n is contained in $c_R^1 \cup \dots \cup c_R^r$. For such an n , decompose c_n as $c_n = c_n^1 \dots c_n^r$ in such a way that for each k , $d_M(\underline{c}_n^k, \underline{c}^k) < R$ and $d_M(\bar{c}_n^k, \bar{c}^k) < R$. Then, for each k , c_n^k converges to c^k , but not with fixed endpoints.

For each n large enough and each $k = 1, \dots, r - 1$, let $\zeta_{k,n}$ denote the geodesic segment joining \bar{c}^k to \bar{c}_n^k . Let $\zeta_{0,n}$ be the point \underline{c} and $\zeta_{r,n}$ be the point \bar{c} . Then, for each n and each $k = 1, \dots, r$, set $\tilde{c}_n^k = \zeta_{k-1,n} c_n^k \zeta_{k,n}^{-1}$. For each k , (\tilde{c}_n^k) tends to c^k with fixed endpoints and, by multiplicativity, $f(c_n) = f(c_n^1 \dots c_n^r) = f(\tilde{c}_n^1 \dots \tilde{c}_n^r) = f(\tilde{c}_n^r) \dots f(\tilde{c}_n^1)$. We have now reduced the problem to the case of paths lying in the domain of a local section and find that $f(c_n)$ tends to $f(c^r) \dots f(c^1) = f(c)$.

We prove now that \mathcal{H}_P is injective. Let us consider two H^1 connections ω and ω' on P which are sent to the same class of $\mathcal{M}(PM, G)$ by the composed mapping $H^1\mathcal{A}(P) \rightarrow H^1\mathcal{A}(P)/H^2\mathcal{J}(P) \xrightarrow{\mathcal{H}_P} \mathcal{M}(PM, G)/\mathcal{F}(M, G)$. We claim that ω and ω' differ by the action of an element of $H^2\mathcal{J}(P)$.

Indeed, let $T = (\tau_c, c \in PM)$ and $T' = (\tau'_c, c \in PM)$ be the holonomies induced by ω and ω' respectively. There exists a gauge transformation $j = (\gamma_m, m \in M)$ such that $T' = j \cdot T$. Let

(U, s) be a local section of P . Then the equality $\langle \omega'_s, \cdot \rangle = j_s \cdot \langle \omega_s, \cdot \rangle$ of multiplicative function on PU holds. By Proposition 2.9, j_s belongs to $H^2(U; G)$. Since this argument is valid for each local section of P , we conclude that j belongs to $H^2\mathcal{J}(P)$. Now, the connections $j \cdot \omega$ and ω' determine the same holonomy on P . According to Proposition 2.12, this implies that they are equal.

2. If P and Q are isomorphic, the statement is straightforward. Let us assume that P and Q are not isomorphic. Then M is necessarily closed. Let f be an element of $\mathcal{M}(PM, G)$ whose class modulo $\mathcal{F}(M, G)$ belongs to the range of \mathcal{H}_P . We claim that $\mathfrak{o}(P)$ can be computed from f . Since $\mathfrak{o}(P) \neq \mathfrak{o}(Q)$, this implies that f does not belong to the range of \mathcal{H}_Q .

Let m be a point of M . Let g denote the genus of M . Let $a_1, b_1, \dots, a_g, b_g$ be $2g$ loops based at m which generate the fundamental group $\pi_1(M, m)$ with the single relation $[b_g^{-1}, a_g^{-1}] \dots [b_1^{-1}, a_1^{-1}] = 1$. Let $L : [0, 1]^2 \rightarrow M$ be a smooth homotopy such that, for each $s \in [0, 1]$, $L(s, \cdot)$ is a smooth loop based at m , $L(0, \cdot)$ is the constant loop at m and $L(1, \cdot)$ is the loop $[b_g^{-1}, a_g^{-1}] \dots [b_1^{-1}, a_1^{-1}]$. The mapping from $[0, 1]$ to G defined by $s \mapsto \tilde{f}(L(s, \cdot))$ is a continuous path starting from 1. Let $\tilde{s} \mapsto \tilde{f}(\tilde{s})$ be the lift starting at 1 in \tilde{G} of this curve. Recall that, if $x, y \in G$ and $\tilde{x}, \tilde{y} \in \tilde{G}$, then $[\tilde{x}, \tilde{y}] = \tilde{x}\tilde{y}\tilde{x}^{-1}\tilde{y}^{-1}$ depends only on x and y . We denote it by $[\widetilde{x, y}]$. Then $\mathfrak{o}(P) = \tilde{f}(1) \left([\widetilde{f(a_1), f(b_1)}] \dots [\widetilde{f(a_g), f(b_g)}] \right)^{-1}$ ([15], Theorem 3.9). \square

To finish this section on Sobolev connections, let us state Uhlenbeck's compactness theorem in the particular case that we are going to use. The original reference for this theorem is [18]. For the case where M has a boundary, and also for a more comprehensive and detailed proof, we refer the reader to [22].

Theorem 2.18 (Compactness theorem) *Let P be a principal G -bundle over M . Let $(\omega_n)_{n \geq 1}$ be a sequence of connections in $H^1\mathcal{A}(P)$ such that $S(\omega_n)$ is uniformly bounded. Then there exists a subsequence $(\omega_{n_k})_{k \geq 1}$, a sequence $(j_k)_{k \geq 1}$ in $H^2\mathcal{J}(P)$ and an element ω of $H^1\mathcal{A}(P)$ such that*

1. $j_k \cdot \omega_{n_k} \rightarrow \omega$ in $H^1\mathcal{A}(P)$,
2. $S(\omega) \leq \liminf_{k \rightarrow \infty} S(\omega_{n_k})$.

3 Large deviations for the Yang-Mills measures

In this section, we prove Theorems 1.15 and 1.16, except for the construction of a connection with minimal energy with prescribed holonomy along the edges of a graph, which is the object of Section 4. We are going to follow a route close to the one followed to construct the measure in [9]. The starting point is the classical short-time estimate of the heat kernel on a compact Riemannian manifold which we apply to G , and, with a minor modification, to its possibly non-compact universal covering \tilde{G} . A large deviation principle for the discrete Yang-Mills measures follows by elementary arguments. An application of the contraction principle produces a large deviation principle for the finite-dimensional distributions associated to families of paths which are contained in some \mathbb{E}^* , where \mathbb{E} is the set of edges of a graph. Just as in the construction of the measure, it is not enough to take the projective limit of these discrete principles: we must first obtain an large deviation principle for all finite dimensional marginals of the holonomy process. For this, we use a standard result on exponential approximations of measures. In identifying the rate function at this stage in terms of the Yang-Mills measure, we make repeated

uses of Uhlenbeck's compactness theorem. Finally, Dawson-Gärtner's theorem yields the large deviation principle for the whole process.

3.1 The discrete Yang-Mills measures

Let N be a Riemannian manifold. Let us denote for all $t > 0$ by $p_t(\cdot, \cdot)$ the heat kernel on N , that is, the kernel of the operator $\exp \frac{t\Delta}{2}$ on the space of square-integrable functions on N . The fundamental estimate is the following. The prototype of this result was proved by Varadhan in [19, 20]. For the form given here the reader may consult [14] or [11].

Theorem 3.1 *Let $p_t(\cdot, \cdot)$ be the heat kernel on a compact Riemannian manifold N . Then, uniformly for all $x, y \in N$, one has*

$$\lim_{t \rightarrow 0} -2t \log p_t(x, y) = d(x, y)^2.$$

When M has a boundary, we deduce directly from this theorem the large deviation principle for the discrete Yang-Mills measure associated to a graph on M with boundary conditions.

Let N_1, \dots, N_p be the connected components of ∂M . Let X_1, \dots, X_p be p conjugacy classes in G . Recall that, if $x \in G$, then $\rho(x)$ denotes the Riemannian distance between 1 and x .

Proposition 3.2 *Let \mathbb{G} be a graph on M . The family of measures $(P_{T;X_1, \dots, X_p}^{\mathbb{G}})_{T>0}$ on $G^{\mathbb{E}^+}$ satisfies a large deviation principle with good rate function*

$$I_{X_1, \dots, X_p}^{\mathbb{E}}(g) = \begin{cases} \sum_{F \in \mathbb{F}} \frac{\rho(h_{\partial F}(g))^2}{2\sigma(F)} & \text{if } h_{N_1}(g) = X_1, \dots, h_{N_p}(g) = X_p, \\ +\infty & \text{otherwise.} \end{cases}$$

Proof – Let S denote the subset of all $g \in G^{\mathbb{E}^+}$ such that the boundary conditions $h_{N_1}(g) = X_1, \dots, h_{N_p}(g) = X_p$ are satisfied. It is a closed subset of $G^{\mathbb{E}^+}$. With the notation of Definition 1.6 and by Theorem 3.1, we have, as T tends to 0, for all $g \in S$,

$$dP_{T;X_1, \dots, X_p}^{\mathbb{G}}(g) = \frac{1}{Z_{T;X_1, \dots, X_p}^{\mathbb{G}}} e^{-\frac{1}{T}(I_{X_1, \dots, X_p}^{\mathbb{E}}(g) + o(1))} d\nu_{X_1}^{N_1} \dots d\nu_{X_p}^{N_p} dg_{int},$$

where $o(1)$ is uniform on S .

According to Proposition 1.7 and to a standard estimation of the supremum of the heat kernel (for instance Theorem V.4.3 of [21], to which we will refer again later), $Z_{T;X_1, \dots, X_p}^{\mathbb{G}} \leq \|p_{T\sigma(M)}\|_{\infty} = O(T^{-\frac{\dim G}{2}})$ is negligible at the exponential scale. The large deviation principle on the subset S follows now from the fact that the measure $d\nu_{X_1}^{N_1} \dots d\nu_{X_p}^{N_p} dg_{int}$ charges every open subset of S . Finally, since S is closed and supports the measures $P_{T;X_1, \dots, X_p}^{\mathbb{G}}$, the large deviation principle holds on $G^{\mathbb{E}^+}$. \square

When M is closed, the Yang-Mills measures on M are defined in terms of the heat kernel on \tilde{G} which may not be compact. However, this possible non-compactness is easy to deal with, since it comes from the presence of a Euclidean direct factor \mathbb{R}^m .

Proposition 3.3 *Uniformly for all $\tilde{g}, \tilde{h} \in \tilde{G}$, one has*

$$\lim_{t \rightarrow 0} -2t \log \tilde{p}_t(\tilde{g}, \tilde{h}) = d_{\tilde{G}}(\tilde{g}, \tilde{h})^2.$$

Proof – Since G is a compact group, its Lie algebra \mathfrak{g} which is also that of \tilde{G} can be written as $\mathfrak{g} = [\mathfrak{g}, \mathfrak{g}] \oplus \mathfrak{z}(\mathfrak{g})$, where $\mathfrak{z}(\mathfrak{g})$ is the center of \mathfrak{g} ([4], Theorem 3.6.2). Accordingly, $\tilde{G} = K \times \mathbb{R}^m$, where K is the subgroup of \tilde{G} with Lie algebra $[\mathfrak{g}, \mathfrak{g}]$ and $m = \dim \mathfrak{z}(\mathfrak{g})$. The group K is compact and simply connected. Let Δ , Δ_K and $\Delta_{\mathbb{R}^m}$ denote the Laplace operators on \tilde{G} , K and \mathbb{R}^m respectively, where K and \mathbb{R}^m are endowed with the induced metric. Observe that Δ_K and $\Delta_{\mathbb{R}^m}$ commute and, independently, that the induced metric on \mathbb{R}^m is a constant Euclidean metric. The scalar product on \mathfrak{g} corresponding to the Riemannian metric on \tilde{G} is invariant under the adjoint action of \tilde{G} on \mathfrak{g} , so that the adjoint action of \mathfrak{g} on itself is antisymmetric and $\mathfrak{z}(\mathfrak{g}) \perp [\mathfrak{g}, \mathfrak{g}]$. This implies the relations $\Delta = \Delta_K + \Delta_{\mathbb{R}^m}$ and $\exp \frac{t\Delta}{2} = \exp \frac{t\Delta_K}{2} \exp \frac{t\Delta_{\mathbb{R}^m}}{2}$. Finally, for all $t > 0$, $k, l \in K$ and $x, y \in \mathbb{R}^m$, and with an obvious notation,

$$\tilde{p}_t((k, x), (l, y)) = p_t^K(k, l) p_t^{\mathbb{R}^m}(x, y).$$

On one hand, as t tends to 0, $-2t \log p_t^K(k, l)$ tends, by Theorem 3.1, to $d_K(k, l)^2$ uniformly. On the other hand, $-2t \log p_t^{\mathbb{R}^m}(x, y) = d_{\mathbb{R}^m}(x, y)^2 - mt \log(2\pi t)$. The result follows now from the identity $d_{\tilde{G}}((k, x), (l, y))^2 = d_K(k, l)^2 + d_{\mathbb{R}^m}(x, y)^2$. \square

Choose $z \in \Pi$. Let \mathbb{G} be a graph on M . Recall the definition of the discrete measures $P_{T,z}^{\mathbb{G}}$ given in Definition 1.5. Let us also introduce, for $\tilde{x} \in \tilde{G}$, the notation $\tilde{\rho}(\tilde{x})$ for the Riemannian distance in \tilde{G} between \tilde{x} and the unit element.

Proposition 3.4 *The family of probability measures $(P_{T,z}^{\mathbb{G}})_{T>0}$ on $G^{\mathbb{E}^+}$ satisfies a large deviation principle with rate function*

$$I_z^{\mathbb{E}}(g) = \min_{z_{\mathbb{F}} \in \Pi_z^{\mathbb{F}}} \sum_{F \in \mathbb{F}} \frac{\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g}) z_F)^2}{2\sigma(F)}, \quad (9)$$

where $\tilde{g} \in \tilde{G}^{\mathbb{E}^+}$ satisfies $\pi(\tilde{g}) = g$.

Remark 3.5 1. *The value of the rate function does not depend on the choice of \tilde{g} for the same reason as the number defined as $D_{T,z}^{\mathbb{G}}(g)$ in Proposition 1.4 does not. Less obvious is the fact, which is part of the last proposition, that the minimum in (9) is attained for some $z_{\mathbb{F}} \in \Pi_z^{\mathbb{F}}$.*
2. *When G is simply connected, the rate function takes the simpler form*

$$I^{\mathbb{E}}(g) = \sum_{F \in \mathbb{F}} \frac{\rho(h_{\partial F}(g))^2}{2\sigma(F)}.$$

In this case, Proposition 3.4 is a direct consequence of Theorem 3.1.

Proof – What makes this proof a bit more difficult than in the case with boundary is the possible presence of an infinite sum in the density $D_{T,z}^{\mathbb{G}}$. We need to truncate this sum and estimate the error we make.

Let us choose a bounded measurable section $G^{\mathbb{E}^+} \rightarrow \tilde{G}^{\mathbb{E}^+}$ of π . Let us simply denote by \tilde{g} the image by this section of $g \in G^{\mathbb{E}^+}$. Set $c = \sup\{\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})) : g \in G^{\mathbb{E}^+}, F \in \mathbb{F}\}$. Let also s and S be two real numbers such that $0 < s < S$ and, for each face F , $s < \sigma(F) < S$. Finally, let $C > 0$ be such that, for all $t > 0$, all $\tilde{x}, \tilde{y} \in \tilde{G}$, $\tilde{\rho}_t(\tilde{x}, \tilde{y}) \leq Ct^{-\frac{\dim G}{2}} \exp -\frac{d_{\tilde{G}}(\tilde{x}, \tilde{y})^2}{Ct}$. Such a constant exists by [21], Theorem V.4.3.

For each integer $k \geq 0$, set $\Lambda_k = \{z_{\mathbb{F}} \in \Pi_z^{\mathbb{F}} | \forall F \in \mathbb{F}, \tilde{\rho}(z_F) < k\}$. Fix $g \in G^{\mathbb{E}^+}$. Then

$$D_{T,z}^G(g) = \sum_{k=0}^{\infty} \left\{ \sum_{z_{\mathbb{F}} \in \Lambda_{k+1} \setminus \Lambda_k} \prod_{F \in \mathbb{F}} \tilde{p}_{T\sigma(F)}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F) \right\}. \quad (10)$$

If $z_{\mathbb{F}}$ belongs to $\Lambda_{k+1} \setminus \Lambda_k$, then there exists a face F such that $\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F) \geq |k - c|$, so

$$\prod_{F \in \mathbb{F}} \tilde{p}_{T\sigma(F)}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F) \leq \left(\frac{C}{(sT)^{\frac{d}{2}}} \right)^{|\mathbb{F}|} e^{-\frac{(k-c)^2}{CST}}.$$

Hence, if L is a non-negative integer, the tail of (10) satisfies

$$\sum_{k=L}^{\infty} \{\dots\} \leq \frac{C}{T^{\frac{d|\mathbb{F}|}{2}}} \sum_{k \geq L} |\Lambda_{k+1}| e^{-\frac{(k-c)^2}{CT}},$$

where C denotes now a constant which varies from line to line. The set $\Pi \subset \tilde{G} \simeq K \times \mathbb{R}^m$ is a sub-lattice of $Z(K) \times R$, where $Z(K)$ is the center of K , which is finite, and R is a discrete additive subgroup of \mathbb{R}^m . Hence, the cardinality $|\Lambda_k|$ is dominated by a power of k . Thus there exists a rational function Q of two variables such that the tail of (10) satisfies

$$\sum_{k=L}^{\infty} \{\dots\} \leq Q(L, \sqrt{T}) e^{-\frac{(L-c)^2}{CT}}.$$

Hence, for each L , the density can be put in the form

$$D_{T,z}^G(g) = \sum_{z_{\mathbb{F}} \in \Lambda_L} \exp \left[-\frac{1}{T} \sum_{F \in \mathbb{F}} \frac{\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F)^2}{2\sigma(F)} + o\left(\frac{1}{T}\right) \right] + \varepsilon(L, T)(g),$$

with $0 \leq \varepsilon(L, T)(g) \leq Q(L, \sqrt{T}) e^{-\frac{(L-c)^2}{CT}}$.

Now, with $g \in G^{\mathbb{E}^+}$ still fixed, the function from $\Pi_z^{\mathbb{F}}$ to \mathbb{R}^+ which sends $z_{\mathbb{F}}$ to $\sum_{F \in \mathbb{F}} \frac{\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F)^2}{2\sigma(F)}$ tends to infinity as $z_{\mathbb{F}}$ tends to infinity. Thus, this function attains its infimum, on a subset $M(g)$ of $\Pi_z^{\mathbb{F}}$ which may not be a singleton. Since $g \mapsto \tilde{g}$ is a bounded mapping, the convergence of the sum above is uniform in g , so that $\cup_{g \in G^{\mathbb{E}^+}} M(g)$ is a bounded set and there exists a positive integer L_0 such that $\cup_{g \in G^{\mathbb{E}^+}} M(g) \subset \Lambda_{L_0}$.

The large deviation principle can now be proved easily. As in the case where M has a boundary, Proposition 1.7 and a classical estimate on the heat kernel imply that $Z_{T,z}^G$ is negligible at the exponential scale. Let $A \subset G^{\mathbb{E}^+}$ be a Borel subset. Then, from the discussion above we deduce that, for $L \geq L_0$,

$$\overline{\lim}_{T \rightarrow 0} T \log P_{T,z}^G(A) \leq \max \left[\overline{\lim}_{T \rightarrow 0} T \log \int_A \varepsilon(L, T)(g) dg, - \inf_{g \in A} \min_{z_{\mathbb{F}} \in \Pi_z^{\mathbb{F}}} \sum_{F \in \mathbb{F}} \frac{\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F)^2}{2\sigma(F)} \right].$$

Since $\overline{\lim}_{T \rightarrow 0} T \log \int_A \varepsilon(L, T)(g) dg \leq -\frac{(L-c)^2}{C}$ tends to $-\infty$ as L tends to infinity, the upper bound of the large deviation principle is proved by taking L large enough. A similar argument for the lower bound finishes the proof. \square

We want to give an expression of the rate functions $I_{X_1, \dots, X_p}^{\mathbb{E}}$ and $I_z^{\mathbb{E}}$ in terms of the Yang-Mills energy. For this, we need to establish a link between H^1 connections and elements of $G^{\mathbb{E}^+}$.

Let J be a subset of PM , for example, the set of edges of a graph, or a set of loops. Any H^1 connection on some G -bundle P over M determines, via the mapping \mathcal{H}_P , an element of $\mathcal{M}(J, G)/\mathcal{F}(M, G)$. Now for every subset K of PM , let K^* denote the set of paths that can be constructed by concatenating elements of K . There is a natural one-to-one correspondence between $\mathcal{M}(K, G)$ and $\mathcal{M}(K^*, G)$. Hence, if K is a subset of PM such that $J \subset K^*$, then any function of $\mathcal{M}(K, G)$ determines a function of $\mathcal{M}(J, G)$ and thus an element of $\mathcal{M}(J, G)/\mathcal{F}(M, G)$. The main example of this situation is the following: K is the set of edges of a graph and J is a set of paths in this graph.

Definition 3.6 *Let J and K be two subsets of PM such that $J \subset K^*$. Two connections of $H^1\mathcal{A}$, or two functions of $\mathcal{M}(K, G)$, or one such connection and one such function are said to agree up to gauge transformation on J , or simply to agree on J , if they induce the same class of $\mathcal{M}(J, G)/\mathcal{F}(M, G)$. We denote this relation by \sim_J .*

From now on, we will alternatively use two sets of assumption, corresponding to the cases with and without boundary. Let us state them once for all and give them a name.

Convention 3.7 1. *(Boundary) means: Assume that M has a non-empty boundary. Assume that N_1, \dots, N_p are the connected components of ∂M . Let X_1, \dots, X_p be p conjugacy classes of G . Let P be a principal G -bundle over M .*

2. *(Closed) means: Assume that M is closed. Let z be an element of Π . Let P be a principal G -bundle over M such that $\mathfrak{o}(P) = z$.*

Let us state the main technical result of this paper.

Proposition 3.8 *Let \mathbb{G} be a graph on M . Let g be an element of $G^{\mathbb{E}^+}$.*

1. *(Boundary) The following equality holds:*

$$I_{X_1, \dots, X_p}^{\mathbb{E}}(g) = \frac{1}{2} \inf\{S(\omega) : \omega \in H^1\mathcal{A}_{X_1, \dots, X_p}(P), \omega \sim_{\mathbb{E}} g\}. \quad (11)$$

2. *(Closed) The following equality holds:*

$$I_z^{\mathbb{E}}(g) = \frac{1}{2} \inf\{S(\omega) : \omega \in H^1\mathcal{A}(P), \omega \sim_{\mathbb{E}} g\}. \quad (12)$$

In both cases, the infima are attained by an element of $W^{1, \infty}\mathcal{A}(P)$, hence continuous and Lipschitz on M , which is smooth outside $\bigcup_{e \in \mathbb{E}} e$.

In (11) and (12), the fact that the left hand side is smaller than the right hand side is a simple consequence of the energy inequality (Proposition 2.5). Besides, the fact that the infima are attained by H^1 connections is a consequence of Uhlenbeck's theorem (Theorem 2.18). Indeed,

from a minimizing sequence one can extract a weakly convergent one and closed constraints on the holonomy are stable by weak H^1 limits. Proving that the minimum is equal to the left hand side is the difficult part. We do this by constructing an explicit minimizer. This is rather long and we postpone the construction until Section 4. Let us state the result here.

Proposition 3.9 *Let \mathbb{G} be a graph on M . Let g be an element of $G^{\mathbb{E}^+}$.*

1. (Boundary) *Assume that $h_{N_1}(g) = X_1, \dots, h_{N_p}(g) = X_p$. Then there exists a connection $\omega \in W^{1,\infty}\mathcal{A}_{X_1, \dots, X_p}(P)$ which is smooth outside $\bigcup_{e \in \mathbb{E}} e$ such that $\omega \sim_{\mathbb{E}} g$ and $S(\omega) = 2I_{X_1, \dots, X_p}^{\mathbb{E}}(g)$.*

2. (Closed) *There exists a connection $\omega \in W^{1,\infty}\mathcal{A}(P)$ which is smooth outside $\bigcup_{e \in \mathbb{E}} e$ such that $\omega \sim_{\mathbb{E}} g$ and $S(\omega) = 2I_z^{\mathbb{E}}(g)$.*

Let us give briefly an idea of what a minimizing connection looks like. The key to the construction is that minimizers of the energy on disks with prescribed holonomy along the boundary are well-known. A connection on a face F with holonomy x along the boundary and minimal energy is gauge-equivalent to a connection of the form $X\lambda$, where X is an element of \mathfrak{g} of minimal norm such that $\exp(\sigma(F)X) = x$ and λ is a smooth 1-form such that $d\lambda = \sigma$. We construct a minimizing connection on M essentially by taking one such minimizer on each face and gluing them all together.

Proof of Proposition 3.8 – 1. Let ω be an H^1 connection which satisfies the boundary conditions and agrees with g on \mathbb{E} . Then, the energy inequality (Proposition 2.5) applied on each face of \mathbb{G} implies that $S(\omega) \geq 2I_{X_1, \dots, X_p}^{\mathbb{E}}(g)$. The reverse inequality follows from Proposition 3.9.

2. Let ω be an H^1 connection on P which agrees with g on \mathbb{E} . It is not enough to apply the energy inequality to ω in this case. Instead, let us choose for each face F a smooth section s_F of P over F . Let us choose a face F . The form ω_F belongs to $H^1\Omega_{\mathfrak{g}}^1(F)$. Let us compute $\langle \omega_F, \partial F \rangle_{\tilde{G}}$, which is a conjugacy class in \tilde{G} and which projects on the conjugacy class $\langle \omega, \partial F \rangle$ of G . The energy inequality applied on each face with the structure group \tilde{G} gives us

$$S(\omega) \geq \sum_{F \in \mathbb{F}} \frac{\tilde{\rho}(\langle \omega_F, \partial F \rangle_{\tilde{G}})^2}{\sigma(F)}.$$

We claim that the right hand side of this inequality is of the form $\sum_{F \in \mathbb{F}} \frac{\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F)^2}{\sigma(F)}$ for some $\tilde{g} \in \tilde{G}^{\mathbb{E}^+}$ and some $z_{\mathbb{F}} \in \Pi_z^{\mathbb{F}}$.

Recall that, if e is an edge, $L(e)$ is the face located on the left of e . For each $e \in \mathbb{E}^+$, set $\tilde{g}_e = \langle \omega_{L(e)}, e \rangle_{\tilde{G}}$. Then $\tilde{g} = (\tilde{g}_e, e \in \mathbb{E}^+)$ belongs to $\tilde{G}^{\mathbb{E}^+}$ and satisfies $\pi(\tilde{g}) = g$. Finally, for each face F , set

$$z_F = \prod_{e \in \mathbb{E}^+ : L(e^{-1}) = F} \langle \omega_{L(e)}, e \rangle_{\tilde{G}} \langle \omega_{L(e^{-1})}, e^{-1} \rangle_{\tilde{G}}.$$

Then it follows from [9], Lemma 1.7, that $z_{\mathbb{F}} = (z_F, F \in \mathbb{F})$ belongs to $\Pi_z^{\mathbb{F}}$. On the other hand, $z_{\mathbb{F}}$ is defined in such a way that, for each face F , $h_{\partial F}^{\tilde{G}}(\tilde{g})z_F = \langle \omega_F, \partial F \rangle_{\tilde{G}}$. Our claim is thus proved and it follows that $S(\omega) \geq 2I_z^{\mathbb{G}}(g)$. The reverse inequality follows as in the case with boundary from Proposition 3.9. \square

3.2 Holonomy along a family of paths in a graph (Contraction principle)

The contraction principle [2, Thm. 4.2.1] allows us to state a large deviation principle for the law of the random holonomy along a finite set of paths in a graph.

Proposition 3.10 *Let \mathbb{G} be a graph on M . Let $J = \{c_1, \dots, c_n\}$ be a finite subset of \mathbb{E}^* .*

1. (Boundary) *The laws of $(H_{c_1}, \dots, H_{c_n})$ under $P_{T;X_1, \dots, X_p}^{\mathbb{G}}$ satisfy as T tends to 0 a large deviation principle on $\mathcal{M}(J, G)$ with rate function*

$$I_{X_1, \dots, X_p}^J(g) = \frac{1}{2} \inf\{S(\omega) : \omega \in H^1 \mathcal{A}_{X_1, \dots, X_p}(P), \omega \sim_J g\}.$$

2. (Closed) *The laws of $(H_{c_1}, \dots, H_{c_n})$ under $P_{T,z}^{\mathbb{G}}$ satisfy as T tends to 0 a large deviation principle on $\mathcal{M}(J, G)$ with rate function*

$$I_z^J(g) = \frac{1}{2} \inf\{S(\omega) : \omega \in H^1 \mathcal{A}(P), \omega \sim_J g\}.$$

Proof – The proof is exactly the same whether or not M has a boundary. We write the proof when M is closed. Changing the names of the probabilities and rate functions gives the proof in the case with boundary.

2. The mapping $h_J = (h_{c_1}, \dots, h_{c_n}) : G^{\mathbb{E}^+} \longrightarrow G^J \simeq G^n$ is continuous. Hence, by contraction of the large deviation principle on $G^{\mathbb{E}^+}$, the laws of $(H_{c_1}, \dots, H_{c_n})$ under $P_{T,z}^{\mathbb{G}}$ satisfy a large deviation principle on G^J with rate function

$$\tilde{I}_z^J(g) = \inf\{I_z^{\mathbb{E}}(k) : k \in G^{\mathbb{E}^+}, h_J(k) = g\}.$$

We claim that this large deviation principle holds on the smaller space $\mathcal{M}(J, G) \subset G^J$. Indeed, $h_J(k) = g$ implies $g \in \mathcal{M}(J, G)$. Hence, the support of \tilde{I}_z^J is contained in the closed subset $\mathcal{M}(J, G)$ of G^J , which supports the laws of $(H_{c_1}, \dots, H_{c_n})$ under $P_{T,z}^{\mathbb{G}}$. The claim follows by [2], Lemma 4.1.5.

Now, on one hand, $h_J(k) = g$ implies $k \sim_J g$. On the other hand, $k \sim_J g$ implies that there exists $j \in \mathcal{F}(M, G)$ such that $h_J(j \cdot k) = g$. Since $I_z^{\mathbb{E}}$ is gauge-invariant, we get the following expression for \tilde{I}_z^J :

$$\begin{aligned} \tilde{I}_z^J(g) &= \inf\{I_z^{\mathbb{E}}(k) : k \in G^{\mathbb{E}^+}, k \sim_J g\} \\ &= \frac{1}{2} \inf\{S(\omega) : \omega \in H^1 \mathcal{A}(P), \exists k \in G^{\mathbb{E}^+}, \omega \sim_{\mathbb{E}} k, k \sim_J g\}. \end{aligned}$$

It appears that $\tilde{I}_z^J(g)$ is the infimum of S over a smaller set of connections than $I_z^J(g)$, so that $\tilde{I}_z^J \geq I_z^J$.

Now, take $g \in G^J$ for which $I_z^J(g) < \infty$, that is, such that there exists an H^1 connection on P which agrees with g on J . For every $\alpha > 0$, there exists a connection ω_α which agrees with g on J and such that $\frac{1}{2}S(\omega_\alpha) \leq I_z^J(g) + \alpha$. This connection induces a certain class of $G^{\mathbb{E}^+} / \mathcal{F}(M, G)$. Let k_α be an element of this class. Then ω_α agrees tautologically with k_α on \mathbb{E} and $k_\alpha \sim_J g$. Hence, $\tilde{I}_z^J(g) \leq \frac{1}{2}S(\omega_\alpha) \leq I_z^J(g) + \alpha$. By letting α tend to 0, we get $\tilde{I}_z^J \leq I_z^J$. \square

For the last part of the proof, we could also have argued that the infimum defining the function I_z^J is attained, as a consequence of Uhlenbeck's compactness theorem. Then, a minimizer agrees on \mathbb{E} with a certain configuration in $G^{\mathbb{E}^+}$ and the inequality $\tilde{I}_z^J \leq I_z^J$ follows.

3.3 Holonomy along arbitrary finite families of paths (Exponential approximation)

The last step at the finite-dimensional level is to prove that Proposition 3.10 holds for an arbitrary finite subset J of PM . To do this, we use an exponential approximation result.

Fix $J = \{p_1, \dots, p_n\}$ an arbitrary finite subset of PM .

Lemma 3.11 *There exist n sequences $(c_i^m)_{m \geq 1}$, $i = 1 \dots n$, of paths such that the following properties hold.*

0. For each $i = 1 \dots n$, $(c_i^m)_{m \geq 1}$ converges to p_i with fixed endpoints.
1. For all $m \geq 1$, there exists a graph \mathbb{G}_m such that c_1^m, \dots, c_n^m belong to \mathbb{E}_m^* .
2. For all $\delta > 0$,

$$\lim_{m \rightarrow \infty} \limsup_{T \rightarrow 0} T \log P_T \left[\max_{1 \leq i \leq n} d(H_{c_i^m}, H_{p_i}) > \delta \right] = -\infty,$$

where P_T stands either for $P_{T; X_1, \dots, X_p}$ or for $P_{T, z}$.

3. For every H^1 connection ω , in $H^1 \mathcal{A}_{X_1, \dots, X_p}(P)$ if M has a boundary, or in $H^1 \mathcal{A}(P)$ if M is closed,

$$\rho(\langle \omega, c_i^m p_i^{-1} \rangle) \leq C [S(\omega) d_\ell(c_i^m, p_i)]^{\frac{1}{2}}$$

for some constant C independent of ω .

Proof – Let us begin with the case where p_1, \dots, p_n are edges. It is proved in [9, Section 2.5.3] that we can find sequences of piecewise geodesic paths $(c_i^m)_{m \geq 1}$, $i = 1 \dots n$, converging to p_1, \dots, p_n with fixed endpoints, such that $c_i^m p_i^{-1}$ bounds for each i a domain diffeomorphic to a disk, of arbitrarily small area. Let σ_i^m denote this area. We assume that $\sigma_i^m < \frac{1}{2} \sigma(M)$ for each i and each m . If M has a boundary, choose $x_1 \in X_1, \dots, x_p \in X_p$. Then, for each bounded non-negative measurable function f , $E_{T; X_1, \dots, X_p} [f(d(H_{p_i}, H_{c_i^m}))]$ is equal to

$$\begin{aligned} & \frac{1}{Z_{T; X_1, \dots, X_p}} \int_{G^{2g+p+1}} f(x) p_{T\sigma_i^m}(x) p_{T(\sigma(M)-\sigma_i^m)}(x^{-1} \prod_{k=1}^g [a_k, b_k] \prod_{l=1}^p y_l^{-1} x_l y_l) \\ & \hspace{15em} da_1 db_1 \dots da_g db_g dy_1 \dots dy_p dx \\ & \leq \frac{1}{Z_{T; X_1, \dots, X_p}} \|p_{\frac{T}{2}\sigma(M)}\|_\infty \int_G f(x) p_{T\sigma_i^m}(x) dx. \end{aligned}$$

If M is closed, then $E_{T, z} [f(d(H_{p_i}, H_{c_i^m}))]$ is equal to

$$\begin{aligned} & \frac{1}{Z_{T, z}} \int_{G^{2g+1}} f(x) \sum_{w \in \Pi} \tilde{p}_{T\sigma_i^m}(\tilde{x}w) \tilde{p}_{T(\sigma(M)-\sigma_i^m)}(\tilde{x}^{-1} \prod_{k=1}^g \widetilde{[a_k, b_k]} w^{-1} z) da_1 db_1 \dots da_g db_g dx \\ & \leq \frac{1}{Z_{T, z}} \|\tilde{p}_{\frac{T}{2}\sigma(M)}\|_\infty \int_G f(x) p_{T\sigma_i^m}(x) dx. \end{aligned}$$

In both cases, by Proposition 1.7 and the estimate of the heat kernel given for example in [21], Theorem V.4.3, we get, for some constant C depending only on G and for every $\delta > 0$,

$$P_T [d(H_{c_i^m}, H_{p_i}) > \delta] \leq C (T\sigma_i^m)^{-\frac{\dim G}{2}} e^{-\frac{C \text{diam}(G)^2}{T} - \frac{\delta^2}{CT\sigma_i^m}}.$$

Hence,

$$\begin{aligned}
\limsup_{T \rightarrow 0} T \log P_T \left[\max_{1 \leq i \leq n} d(H_{c_i^m}, H_{p_i}) > \delta \right] &\leq \limsup_{T \rightarrow 0} T \log \sum_{i=1}^n P_T [d(H_{c_i^m}, H_{p_i}) > \delta] \\
&= \max_{1 \leq i \leq n} \limsup_{T \rightarrow 0} T \log P_T [d(H_{c_i^m}, H_{p_i}) > \delta] \\
&\leq C \text{diam}(G)^2 - \min_{1 \leq i \leq n} \frac{\delta^2}{C \sigma_i^m}.
\end{aligned}$$

Since the sequences $(\sigma_i^m)_m$ converge to 0, the limit as m tends to infinity of this expression is equal to $-\infty$.

Let ω be an H^1 connection. By the energy inequality (Proposition 2.4),

$$\rho(\langle \omega, c_i^m p_i^{-1} \rangle)^2 \leq \sigma_i^m S(\omega).$$

Now, the domain bounded by $p_i^{-1} c_i^m$ is contained in a tube around p_i of width $d_\infty(p_i, c_i^m) \leq d_\ell(p_i, c_i^m)$. Thus, there exists a constant K , depending on the paths p_1, \dots, p_n , such that $\sigma_i^m \leq K d_\ell(p_i, c_i^m)$. Finally, we get

$$\rho(\langle \omega, c_i^m p_i^{-1} \rangle)^2 \leq K d_\ell(p_i, c_i^m) S(\omega)$$

and property 3 holds.

If the paths p_1, \dots, p_n are not edges, let us write them in some way as concatenations of edges. Let $\{q_1, \dots, q_r\}$ be the set of distinct edges that have been used in at least one of the decompositions. We apply the arguments above to this new set of paths. We find r sequences $(d_j^m)_{m \geq 1}$, $j = 1 \dots r$ satisfying properties 0-3. Let us make the further assumption that $\ell(d_j^m) \geq \ell(q_j)$ for all j and m .

Assume for instance that $p_1 = q_{i_1} \dots q_{i_s}$, where $1 \leq i_1, \dots, i_s \leq r$. Then, for all $m \geq 1$,

$$d(H_{p_1}, H_{d_{i_1}^m \dots d_{i_s}^m}) \leq \sum_{j=1}^s d(H_{q_{i_j}}, H_{d_{i_j}^m}).$$

Set $c_1^m = d_{i_1}^m \dots d_{i_s}^m$ and define the others c_i^m , $i = 2 \dots n$ in a similar fashion. Let N be the largest number of non necessarily distinct edges that it is necessary to concatenate in order to get one of the paths p_i . We have

$$P_T \left[\max_{1 \leq i \leq n} d(H_{c_i^m}, H_{p_i}) > \delta \right] \leq P_T \left[\max_{1 \leq j \leq r} d(H_{d_j^m}, H_{q_j}) > \frac{\delta}{N} \right]$$

and property 2 follows.

Let us prove property 3 for $i = 1$. By applying the special case of property 3 that we have proved above to q_{i_1}, \dots, q_{i_s} , we find a constant K' such that

$$\begin{aligned}
\rho(\langle \omega, c_1^m p_1^{-1} \rangle)^2 &\leq \left(\sum_{j=1}^s \rho(\langle \omega, d_{i_j}^m q_{i_j}^{-1} \rangle) \right)^2 \\
&\leq K' S(\omega) s \sum_{j=1}^s d_\ell(d_{i_j}^m, q_{i_j}).
\end{aligned}$$

Since $\ell(q_i) \leq \ell(d_i^m)$, we have for each j the inequality $|\ell(d_{i_j}^m) - \ell(q_{i_j})| \leq |\ell(c_i^m) - \ell(p_i)|$ and the last term is bounded above by $K's^2 S(\omega) d_\ell(c_1^m, p_1)$. Finally, $s \leq N$ and property 3 follows. \square

Proposition 3.12 *Let $J = \{p_1, \dots, p_n\}$ be a finite subset of PM .*

1. (Boundary) *The laws of $(H_{p_1}, \dots, H_{p_n})$ under $P_{T; X_1, \dots, X_p}$, satisfy, as T tends to 0, a large deviation principle on $\mathcal{M}(J, G)$ with rate function*

$$I_{X_1, \dots, X_p}^J(g) = \frac{1}{2} \inf \{S(\omega) : \omega \in H^1 \mathcal{A}_{X_1, \dots, X_p}(P), \omega \sim_J g.\}$$

2. (Closed) *The laws of $(H_{p_1}, \dots, H_{p_n})$ under $P_{T, z}$, satisfy, as T tends to 0, a large deviation principle on $\mathcal{M}(J, G)$ with rate function*

$$I_z^J(g) = \frac{1}{2} \inf \{S(\omega) : \omega \in H^1 \mathcal{A}(P), \omega \sim_J g.\}$$

Proof – Here again, the proof is exactly the same with and without boundary. We drop the subscripts that usually indicate in which context we are.

Let $(c_i^m)_{m \geq 1}$, $i = 1 \dots n$, be given by Lemma 3.11. For each $m \geq 1$, denote by J_m the set of paths $\{c_1^m, \dots, c_n^m\}$. By Proposition 3.10, the laws of $(H_{c_1^m}, \dots, H_{c_n^m})$ under P_T satisfy a large deviation principle in G^J with rate function I^{J_m} .

By a standard result on exponential approximations [2, Theorem 4.2.16], property 2 of Lemma 3.11 ensures that the laws of $(H_{p_1}, \dots, H_{p_n})$ under P_T , $T > 0$ satisfy a large deviation principle on G^J with rate function

$$\hat{I}^J(g) = \sup_{\delta > 0} \liminf_{m \rightarrow \infty} \inf_{h \in B(g, \delta)} I^{J_m}(h).$$

Here, $B(g, \delta)$ denotes the open ball of radius δ around g in G^J . The proof is completed by the next lemma, after noticing that the domain $\{I^J < +\infty\}$ is contained in the closed subset $\mathcal{M}(J, G)$ of G^J . \square

Lemma 3.13 *We keep the preceding notation. Then*

$$\sup_{\delta > 0} \liminf_{m \rightarrow \infty} \inf_{h \in B(g, \delta)} \inf_{\omega \sim_{J_m} h} S(\omega) = \inf_{\omega \sim_J g} S(\omega).$$

In the case with boundary, the two last infima are taken over all H^1 connections which satisfy the boundary conditions. In the closed case, they are taken over all H^1 defined on a principal G -bundle over M which belongs to the correct isomorphism class.

Proof – For each $m \geq 1$ and $\delta > 0$, define the set

$$O_{m, \delta} = \{\omega \in H^1 \mathcal{A} : \exists h \in B(g, \delta), \omega \sim_{J_m} h\}.$$

Since the holonomy along a fixed path depends continuously on the connection in the H^1 topology, these are open subsets of $H^1 \mathcal{A}_{X_1, \dots, X_p}(P)$ or $H^1 \mathcal{A}(P)$.

Now, let ω be an H^1 connection such that $\omega \sim_J g$. According to Proposition 2.16, the holonomy induced by ω is continuous with fixed endpoints. Hence, for every $\delta > 0$, ω belongs to $O_{m,\delta}$ for m large enough.

Choose $\alpha > 0$ and an H^1 connection ω_0 such that $\omega_0 \sim_J g$ and such that $S(\omega_0) \leq \inf\{S(\omega) : \omega \sim_J g\} + \alpha$. Choose $\delta > 0$. By the observation above,

$$\liminf_{m \rightarrow \infty} \inf_{O_{m,\delta}} S \leq S(\omega_0).$$

By letting α , then δ tend to 0, we get the inequality $\hat{I}^J \leq I^J$.

Assume that this inequality is strict. Then, for some $\alpha > 0$, for all $\delta > 0$, one has

$$\liminf_{m \rightarrow \infty} \inf_{O_{m,\delta}} S \leq \inf_{\omega \sim_J g} S(\omega) - \alpha = s - \alpha,$$

where we have set $s = \inf_{\omega \sim_J g} S(\omega)$. Let us fix $\delta > 0$. We can construct an increasing sequence of integers $(m_k)_{k \geq 1}$ and a sequence of H^1 connections $(\omega_{m_k})_{k \geq 1}$ with $\omega_{m_k} \in O_{m_k,\delta}$ and $\sup_k S(\omega_{m_k}) \leq s - \alpha/2$. From this sequence with bounded energy, we can, by Uhlenbeck's theorem, extract a subsequence which is gauge-equivalent to a weakly convergent sequence of connections. Thus, there exists a subsequence $(\omega_r)_{r \geq 1}$ of $(\omega_{m_k})_{k \geq 1}$, a sequence $(j_r)_{r \geq 1}$ in $H^2\mathcal{J}(P)$ and an H^1 connection ω such that $j_r \cdot \omega_r \rightharpoonup \omega$ and $S(\omega) \leq s - \alpha/2$. Let $(N_r)_{r \geq 1}$ denote an increasing sequence such that $\omega_r \in O_{N_r,\delta}$.

For every path c , the holonomy along c of $j_r \cdot \omega_r$ converges to that of ω . This holds in particular for the paths p_1, \dots, p_n . Moreover, since $S(\omega_r)$ is bounded independently of r , by property 3 of Lemma 3.11, the distance $\rho(\langle j_r \cdot \omega_r, c_i^{N_r} p_i^{-1} \rangle)$ tends to 0 as r tends to infinity. Hence, for each $i = 1 \dots n$, the holonomy of $j_r \cdot \omega_r$ along $c_i^{N_r}$ converges as r tends to infinity to the holonomy of ω along p_i . Hence, $\omega \in O_{\infty,2\delta}$, where we take the convention $J_\infty = J$.

For every $\delta > 0$, we are thus able to construct a connection ω in $O_{\infty,2\delta}$ such that $S(\omega) \leq s - \alpha/2$. For each $n \geq 1$, let us do this construction with $\delta = 1/2n$. This produces a sequence (ω_n) , from which we may again extract a subsequence gauge-equivalent to a weakly H^1 convergent one, with limit ω_* . This limit satisfies both $S(\omega_*) \leq s - \alpha/2$ and $\omega_* \sim_J g$. This contradicts the definition of s . \square

3.4 The Yang-Mills measures (Projective limit)

As explained in [9], Section 2.10.2, the probability space $(\mathcal{M}(PM, G), \mathcal{C}, P_T)$ is the projective limit of the spaces $(\mathcal{M}(J, G), \mathcal{C}_J, P_T^J)$, where J spans the set of finite subsets of PM and P_T^J denotes the distribution of the holonomy along the paths of J under P_T . A straightforward application of Dawson-Gärtner's theorem ([2], Theorem 4.6.1) gives the following result.

Proposition 3.14 1. (Boundary) The probability measures $(P_{T;X_1,\dots,X_p})_{T>0}$ satisfy, as T tends to 0, a large deviation principle on $\mathcal{M}(PM, G)$ with rate function

$$\tilde{I}_{X_1,\dots,X_p}^{\text{YM}}(f) = \sup_{J \subset PM, |J| < \infty} I_{X_1,\dots,X_p}^J(f) = \frac{1}{2} \sup_{J \subset PM, |J| < \infty} \inf\{S(\omega) : \omega \in H^1\mathcal{A}_{X_1,\dots,X_p}(P), \omega \sim_J f\}.$$

2. (Closed) The probability measures $(P_{T,z})_{T>0}$ satisfy, as T tends to 0, a large deviation principle on $\mathcal{M}(PM, G)$ with rate function

$$\tilde{I}_z^{\text{YM}}(f) = \sup_{J \subset PM, |J| < \infty} I_z^J(f) = \frac{1}{2} \sup_{J \subset PM, |J| < \infty} \inf\{S(\omega) : \omega \in H^1\mathcal{A}(P), \omega \sim_J f\}.$$

The proof of Theorems 1.15 and 1.16 will be complete after we have proved that $\tilde{I}^{\text{YM}} = I^{\text{YM}}$ in both cases. Since the proof of this equality is the same with and without boundary, we drop again the subscripts. Let us start with an easy lemma.

Lemma 3.15 *The inequality $\tilde{I}^{\text{YM}} \leq I^{\text{YM}}$ holds.*

Proof – Let $f \in \mathcal{M}(PM, G)$. If $I^{\text{YM}}(f) = +\infty$, then it is certainly true that $\tilde{I}^{\text{YM}}(f) \leq I^{\text{YM}}(f)$. Let us assume that $I^{\text{YM}}(f) < \infty$. Then there exists an H^1 connection ω which agrees with f on PM . In particular, it agrees with f on J for every subset J of PM . Hence, $\tilde{I}^{\text{YM}}(f) \leq I^{\text{YM}}(f)$. \square

Let us define another function $\hat{I}^{\text{YM}} : \mathcal{M}(PM, G) \longrightarrow [0, +\infty]$ by

$$\hat{I}^{\text{YM}}(f) = \sup_{\mathbb{G} \text{ graph}} I^{\mathbb{E}}(f).$$

Since \hat{I}^{YM} is a supremum over a smaller class of subsets of PM than \tilde{I}^{YM} , the inequality $\hat{I}^{\text{YM}} \leq \tilde{I}^{\text{YM}}$ holds. According to Lemma 3.15, it is enough to prove that $\hat{I}^{\text{YM}} = I^{\text{YM}}$, or even that $I^{\text{YM}} \leq \hat{I}^{\text{YM}}$.

Proposition 3.16 *The inequality $I^{\text{YM}} \leq \hat{I}^{\text{YM}}$ holds on $\mathcal{M}(PM, G)$.*

Proof – Consider $f \in \mathcal{M}(PM, G)$. Assume that $\hat{I}^{\text{YM}}(f) < \infty$, otherwise there is nothing to prove. Let l be a simple loop in PM . There exists a graph, say \mathbb{G} , such that l belongs to \mathbb{E}^* . Assume that l bounds a domain V diffeomorphic to an open disk. Then we may assume that V is a face of \mathbb{G} .

If M has a boundary, then by definition of \hat{I}^{YM} ,

$$\rho(f(l))^2 = \rho(f(\partial V))^2 \leq 2\hat{I}^{\text{YM}}(f)\sigma(V). \quad (13)$$

If M is closed, we can only say that there exists a lift $\tilde{f}(\tilde{l})$ of $f(l)$ to \tilde{G} and $z \in \Pi$ such that $\tilde{\rho}(\tilde{f}(\tilde{l})z)^2 \leq 2\hat{I}^{\text{YM}}(f)\sigma(V)$. However, notice that, if $x \in G$ and $\tilde{x} \in \tilde{G}$ satisfy $\pi(\tilde{x}) = x$, then $\rho(x) = \min_{z \in \Pi} \tilde{\rho}(\tilde{x}z)$. So, (13) holds even when M is closed. In both cases, Proposition 2.7 allows us to deduce that f is continuous with fixed endpoints.

Now let m be a point in M . There exists a countable dense subset of the space $L_m M$ of loops based at m . For example, consider a countable dense subset of M . Then the set $\Lambda_m M$ of piecewise geodesic loops based at m and joining a finite number of these points is countable and dense in $L_m M$. Let $(\zeta_n)_{n \geq 1}$ be a sequence which exhausts $\Lambda_m M$. For each $n \geq 1$, there exists a graph, say \mathbb{G}_n , such that $\zeta_i \in \mathbb{E}_n^*$ for each $i = 1 \dots n$.

For each $n \geq 1$, let ω_n be an H^1 connection which agrees with f on \mathbb{G}_n and such that $S(\omega_n) = 2I^{\mathbb{E}_n}(f)$. Such a connection exists by Proposition 3.9. Uhlenbeck's compactness theorem allows us to extract a weakly convergent subsequence of the sequence $(\omega_n)_{n \geq 1}$, up to gauge transformations. The weak limit of this subsequence has an energy at most equal to $2\hat{I}_{YM}(f)$ and it agrees with f on ζ_1, \dots, ζ_n for each n . Since G is compact, this implies that ω and f agree on $\Lambda_m M$. Finally, since both are continuous on $L_m M$, they agree on $L_m M$ itself. As pointed out in the remark 2.17, this is equivalent to saying that ω and f agree on PM .

Since there exists an H^1 connection ω which agrees with f , $I^{\text{YM}}(f)$ is finite. It is equal to $\frac{1}{2}S(\omega)$ and we have observed that $\frac{1}{2}S(\omega) \leq \hat{I}^{\text{YM}}(f)$. The result is proved. \square

4 Connections that minimize the Yang-Mills energy under holonomy constraints

4.1 The main result

The purpose of this section is to prove Proposition 3.9. We deduce it from the next proposition, in which we assume that M is closed. Once for all, let us choose a simple graph $\mathbb{G} = (\mathbb{V}, \mathbb{E}, \mathbb{F})$ on M . Let \mathbb{E}^+ be an orientation of \mathbb{V} , which satisfies the properties explained in Lemma 1.2.

Proposition 4.1 *Let $g = (g_e)_{e \in \mathbb{E}^+}$ be an element of $G^{\mathbb{E}^+}$. Let $\tilde{g} = (\tilde{g}_e)_{e \in \mathbb{E}^+}$ be an element of $\tilde{G}^{\mathbb{E}^+}$ such that $g = \pi(\tilde{g})$. Let z be an element of Π . Let $z_{\mathbb{F}} = (z_F)_{F \in \mathbb{F}}$ be an element of $\Pi_z^{\mathbb{F}}$. There exists a principal G -bundle P over M and a connection ω on P such that the following properties hold.*

1. $\mathfrak{o}(P) = z$.

2. ω belongs to $W^{1,\infty} \mathcal{A}(P)$ and, for each open face F of \mathbb{G} , the restriction of ω to $P|_F$ is smooth.

3. ω and g agree on \mathbb{E} .

4. For each face F of \mathbb{G} , $S_F(\omega) = \frac{\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F)^2}{\sigma(F)}$.

Let us explain how this result implies Proposition 3.9.

Proof of Proposition 3.9 – 1. Let \overline{M} be a closure of M , that is, a closed surface in which M is embedded in such a way that $\overline{M} \setminus M$ is a disjoint union of p disks. Then \mathbb{G} is still a simple graph on \overline{M} , it only has p more faces. Let $\tilde{g} \in \tilde{G}^{\mathbb{E}^+}$ be such that $\pi(\tilde{g}) = g$. For each face F of \mathbb{G} contained in M , let z_F be an element of Π such that $\tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F) = \rho(h_{\partial F}(g))$. Such a z_F exists because, for each $\tilde{x} \in \tilde{G}$, $\rho(\pi(\tilde{x})) = \min_{z \in \Pi} \tilde{\rho}(\tilde{x}z)$. For the faces of \mathbb{G} not contained in M , choose z_F arbitrarily. Then set $z = \prod_F z_F$. Proposition 4.1 produces a bundle Q over \overline{M} and a connection η on Q . A bundle P is given by assumption over M . Since M has a boundary, P is trivial, and so is the restriction of Q to M . Let $\varphi : P \rightarrow Q|_M$ be a bundle isomorphism. Set $\omega = \varphi^* \eta$. Then, since ω agrees with g on \mathbb{E} and g satisfies the correct boundary conditions, ω belongs to $W^{1,\infty} \mathcal{A}_{X_1, \dots, X_p}(P)$. It is smooth outside $\cup_{e \in \mathbb{E}} e$ because η is. Finally, the choice of $(z_F)_{F \in \mathbb{F}}$ guarantees that $S(\omega)$ is equal to $I_{X_1, \dots, X_p}^{\mathbb{E}}(g)$.

2. In this case, the result is straightforward: it suffices to choose $\tilde{g} \in \tilde{G}^{\mathbb{E}^+}$ such that $\pi(\tilde{g}) = g$ and then $z_{\mathbb{F}}$ which minimizes the right hand side of (9). \square

4.2 An open covering of M

We begin by constructing an open covering of M which is nicely adapted to \mathbb{G} . Recall that M is endowed with a Riemannian metric. This metric allows us to define tubular neighbourhoods around embedded submanifolds of M , see for example [5].

Up to this point, we have always called *faces* the closed faces of \mathbb{G} . In this section, we change this convention and decide to call faces the *open faces* of M .

Lemma 4.2 *There exist two real numbers $R, L > 0$ such that the following properties hold.*

1. The balls $B(v, R), v \in \mathbb{V}$ are diffeomorphic to disks and pairwise disjoint. Moreover, an edge e meets the ball $B(v, R)$ if and only if $v \in \{\underline{e}, \bar{e}\}$. In this case, $e \cap B(v, R)$ is connected.
2. Let $v \in \mathbb{V}$. Let f_1, \dots, f_k be the edges sharing v as their starting point, indexed in their cyclic order around v induced by the orientation of M . It is possible to choose polar normal coordinates (r, θ) in $B(v, R)$ such that there exist k smooth functions $\theta_1, \dots, \theta_k : [0, R) \rightarrow [0, 2\pi)$ and $2k$ real numbers $0 \leq \theta_1^- < \theta_1^+ < \dots < \theta_k^- < \theta_k^+ < 2\pi$, such that, for all $j = 1 \dots k$, $(r, \theta_j(r))$ is a parametrization of f_j inside $B(v, R)$ and $\theta_j^- < \inf_{[0, R)} \theta_j < \sup_{[0, R)} \theta_j < \theta_j^+$. We call the sector $\{(r, \theta) : \theta_j^- < \theta < \theta_j^+\}$ the angular sector of f_j at v .
Moreover, let $(r, \theta) \mapsto \sigma(r, \theta)$ be the smooth density of the measure σ with respect to $r dr d\theta$ on $B(v, R)$. Then any partial derivative of any order of σ is uniformly bounded on $D(0, R)$.
For each edge e , let e° denote the intersection of e with the subset $M \setminus \bigcup_{v \in \mathbb{V}} \overline{B(v, R/2)}$. Let T_e denote the tubular neighbourhood of radius L around e° .
3. For each edge e , the tubular neighbourhood of radius L around e° exists. It is denoted by T_e . If $e, e' \in \mathbb{E}$ are not equal nor inverse of each other, then T_e and $T_{e'}$ are disjoint.
4. Let v be a vertex and e an edge. The tube T_e meets $B(v, R)$ only if v is an endpoint of e .
5. Let e be an edge. There exists a coordinate chart $T_e \rightarrow (-3, 3) \times (-1, 1)$ with coordinates (x, y) such that $e \cap T_e = \{y = 0\}$, $L(e) \cap T_e = \{y > 0\}$, $B(\underline{e}, 3R/4) \cap T_e = \{x < -2\}$, $B(\underline{e}, R) \cap T_e = \{x < -1\}$, $B(\bar{e}, R) \cap T_e = \{x > 1\}$ and $B(\bar{e}, 3R/4) \cap T_e = \{x > 2\}$.
Moreover, let $(x, y) \mapsto \sigma(x, y)$ be the smooth density of the measure σ with respect to $dx dy$ on T_e . Then any partial derivative of any order of σ is uniformly bounded on $(-3, 3) \times (-1, 1)$.

Proof – Let R_{inj} be the injectivity radius of M , so that any ball of radius smaller than R_{inj} is diffeomorphic to a disk. Set $R_1 = R_{inj} \wedge \frac{1}{2} \inf_{v \neq w} d(v, w)$, where the infimum is taken over all pairs of distinct vertices. The balls $B(v, R), v \in \mathbb{V}$ are diffeomorphic to disks and pairwise disjoint as soon as $R \leq R_1$.

Let us choose in each ball $B(v, R_1)$ a system of normal polar coordinates. This amounts to choosing the initial speed of the geodesic of equation $\{\theta = 0\}$. We choose it in such a way that no edge starting from v is tangent to this geodesic at v .

Let v be a vertex and e an edge such that $\underline{e} = v$. Since e is a segment of an embedded submanifold, it can be parametrized near v and inside $B(v, R_1)$ as $e(s) = (r(s), \theta(s))$, where $s \geq 0$ and $e(0) = v$. Then $\dot{r}(0) > 0$. Choose $s_v(e) > 0$ such that $\dot{r}(s) > 0$ for all $s \in [0, s_v(e)]$. Set $R_2 = \inf\{r(s_v(e)) : v \in \mathbb{V}, \underline{e} = v\}$.

Let e be an edge. Set $\tilde{e} = e \cap (B(\underline{e}, R_2) \cup B(\bar{e}, R_2))^c$. Finally, set $R_3 = \frac{1}{2} \inf d(\tilde{e}, v)$, where the infimum is taken over all edges and all vertices. Observe that $R_3 < R_2$. Take $R \leq R_3$. Then the balls $B(v, R)$ satisfy the first point of the lemma. Moreover, we can say that, if e is incident to v , it crosses the circle of radius R around v transversally.

Let us consider a vertex v and an edge e incident at v . Once again, because e is a segment of an embedded submanifold, the local parametrization $(r(s), \theta(s))$ of e defined above is such that $s \mapsto \theta(s)$ can be extended by continuity at $s = 0$.

Let us denote by f_1, \dots, f_k be the edges starting at v , given in their cyclic order around v , and $(r_1, \theta_1), \dots, (r_k, \theta_k)$ their local parametrizations. We assume that the edges are indexed in such a

way that $0 < \theta_1(0) < \dots < \theta_k(0) < 2\pi$. Set $\delta = \theta_1(0) \wedge (2\pi - \theta_k(0)) \wedge \inf_{1 \leq i \leq k-1} |\theta_{i+1}(0) - \theta_i(0)|$. Choose $s'_v > 0$ such that, for all $i = 1 \dots k$ and all $s \in [0, s'_v]$, $|\theta_i(s) - \theta_i(0)| < \delta/4$. Finally, set $R_4 = \inf r_i(s'_v)$, where v runs over \mathbb{V} and r_i is the local parametrization of an edge incident to v . Then, any R such that $0 < R \leq R_4$ satisfies the two first points of the lemma, except maybe for the boundedness condition on the derivatives of σ . Let us choose $R = R_4/2$ to make sure that it holds. This allows us to define e° for each edge e . In fact, let us temporarily consider e^\bullet which is the larger portion of e defined by $e^\bullet = e \cap (M \setminus \bigcup_{v \in \mathbb{V}} \overline{B(v, R/4)})$.

For each edge e , let L_e be the largest width of a tube around e^\bullet . Set $L_1 = \inf_e L_e \wedge \frac{1}{2} \inf_{e \neq f} d(e^\bullet, f^\bullet)$. Then any positive L smaller than L_1 satisfies the third point. Now set $L_2 = \frac{1}{2} \inf_{e,v} d(e^\bullet, B(v, R))$, where the infimum runs over all pairs (e, v) with $v \notin \{\underline{e}, \bar{e}\}$. Any positive L smaller than $L_3 = L_1 \wedge L_2$ satisfies the fourth point.

Let e be an edge. Let T_e be a tube around e^\bullet , endowed with Fermi coordinates (t, s) , such that e^\bullet is defined by the equation $s = 0$. Consider a curve γ which crosses e^\bullet only once and transversally. Near its intersection point with e^\bullet , γ can be parametrized as $(t(\tau), s(\tau))$, with $\tau = 0$ corresponding to the intersection point. Then, $\dot{s}(0) \neq 0$. Hence, there exists a positive number $L'_e(\gamma)$ such that the portion of γ contained in a tube of width smaller than $L'_e(\gamma)$ is the graph in Fermi coordinates of a smooth function $s \mapsto T(s)$, which crosses the boundary of the tube transversally. Then, the domain $\{(s, t) : t \leq T(s)\}$ is diffeomorphic to a rectangle by a diffeomorphism which sends e^\bullet to a segment parallel to an edge.

By applying for each edge e this argument to the circles of centers \underline{e} and \bar{e} and radii R and $3R/4$, we find a positive width L'_e such that $L_4 = L_3 \wedge \inf_e L'_e$ satisfies the fifth and sixth points, excepts perhaps for the boundedness condition of the derivatives of σ . For each edge e , the tube of radius $L = L_4/2$ around e° is contained, as well as its closure, in the tube of radius L around e^\bullet and satisfies all the required properties. \square

Let R and L be two positive numbers given by this lemma. We define a collection of open subsets of M as follows. For each vertex v , set $U_v = B(v, R)$. For each edge e , set $U_e = B(\underline{e}, 3R/4) \cup T_e \cup B(\bar{e}, 3R/4)$. For each face F , set $U_{\partial F} = \cup_{L(e)=F} U_e$ and $U_F = F \cup U_{\partial F}$. Finally, set $U_{\mathbb{G}} = \cup_{F \in \mathbb{F}} U_{\partial F} = \cup_{e \in \mathbb{E}} U_e$.

4.3 Around the singularities

The connections we are going to construct are singular on the subset $\cup_{e \in \mathbb{E}} e$ of M . This set contains two kinds of points: the vertices of \mathbb{G} and the points which are interior to an edge. We treat these two cases separately.

Proposition 4.3 *Let v be a vertex of \mathbb{G} . For each face F such that v lies on ∂F , let X_F be an element of \mathfrak{g} . There exists $\omega \in W^{1,\infty} \Omega_{\mathfrak{g}}^1(U_v)$ such that the following properties hold.*

1. *If e is an edge starting at v , then ω vanishes in all directions on $U_v \cap e$.*
2. *If F is a face such that v lies on ∂F , then ω is smooth on $U_v \cap F$. Moreover, there exists $\lambda \in \Omega^1(U_v \cap F)$ such that $d\lambda = \sigma$ and $\omega = X_F \lambda$ on $U_v \cap F$. In particular, the curvature of ω is equal to $X_F \sigma$ on $U_v \cap F$.*

Proof – Let us denote, as in Lemma 4.2, by f_1, \dots, f_k be the edges starting at v , given in their cyclic order around v , and $(r_1, \theta_1), \dots, (r_k, \theta_k)$ their local parametrizations. Pick an integer i between 1 and k . Set $X_i = X_F$, where F is the face sitting between f_i and f_{i+1} , with

the convention $f_{k+1} = f_1$. Let $\sigma(r, \theta)$ be the smooth positive function on $D(0, R)$ such that $\sigma = \sigma(r, \theta)rdrd\theta$. Define, for $r \in [0, R)$,

$$\beta_i(r) = X_i \int_{\theta_i(r)}^{\theta_{i+1}(r)} \sigma(r, \theta) d\theta$$

and

$$\alpha_i(r) = \frac{2}{r^2} \int_0^r u\beta_i(u) du.$$

Let $\varphi_i : [0, 2\pi) \rightarrow [0, \infty)$ be a smooth nonnegative function such that $\int_0^{2\pi} \varphi_i(\theta) d\theta = 1$ and such that the support of φ_i is contained in $(\theta_i^+, \theta_{i+1}^-)$. Set

$$a(r, \theta) = \varphi_i(\theta)\alpha_i(r), \quad \theta_i(r) \leq \theta < \theta_{i+1}(r),$$

and define a function b on $B(v, R)$ by setting $b(r, \theta_1(r)) = 0$ and

$$\frac{\partial b}{\partial \theta} = \sigma(r, \theta)X_i - \varphi_i(\theta)\beta_i(r), \quad \theta_i(r) \leq \theta < \theta_{i+1}(r).$$

After doing this for each $i = 1 \dots k$, we have defined two functions a and b on U_v . Finally, set

$$\omega = \frac{r^2}{2}a(r, \theta)d\theta - rb(r, \theta)dr.$$

We claim that ω belongs to $W^{1, \infty}\mathcal{A}(U_v)$. To see this, let us write ω in Cartesian coordinates corresponding to our choice of polar coordinates:

$$\omega = \left(-\frac{y}{2}a(r, \theta) - xb(r, \theta)\right) dx + \left(\frac{x}{2}a(r, \theta) - yb(r, \theta)\right) dy.$$

Recall that the function σ is bounded as well as its derivatives on $D(0, R)$. For each $i = 1 \dots k$, the function β_i is continuous and bounded on $[0, R)$ and so is α_i . Both are also smooth on $(0, R)$. Hence, a is bounded on U_v and smooth except at v . Observe that $\int_{\theta_i(r)}^{\theta_{i+1}(r)} \sigma(r, \theta)X_i - \varphi_i(\theta)\beta_i(r) d\theta = 0$, so that b is continuous and bounded on $U_v - \{v\}$. It is also smooth in U_v outside \mathbb{G} .

So far, we have proved that the components of ω belongs to L^∞ . Now, along any segment in U_v parallel to one of the coordinate axes and which does not contain v , a is smooth and b is continuous and piecewise smooth with bounded derivative. In particular, both are absolutely continuous and so are the components of ω . Let us show that the almost-everywhere defined derivatives of these components are uniformly bounded on U_v . For example,

$$\frac{\partial}{\partial x}(ya)(r, \theta) = \frac{2}{r^2} \int_0^r u\beta_i(u) du \left(-\frac{y^2}{r^2}\varphi'(\theta) - \frac{2xy}{r^2}\varphi(\theta)\right) + \frac{2xy}{r^2}\varphi_i(\theta)\beta_i(r),$$

so that

$$\|\partial_x(ya)\|_{L^\infty} \leq 10 \sup_{1 \leq i \leq k} (\|\beta_i\|_\infty(\|\varphi_i\|_\infty + \|\varphi_i'\|_\infty)) < \infty.$$

A similar computation shows that the partial derivatives of xa, ya, xb, yb are all uniformly bounded on U_v . Finally, the components of ω are L^∞ , absolutely continuous along almost

all segments parallel to one of the coordinate axes inside U_v , with partial derivative belonging to $L^\infty(U_v)$. This implies that the components of ω belong to $W^{1,\infty}$.

The fact that ω vanishes along each edge starting at v comes from the fact that a vanishes in the angular sector of each edge and $b(r, \theta_i(r))$ vanishes identically for each $i = 1 \dots k$.

Let F be a face incident to v . By construction, ω is smooth inside F and it can be written as $X_F(v)$ times a real-valued 1-form. It is readily checked that the differential of this form is σ itself: a and b have been designed for that purpose. The statement on the curvature of ω follows immediately. \square

Proposition 4.4 *Let e be an edge of \mathbb{G} . For each face F bounded by e , let X_F be an element of \mathfrak{g} . There exists $\omega \in W^{1,\infty}\Omega_{\mathfrak{g}}^1(T_e)$ such that the following properties hold.*

1. *The form ω vanishes in all directions on $T_e \cap e$.*
2. *If F is a face bounded by e , then ω is smooth on $T_e \cap F$. Moreover, there exists $\lambda \in \Omega^1(T_e \cap F)$ such that $d\lambda = \sigma$ and $\omega = X_F\sigma$ on $T_e \cap F$.*

Let \tilde{g} be an element of \tilde{G} . There exists a smooth function $j : T_e \rightarrow \tilde{G}$ which is identically equal to 1 on $T_e \cap U_{\underline{e}}$ and identically equal to \tilde{g}^{-1} on $T_e \cap U_{\bar{e}}$. Then the following properties hold.

3. *$\langle j \cdot \omega, e^\circ \rangle_{\tilde{G}} = \tilde{g}$.*
4. *The form $j \cdot \omega$ vanishes in all directions on $T_e \cap (U_{\underline{e}} \cup U_{\bar{e}}) \cap e$.*
5. *If F is a face bounded by e , then the curvature of $j \cdot \omega$ is equal to $X_F\sigma$ on $T_e \cap F \cap U_{\underline{e}}$, and it is equal to $\text{Ad}(\tilde{g})X_F\sigma$ on $T_e \cap F \cap U_{\bar{e}}$.*

Proof – Let $(x, y) \in (-3, 3) \times (-1, 1)$ be the local coordinates on T_e given by Lemma 4.2. Let $\sigma(x, y)$ be the smooth positive function on $(-3, 3) \times (-1, 1)$ such that the equality $\sigma = \sigma(x, y)dx dy$ holds. Set

$$\omega(x, y) = -(X_{L(e)}\mathbf{1}_{y \geq 0} + X_{L(e^{-1})}\mathbf{1}_{y \leq 0}) \left(\int_0^y \sigma(x, t) dt \right) dx.$$

Since $(x, y) \mapsto \sigma(x, y)$ is bounded on $(-3, 3) \times (-1, 1)$ as well as its derivatives, ω is Lipschitz, hence $W^{1,\infty}$. It is also smooth on T_e outside e and satisfies property 2.

Let $\psi : (-3, 3) \rightarrow \tilde{G}$ be a smooth mapping such that $\psi(x) = 1$ whenever $x \leq -1$ and $\psi(x) = \tilde{g}^{-1}$ whenever $x \geq 1$. Finally, set $j(x, y) = \psi(x)$. Property 3, 4 and 5 are straightforward. \square

4.4 On a neighbourhood of the graph

We want to combine the two constructions presented above to get an element of $W^{1,\infty}\Omega_{\mathfrak{g}}^1(U_{\mathbb{G}})$. For this, we need to choose a configuration in $\tilde{G}^{\mathbb{E}^+}$ and several elements of \mathfrak{g} . To begin with, a configuration \tilde{g} is given by assumption in Proposition 4.1. Then, for each face $F \in \mathbb{F}$, let us choose a vertex $o(F)$ on the boundary of F and call ∂F the loop based at $o(F)$ going once around F with positive orientation. If v is a vertex on ∂F other than $o(F)$, denote by $\partial F_{o \rightarrow v}$ the portion of ∂F going from $o(F)$ to v . Let us decide that $\partial F_{o \rightarrow o}$ is the path ∂F itself. For each vertex v on the boundary of F , set $\tilde{x}_F(v) = h_{\partial F_{o \rightarrow v}}^{\tilde{G}}(\tilde{g})$. We denote $\tilde{x}_F(o(F)) = h_{\partial F}^{\tilde{G}}(\tilde{g})$ simply by \tilde{x}_F .

Choose an element $X_F \in \mathfrak{g}$ of minimal norm such that $\widetilde{\exp}(\sigma(F)X_F) = \tilde{x}_F z_F$, where $\widetilde{\exp} : \mathfrak{g} \rightarrow \widetilde{G}$ is the exponential map. Observe that $\tilde{\rho}(\tilde{x}_F z_F) = \sigma(F)\|X_F\|$. Finally, for each v on the boundary of F , set $X_{F,v} = \text{Ad}(\tilde{x}_F(v))X_F$. By definition, the following compatibility condition is satisfied for each edge e :

$$\forall e \in \mathbb{E}^+, \forall F \in \{L(e), L(e^{-1})\}, X_{F,\bar{e}} = \text{Ad}(\tilde{g}_e)X_{F,\underline{e}}. \quad (14)$$

Once these choices are made, Proposition 4.3 provides us with a collection of 1-forms $(\omega_v, v \in \mathbb{V})$ and Proposition 4.4 with a collection $(\omega_e, e \in \mathbb{E}^+)$. We prove now that it is possible to let a gauge transformation act on each form ω_e in such a way that it coincides with $\omega_{\underline{e}}$ and $\omega_{\bar{e}}$ on the domains $B(\underline{e}, 3R/4) \cap T_e$ and $B(\bar{e}, 3R/4)$ respectively.

Proposition 4.5 *Let $e \in \mathbb{E}^+$ be an edge. There exists $j_e \in W^{2,\infty}(T_e; \widetilde{G})$ such that the following properties hold.*

1. *The forms $j_e \cdot \omega_e$ and $\omega_{\underline{e}}$ (resp. $\omega_{\bar{e}}$) coincide on $B(\underline{e}, 3R/4) \cap T_e$ (resp. $B(\bar{e}, 3R/4) \cap T_e$).*
2. *j_e is smooth outside e and identically equal to 1 on e .*

Proof – Let $(x, y) \in (-3, 3) \times (-1, 1)$ be the local coordinates on T_e given by Lemma 4.2. The forms ω_e and $\omega_{\underline{e}}$ are both defined on $(-3, -1) \times (-1, 1)$. We are going to apply Lemma 2.8 and Proposition 2.9 to these two forms on this domain.

First, let us choose on M an auxiliary Riemannian metric for which e is a geodesic. Let l be a loop based at $(-2, 0)$ and contained in $(-3, -1) \times [0, 1)$. Then, according to Propositions 4.3 and 4.4, there exist two smooth 1-forms λ and λ' such that $d\lambda = d\lambda' = \sigma$, $\omega_{\underline{e}} = X_{L(e), \underline{e}}\lambda$ and $\omega_e = X_{L(e), \underline{e}}\lambda'$ on $(-3, -1) \times [0, 1)$. Since this domain is simply connected, we conclude that $\langle \omega_{\underline{e}}, l \rangle_{\widetilde{G}} = \widetilde{\exp}(X_{L(e), \underline{e}} \int_l \lambda) = \widetilde{\exp}(X_{L(e), \underline{e}} \int_l \lambda') = \langle \omega_e, l \rangle_{\widetilde{G}}$. Hence, Lemma 2.8 shows that there exists $j_+ : (-3, -1) \times [0, 1) \rightarrow \widetilde{G}$ which transforms the holonomy of ω into that of ω' on this domain.

Replacing $L(e)$ by $L(e^{-1})$ and $(-3, -1) \times [0, 1)$ by $(-3, -1) \times (-1, 0]$, we find that the equality $\langle \omega_{\underline{e}}, l \rangle_{\widetilde{G}} = \langle \omega_e, l \rangle_{\widetilde{G}}$ holds also if l is contained in $(-3, -1) \times (-1, 0]$. Lemma 2.8 gives similarly $j_- : (-3, -1) \times [0, 1) \rightarrow \widetilde{G}$. Moreover, $j_+(-2, 0) = j_-(-2, 0) = 1$.

Finally, both ω and ω_e vanish in all directions on $e \cap U_{\underline{e}} \cap T_e$, so that if c is a segment contained in e , both $c^*\omega_e$ and $c^*\omega_{\underline{e}}$ are equal to 0. Hence, both j_+ and j_- are identically equal to 1 on e . They combine to give a function $j_{\underline{e}} : (-3, -1) \times (-1, 1) \rightarrow \widetilde{G}$ which transforms the holonomy of ω into that of ω' , as long as one restricts oneself to paths which are finite concatenations of paths which stay on either side of e . Since e is geodesic for the auxiliary metric on M , piecewise geodesic paths for this metric have this property. Hence, the assumptions of Proposition 2.9 are satisfied on $(-3, -1) \times (-1, 1)$ by ω , ω' and j , with $k = 1$ and $p = \infty$. Hence, $j_{\underline{e}}$ belongs to $W^{2,\infty}(U; \widetilde{G})$. According to the remark made immediately after Proposition 2.9, $j_{\underline{e}}$ is smooth outside e . It is also identically equal to 1 on e by construction.

At the vertex \bar{e} , the compatibility conditions $X_{L(e^{\pm 1}), \bar{e}} = \text{Ad}(\tilde{g}_e)X_{L(e^{\pm 1}), \underline{e}}$ and the same arguments as above imply that the forms ω_e and $\omega_{\bar{e}}$ restricted to $U_{\bar{e}} \cap T_e$ also satisfy the assumptions of Lemma 2.8. Thus, we find in the same way $j_{\bar{e}} \in W^{2,\infty}(U_{\bar{e}} \cap T_e; \widetilde{G})$, smooth outside e , identically equal to 1 on e and such that $j_{\bar{e}} \cdot \omega_e = \omega_{\bar{e}}$ on $U_{\bar{e}} \cap T_e$.

There remains to extend $j_{\underline{e}}$ and $j_{\bar{e}}$ to an element of $W^{2,\infty}(T_e; \widetilde{G})$. The functions $j_{\underline{e}}$ and $j_{\bar{e}}$, which are in particular Lipschitz, extend respectively to continuous functions on $[-3, -1] \times [-1, 1]$ and $[1, 3] \times [-1, 1]$. We start by interpolating them by a continuous function $j_0 : [-3, 3] \times$

$[-1, 1] \longrightarrow \tilde{G}$, for example by setting, for $-1 \leq x \leq 0$ and $0 \leq y \leq 1$, $j_0(x, y) = j_{\underline{e}}(-1, (y - x - 1)^+)$ and defining j_0 similarly on the three other quarters of $[-1, 1]^2$. We still have $j_0(x, 0) = 1$ for every $x \in [-3, 3]$.

Let us embed \tilde{G} in the linear space $\mathbb{M}_N(\mathbb{R})$ of $N \times N$ real matrices for some $N \geq 1$. Let us endow this space of matrices with a Euclidean scalar product. As a smooth Riemannian submanifold of $\mathbb{M}_N(\mathbb{R})$, G admits a tubular neighbourhood. In particular, there exists $\varepsilon > 0$ such that, if G^ε denotes the set of matrices at a Euclidean distance smaller than ε to G , then there exists a smooth mapping $pr_G : G^\varepsilon \longrightarrow G$ which is the identity when restricted to G . Since the range of j_0 is bounded, we can also assume that ε is small enough to guarantee that any matrix closer than ε to the range of j_0 is invertible.

Take $\varepsilon' > 0$ and let $J : [-3, 3] \times [-1, 1] \longrightarrow \mathbb{M}_N(\mathbb{R})$ be a smooth function such that $\|J - j_0\|_\infty < \varepsilon'$. Such a J can be constructed for instance by smoothing j_0 by convolution. Replacing J by $J(x, y)J(x, 0)^{-1}$, and provided ε' is small enough, we can assume that $\|J - j_0\|_\infty < \varepsilon$ and $J(x, 0) = 1$ for all $x \in [-3, 3]$.

Let now $\varphi : [-3, 3] \longrightarrow [0, 1]$ be a smooth function such that $\varphi(x) = 0$ if $|x| \geq 2$ and $\varphi(x) = 1$ if $|x| \leq 1$. Set $j_e(x, y) = pr_G[(1 - \varphi(x))j_0(x, y) + \varphi(x)J(x, y)]$. The function defined in this way belongs to $W^{2,\infty}((-3, 3) \times (-1, 1))$, is smooth outside the segment $x = 0$ and equal to 1 identically on this segment. Moreover, it coincides respectively with $j_{\underline{e}}$ and $j_{\bar{e}}$ on $(-3, -2) \times (-1, 1)$ and $(2, 3) \times (-1, 1)$. Hence, j_e has the desired properties. \square

Proposition 4.5 applied to each form ω_e produces a new collection $(j_e \cdot \omega_e, e \in \mathbb{E}^+)$ of forms which, together with $(\omega_v, v \in \mathbb{V})$, determine a unique element $\omega_{\mathbb{G}} \in W^{1,\infty}\Omega_{\mathfrak{g}}^1(U_{\mathbb{G}})$. Let us summarize the properties of $\omega_{\mathbb{G}}$.

Proposition 4.6 *There exists an element $\omega_{\mathbb{G}} \in W^{1,\infty}\Omega_{\mathfrak{g}}^1(U_{\mathbb{G}})$ such that the following conditions hold.*

1. *For each edge $e \in \mathbb{E}^+$, $\langle \omega_{\mathbb{G}}, e \rangle_{\tilde{G}} = \tilde{g}_e$. Moreover, $\omega_{\mathbb{G}}$ vanishes in all directions on e inside $B(\underline{e}, 3R/4) \cup B(\bar{e}, 3R/4)$.*
2. *For each face $F \in \mathbb{F}$, $\omega_{\mathbb{G}}$ is smooth on $U_{\mathbb{G}} \cap F$.*
3. *For each face F and each edge e bounding F , $\omega_{\mathbb{G}}$ is gauge-equivalent in $T_e \cap F$ to $X_{F,v}\lambda$, where v is any vertex on the boundary of F and λ is a smooth 1-form such that $d\lambda = \sigma$.*
4. *For each face F and each vertex v on the boundary of F , $\omega_{\mathbb{G}}$ is equal to $X_{F,v}\lambda$ on $B(v, 3R/4) \cap F$, where λ is a smooth 1-form such that $d\lambda = \sigma$.*

4.5 The principal bundle

In this paragraph, we construct the principal bundle P on which the minimizing connection is going to be defined. For this, we start by proving that there exists a family $(z_e)_{e \in \mathbb{E}^+}$ of elements of Π indexed by \mathbb{E}^+ such that, for each face F , one has

$$z_F = \prod_{e \in \mathbb{E}^+ : L(e^{-1}) = F} z_e. \quad (15)$$

This is a simple consequence of the property 1 of Lemma 1.2. Indeed, the subsets $\{e \in \mathbb{E}^+ : L(e^{-1}) = F\}$ are non-empty and form, as F spans \mathbb{F} , a partition of \mathbb{E}^+ . Hence, the mapping $\Pi^{\mathbb{E}^+} \longrightarrow \Pi^{\mathbb{F}}$ defined by $(z_e)_{e \in \mathbb{E}^+} \mapsto (\prod_{e \in \mathbb{E}^+ : L(e^{-1}) = F} z_e)_{F \in \mathbb{F}}$ is onto.

Let us choose $(z_e)_{e \in \mathbb{E}^+}$ such that (15) holds. Now, for each $z \in \Pi$, let us choose a smooth curve $\zeta_z : (-3, 3) \rightarrow \tilde{G}$ such that $\zeta_z(t) = 1$ if $t \leq -1$ and $\zeta_z(t) = z$ if $t \geq 1$. Pick $e \in \mathbb{E}^+$. Consider as usual the coordinates $(-3, 3) \times (-1, 1)$ on T_e given by Lemma 4.2. Define $\psi_e : U_e \rightarrow G$ by setting $\psi_e(m) = 1$ if $m \in B(\underline{e}, 3R/4) \cup B(\bar{e}, 3R/4)$ and $\psi_e(m) = \pi(\zeta_{z_e}(x))$ if $m \in T_e$ and $m = (x, y)$. Finally, extend ψ_e on $U_{\mathbb{G}}$ by setting $\psi_e(m) = 1$ if $m \notin U_e$. Observe that, if $e \neq e'$, then the subsets $\{m \in U_{\mathbb{G}} : \psi_e(m) \neq 1\}$ and $\{m \in U_{\mathbb{G}} : \psi_{e'}(m) \neq 1\}$ are disjoint, so that $\psi_e \psi_{e'} = \psi_{e'} \psi_e$ everywhere on $U_{\mathbb{G}}$. For each face F , define $\psi_F : U_{\partial F} \rightarrow G$ by setting

$$\psi_F = \prod_{e \in \mathbb{E}^+ : L(e^{-1}) = F} \psi_e.$$

Finally, if F and F' are two faces which share at least one common vertex, then define $\psi_{FF'} : U_{\partial F} \cap U_{\partial F'} \rightarrow G$ by $\psi_{FF'} = \psi_F^{-1} \psi_{F'}$. Since for every pair (F, F') of faces, $U_{\partial F} \cap U_{\partial F'} = U_F \cap U_{F'}$, the collection $(\psi_{FF'})_{F, F' \in \mathbb{F}}$ is a G -valued Čech 1-cochain on M , actually a 1-cocycle.

At first sight, it may seem that this cocycle is actually a coboundary. In fact, the equality $\psi_{FF'} = \psi_F^{-1} \psi_{F'}$ is misleading. It holds on $U_F \cap U_{F'}$ which happens to be equal to $U_{\partial F} \cap U_{\partial F'}$ and the point is that, in general, neither ψ_F nor $\psi_{F'}$ can be extended to smooth or continuous G -valued functions on U_F or $U_{F'}$.

Definition 4.7 *Let P be the principal G -bundle over M defined by the covering $(U_F)_{F \in \mathbb{F}}$ of M and the transition functions $(\psi_{FF'})_{F, F' \in \mathbb{F}}$.*

The bundle P can be described as follows. Consider the disjoint union $\bigsqcup_{F \in \mathbb{F}} (U_F \times G)$. An element of this union is denoted by $(m, g)_F$. Declare $(m, g)_F$ and $(m', g')_{F'}$ to be equivalent if $m = m'$ and $g = \psi_{FF'}(m)g'$. Let \sim denote this equivalence relation. Then P is the manifold $\bigsqcup_{F \in \mathbb{F}} (U_F \times G) / \sim$ on which G acts by right multiplication on the second factor.

In fact, P constructed in this way is endowed with a family of local sections. Indeed for each face F , there is a smooth section s_F of P over U_F which sends each point m to the class of $(m, 1)_F$. If two faces F and F' share at least one common vertex, then s_F and $s_{F'}$ are related on $U_F \cap U_{F'}$ by $s_{F'} = s_F \psi_{FF'}$. In particular, if v is a vertex of \mathbb{G} and if $v \in U_F \cap U_{F'}$, then $s_F(v) = s_{F'}(v)$.

We turn now to the construction of the connection on P .

4.6 Inside the faces

Let $\delta > 0$ be such that, for each face F , the open subset $U_{\partial F}^\delta = \{m \in M : d(m, \partial F) < \delta\}$ of M is contained in $U_{\partial F}$. Such a δ exists because the boundaries of the faces of \mathbb{G} are compact subsets of M . Define, for each face F , $U_F^\delta = F \cup U_{\partial F}^\delta$. Define also $U_{\mathbb{G}}^\delta = \cup_{F \in \mathbb{F}} U_{\partial F}^\delta$. This open subset of M is contained in $U_{\mathbb{G}} = \cup_{F \in \mathbb{F}} U_{\partial F}$.

The domains of the local sections (s_F, U_F^δ) cover M . Hence, according to the remark 1.11, a $W^{1, \infty}$ connection on P is specified by the data of a collection $(\eta_F)_{F \in \mathbb{F}}$, where for each F , $\eta_F \in W^{1, \infty} \Omega_{\mathfrak{g}}^1(U_F^\delta)$ and, on $U_F^\delta \cap U_{F'}^\delta$, $\eta_{F'} = \psi_{FF'} \cdot \eta_F$. In order to construct such a family, let us start by restricting the form $\omega_{\mathbb{G}}$ given by Proposition 4.6 to each one of the open subsets $U_{\partial F}$. In this way, we get a collection $(\omega_{\partial F}^0)_{F \in \mathbb{F}}$ of 1-forms. For each face F , set

$$\omega_{\partial F} = \psi_F \cdot \omega_{\partial F}^0.$$

Observe that, if F and F' are two faces, then, on $U_F \cap U_{F'}$, one has $\omega_{\partial F'} = \psi_{FF'} \cdot \omega_{\partial F}$. Hence, these locally defined 1-forms almost define a connection on P . There remains only to extend each $\omega_{\partial F}$ to a 1-form defined on U_F .

For each face F , let $\lambda_F \in \Omega^1(F)$ be such that $d\lambda_F = \sigma$. Consider the element $X_F = X_{o(F),F}$ of \mathfrak{g} . We apply a gauge transformation to the form $X_F \lambda_F$ in order to extend $\omega_{\partial F}$ inside F . However, as in Proposition 4.5, we get two forms which do not coincide on $F \cap U_{\partial F}$, but on a smaller domain, that we arrange to contain $F \cap U_{\partial F}^\delta$.

Proposition 4.8 *Let F be a face of \mathbb{G} . There exists a smooth G -valued function $j_F : F \rightarrow G$ such that $j_F \cdot (X_F \lambda_F) = \omega_{\partial F}$ on $F \cap U_{\partial F}^\delta$.*

We begin by proving the following result.

Lemma 4.9 *Let F be a face of \mathbb{G} . Let m be a point of $F \cap B(o(F), 3R/4)$. Let l be a smooth loop based at m contained in $F \cap U_{\partial F}$. Then $\langle \omega_{\partial F}, l \rangle_{\tilde{G}} = \widetilde{\exp}(X_F \int_l \lambda_F)$.*

Proof – Let $*\Omega_{\partial F}$ be the unique element of $C^\infty(F \cap U_{\partial F})$ such that the curvature of $\omega_{\partial F}$ is equal to $*\Omega_{\partial F} \sigma$ on this domain.

We begin by treating the case where l is homotopic to a constant loop. Let $L : [0, 1]^2 \rightarrow F \cap U_{\partial F}$ be a smooth homotopy $(s, t) \mapsto L(s, t) = l_s(t)$ such that for each $s \in [0, 1]$, l_s is a smooth loop based at m , l_0 is the constant loop and l_1 is just l . Let us pull $\omega_{\partial F}$ back by this smooth homotopy and work on $[0, 1]^2$. For each $s \in [0, 1]$, set $h_s = \langle \omega_{\partial F}, l_s \rangle_{\tilde{G}} = \langle L^* \omega_{\partial F}, \{s\} \times [0, 1] \rangle_{\tilde{G}}$. Then, according to [6], Theorem 2.2, or Proposition A-1, the mapping $s \mapsto h_s$ satisfies the following differential equation:

$$h_s^{-1} \frac{\partial h_s}{\partial s} = \int_0^1 \langle \omega_{\partial F}, l_s([0, t]) \rangle_{\tilde{G}}^{-1} * \Omega_{\partial F}(l_s(t)) \langle \omega_{\partial F}, l_s([0, t]) \rangle_{\tilde{G}} \Phi^* \sigma_{(s,t)}(\partial_s, \partial_t) dt. \quad (16)$$

We claim that $\langle \omega_{\partial F}, l_s([0, t]) \rangle_{\tilde{G}}^{-1} * \Omega_{\partial F}(l_s(t)) \langle \omega_{\partial F}, l_s([0, t]) \rangle_{\tilde{G}}$ is identically equal to $*\Omega_{\partial F}(m)$. More generally, we claim that, whenever a path c in $F \cap U_{\partial F}$ starts at m and finishes at some point n , then

$$*\Omega_{\partial F}(m) = \langle \omega_{\partial F}, c \rangle_{\tilde{G}}^{-1} * \Omega_{\partial F}(n) \langle \omega_{\partial F}, c \rangle_{\tilde{G}}.$$

Indeed, this relation is true if $\omega_{\partial F}$ is replaced by a connection of the form $X\lambda$, where $X \in \mathfrak{g}$ and λ satisfies $d\lambda = \sigma$. Moreover, this relation is gauge-invariant: if it is true for some connection, it is also satisfied by the image of this connection by any gauge transformation. Finally, the relation is multiplicative: if it holds for two paths which one can concatenate, then it holds for their concatenation. Now the result follows from the fact that any path in $F \cap U_{\partial F}$ can be written as a concatenation of finitely many shorter paths, each of which is contained in a domain where $\omega_{\partial F}$ is gauge-equivalent to a connection of the form $X\lambda$.

Now, (16) implies that

$$h_1 = \exp \left(*\Omega_{\partial F}(m) \int_0^1 \left\{ \int_0^1 L^* \sigma_{(s,t)}(\partial_s, \partial_t) dt \right\} ds \right).$$

On the other hand,

$$\frac{\partial}{\partial s} \int_{l_s} \lambda_F = \int_0^1 \partial_s [(L^* \lambda_F)_{(s,t)}(\partial_t)] dt = \int_0^1 d(L^* \lambda_F)_{(s,t)}(\partial_s, \partial_t) dt = \int_0^1 L^* \sigma_{(s,t)}(\partial_s, \partial_t) dt,$$

because $d_{(s,0)}L(\partial_s) = d_{(s,1)}L(\partial_s) = 0$ for all $s \in [0, 1]$. Finally, since $*\Omega_{\partial F}(m) = X_F$, we get $\langle \omega_{\partial F}, l \rangle_{\tilde{G}} = \widetilde{\text{exp}}(X_F \int_l \lambda_F)$.

Let us drop the assumption that l is homotopic to a constant loop. Let γ be the geodesic segment from m to $o(F)$. Set $w = \gamma \partial F \gamma^{-1}$. Then w generates the fundamental group $\pi_1(U_{\partial F}, m)$, which is isomorphic to \mathbb{Z} . Let r be the unique integer such that lw^{-r} is homotopic to a constant loop. We cannot apply the discussion above to lw^{-r} because it is not contained in $F \cap U_{\partial F}$. So, let $(w_n)_{n \geq 0}$ be a sequence of simple loops based at m such that w_n converges to w and, for each $n \geq 0$, lw_n^{-r} is contained in $F \cap U_{\partial F}$ and homotopic to a constant loop. Then $\langle \omega_{\partial F}, l \rangle_{\tilde{G}} = \langle \omega_{\partial F}, w_n \rangle_{\tilde{G}}^r \langle \omega_{\partial F}, lw_n^{-r} \rangle_{\tilde{G}}$.

On one hand, $\langle \omega_{\partial F}, w_n \rangle_{\tilde{G}}$ converges to $\langle \omega_{\partial F}, w \rangle_{\tilde{G}}$ because w_n converges to w with fixed endpoints. The holonomy of $\omega_{\partial F}$ along ∂F is equal to $h_{\partial F}^{\tilde{G}}(\tilde{g})z_F = \widetilde{\text{exp}}(\sigma(F)X_F)$ and its holonomy along γ commutes to X_F because, on $B(o(F), 3R/4)$, $\omega_{\partial F}$ takes its values in $\mathbb{R}X_F$. So, $\langle \omega_{\partial F}, w \rangle_{\tilde{G}} = \widetilde{\text{exp}}(\sigma(F)X_F)$.

On the other hand, by the discussion of the homotopically trivial case, $\langle \omega_{\partial F}, lw_n^{-r} \rangle_{\tilde{G}} = \widetilde{\text{exp}}(X_F \int_l \lambda_F) \widetilde{\text{exp}}(-rX_F \int_{w_n} \lambda_F)$. Since $d\lambda_F = \sigma$, $\int_{w_n} \lambda_F$ is equal to the area enclosed by w_n . Since w_n converges uniformly to $w = \gamma \partial F \gamma^{-1}$, this area tends to $\sigma(F)$. The result follows. \square

Proof of Proposition 4.8 – Let F be a face of \mathbb{G} . Both forms $\omega_{\partial F}$ and $X_F \lambda_F$ are smooth on $F \cap U_{\partial F}$. We have just proved that they have the same holonomy in \tilde{G} along all smooth loops based at m . Hence, by Lemma 2.8 and Proposition 2.9, there exists a smooth mapping $j_F^0 : F \cap U_{\partial F} \rightarrow \tilde{G}$ such that $j_F^0 \cdot (X_F \lambda_F)$ coincides with $\omega_{\partial F}$ on $F \cap U_{\partial F}$. Consider a diffeomorphism between F and the unit disk in \mathbb{R}^2 with polar coordinates (r, θ) such that $F \cap U_{\mathbb{G}}^\delta \subset \{r > \frac{3}{4}\} \subset \{r > \frac{1}{2}\} \subset F \cap U_{\partial F}$. Consider the restriction of j_F^0 to $\{r > \frac{1}{2}\}$. In order to extend j_F^0 , we start by embedding \tilde{G} in a vector space of matrices. The possible non compactness of \tilde{G} is, as always in this paper, not a problem because \tilde{G} is the direct product of a compact group and a group isomorphic to $(\mathbb{R}^m, +)$ for some $m \geq 0$. We may even assume that, for some $\varepsilon > 0$, and with the same notation as in the proof of Proposition 4.5, there is a smooth mapping $pr_{\tilde{G}} : \tilde{G}^\varepsilon \rightarrow \tilde{G}$ which restricts to the identity on \tilde{G} . Now, since \tilde{G} is simply connected, we can extend j_F^0 by continuity on $D(0, 1)$. Since j_0 is uniformly continuous on $D(0, \frac{7}{8})$, we can also approximate it by a smooth matrix-valued function J defined on $D(0, \frac{7}{8})$ and such that $\|J - j_0\|_\infty < \varepsilon$. Finally, let $\varphi : [0, 1] \rightarrow [0, 1]$ be a smooth function such that $\varphi(r) = 1$ if $r \leq \frac{1}{2}$ and $\varphi(r) = 0$ if $r \geq \frac{3}{4}$. We define $j_F : D(0, 1) \rightarrow G$ in polar coordinates by $j_F(r, \theta) = \pi[pr_{\tilde{G}}((1 - \varphi(r))j_0(r, \theta) + \varphi(r)J(r, \theta))]$. The forms $j_F \cdot (X_F \lambda_F)$ and $\omega_{\partial F}$ coincide on the domain of equation $\{r > \frac{3}{4}\}$, which contains $F \cap U_{\mathbb{G}}^\delta$. \square

We can finish the proof of the existence of a minimizing connection.

Proof of Proposition 4.1 – Let P be the principal G -bundle over M constructed in Section 4.5. It satisfies $\mathfrak{o}(P) = \prod_{e \in \mathbb{E}^+} z_e = \prod_{F \in \mathbb{F}} z_F = z$ (see Appendix A of [8]).

For each face F , let us call ω_F the element of $W^{1, \infty} \Omega_{\mathfrak{g}}^1(U_F^\delta)$ which is equal to $j_F \cdot (X_F \lambda_F)$ on F and to $\omega_{\partial F}$ on $U_{\partial F}^\delta$. If F and F' are two faces which share at least a common vertex, then, on $U_F^\delta \cap U_{F'}^\delta = U_{\partial F}^\delta \cap U_{\partial F'}^\delta$, we have $\omega_{F'} = \omega_{\partial F'} = \psi_{FF'} \cdot \omega_{\partial F} = \psi_{FF'} \cdot \omega_F$. Hence, the forms $(\omega_F)_{F \in \mathbb{F}}$ determine a connection on P . Let ω denote this connection. By construction, ω belongs to $W^{1, \infty} \mathcal{A}(P)$ and it is smooth outside $\cup_{e \in \mathbb{E}} e$.

Let v be a vertex of \mathbb{G} . We have observed at the end of Section 4.5 that the local sections

s_F determine without ambiguity a point $s_F(v)$, which we denote by $p_v \in P_v$. Let us compute the holonomy of ω along the edges of \mathbb{G} with respect to the reference points p_v .

Let $(\tau_c, c \in PM)$ be the holonomy induced by ω . Let e be an edge of \mathbb{E}^+ . By definition of ω , we have for each face F the equality $\omega_{s_F} = \omega_F$. In particular, $\omega_{s_{L(e)}} = \omega_{L(e)}$. Hence, $\tau_e(p_e) = p_e \langle \omega_{L(e)}, e \rangle$. Since $\psi_{L(e)}$ is identically equal to 1 in a neighbourhood of e , $\langle \omega_{L(e)}, e \rangle = \langle \omega_{\mathbb{G}}, e \rangle$. By Proposition 4.6, this is equal to g_e . Since this holds for every $e \in \mathbb{E}^+$, ω agrees with $(g_e)_{e \in \mathbb{E}^+}$ on \mathbb{E} .

Finally, let F be a face of \mathbb{G} . Inside this face, ω_F is gauge-equivalent to $X_F \lambda_F$, so that its curvature is equal to $X_F \sigma$. Hence, $S_F(\omega) = \|X_F\|^2 \sigma(F)$. The assertion on $S_F(\omega)$ follows now from the identity $\sigma(F) \|X_F\| = \tilde{\rho}(h_{\partial F}^{\tilde{G}}(\tilde{g})z_F)$. \square

Appendix: A proof of the energy inequality

In this appendix, we give a proof of the energy inequality (Proposition 2.4), which consists basically in applying Cauchy-Schwarz inequality in the right context. The proof of a more general instance of this result can be found in [16]. However, we are not aware of a such a compact proof as the one we give here.

Let D be the closed unit disk in \mathbb{R}^2 endowed with a smooth volume form σ . Pick ω in $\Omega_{\mathfrak{g}}^1(D)$. Let $\Omega \in \Omega_{\mathfrak{g}}^2(D)$ be its curvature and $*\Omega$ the unique smooth \mathfrak{g} -valued function on D such that $\Omega = *\Omega\sigma$.

For all $s, t \in [0, 1]$, set

$$\gamma(s, t) = (1 - s + s \cos(2\pi t), s \sin(2\pi t)) \in \mathbb{R}^2.$$

For every $s \in [0, 1]$, $t \mapsto \gamma(s, t)$ is a smooth loop based at the point $(1, 0)$. The mapping γ realizes a smooth homotopy between the constant loop at $(1, 0)$ and the boundary of D starting at $(1, 0)$ with the usual orientation.

Set $h(s, t) = \langle \omega, \gamma(s, \cdot) \rangle$. By standard results on ordinary differential equations, h is a smooth mapping from $[0, 1]^2$ to G . Let us denote by h_s and h_t the partial derivatives of h with respect to s and t . Similarly, let γ_s and γ_t denote the partial derivatives of γ with respect to s and t .

Proposition A-1 For all $s \in [0, 1]$,

$$h(s, 1)^{-1} h_s(s, 1) = \int_0^1 h(s, t)^{-1} * \Omega(\gamma_t(s, t), \gamma_s(s, t)) h(s, t) dt.$$

Proof – For every $s \in [0, 1]$, consider the vector field X_s defined on $[0, 1] \times G$ by

$$X_s(u, g) = (1, -\omega(\gamma_t(s, u))g).$$

Let $(\Phi_s^t)_{t \geq 0}$ denote the flow of X_s . The group G acts on $[0, 1] \times G$ by multiplication on the right on the second factor. Since X_s is G -invariant, its flow is also invariant. By definition of $h(s, t)$, one has

$$\Phi_s^t(u, g) = (u + t, h(s, u + t)h(s, u)^{-1}g). \quad (\text{A-17})$$

In particular, $\Phi_s^1(0, 1) = (1, h(s, 1))$. Since X_s depends smoothly on s , a classical result (see [4], Thm. B.3 for example) asserts that

$$\frac{\partial}{\partial s} \Phi_s^1(0, 1) = \int_0^1 d_{\Phi_s^t(0,1)} \Phi_s^{1-t} \left[\frac{\partial X_s}{\partial s}(\Phi_s^t(0, 1)) \right] dt.$$

The term inside the brackets is equal to $(0, -\frac{\partial}{\partial s} [\omega(\gamma_t(s, t))] h(s, t))$. By differentiating (A-17), one gets

$$\begin{aligned} d_{\Phi_s^t(0,1)} \Phi_s^{1-t} \left[(0, -\frac{\partial}{\partial s} [\omega(\gamma_t(s, t))] h(s, t)) \right] = \\ - (0, h(s, 1) h(s, t)^{-1} \frac{\partial}{\partial s} [\omega(\gamma_t(s, t))] h(s, t)). \end{aligned}$$

Hence,

$$h(s, 1)^{-1} h_s(s, 1) = - \int_0^1 h(s, t)^{-1} \frac{\partial}{\partial s} [\omega(\gamma_t(s, t))] h(s, t) dt. \quad (\text{A-18})$$

Since the vector fields $\partial_t \gamma$ and $\partial_s \gamma$ commute, one has at every point

$$-\frac{\partial}{\partial s} [\omega(\gamma_t)] = d\omega(\gamma_t, \gamma_s) - \frac{\partial}{\partial t} [\omega(\gamma_s)]. \quad (\text{A-19})$$

Let us compute $-\int_0^1 h(s, t)^{-1} \frac{\partial}{\partial t} [\omega(\gamma_s)] h(s, t) dt$ by integration by parts. The boundary terms vanish, because the loops $t \mapsto \gamma(s, t)$ share the same basepoint. There remains

$$- \int_0^1 h(s, t)^{-1} \frac{\partial}{\partial t} [\omega(\gamma_s)] h(s, t) dt = \int_0^1 h(s, t)^{-1} [\omega(\gamma_t), \omega(\gamma_s)] h(s, t) dt, \quad (\text{A-20})$$

where the last bracket is the Lie bracket of \mathfrak{g} . Combining (A-18), (A-19) and (A-20), we find what we want, that is,

$$h(s, 1)^{-1} h_s(s, 1) = \int_0^1 h(s, t)^{-1} * \Omega(\gamma_t(s, t), \gamma_s(s, t)) h(s, t) dt.$$

□

Corollary A-2 (Energy inequality) *The following inequality holds:*

$$\rho(\langle \omega, \partial D \rangle)^2 \leq S_D(\omega) \sigma(D).$$

Proof – Since conjugation preserves the norm in \mathfrak{g} ,

$$\|h_s(s, 1)\| \leq \int_0^1 \|\Omega(\gamma_t(s, t), \gamma_s(s, t))\| dt.$$

By Cauchy-Schwarz inequality,

$$\begin{aligned} \left| \int_0^1 \|\Omega(\gamma_t(s, t), \gamma_s(s, t))\| dt \right|^2 &= \left| \int_0^1 \| * \Omega(\gamma(s, t)) \| |\sigma(\gamma_t(s, t), \gamma_s(s, t))| dt \right|^2 \\ &\leq \int_0^1 \| * \Omega(\gamma(s, t)) \|^2 |\sigma(\gamma_t(s, t), \gamma_s(s, t))| dt \int_0^1 |\sigma(\gamma_t(s, t), \gamma_s(s, t))| dt. \end{aligned}$$

Set $A(r) = \int_0^r ds \int_0^1 |\sigma(\gamma_t(s, t), \gamma_s(s, t))| dt$. It is the area enclosed by the path $\gamma(r, \cdot)$. The function A is a diffeomorphism of $[0, 1]$ onto $[0, \sigma(D)]$. Reparametrize $s \mapsto h(s, 1)$ by setting

$$k(u) = h(A^{-1}(\sigma(D)u), 1).$$

The path $u \mapsto k(u)$ is a path in G from 1 to $h(1, 1)$. Hence,

$$\rho(h(1, 1))^2 \leq \ell(k)^2 = \left| \int_0^1 \|\dot{k}(u)\| du \right|^2 \leq \int_0^1 \|\dot{k}(u)\|^2 du.$$

The last term is the energy of the path k , which can be estimated as follows.

$$\begin{aligned} \int_0^1 \|\dot{k}(u)\|^2 du &= \int_0^1 \|h_s(A^{-1}(\sigma(D)u), 1)\|^2 \frac{\sigma(D)^2}{A'(A^{-1}(\sigma(D)u))^2} du \\ &= \int_0^1 \|h_s(s, 1)\|^2 \frac{\sigma(D)}{A'(s)} ds \\ &\leq \sigma(D) \int_0^1 \left(\int_0^1 \|\ast \Omega(\gamma(s, t))\|^2 |\sigma(\gamma_t(s, t), \gamma_s(s, t))| dt \right) ds \\ &= \sigma(D) \int_D \|\ast \Omega\|^2 \sigma \\ &= \sigma(D) S_D(\omega). \end{aligned}$$

This implies the result because $h(1, 1) = \langle \omega, \partial D \rangle$. □

References

- [1] Robert A. Adams. *Sobolev spaces*. Academic Press, New York-London, 1975. Pure and Applied Mathematics, Vol. 65.
- [2] Amir Dembo and Ofer Zeitouni. *Large deviations techniques and applications*. Jones and Bartlett Publishers, Boston, MA, 1993.
- [3] Bruce K. Driver. Two-dimensional Euclidean quantized Yang-Mills fields. In *Probability models in mathematical physics (Colorado Springs, CO, 1990)*, pages 21–36. World Sci. Publishing, Teaneck, NJ, 1991.
- [4] J. J. Duistermaat and J. A. C. Kolk. *Lie groups*. Universitext. Springer-Verlag, Berlin, 2000.
- [5] Alfred Gray. *Tubes*. Addison-Wesley Publishing Company Advanced Book Program, Redwood City, CA, 1990.
- [6] Leonard Gross. A Poincaré lemma for connection forms. *J. Funct. Anal.*, 63(1):1–46, 1985.
- [7] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of differential geometry. Vol. I*. John Wiley & Sons Inc., New York, 1996. Reprint of the 1963 original.

- [8] H. Blaine Lawson, Jr. and Marie-Louise Michelsohn. *Spin geometry*, volume 38 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1989.
- [9] Thierry Lévy. Yang-Mills measure on compact surfaces. *Mem. Amer. Math. Soc.*, 166(790):xiv+122, 2003.
- [10] Thierry Lévy. Discrete and continuous Yang-Mills measure for non-trivial bundles over compact surfaces. *Preprint*, math-ph/0501014, 2005.
- [11] Peter Li and Shing-Tung Yau. On the parabolic kernel of the Schrödinger operator. *Acta Math.*, 156(3-4):153–201, 1986.
- [12] A. A. Migdal. Recursion equations in gauge field theories. *Sov. Phys. JETP*, 42(3):413–418, 1975.
- [13] Shigeyuki Morita. *Geometry of differential forms*, volume 201 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 2001.
- [14] James R. Norris. Heat kernel asymptotics and the distance function in Lipschitz Riemannian manifolds. *Acta Math.*, 179(1):79–103, 1997.
- [15] Ambar Sengupta. Gauge theory on compact surfaces. *Mem. Amer. Math. Soc.*, 126(600):viii+85, 1997.
- [16] Ambar Sengupta. A Yang-Mills inequality for compact surfaces. *Infin. Dimens. Anal. Quantum Probab. Relat. Top.*, 1(1):1–16, 1998.
- [17] Norman Steenrod. *The topology of fibre bundles*. Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1999.
- [18] Karen K. Uhlenbeck. Connections with L^p bounds on curvature. *Comm. Math. Phys.*, 83(1):31–42, 1982.
- [19] S. R. S. Varadhan. Diffusion processes in a small time interval. *Comm. Pure Appl. Math.*, 20:659–685, 1967.
- [20] S. R. S. Varadhan. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Comm. Pure Appl. Math.*, 20:431–455, 1967.
- [21] N. Th. Varopoulos, L. Saloff-Coste, and T. Coulhon. *Analysis and geometry on groups*. Cambridge University Press, Cambridge, 1992.
- [22] Katrin Wehrheim. *Uhlenbeck compactness*. EMS Series of Lectures in Mathematics. European Mathematical Society (EMS), Zürich, 2004.
- [23] Edward Witten. On quantum gauge theories in two dimensions. *Comm. Math. Phys.*, 141(1):153–209, 1991.