Profinite complexes of curves, their automorphisms and anabelian properties of moduli stacks of curves

M.Boggi and P.Lochak

Abstract

Let $\mathcal{M}_{g,[n]}$, for $2g-2+n>0$, be the Deligne-Mumford (briefly D-M) moduli stack of smooth curves of genus g labeled by n unordered distinct points. The main result of the paper is that a finite, connected étale cover $\mathcal{M}^\lambda$ of $\mathcal{M}_{g,[n]}$, defined over a sub-p-adic field k, is "almost" anabelian in the sense conjectured by Grothendieck for curves and their moduli spaces.

The precise result is the following. Let $\pi_1(\mathcal{M}^\lambda)$ be the geometric algebraic fundamental group of $\mathcal{M}^\lambda$ and let $Out^*(\pi_1(\mathcal{M}^\lambda))$ be the group of its exterior automorphisms which preserve the conjugacy classes of elements corresponding to simple loops around the Deligne-Mumford boundary of $\mathcal{M}^\lambda$ (this is the "$*$-condition" motivating the "almost" above). Let us denote by $Out^*_{G_k}(\pi_1(\mathcal{M}^\lambda))$ the subgroup consisting of elements which commute with the natural action of the Galois group G_k of $\overline{k}$ over k. Let us assume, moreover, that the generic point of the D-M stack $\mathcal{M}^\lambda$ has trivial automorphisms group. Then, there is a natural isomorphism:

$$Aut_k(\mathcal{M}^\lambda) \cong Out^*_{G_k}(\pi_1(\mathcal{M}^\lambda)).$$

This partially extends to moduli spaces of curves the anabelian properties proved by Mochizuki for hyperbolic curves over sub-p-adic fields (see [Mo]).

AMS Classification: 11R32, 14D22, 57M99.

1. Introduction

The original motivation of this paper was to provide a geometric counterpart to the rather technical and sometimes dry results contained in [B1]. In particular, profinite curve complexes, whose theory was initiated there, seemed an adequate tool to track down some of the anabelian properties which were conjectured to hold for moduli spaces of smooth curves. We include in this set any representable finite connected étale cover $\mathcal{M}^\lambda$ of the D-M stack $\mathcal{M}_{g,[n]}$ of smooth curves labeled by n unordered distinct points, for $2g-2+n>0$, i.e., according to the terminology adopted in [B1], any *level structure* over $\mathcal{M}_{g,[n]}$. The level structure is *Galois* if $\mathcal{M}^\lambda \to \mathcal{M}_{g,[n]}$ is a Galois cover. In particular, observe that the stack $\mathcal{M}_{g,n}$ of smooth curves labeled by n ordered distinct points is a Galois level structure over $\mathcal{M}_{g,[n]}$ with Galois group the symmetric group $\mathcal{S}_n$. The brackets ($[n]$) are used (and will be used) to make clear that the points are unordered.

Since there is not yet a well established definition of "anabelian", let us make more precise what we mean. After the groundbreaking work of Mochizuki (see [Mo]), who proved the anabelian conjecture of Grothendieck for hyperbolic curves (formulated in [G]) and actually something more, it seems natural to propose the following definition.

Let k be a sub-p-adic field, i.e. a subfield of a finitely generated extension of $\mathbb{Q}_p$. Consider the functor Φ which associates to a smooth irreducible X D-M stack of finite type over k its *geometric* algebraic fundamental group, i.e. the algebraic fundamental group of $X_{\overline{k}} := X \otimes \overline{k}$, endowed with the natural action of the Galois group G_k. To a k-morphism $Y \to X$ of schemes of finite type over k, is then associated the exterior G_k-equivariant homomorphism of geometric fundamental groups. Let $Hom_k^{dom}(Y, X)$ denote the set of representable dominant k-morphisms and $Hom_{G_k}^{op}(\pi_1(Y_{\overline{k}}), \pi_1(X_{\overline{k}}))^{ext}$ the set of exterior G_k-equivariant open homomorphisms between their fundamental groups.

A smooth D-M stack X of finite type over k is *anabelian* if, for every smooth D-M stack Y of finite type over k, the functor Φ establishes a bijection:

$$Hom_k^{dom}(Y, X) \cong Hom_{G_k}^{op}(\pi_1(Y_{\overline{k}}), \pi_1(X_{\overline{k}}))^{ext}.$$

The first test of anabelianity is then the verification of the above property for $Y = X$. In the only case in which anabelianity has been fully proved, i.e. hyperbolic curves, this test actually proved to be decisive. Let us remark that a dominant endomorphism of a hyperbolic curve is necessarily an automorphism. The same holds for an endomorphism of a level structure $\mathcal{M}^\lambda$ (by Royden theorem). Likewise, an open endomorphism of the fundamental group of a hyperbolic curve is necessarily an automorphism. The same holds for the topological fundamental group of a level structure (by Ivanov's results in [2]) and presumably also for its geometric algebraic fundamental group. In the case of hyperbolic curves, but it is reasonable to guess this holds for all anabelian varieties, the test on endomorphisms reduces to the form:

$$Aut_k(X) \cong Out_{G_k}(\pi_1(X_{\overline{k}})).$$

The original purpose of the paper was to prove this identity for the moduli stacks $\mathcal{M}^\lambda$. Of course, one cannot expect the above identity to hold when the generic point of X has a non-trivial automorphism. In fact, let A be the group of generic automorphisms of the stack X and denote by $X/\!\!/A$ the stack obtained from X rigidifying with respect to A. By the theory of fundamental groups of stacks (see [No]), there is a short exact sequence of fundamental groups:

$$1 \to A \to \pi_1(X_{\overline{k}}) \to \pi_1(X_{\overline{k}}/\!\!/A) \to 1.$$

Therefore, a non-trivial $f \in A$ determines a stack automorphism of X acting by conjugation on the fundamental group. In particular, this is what happens with the moduli stacks $\mathcal{M}_2$, $\mathcal{M}_{1,1}$, $\mathcal{M}_{1,[2]}$ and $\mathcal{M}_{0,[4]}$.

A series of remarks are in order at this point. For any field k of characteristic zero, the algebraic fundamental group of $X_{\overline{k}}$ is isomorphic to the profinite completion of the topological fundamental group of $X_{\mathbb{C}}$ (see [N2]). This has already some interesting consequences. Both for X an hyperbolic curve or a level structure $\mathcal{M}^\lambda$ over $\mathcal{M}_{g,[n]}$, whose generic point has trivial automorphisms group, from classical results (Hurwitz and Royden theorems, respectively), it follows that the functor Φ induces a monomorphism:

$$\Phi : Aut_k(X) \hookrightarrow Out_{G_k}(\pi_1(X_{\overline{k}})).$$

So, the above test of anabelianity simply demands whether it is possible to reconstruct an automorphism of X from a given Galois equivariant outer automorphism of $\pi_1(X_{\overline{k}})$. In case X is not compact, for instance when X is a punctured curve or any moduli space of smooth curves, there is a necessary geometric condition the given automorphism must satisfy in order this to be possible. Let $\overline{X}$ be a toroidal compactification of X such that any automorphism of X extends to $\overline{X}$ (for a moduli stack of smooth curves, this is provided by the Deligne- Mumford compactification). Then, it is clear that $f \in Im\Phi$ only if the automorphism f preserve the set of conjugacy classes in $\pi_1(X_{\overline{k}})$ of elements corresponding to small loops in $X_{\mathbb{C}}$ around the boundary of $\overline{X}_{\mathbb{C}}$. This is the so-called *inertia preserving* condition (or $*$-condition). In the sequel, Aut^* (resp. Out^*) will denote the subgroup of elements of Aut (resp. Out) satisfying the $*$-condition.

An argument by Nakamura (see [N]) shows that the $*$-condition is satisfyed in the case of a punctured hyperbolic curve. The argument is based on the Weil-Deligne theory of weights for the ℓ-adic cohomology of an algebraic variety (actually, it only considers first cohomology groups). The same argument fails to work for moduli stacks of curves. The reason is that, in general, small loops in $\mathcal{M}^\lambda$ around the Deligne-Mumford boundary determine in the first homology group torsion elements and hence cannot be detected by ℓ-adic cohomology.

At the moment, we are not able to fill this gap even though there are serious reasons, some will be exposed below, to guess that the $*$-condition is verified and not only by the Galois equivariant elements. The result we are able to prove here is that, for any level structure $\mathcal{M}^\lambda$ over $\mathcal{M}_{g,[n]}$, defined over a sub-p-adic field k, such that its generic point has trivial automorphisms group the functor Φ induces an isomorphism:

$$Aut_k(\mathcal{M}^\lambda) \cong Out^*_{G_k}(\pi_1(\mathcal{M}^\lambda_{\overline{k}})).$$

In case the level structure $\mathcal{M}^\lambda$ over $\mathcal{M}_{g,[n]}$ has a generic non-trivial automorphism, the homorphism of groups induced by Φ is still surjective but not anymore injective. In fact, if this is the case, $Out^*_{G_k}(\pi_1(\mathcal{M}^\lambda_{\overline{k}}))$ is isomorphic to the automorphism group of the stack obtained from $\mathcal{M}^\lambda$ by rigidification, i.e. stripping off the generic automorphisms group (see [R] for a detailed description of this procedure). The latter case may occur only if either $(g,n) \in \{(1,1),(1,2),(2,0)\}$ and Γ^λ contains the hyperelliptic involution (which, in these cases, is both the generator of the center of $\Gamma_{g,[n]}$ and of the generic automorphism of $\mathcal{M}_{g,[n]}$), or $(g,n) = (0,4)$ and Γ^λ intersect non-trivially the Klein subgroup of $\Gamma_{0,[4]}$ (which is isomorphic to $\mathbb{Z}/2 \times \mathbb{Z}/2$).

The proof is via Teichmüller theory. Let us then briefly recall some well known facts. From the point of view of moduli spaces of curves, Teichmüller theory is the study of the geometry of the universal cover, usually denoted $\mathcal{T}_{g,n}$ and called the Teichmüller space, and of the topological fundamental group (the Teichmüller group) of the topological stack underlying the moduli stack $\mathcal{M}_{g,[n]} \otimes \mathbb{C}$. By Teichmüller theory, the latter group is isomorphic to the orientation preserving mapping class group of the differential n-punctured, genus g Riemann surface $S_{g,n}$. In order to have a natural identification of the two groups, it is enough to fix a conformal structure on $S_{g,n}$, i.e. a base point on the Teichmüller space $\mathcal{T}_{g,n}$, compatible with the choice of base point on $\mathcal{M}_{g,n}$. The mapping class group of $S_{g,n}$ (also called Teichmüller modular group) is usually denoted by $\Gamma_{g,[n]}$ while the notation $\Gamma_{g,n}$ is reserved to its normal subgroup corresponding to the $\mathcal{S}_n$-cover $\mathcal{M}_{g,n} \to \mathcal{M}_{g,[n]}$.

A standard set of generators for $\Gamma_{g,[n]}$ is provided by all Dehn twists plus the half Dehn twists along simple curves bounding a disc containing two punctures of $S_{g,n}$ (the latter are usually called *braids*). Geometrically, in $\pi_1^{top}(\mathcal{M}_{g,[n]})$ they all correspond to small loops around the D-M boundary.

As usual, we identify the topological fundamental group of a level structure $\mathcal{M}^\lambda$ with the corresponding subgroup Γ^λ of $\Gamma_{g,[n]}$ and call it *a level*. Its profinite completion is denoted by $\hat{\Gamma}^\lambda$. Small loops around the D-M boundary of $\mathcal{M}^\lambda$ then correspond to suitable powers of Dehn twists or braids. So, the $*$-condition stated above translates into the condition that an element of $Aut(\hat{\Gamma}^\lambda)$ must preserve the set of conjugacy classes of powers of Dehn twists and braids. From the classification of mapping classes in the Teichmüller group, it follows that this condition is superfluous in the discrete case, since it is satisfied by all automorphisms of the groups Γ^λ. This is why we expect the $*$-condition to be superfluous in the profinite case as well. Unfortunately, the lack of any satisfactory classification of elements of $\hat{\Gamma}_{g,[n]}$ jeopardized all efforts in this direction.

As mentioned above, the most important tools at our disposal for the proof are profinite curve complexes. The interest in these objects go certainly beyond the purposes of this paper. In fact, they represent a natural generalization of "dessin d'enfants" and have already revealed to be very useful in the study of the Grothendieck-Teichmüller group as the results announced in [Lo] show.

Let us introduce first the discrete curve complexes $C(S_{g,n})$ and $C_P(S_{g,n})$ (they will be treated more extensively in Section 2).

The complex of curves $C(S_{g,n})$ is defined to be the simplicial complex whose simplices are given by sets of distinct, non-trivial, isotopy classes of circles on $S_{g,n}$, such that they admit a set of disjoint representatives none of them bounding a disc with a single puncture. It is a simplicial complex of dimension $3g-3+n-1$ (the modular dimension minus one).

The pants complex $C_P(S_{g,n})$ is the two dimensional complex whose vertices are given by the simplices of highest dimension in $C(S_{g,n})$ (corresponding to pants decomposition of $S_{g,n}$). Given two vertices $\underline{s}, \underline{s}' \in C_P(S_{g,n})$, they are connected by an edge if and only if $\underline{s}$ and $\underline{s}'$ have $3g-3+n-1$ curves in common, so that up to relabeling (and of course isotopy) $s_i = s'_i$, $i = 1, \ldots, 3g-3+n-1$, whereas s_0 and s'_0 differ by an *elementary move*, which means the following. Cutting $S_{g,n}$ along the s_i's, $i > 0$, there remains a surface Σ of modular dimension 1, so Σ is of type $(1,1)$ or $(0,4)$. Then s_0 and s'_0, which are supported on Σ, should intersect in a minimal way, that is they should have intersection number 1 in the first case, and 2 in the second case.

In a sense, their algebraic analogue are more natural. Let $C^\lambda(S_{g,n})$ be the nerve of the Deligne-Mumford boundary of the level structure $\mathcal{M}^\lambda$ (taken for technical reason in the category of simplicial sets). Then, in the category of simplicial sets, $C^\lambda(S_{g,n})$ realizes the quotient of the complex of curves $C(S_{g,n})$ by the natural action of the fundamental group Γ^λ of the level structure $\mathcal{M}^\lambda$. The family of simplicial finite sets $\{C^\lambda(S_{g,n})\}$, for Γ^λ varying in the tower of all levels of $\Gamma_{g,[n]}$, forms an inverse system and its inverse limit $\hat{C}(S_{g,n})$, taken in the category of simplicial profinite sets, is what we call the *profinite complex of curves*.

The quotient $C_P^\lambda(S_{g,n}) := C_P(S_{g,n})/\Gamma^\lambda$, for Γ^λ containing an abelian level of order ≥ 3, may be taken in the category of simplicial complexes and has an even more natural geometric interpretation. Let $\mathcal{F}^\lambda$ be the locus in $\overline{\mathcal{M}}^\lambda$ of points parametrizing curves with at least $3g-3+n-1$ singular points. Then, $\mathcal{F}^\lambda \otimes \mathbb{C}$ is itself a stable curve whose irreducible components are smooth modular curves intersecting precisely in the vertices of the natural triangulations of which they are endowed as covers of $\overline{\mathcal{M}}_{0,4} \cong \mathbb{P}^1$ ramifying only above $\{0, 1, \infty\}$. Thus, if we give to each edge of

$C_P^\lambda(S_{g,n})$ lenght one, it gets identified with the above triangulation of $\mathcal{F}^\lambda \otimes \mathbb{C}$. *The profinite pants complex* $\hat{C}_P(S_{g,n})$ is then defined as the inverse limit of all the $C_P^\lambda(S_{g,n})$.

Let $Aut(\hat{C}(S_{g,n}))$ and $Aut(\hat{C}_P(S_{g,n}))$ denote the respective groups of continuous simplicial automorphisms. There are natural inclusions $Aut^*(\hat{\Gamma}^\lambda) \leq Aut(\hat{C}(S_{g,n}))$ and $Aut(\hat{C}_P(S_{g,n})) \leq Aut(\hat{C}(S_{g,n}))$ (Corollary 4.9 and Proposition 4.11, respectively). The key technical result is then the following. Let $Aut^+(\hat{C}_P(S_{g,n}))$ be the index two subgroup of orientation preserving automorphisms of $Aut(\hat{C}_P(S_{g,n}))$. For $2g-2+n>0$ and $(g,n) \neq (1,2)$ (Theorem 4.15, however, takes care of this case as well), there is a natural isomorphism:

$$Aut^+(\hat{C}_P(S_{g,n})) \cong Inn(\hat{\Gamma}_{g,[n]}).$$

Let $Aut^*_{G_k}(\hat{\Gamma}^\lambda)$ be the subgroup of those automorphisms which, modulo inner automorphisms, commute with the Galois action. Let k' be a finite extension of k such that all automorphisms of $\mathcal{M}^\lambda_{\bar{k}}$ are already defined over k'. Let us assume, moreover, that the level Γ^λ is *Galois*, i.e. Γ^λ is a normal subgroup of $\Gamma_{g,[n]}$. The next move then is to prove that, for $(g,n) \neq (1,2)$, there is a natural isomorphism:

$$Aut^*_{G_{k'}}(\hat{\Gamma}^\lambda) \cong Aut^+(\hat{C}_P(S_{g,n})).$$

For $(g,n) = (1,2)$, a similar result holds but further restrictions must be imposed on $Aut(\hat{C}_P(S_{g,n}))$ (like in Theorem 4.15) .

The proof of this isomorphism start with the remark that $Aut^*(\hat{\Gamma}^\lambda)$ already acts on the vertices of $\hat{C}_P(S_{g,n})$. So, one needs to prove that Galois compatible automorphisms also preserve the edges of $\hat{C}_P(S_{g,n})$. Modulo inner automorphisms, it is possible to assume that the given automorphism preserve some $3g-3+n-2$-simplex in $\hat{C}(S_{g,n})$. In this way, the problem is reduced to the case of modular dimension one, i.e. of hyperbolic curves, which is dealt thanks to Mochizuki's Theorem.

The composition of the two above isomorphism, for all $2g-2+n>0$, gives then a natural isomorphism:

$$Aut^*_{G_{k'}}(\hat{\Gamma}^\lambda) \cong Inn(\hat{\Gamma}_{g,[n]}).$$

Finally, for a Galois level $\Gamma^\lambda \leq \Gamma_{g,[n]}$, we get:

$$Out^*_{G_{k'}}(\hat{\Gamma}^\lambda) \cong \hat{\Gamma}_{g,[n]}/(\hat{\Gamma}^\lambda \cdot Z),$$

where Z denotes the center of $\hat{\Gamma}_{g,[n]}$. Now, the finite group $\hat{\Gamma}_{g,[n]}/(\hat{\Gamma}^\lambda \cdot Z)$ is naturally isomorphic to the automorphism group of the rigidified stack associated to $\mathcal{M}^\lambda_{\bar{k}}$. By descent, the above isomorphism implies the anabelian claim we made for any sub-p-adic field of definition k for $\mathcal{M}^\lambda$ and also for the level structure which are not Galois.

Prior to the results of the present paper, only the case $g=0$ had been studied, especially in [N] and more recently in [MT]; see also [HS] for a result with a different flavor involving the Grothendieck-Teichmüller group, and Proposition 4.14 below. The result in [N] is proved in a pronilpotent setting but, as noticed in [IN], it can readily be transfered to the (full) profinite case, using the anabelian theorem for affine hyperbolic curves due to A.Tamagawa. Then one gets (see [IN]) the statement above for $g=0$ and $k=\mathbb{Q}$ or more generally a number field (nowadays, this could be jazzed up to sub-p-adic fields using [Mo]) but *without* the inertia preserving condition, that is: $Out_{G_{\mathbb{Q}}}(\hat{\Gamma}_{0,[n]})$ is trivial for any n (> 2). It was shown indeed in [N] that in genus 0, any Galois invariant automorphism is inertia preserving. It should be added that the case $g = 0$ is really special from the anabelian viewpoint. Indeed the schemes $\mathcal{M}_{0,n}$ (as well as the stacks $\mathcal{M}_{0,[n]}$) are "anabelian" according to any sensible definition of that word. This is because $\mathcal{M}_{0,n}$ is built up by iterating the natural fibration $\mathcal{M}_{0,k+1} \to \mathcal{M}_{0,k}$, whose fibre is an hyperbolic curve (here a marked projective line). So the affine scheme $\mathcal{M}_{0,n}$ is globally an "Artin good neighborhood" and these are in some sense the paradigm for higher dimensional anabelian varieties. This iterative structure (which is group theoretically elucidated in [MT]) allows for a dimensional induction. We also use that kind of induction in the case of higher genus, but in a completely different way, taking advantage of the local structure of complexes of curves.

Acknowledgments: It is a pleasure to thank Leila Schneps for many useful conversations around the themes discussed in the paper, and for her help in proving Proposition 4.14. The second author (P.L.) warmly thanks the Mathematics Department of the University of Tel Aviv for a very interesting and pleasant stay during which the redaction of the paper was completed.

2. Discrete complexes of curves and their automorphisms

In this section, after recalling the necessary definitions, we will review the results on automorphisms of (discrete) complexes of curves with an eye on the profinite setting. Actually we also reprove the results of [M] in a different way, which will enable us to extend them to the profinite case in §4. They will appear as a consequence of a result (Theorem 2.10 below) which itself will be extended to the profinite case in §4 and may be of independent interest. Clearly the setting in this section is topological, occasionally complex analytic.

Let us first fix some classical notation. Concerning curves (and topological surfaces), we start with a finite hyperbolic type, that is a pair (g, n) of positive integers with $2g - 2 + n > 0$. Given such a type, we let $S \cong S_{g,n}$ denote a – unique up to diffeomorphism – differentiable surface of genus g with n punctures, where the points are not ordered. Its Euler characteristic is of course $\chi(S_{g,n}) = 2 - 2g - n < 0$. We write $\mathcal{T}(S)$, $\mathcal{M}(S)$, $Mod(S)$ for the Teichmüller space, moduli space and (extended) mapping class group attached to S. The Teichmüller space $\mathcal{T}(S)$ is noncanonically identified with the Teichmüller space $\mathcal{T}_{g,n}$ previously defined. In particular $\mathcal{M}(S)$ is – again noncanonically – identified with $\mathcal{M}_{g,[n]}$, the moduli stack of curves of the given type, labeled with n unordered distinct points, which we regard either as a complex analytic stack or as an algebraic stack according to the context. It has dimension $d(S) = d_{g,n} = 3g - 3 + n$, which we call the modular dimension of S (or of the given type). We will often drop any mention of the type when it is clear from the context.

We let $Mod^+(S) \subset Mod(S)$ denote the index 2 subgroup of orientation preserving diffeomorphisms of S, up to isotopies. More generally an upper $+$ will mean "orientation preserving" in the sequel. We usually write $\Gamma(S) = Mod^+(S)$ and call it the (Teichmüller) modular group. $\Gamma(S)$ is (noncanonically) isomorphic to $\Gamma_{g,[n]}$, which is instead identified to the topological fundamental group of the complex analytic stack $\mathcal{M}_{g,[n]}$.

We will mainly use two types of complexes of curves, which are both classical, and which we denote $C(S)$ and $C_P(S)$. The complex $C(S)$ is the original one introduced by W.J.Harvey; we will need in particular the results contained in [I1] and [L] to which we refer for further detail and references. Here we recall that $C(S)$ is a simplicial complex of dimension $d(S)-1$; it is not locally finite. A k-simplex of $C(S)$ is defined by a multicurve $\underline{s}=(s_0,\ldots,s_k)$, that is a set of $k+1$ isotopy classes of loops which are distinct, nontrivial, not bounding a punctured disc and have pairwise intersection number 0. The intersection number of two isotopy classes s and s' is defined as the minimum of the intersection numbers of representatives of s and s'. The face and degeneracy operators are defined by adding and deleting loops respectively. We will write $C^{(k)}(S)$ for the k-dimensional skeleton of $C(S)$. For any simplex, there exists in fact a set of $k+1$ representatives which are disjoint simple closed curves on S; namely endow S with a Poincaré metric (constant negative curvature) and pick as a representative of s_i the unique geodesic in that isotopy class. It is a deep and fundamental result (N.V.Ivanov, J.Harer) that $C(S)$ has the homotopy type of a wedge of spheres; we refer to these authors or to [B1] for a precise statement.

Consider now the group of simplicial automorphisms $Aut(C(S))$ of the complex of curves. There is a natural map $Mod(S)\to Aut(C(S))$ induced by letting a diffeomorphism act on loops, up to isotopy. The elements of the center lie in the kernel of that map because they commute with twists, so there is an induced map $\theta: Inn(Mod(S)))\to Aut(C(S))$. Assume now that $C(S)$ is connected, that is $d(S)>1$. We will return to the cases of modular dimension 1 (i.e. types $(0,4)$ and $(1,1)$) at the end of the section. Then it is not too difficult to show that θ is injective. A deep fundamental fact is that θ is also surjective for $(g,n)\neq(1,2)$. We state this as a theorem whose surjectivity part in i) is due to N.V.Ivanov ([I1])and F.Luo ([L]), as well as M.Korkmaz in the cases of low genera. We will comment on ii) after the statement.

Theorem 2.1: *Let S be a hyperbolic surface of type (g,n) with $d(S)>1$; then:*
i) the natural map $\theta: Inn(Mod(S))\to Aut(C(S))$ is an isomorphism except if $(g,n)=(1,2)$, in which case it is injective but not surjective; in fact θ maps $Inn(Mod(S_{1,2}))$ onto the strict subgroup of the elements $Aut(C(S_{1,2}))$ which globally preserve the set of vertices representing nonseparating curves;
ii) $Aut(C^{(1)}(S))=Aut(C(S))$.

Of course, if the type is different from $(1,2)$ and $(2,0)$, $Mod(S)$ is centerfree and θ provides an isomorphism between $Mod(S)$ and $Aut(C(S))$. Item ii) is easy and well-known and was added because of the parallel statement for the pants complex (see Theorem 2.5 below) and the fact that it will remain valid in the profinite setting. We include the short proof here because similar arguments will appear below. There is a natural map $Aut(C(S))\to Aut(C^{(1)}(S))$ which to an automorphism of the complex associates its restriction to the 1-skeleton. Now $C(S)$ is a flag complex, that is a simplex is determined by its boundary, and by induction it is determined by its vertices. This implies that the above restriction map is injective, indeed the restriction to the set of vertices (i.e. $C^{(0)}(S)$) is already injective. To prove surjectivity it is enough to give a graph theoretic characterization of the higher dimensional simplices of $C(S)$ and this is easily available: a moment contemplation will confirm that the k-dimensional simplices are in one-to-one correspondence with the complete subgraphs of $C^{(1)}(S)$ with $k+1$ vertices, i.e. subgraphs such that any two vertices are connected by an edge. This characterization proves ii). $\square$

Let us also add that the odd looking case of type $(1,2)$ is actually easy to understand. It comes from the fact that $C(S_{1,2})$ and $C(S_{0,5})$ are isomorphic whereas $\Gamma_{1,[2]}/Z(\Gamma_{1,[2]})$ maps into $\Gamma_{0,[5]}$ as a subgroup of index 5; indeed θ maps $Inn(Mod(S_{1,2}))$ injectively onto an index 5 subgroup of $Aut(C(S_{1,2}))$. See [L] for a short geometric discussion.

N.V.Ivanov showed how to use the description of $Aut(C(S))$ afforded by Theorem 2.1 in order study the action of $\Gamma(S)$ on Teichmüller space. He recovers in this way (see [I1]) the basic result of H.Royden (see [EK]) about automorphisms of Teichmüller spaces:

Theorem 2.2: *If $d(S) > 1$, any complex automorphism of $\mathcal{T}(S)$ is induced by an element of $Mod(S)$.*

As N.V.Ivanov again showed, Theorem 2.1 also has immediate bearing on the automorphisms of modular groups. This is because any inertia preserving automorphism of $\Gamma(S)$, i.e. any element of $Aut^*(\Gamma(S))$ clearly induces an element of $Aut(C(S))$. Now in the discrete setting, we have the following:

Theorem 2.3: *All automorphisms of $\Gamma(S)$ are inertia preserving, that is:*

$$Aut^*(\Gamma(S)) = Aut(\Gamma(S)).$$

This result, which is essentially due to N.V.Ivanov (see [I2] and reference therein) rests on a group theoretic characterization of twists inside $\Gamma(S)$. It is rarely stated independently or emphasized but we would like to stress it in view of the profinite case (we also refer to [McC] for a nice proof based on the notion of "stable rank"). This is because first we do not know how to prove the profinite analog, which is unfortunate, and second because in the profinite setting this would feature a rather striking and precise analog of the so-called "local correspondence" in birational anabelian geometry. We hope to return to these questions elsewhere.

Armed with Theorem 2.3, it is easy to use Theorem 2.1 in order to study the automorphisms of $\Gamma(S)$ (compare Proposition 4.11 below). Actually it is no more difficult to study morphisms between finite index subgroups, as in [I1] (Theorem 2); we state this as:

Theorem 2.4: *Assume $d(S) > 1$ and $\Gamma = \Gamma(S)$ has trivial center; let $\Gamma_1, \Gamma_2 \subset \Gamma$ be two finite index subgroups. Then any isomorphism ϕ between Γ_1 and Γ_2 is induced by an element of $Mod(S)$, namely there exists $\gamma \in Mod(S)$ such that $\phi(\gamma_1) = \gamma^{-1}\gamma_1\gamma$ for any $\gamma_1 \in \Gamma_1$. In particular $Out(\Gamma(S)) \simeq \mathbb{Z}/2$.*

As usual one can study the two cases with nontrivial center, that is $(1,2)$ and $(2,0)$ in detail; see [McC] for the latter one.

Turning now to the *pants complex* $C_P(S)$, we note that it was briefly introduced in the appendix of the classical 1980 paper by A.Hatcher and W.Thurston (see [HLS] or [M]) and first studied in [HLS] where it is shown to be connected and simply connected. It is a two dimensional, not locally finite complex whose vertices are given by the pants decomposition (i.e. maximal multicurves) of S; they correspond to the simplices of highest dimension ($= d(S)-1$) of $C(S)$. Given two vertices $\underline{s}, \underline{s}' \in C_P(S)$, they are connected by an edge if and only if $\underline{s}$ and $\underline{s}'$ have $d(S)-1$ curves in common, so that up to relabeling (and of course isotopy) $s_i = s_i'$, $i = 1, \ldots, d(S)-1$, whereas s_0 and s_0' differ by an *elementary move*, which means the following. Cutting S along the s_i's, $i > 0$, there remains a surface Σ of modular dimension 1, so Σ is of type $(1,1)$ or $(0,4)$. Then s_0 and s'_0, which are supported on Σ, should intersect in a minimal way, that is they should have intersection number 1 in the first case, and 2 in the second case.

We have thus defined the 1-skeleton $C_P^{(1)}(S)$ of $C_P(S)$ which, following [M], we call the *pants graph* of S. We will not give here the definition of the 2-cells of $C_P(S)$ (see [HLS] or [M]), as we will actually not use it. They describe certain relations between elementary moves, that is they can be considered as elementary homotopies; as mentioned above pasting them in makes $C_P(S)$ simply connected (cf. [HLS]).

Here we will use and reprove the recent analog of Theorem 2.1 proved by D.Margalit in [M]. Again we first restrict to the case $d(S) > 1$ and will comment on the one dimensional cases below. As in the case of $C(S)$, there is a natural map $\theta_P : Inn(Mod(S)) \to Aut(C_P(S))$ and one inquires once again about its kernel and range. The result of D.Margalit reads as follows:

Theorem 2.5: *Let S be a hyperbolic surface of type (g,n) with $d(S) > 1$; then:*
i) the natural map $\theta_P : Inn(Mod(S)) \to Aut(C_P(S))$ is an isomorphism;
ii) $Aut(C_P^{(1)}(S)) = Aut(C_P(S))$.

So the statement completely parallels that of Theorem 2.1 for $C(S)$, with the only difference that type $(1,2)$ is not exceptional here. Also item ii) is much harder to prove that in Theorem 2.1. So up to a small kernel in small dimensions, the mapping class group $\Gamma(S)$ can be recovered as the group of orientation preserving automorphisms of the pants *graph* $C_P^{(1)}(S)$. As mentioned above, in the sequel we will not have to make use nor even define the profinite analog of the full pants complex; therefore *from here on $C_P(S)$ will refer by default to the pants graph*, i.e. the 1-skeleton of the pants complex. We remark however that the use of the full two dimensional pants complex was a crucial feature of [HLS] because for instance its faces are intimately connected with the defining relations of the Grothendieck-Teichmüller group. So it *is* remarkable that for our present purpose (and in Theorem 2.5 above) one can forget about the faces, which are actually encoded in the 1-skeleton as detailed in [M].

Let us return to the one dimensional cases $(0,4)$ and $(1,1)$. Then the (two dimensional) pants complexes $C_P(S_{0,4})$ and $C_P(S_{1,1})$ both coincide with the classical Farey tesselation (see e.g. [M]). The curves complexes $C(S_{0,4})$ and $C(S_{1,1})$ are 0-dimensional according to the general definition but by convention and in order to add structure, they are usually redefined as coinciding with the Farey tesselation as well. We denote the latter simply by F and do not notationally distinguish it from its 1-skeleton, namely the Farey graph; both objects carry the same information and can easily be recovered from each other. In other words all four complexes are just identified with the Farey tesselation or graph F. It will appear below that the convention just mentioned about $C(S)$ for $d(S) = 1$ is not necessarily the right one, but that does not matter here. In any case, passing to automorphisms, we have: $Aut(F) = PGL_2(\mathbb{Z})$ with the index 2 orientation preserving subgroup $Aut^+(F) = PSL_2(\mathbb{Z})$.

Before we move on, let us summarize the situation in the discrete case again, with a view towards the profinite setting. We focus on $\Gamma(S) = Mod^+(S)$ which is the natural group in the algebro-geometric setting. So we have a tautological short exact sequence:

$$1 \to \Gamma(S) \to Mod(S) \to \mathbb{Z}/2 \to 1,$$

where here and below one should think of $\mathbb{Z}/2$ as the Galois group $Gal(\mathbb{C}/\mathbb{R})$. Now consider for simplicity a surface S with $d(S) > 1$ and not of type $(1, 2)$, i.e. S is of type $(0, 5)$ or $d(S) > 2$. Theorem 2.1 states that there is a short exact sequence:

$$1 \to Inn(\Gamma(S)) \to Aut(C(S)) \to \mathbb{Z}/2 \to 1. \tag{1}$$

Theorem 2.5 asserts in particular that $Aut(C_P(S))$ is isomorphic to $Aut(C(S))$ and in fact the proof of Theorem 2.5 in [M] consists in constructing an explicit map between these groups (see also Theorem 2.13 below). Now about $Aut(\Gamma(S))$; we first know that any automorphism permutes cyclic subgroups generated by twists (Theorem 2.3). Using this it is easy to prove Theorem 2.4, so in particular there is an exact sequence:

$$1 \to Inn(\Gamma(S)) \to Aut(\Gamma(S)) \to \mathbb{Z}/2 \to 1. \tag{2}$$

By Theorem 2.3 there is a natural map $Aut(\Gamma(S)) \to Aut(C(S))$ and it is an isomorphism. This is a close analog of Tits rigidity theorem which states that (with suitable assumptions) any automorphism of a building is induced by an automorphism of the corresponding algebraic group. We remark that actually N.V.Ivanov proves more, namely he determines the automorphism group of any finite index subgroup of $\Gamma(S)$ as in Theorem 2.4, which for our purpose is quite significant. Finally we remark that all the above exact sequences are split because of the existence of a reflection, that is an orientation reversing order 2 diffeomorphism of the surface S.

Perhaps the main surprise of the profinite case will consist in the fact that the automorphism groups of the profinite completions of $C(S)$ and $C_P(S)$ are very different. In trying to make this more natural, we came across the following considerations and results which will later (in §4) be adapted to the profinite case. As mentioned above, even in the discrete case, they may present some independent interest. Given a hyperbolic surface S as above, we associate to it three (discrete) graphs. First $C^{(1)} = C^{(1)}(S)$ is just the 1-skeleton of $C(S)$; second $C_P = C_P(S)$ will denote as above the pants graph. The third graph $C_*(S)$ is defined as follows. Its vertices are again the pants decompositions (maximal multicurves) of S and two vertices are connected by an edge if and only if the associated multicurves differ by exactly one curve. In other words the two vertices are represented by multicurves $\underline{s} = (s_i)_i$ and $\underline{s}' = (s'_i)_i$ where $i = 0, \dots, d(S) - 1$ and up to relabeling we can assume that $s_i = s'_i$ for $i > 0$, whereas s_0 and s'_0 lie on a surface of type $(0, 4)$ or $(1, 1)$. Comparing to C_P, we have just dropped the condition of minimal intersection.

As mentioned in the proof of Theorem 2.1 ii), it is well-known that $C(S)$ is easily reconstructed from $C^{(1)}(S)$: the cells of $C(S)$ are given by the complete subgraphs of $C^{(1)}(S)$. In particular $Aut(C^{(1)}(S)) = Aut(C(S))$ and the same will hold true in the profinite case. Here we will not mention $C^{(1)}$ anymore and concentrate mostly on C_*, emphasizing the relevance of that graph which has never been studied in detail. We remark that from the point of view of topology "in the large", these three graphs are not so different. In particular, when endowed with their respective graph metrics (every edge has unit length) all three are quasi-isometric, in fact quasi-isometric to $C(S)$, so by the main result of [MM], all three are hyperbolic. However this is not relevant for our present purpose.

From their definitions, the graphs $C_*(S)$ and $C_P(S)$ the graphs $C_*(S)$ and $C_P(S)$ have the same sets of vertices, namely the pants decompositions of the surface S; we denote it $V(S)$. We write $E(S)$ (resp. $E_P(S)$) for the set of edges of $C_*(S)$ (resp. $C_P(S)$). We may view $E_P(S)$ as a subset of $E(S)$ and $C_P(S) \subset C_*(S)$ as a subgraph of $C_*(S)$ with the same vertices. If S is connected of dimension 0, it is of type $(0,3)$ (a "pair of pants"); by convention, $C_P(S) = C_*(S)$ is reduced to a point (with no edge; note that usually one defines $C(S_{0,3}) = \emptyset$). If S is connected and $d(S) = 1$, it is of type $(0,4)$ or $(1,1)$. In both cases $C_P(S) = F$ coincides with the Farey graph. On the other hand, it is easily checked that $C_*(S)$ is the complete graph associated to F, which we denote by G. This is simply because two curves on a surface of (modular) dimension 1 always intersect. Finally if $d(S) > 1$, $C_*(S)$ is nothing but the 1-skeleton of $C(S)^*$, the complex dual to $C(S)$. For that reason, when $d(S) = 1$, it becomes natural to define $C(S)$ as the dual of G, which is not the usual convention but seems to be the right one here.

It is useful to be able to deal with non connected surfaces and C_* and C_P turn out to be particularly easy to deal with in this respect. Let $S = S' \coprod S''$ be given as the disjoint sum of S' and S'', which themselves need not be connected. First modular dimension is additive: $d(S) = d(S') + d(S'')$. Then it is easy to describe $C_*(S)$ and $C_P(S)$ in terms of the graphs associated to S' and S''. For the vertices we get: $V(S) = V(S') \times V(S'')$; and for the edges of $C_*(S)$: $E(S) = E(S') \times V(S'') \coprod V(S') \times E(S'')$. Simply change E into E_P for the case of C_P. These prescriptions immediately generalize to an arbitrary number r of not necessarily connected pieces. If $S = \coprod_i S_i$, $d(S) = \Sigma_i d(S_i)$, $V(S) = \prod_i V(S_i)$ and $E(S) = \coprod_i V(S_1) \times \ldots \times E(S_i) \times \ldots \times \ldots V(S_r)$; replace again E with E_P when dealing with C_P.

From now on, surfaces are assumed to be hyperbolic (i.e. all connected components are hyperbolic) but not necessarily connected unless this is explicitly mentioned. For simplicity we will also state or prove some statements for C_* only, when they involve only one of the two graphs C_* and C_P. The transposition to C_P is usualy quite literal and will be left to the reader.

Given a surface S, a *subsurface* T is defined as $T = S \setminus \sigma$ where $\sigma \in C(S)$; we denote it S_σ and it is nothing but S cut or slit along the multicurve representing σ. In this definition, the curves are defined as usual up to isotopy. There is a natural inclusion $C_*(S_\sigma) \subset C_*(S)$; in fact $C_*(S_\sigma)$ is the full subgraph of $C_*(S)$ whose vertices correspond to those pants decompositions of S which include σ (idem C_P). For $\sigma \in C(S)$, we let $|\sigma|$ denote the number of curves which constitute σ. So $|\sigma| = dim(\sigma) + 1$ if $dim(\sigma)$ denotes the dimension of the simplex $\sigma \in C(S)$. The quantity $|\sigma|$ turns out to be more convenient in our context; in particular $d(S_\sigma) = d(S) - |\sigma|$. We include throughout the case of an empty cell (dimension -1): $S_\emptyset = S$. For example if σ is a maximal multicurve (pants decomposition), S_σ is a disjoint union of pants and $C_*(S_\sigma)$ is reduced to a point. We say that two simplices $\rho, \sigma \in C(S)$ are compatible if the curves which compose ρ and σ do not intersect properly, that is they are either disjoint or coincide. Complex theoretically it means that ρ and σ lie in the closure of a common top dimensional simplex of $C(S)$. If ρ and σ are compatible, we define their unions and intersections $\rho \cup \sigma, \rho \cap \sigma \in C(S)$ in the obvious way.

The next statement is both obvious and important:

Lemma 2.6: *If $\rho, \sigma \in C(S)$ are compatible simplices $C_*(S_\rho) \cap C_*(S_\sigma) = C_*(S_{\rho \cup \sigma})$. If they are not compatible, this intersection is empty.* $\square$

Here all graphs $C_*(S_\tau)$ ($\tau \in C(S)$) are considered as subgraphs of $C_*(S)$. This lemma has a number of equally obvious consequences. For instance $C_*(S_\rho) \subset C_*(S_\sigma)$ if and only if $\sigma \subset \rho$.

Let us now return to the connections between C_* and C_P. The inclusion $C_P \subset C_*$ can be made more precise:

Lemma 2.7: *$C_*(S)$ is obtained from $C_P(S)$ by replacing every maximal copy of F inside $C_P(S)$ by a copy of G.* $\square$

A maximal copy of F is a subgraph of $C_P(S)$ which is isomorphic to F and is not properly contained in another such subgraph. Note that the operation described in the lemma is not reversible; one cannot recognize C_P inside C_* without additional information. This has an easy but again important consequence in terms of automorphisms:

Lemma 2.8:

$$Aut(C_P(S)) \subset Aut(C_*(S))$$

An automorphism of C_P determines a permutation of the common vertex set V which determines an automorphism of C_* provided it is compatible with the edges of C_*. Lemma 2.7 and the fact that G is a complete graph ensure that this is always the case. $\square$

So any automorphism of $C_P(S)$ determines an automorphism of $C_*(S)$ because both graphs share the same set of vertices and automorphisms of flag complexes are determined by their effect on the vertices. However *a priori* only certain automorphisms of $C_*(S)$ will preserve the additional structure given by the edges of $C_P(S)$, inducing an automorphism of this subgraph. In dimension 1, $Aut(G)$ is nothing but the permutation group on its vertices. Any automorphism of F determines a unique automorphism of G by looking at its effect on the vertices, but $Aut(F) \simeq PGL_2(\mathbb{Z})$ is certainly much smaller than $Aut(G)$. In the discrete case a kind of rigidification occurs for $d(S) > 1$ but this will not be so in the profinite case.

The (semi)local structure of C_* and C_P is not so mysterious. Indeed we have:

Lemma 2.9: *Let $v \in V(S)$ be a vertex of $C_*(S)$ and $C_P(S)$, with $d(S) = k \geq 0$. Then v lies at the intersection of exactly k (maximal) copies of G (resp. F) in $C_*(S)$ (resp. $C_P(S)$). For any two copies G_i, G_j $(i \neq j)$ one has $G_i \cap G_j = \{v\} \subset C_*(S)$ and two vertices $w_i \in G_i$ $w_j \in G_j$ with $w_i \neq v$, $w_j \neq v$ are* not *joined by an edge in $C_*(S)$. The same statement holds for $C_P(S)$ (with G replaced by F).*

Let v be given as a pants decomposition $v = (\alpha_1, \ldots, \alpha_k)$. The main point here is that any triangle (complete graph on three vertices) of C_* or C_P is obtained by varying one of the α_i's keeping all the other curves α_j fixed. This in turn depends only on the already mentioned (and obvious) fact that two curves on a surface of dimension 1 always intersect. So we get k copies of G inside C_* which are indexed by the curves appearing in v. The rest of the statement and the transposition to C_P are easily verified.

Note that this shows that $d(S)$ can be read off graph theoretically on C_* or C_P. In fact it can be detected locally around any vertex v. To this end one can look for a family $(w_i)_{i \in I}$ of vertices of $C_*(S)$ (or $C_P(S)$) such that each w_i is connected to v by an edge and no two distinct w_i's are connected. Then $d(S)$ is the maximal possible number of such vertices (i.e. the maximal cardinal of the index set I). $\square$

We now would like to reconstruct $C(S)$ from $C_*(S)$, hence also from $C_P(S)$ by Lemma 2.7. One way to do this is to set up a correspondence between the subgraphs of $C_*(S)$ which are graph theoretically isomorphic to some $C_*(S_\sigma)$ ($\sigma \in C(S)$) and the subsurfaces of S. This correspondence, which we will later extend to the profinite setting is also interesting by itself. A precise wording goes as follows:

Theorem 2.10: *Let $C \subset C_*(S)$ be a subgraph which is (abstractly) isomorphic to $C_*(\Sigma)$ for a certain surface Σ and is maximal with this property. Then there exists a unique $\sigma \in C(S)$ such that $C = C_*(S_\sigma)$.*

Let us start with some remarks and reductions. First we note that the word "maximal" is indeed necessary. For instance there are proper subgraphs of F (resp. G) which are isomorphic to F (resp. G). Second, implicit in the statement is the fact that any $C_*(S_\sigma) \subset C_*(S)$ does indeed answer the problem, namely it is maximal in its isomorphy class. Assume on the contrary that we have a nested sequence $C_*(S_\sigma) \subset C \subset C_*(S)$ where $d(S_\sigma) = k$, C is isomorphic to $C_*(S_\sigma)$ and the first inclusion is strict. Since C is connected, we can find a vertex $w \in C \setminus C_*(S_\sigma)$ which is connected by an edge to a vertex $v \in C_*(S_\sigma)$. Since S_σ has dimension k, we can find k vertices $w_i \in C_*(S_\sigma)$ as in the proof of Lemma 2.9 (with respect to v). But $w \in C$, is connected to v and it is easy to check that it is not connected to any of the w_i. In other words we have actually found $k+1$ vertices which are connected to v and no two of which are connected, which contradicts the fact that C is isomorphic to $C_*(S_\sigma)$.

Having justified the statement, we can turn to the proof of Theorem 2.10, remarking that uniqueness is clear: obviously $C_*(S_\sigma)$ coincides with $C_*(S_\tau)$ ($\sigma, \tau \in C(S)$) if and only if $\sigma = \tau$; this is also a very particular case of Lemma 2.6. From Lemma 2.9 we can now define $d(C) = d(\Sigma)$, which determines $|\sigma|$ (assuming the existence of σ) since $d(S_\sigma) = d(C) = d(S) - |\sigma|$. Next the result is true if $d(\Sigma) = 0$ because then $C_*(\Sigma)$ is just a point and so is C. Hence it does correspond to a vertex of $C_*(S)$, in other words to an actual pants decomposition of S. We will prove the result by induction on $k = d(\Sigma)$ but it is actually useful and enlightening to prove the case $k = 1$ directly. This is easy and essentially well-known in a different context. Much as in Lemma 2.9, the point is that any complete graph, and in fact any triangle inside $C_*(S)$ (or $C_P(S)$) determines a unique subsurface Σ with $d(\Sigma) = 1$. This sets up a one-to-one correspondence between subsurfaces of S of dimension 1 and maximal complete subgraphs of $C_*(S)$.

Now let $k > 1$, assume the result has been proved for $d(C) < k$ and consider a graph $C \subset C_*(S)$ as in the statement, with $d(C) = k$. We fix an isomorphism $C \xrightarrow{\sim} C_*(\Sigma)$. Changing notation slightly for convenience, we are looking for a subsurface $T \subset S$, defined by a cell of $C(S)$, and such that $C = C_*(T)$. Note that it may happen that the surfaces Σ and T (assuming the existence of the latter) are not of the same type because of the well-known exceptional low-dimensional isomorphisms between complexes of curves. One will have $C_*(\Sigma) \simeq C_*(T)$ and indeed, as a consequence of the result itself, $C(\Sigma) \simeq C(T)$, so for instance Σ could be of type $(0, 6)$ and T of type $(2, 0)$.

We may now consider subsurfaces of Σ and transfer the information to $C \subset C(S)$. Namely for any $\sigma \in C_*(\Sigma)$, we denote by $C_\sigma \subset C$ the subgraph corresponding to $C_*(\Sigma_\sigma)$ under the fixed isomorphism $C \simeq C_*(\Sigma)$. Actually, forgetting about this isomorphism, we just write $C_\sigma =$

$C_*(\Sigma_\sigma) \subset C \subset C_*(S)$. By the induction hypothesis, for any $\sigma \in C(\Sigma)$, $\sigma \neq \emptyset$, here corresponds to C_σ a unique subsurface $S_{(\sigma)} \in S$. Beware of the fact that σ now runs over the cells of $C(\Sigma)$, not of $C(S)$ and this is the reason of the added brackets. In these terms we are trying to extend this correspondence to $\sigma = \emptyset$, i.e. find $T = S_{(\emptyset)}$.

In order to show the existence of T, it is actually enough to show that there exists a k-dimensional subsurface of S, call it precisely T, such that any $S_{(\sigma)}$ with $\sigma \in C(\Sigma)$ not empty is contained in T. Indeed, the corresponding C_σ's form a covering of C. So assuming the existence of such a subsurface T, we find that $C \subset C_*(T)$; these two subgraphs being isomorphic and C being maximal by assumption, they coincide. In order to prove the existence of T, we can now restrict attention to the largest possible $S_{(\sigma)}$'s, i.e. to the case $|\sigma| = 1$, which simply means that σ consists of a single loop.

We are thus reduced to proving that there exists a k-dimensional subsurface $T \subset S$ such that for any loop α on Σ $S_{(\alpha)}$ is contained in T. Now $C(\Sigma)$ is connected because $k > 1$ and this can be used as follows. If α and β are two non intersecting curves on Σ, Σ_α and Σ_β are two subsurfaces of Σ of dimension $k-1$ intersecting along the subsurface $\Sigma_{\alpha\cup\beta}$ of dimension $k-2$, where $\alpha \cup \beta$ is considered as a simplex of $C(\Sigma)$. Informally speaking for the time being, the union $S_{(\alpha)} \cup S_{(\beta)}$ has dimension k and this is the natural candidate for T. In other words the latter, if it exists, is determined by any two non intersecting loops of Σ. Returning to the formal proof, let γ and δ be two arbitrary loops on Σ. There exists a path in the 1-skeleton of $C(\Sigma)$ connecting γ to δ. It is given as a finite sequence $\gamma, \alpha_1, \ldots, \alpha_n, \delta$ of loops such that α_1 does not intersect γ, α_n does not intersect δ and for $1 < i < n$, α_i does not intersect α_{i-1} and α_{i+1}. Using the existence of such a chain, we are reduced to the following situation. Let α, β, and γ be three loops on Σ such that $\alpha \cap \beta = \beta \cap \gamma = \emptyset$; there remains again to show that $S_{(\alpha)}$, $S_{(\beta)}$ and $S_{(\gamma)}$ are contained in a common k-dimensional subsurface T, and this will complete the proof of the result.

We can write $S_{(\alpha)} = S_\rho$, $S_{(\beta)} = S_\sigma$, $S_{(\gamma)} = S_\tau$, for certain simplices $\rho, \sigma, \tau \in C(S)$ with $|\rho| = |\sigma| = |\tau| = d(S) - k + 1$. Moreover because $\alpha \cap \beta = \emptyset$ (resp. $\beta \cap \gamma = \emptyset$) ρ and σ (resp. σ and τ) are compatible simplices. So we can consider $\rho \cap \sigma$ and $\sigma \cap \tau$, with $|\rho \cap \sigma| = |\sigma \cap \tau| = d(S) - k$. The corresponding surfaces $S_{\rho\cap\sigma}$ and $S_{\sigma\cap\tau}$ are both subsurfaces of S of dimension k. There remains only to show that they coincide: $S_{\rho\cap\sigma} = S_{\sigma\cap\tau} (= T)$. We reason much as above, when proving that a subcomplex of type $C_*(S_\sigma) \subset C_*(S)$ is maximal in its isomorphy class. The complexes $C_{\rho\cap\sigma}$ and $C_{\sigma\cap\tau}$ are two subcomplexes of dimension k inside C which is also of dimension k, and they are maximal such complexes, being attached to subsurfaces of S. This forces them to coincide – and in fact coincide with the whole of C. More formally, assume the contrary, that is $S_{\rho\cap\sigma}$ and $S_{\sigma\cap\tau}$ are distinct. Then, breaking the symmetry for a moment and relabeling if necessary, we can choose as above two vertices $v \in C_{\rho\cap\sigma}$ and $w \in C_{\sigma\cap\tau} \setminus C_{\rho\cap\sigma}$ which are connected by an edge. Then again pick a maximal family (w_i) of k vertices in $C_{\rho\cap\sigma}$ which are connected to v and are not mutually connected. Adding in the vertex w we get a family of $k+1$ vertices with the same properties, which contradicts the fact that $d(C) = k$ and finishes the proof. □

We now move to some consequences of this result. First one has:

Corollary 2.11: *$C(S)$ can be (graph theoretically) reconstructed from $C_*(S)$.*

In fact starting from $C_*(S)$, one builds a complex by considering subgraphs C as in the statement of the theorem, with the inclusion map as a boundary operator. The result ensures that this simplicial complex is isomorphic to the curve complex $C(S)$. □

One then immediately gets:

Corollary 2.12: *$Aut(C_*(S)) = Aut(C(S))$.* □

Now taking Lemma 2.8 into account, this yields:

Theorem 2.13: *There is a natural injective map: $Aut(C_P(S)) \hookrightarrow Aut(C(S))$.* □

Although this statement is obtained here as a consequence of Theorem 2.10, we state it as an independent result in order to stress its importance and because in section 4 we will obtain its profinite counterpart. It is the keypoint in deducing Theorem 2.5 from Theorem 2.1 (see also [M]) from which everything else follows easily. Namely since $Aut(C(S)) = Mod(S)$ (Theorem 2.1), Theorem 2.13 implies that $Aut(C_P(S)) = Mod(S)$ as well, where we recall that $C_P(S)$ actually denotes the pants graph. This in turn immediately implies that the automorphism group of the full (2 dimensional) pants complex also coincides with $Mod(S)$, so in particular the automorphism group of the full pants complex coincides with the automorphism group of its 1-skeleton (the pants graph). For the fact that here type $(1, 2)$ is no exception, see the last page of [M].

3. Profinite complexes of curves

In this paragraph we introduce the profinite analog of $C(S)$, giving several descriptions and proving their equivalence. This can be seen as a continuation of the fundational work started in [B1] but it falls short of completing it. More remains to be unearthed, although we will prove more than enough for our future needs. We will also introduce the profinite pants graph but still lack a satisfactory description of it as will be made clearer in the remarks at the end of this section.

Let us first recall the main definitions, refering to [B1] for more detail. We also refer to [S] for the general framework of topology for profinite or more generally compact completely disconnected spaces. We fix a hyperbolic surface S of type (g, n) as usual and often drop the mention of the surface or the type from the notation. We index the inverse system of the finite index subgroups of $\Gamma = \Gamma(S) \simeq \Gamma_{g,[n]}$ by the set Λ, so that for any $\lambda \in \Lambda$ we have a subgroup Γ^λ and by definition:

$$\hat{\Gamma} = \varprojlim_{\lambda \in \Lambda} \Gamma/\Gamma^\lambda.$$

Here we work for simplicity with the full profinite completion only but the theory extends to other completions as well, i.e. inverse subsystems of Λ. Note that for most purposes one can restrict consideration to the normal or even characteristic subgroups Γ^λ since both types define cofinal inverse subsystems. To any Γ^λ there corresponds a finite étale cover $\mathcal{M}^\lambda \to \mathcal{M}$. In particular, for $m \geq 2$ a positive integer, the abelian level $\mathcal{M}(m)$ is defined by the subgroup $\Gamma(m)$ which is the kernel of the natural map $\Gamma \to Sp_{2g}(\mathbb{Z}/m)$, that is $\Gamma(m)$ is the group of diffeomorphisms of S which fix the homology of the associated unmarked or compact surface modulo m. For $\lambda, \mu \in \Lambda$ we write $\mu \geq \lambda$ if $\Gamma^\mu \subseteq \Gamma^\lambda$ i.e. if $\mathcal{M}^\mu$ is a covering of $\mathcal{M}^\lambda$, and we say that $\mathcal{M}^\mu$ (resp. Γ^μ) dominates $\mathcal{M}^\lambda$ (resp. Γ^λ).

We regard the complex $C(S)$ as a simplicial set, denoted $C(S)_\bullet$; for $k \geq 0$, $C(S)_k$ denotes the set of its k-simplices and $C(S)^{(k)}$ its k-skeleton. The complex $C(S)$ is naturally equipped with a

Γ-action, which is geometric in the sense of [B1]. In fact, when restricted to a Γ^λ which dominates $\Gamma(m)$ for some $m > 2$, the action is simplicial, that is it commutes with the face and degeneracy operators (cf. [B1], §5). This makes it possible to define the profinite completion as the inverse limit:

$$\hat{C}(S)_\bullet = \varprojlim_{\lambda\in\Lambda} C(S)_\bullet/\Gamma^\lambda.$$

So we regard $\hat{C}(S)_\bullet$ as a simplicial object in the category of profinite sets but usually denote it simply $\hat{C}(S)$. There is a canonical inclusion $C(S) \hookrightarrow \hat{C}(S)$ with dense image, and a natural continuous action of $\hat{\Gamma}$ on $\hat{C}(S)$.

In a completely similar fashion, we define $\hat{C}_P(S)$ as the inverse limit:

$$\hat{C}_P(S)_\bullet = \varprojlim_{\lambda\in\Lambda} C_P(S)_\bullet/\Gamma^\lambda$$

and regard it as a simplicial object in the category of profinite sets, in fact as a prograph, usually denoted $\hat{C}_P(S)$. There is again a canonical inclusion $C_P(S) \hookrightarrow \hat{C}_P(S)$ with dense image, and a natural continuous action of $\hat{\Gamma}$ on $\hat{C}_P(S)$. Finally, as in the discrete case, there is a one-to-one correspondence between the vertices of $\hat{C}_P(S)$ (i.e. $\hat{C}_P(S)_0$) and the simplices of $\hat{C}(S)$ of maximal dimension (i.e. $\hat{C}(S)_{d-1}$, $d = d(S)$). A deep additional information is however contained in the edges of $\hat{C}_P(S)$, which we do not know at the moment how to decipher in a satisfactory way.

We now concentrate on alternative, more geometric and managable descriptions of $\hat{C}(S)$, to which end we introduce more geometric objects which are also important for their own sake. The point is that a k-simplex of the discrete complex $C(S)$ can be described in at least two equivalent ways. It can be seen as a set of $k+1$ non intersecting simple closed curves, that is as a multicurve $\underline{s} = (s_0, \ldots, s_k)$, or as the set of $k+1$ commuting twists along the s_i. Moreover, as already noted, given a loop s, it determines a twist $\tau_s \in \Gamma$, provided an orientation of the reference surface S has been fixed, which we assume from here on. The above considerations are elementary in the discrete case; the point will be to translate them in the profinite setting and show that they are in substance still valid in that case but now reflect deep properties of the profinite complex $\hat{C}(S)$.

Let $\mathcal{S} = \mathcal{S}(S)$ denote the set of loops on S, that is of isotopy classes of simple closed curves; let $\mathcal{T} = \mathcal{T}(S) \subset \Gamma = \Gamma(S)$ denote the set of twists in Γ. The group Γ naturally acts on $\mathcal{S}$ and it acts on $\mathcal{T}$ by conjugation. There is a natural map $\tau : \mathcal{S} \to \mathcal{T}$ sending $s \in \mathcal{S}$ to the twist $\tau_s \in \mathcal{T}$ along s. Moreover τ is one-to-one and Γ-equivariant, that is for $\gamma \in \Gamma$ and $s \in \mathcal{S}$: $\tau_{\gamma\cdot s} = \gamma\tau_s\gamma^{-1}$. Now define the profinite set $\hat{\mathcal{S}}$ of profinite loops as:

$$\hat{\mathcal{S}} = \varprojlim_{\lambda\in\Lambda} \mathcal{S}/\Gamma^\lambda,$$

noting that there is indeed a finite number of Γ^λ-orbits of loops for any $\lambda \in \Lambda$. The natural map $\mathcal{S} \to \hat{\mathcal{S}}$ is a monomorphism with dense image and $\hat{\mathcal{S}}$ comes equipped with a natural $\hat{\Gamma}$-action extending the Γ-action on $\mathcal{S}$.

On the other hand let us denote by $\bar{\mathcal{T}}$ the closure of $\mathcal{T} \subset \Gamma$ in $\hat{\Gamma}$ and call it the set of profinite twists. We first have the following easy description:

Lemma 3.1: $\bar{\mathcal{T}} = \{\gamma\tau_s\gamma^{-1}; s \in \mathcal{S}, \gamma \in \hat{\Gamma}\}$.

As mentioned above $\mathcal{S}$ consists of finally many Γ-orbits so that we can write: $\mathcal{S} = \coprod_{s\in F} \Gamma \cdot s$, where $s \in \mathcal{S}$ runs over a finite set F of representatives of these orbits. In the same way, we have that $\mathcal{T} = \coprod_{s\in F} \Gamma \cdot \tau_s$ where the action of Γ is of course by conjugation. The lemma states that $\bar{\mathcal{T}} = \coprod_{s\in F} \hat{\Gamma} \cdot \tau_s \subset \hat{\Gamma}$. In fact this is a closed set because $\hat{\Gamma}$ is compact, the action is continuous and F is finite. Finally $\mathcal{T}$ is dense in that set, which completes the proof of the lemma. □

The above, using only loops and twists, is well adapted to the study of the vertices of the curves complex and we will eventually come back to it. However, in order to study the higher dimensional skeleta, it is useful to introduce a group theoretic setting. Namely let $\mathcal{G} = \mathcal{G}(\hat{\Gamma})$ denote the set of closed subgroups of $\hat{\Gamma}$. This is again a profinite set since it can be written as:

$$\mathcal{G} = \varprojlim_{\lambda\in\Lambda} \mathcal{G}(\Gamma/\Gamma^{\lambda}),$$

where Γ^{λ} runs over the normal subgroups of finite index in Γ and $\mathcal{G}(\Gamma/\Gamma^{\lambda})$ denotes the finite set of the subgroups of Γ/Γ^{λ}. We also have an action of $\hat{\Gamma}$ on $\mathcal{G}$ by conjugation.

Let now $\sigma \in C(S)$ be a k-simplex, determined by a multicurve $\underline{s} = (s_0, \ldots, s_k)$, that is by the twists τ_i along the s_i. We define a map $C(S) \to \mathcal{G}$ by sending $\sigma \in C(S)_k$ to $G_\sigma \subset \mathcal{G}$, the group topologically generated by the τ_i in $\hat{\Gamma}$, which is free abelian; in fact G_σ is the closed subgroup $\hat{\mathbb{Z}}\tau_0 \oplus \ldots \oplus \hat{\mathbb{Z}}\tau_k \subset \hat{\Gamma}$ but this is far from obvious (cf. [B1] and [PdJ]). We now take the closure $\overline{C}(S)_k$ of the image of $C(S)_k$ in $\mathcal{G}$. One can get a concrete view of $\overline{C}(S)_k$ much in the same way as we did for $\bar{\mathcal{T}}$ in the proof of lemma 3.1. Choose a finite set F_k of representatives of the k-simplices under the action of Γ; then we can write:

$$\overline{C}(S)_k = \coprod_{\sigma\in F_k} \hat{\Gamma} \cdot \sigma,$$

where $\hat{\Gamma}$ acts by conjugation. This is again because the right-hand side is closed in $\mathcal{G}$ and contains the left-hand side as a dense subset. This way of writing displays the action of $\hat{\Gamma}$ on $\overline{C}(S)_k$. It also makes it clear that, varying k, we can extend the face and degeneracy operators and construct the simplicial profinite set $\overline{C}(S)_\bullet$, which we call the group theoretic $\hat{\Gamma}$-completion of $C(S)$ and which comes equipped with a continuous geometric action of $\hat{\Gamma}$. We usually denote it simply $\overline{C}(S)$. By the universal property of the completion $\hat{C}(S)$, there is a natural continuous $\hat{\Gamma}$-equivariant morphism $\hat{C}(S) \to \overline{C}(S)$ which is onto since the image is dense and closed.

One of the main results of [B1] now reads:

Theorem 3.2 (see [B1], §7): *The natural map $\hat{C}(S) \to \overline{C}(S)$ is a $\hat{\Gamma}$-equivariant isomorphism of simplicial profinite sets.*

The real content of this result may not be clear at first encounter, so it may be useful to elucidate it in the case of vertices. First and by the very definitions: $\hat{C}(S)_0 = \hat{\mathcal{S}}$. Next recall the bijective map $\tau : \mathcal{S} \to \mathcal{T}$ between loops and twists. By the universal property of $\hat{C}(S)$ it extends to a map (with the same name) $\tau : \hat{\mathcal{S}} \to \bar{\mathcal{T}}$, the latter set being a profinite set with $\hat{\Gamma}$-action. This map is onto by the usual argument that the image is both dense and closed. Moreover we have a map $\gamma : \bar{\mathcal{T}} \to \overline{C}(S)_0$ which can be constructed as follows; consider $\bar{\mathcal{T}} \subset \hat{\Gamma}$ and map any element to the closed subgroup it generates. Then $\mathcal{T}$ is dense in $\bar{\mathcal{T}}$ and the image of $\mathcal{T}$ is contained in $\overline{C}(S)_0$; thus the image of $\bar{\mathcal{T}}$ is contained in $\overline{C}(S)_0$ (one could also use the explicit decomposition of $\bar{\mathcal{T}}$ and $\overline{C}(S)_0$ in $\hat{\Gamma}$-orbits). The map γ is onto as usual. Now we know from the theorem that the composite map $\gamma \circ \tau$ is an isomorphism, so both maps τ and γ are in fact isomorphisms. In particular we have proved:

Corollary 3.3: *The map τ is a $\hat{\Gamma}$-equivariant homeomorphism from $\hat{\mathcal{S}}$ to $\bar{\mathcal{T}}$.* □

Note that this shows that (indeed is equivalent to):

$$\bar{\mathcal{T}} \simeq \hat{\mathcal{T}} = \varprojlim_{\lambda \in \Lambda} \mathcal{T}/\Gamma^{\lambda}.$$

Choosing representatives $s \in F \subset \mathcal{S}$ for the Γ-action on $\mathcal{S}$, hence also for the $\hat{\Gamma}$-action on $\hat{\mathcal{S}}$, we see that the map τ can be described quite explicitly by: $\tau(\gamma \cdot s) = \gamma \tau_s \gamma^{-1}$, for any $s \in F \subset \mathcal{S}$ and $\gamma \in \hat{\Gamma}$. This completely determines the extension of τ to $\hat{\mathcal{S}}$ and the same formula remains valid for any $s \in \hat{\mathcal{S}}$ and any $\gamma \in \hat{\Gamma}$.

What Theorem 3.2 tells us in that case is that $\tau : \hat{\mathcal{S}} \to \bar{\mathcal{T}}$ is injective. This can actually be proved directly, starting from Lemma 3.1 and the reader may find it entertaining to look for a direct proof. Here we will translate this injectivity in terms of ordinary topology. What it says is that for any $s \in \mathcal{S}$ – we can indeed reduce to the case of a discrete curve – if $\tau_{\gamma \cdot s} = \tau_s \in \hat{\Gamma}$, then $\gamma \cdot s = s \in \hat{\mathcal{S}}$. Assume the first equality and write $\gamma \cdot s = \varprojlim_{\lambda} (\gamma_\lambda \cdot s)$ with $\gamma_\lambda \in \Gamma$ and $\gamma_\lambda \cdot s \in \mathcal{S}/\Gamma^{\lambda}$. Then for any $\lambda \in \Lambda$ and $\lambda' \in \Lambda$ large enough depending on λ, we find that $\gamma_{\lambda'} \tau_s \gamma_{\lambda'}^{-1} \in \Gamma^{\lambda} \tau_s$. Injectivity means that we can actually pick $\gamma_{\lambda'} \in \Gamma^{\mu}$ with $\mu \to \infty$ as $\lambda \to \infty$. This can be rephrased as the following statement:

Corollary 3.4: *For any $s \in \mathcal{S}$ and $\lambda \in \Lambda$ there exists $\mu \in \Lambda$ such that if $\tau_{\gamma \cdot s} \in \Gamma^{\mu} \tau_s$ for some $\gamma \in \Gamma$, one can find $\delta \in \Gamma^{\lambda}$ such that $\tau_{\delta \cdot s} = \tau_{\gamma \cdot s}$.* □

In other words, assume that $\gamma \tau_s \gamma^{-1} \in \Gamma^{\mu} \tau_s$, then $\gamma Z(\tau_s) \subset \Gamma$ intersects Γ^{λ}, where $Z(\tau_s) \subset \Gamma$ is the centralizer of τ_s.

We highlighted this technical looking statement partly because it is purely topological, dealing with discrete twists and discrete modular groups. It asserts that if for some (ordinary) twist τ and $\gamma \in \Gamma$, $\gamma \cdot \tau$ is close to τ, then one can find δ close to 1 in the sense of profinite topology, such that $\delta \cdot \tau = \gamma \cdot \tau$. One can speculate that such properties point to profinite (or nonarchimedean) analogs of the hyperbolicity properties of the complexes of curves, as recently investigated in the wake of [MM].

Returning to protwists, we proved above that $\bar{\mathcal{T}} \simeq \overline{C}(S)_0$, and we now know that $\bar{\mathcal{T}} \simeq \hat{\mathcal{T}}$. What the equality $\hat{\mathcal{T}} = \overline{C}(S)_0$ tells us is that a protwist is determined by the group it generates, just as in the discrete case. In other words we have the following:

Corollary 3.5: *Let $\tau, \tau' \in \hat{\mathcal{T}}$ and assume the free procyclic groups they generate in $\hat{\Gamma}$ coincide: $\langle \tau \rangle = \langle \tau' \rangle$. Then $\tau = \tau'$.* □

Although we stated this corollary explicitly, we are actually going to prove a stronger statement below. Recall that given a simplex $\sigma \in \hat{C}(S)$, we denote by G_σ its image in $\overline{C}(S)$, that is the (pro)free abelian group generated by the (pro)twists defined by the vertices of σ. With this notation Theorem 3.2, namely the injectivity of the canonical map $\hat{C}(S) \to \overline{C}(S)$, says that G_σ is determined by σ: if $G_\sigma = G_{\sigma'} \subset \hat{\Gamma}$, then $\sigma = \sigma' \in \hat{C}(S)$.

In order to improve on these results our technical tool will consist in a weighted version of Theorem 3.2. In order to state it properly, we first have to make sure that the topological type of a protwist $\tau \in \hat{\mathcal{T}} \subset \hat{\Gamma}$ is well-defined. We have seen above that such a twist is conjugate in $\hat{\Gamma}$ to a twist $\tau^{(0)} \in \mathcal{T} \subset \Gamma$. That (discrete) twist has a topological type; namely if $\tau^{(0)} = \tau_\gamma$ with $\gamma \in \mathcal{S}$, it is defined to be the type of the possibly disconnected surface $S \setminus \gamma$, that is S cut along γ. We define the topological type of τ to be that of $\tau^{(0)}$, which makes sense, thanks to the following:

Lemma 3.6: *Two discrete twists $\tau, \tau' \in \mathcal{T} \subset \Gamma$ are conjugate in $\hat{\Gamma}$ if and only if they have the same topological type. Consequently the topological type of a protwist is well-defined.*

We need only prove the only if part. But by [B1] (§7), if τ and τ' have different topological types, their centralizers $Z_{\hat{\Gamma}}(\tau)$ and $Z_{\hat{\Gamma}}(\tau')$ are not isomorphic. So they cannot be conjugate – otherwise their centralizers would be as well. □

Now, consider again the finite set F of the orbits of the vertices of $C(S)$ under the action of Γ, which is also the set of $\hat{\Gamma}$-orbits of the vertices of $\hat{C}(S)$. This set F simply enumerates the various topological types of the twists and thus also of the protwists. It is useful to keep in mind that it also enumerates the components of the boundary $\partial\mathcal{M}$ of $\overline{\mathcal{M}}$. Let $w : F \to \mathbb{Z}_+^*$ be a (non degenerate) weight function, assigning a strictly positive integer to any element of F. This also defines $w_s \in \mathbb{Z}_+^*$ for $s \in \mathcal{S}$ and more generally $w_\sigma \in (\mathbb{Z}_+^*)^{k+1}$ for any simplex $\sigma \in \hat{C}(S)_k$. We write w_s (resp. w_σ) or $w(s)$ (resp. $w(\sigma)$) indifferently.

Given a weight function w as above, we define $\overline{C}^w(S)$ as follows. Let first $\sigma \in C(S)$ be a discrete k-simplex, corresponding as usual to a multicurve $\underline{s} = (s_i)$ and twists τ_i. Then assign to σ the group $G_\sigma^w \subset \mathcal{G}$ topologically generated by the powers $\tau_i^{w_i}$ with $w_i = w(s_i)$; again $G_\sigma^w \simeq \hat{\mathbb{Z}}^{k+1}$. This determines an injective map $i_w : C(S) \to \mathcal{G}$ and we define $\overline{C}^w(S)$ to be the closure of the image $i_w(C(S))$. Clearly $\overline{C}(S)$ corresponds to using the weight function which is constant equal to 1. The complex $\overline{C}^w(S)_\bullet$ is again endowed with a continuous $\hat{\Gamma}$-action, hence by the universality of $\hat{C}(S)$ a natural surjective map $p : \hat{C}(S) \to \overline{C}^w(S)$.

Let us make two useful preliminary remarks. First the stabilizer of a simplex $\sigma \in \overline{C}^w(S)$ is by definition the normalizer $N_{\hat{\Gamma}}(G_\sigma^w)$ of the group G_σ^w in $\hat{\Gamma}$. Second we note that the complexes $\overline{C}^w(S)$ are flag complexes by definition: a simplex is uniquely determined by its faces, hence (by induction) by its vertices. The profinite complex $\hat{C}(S)$ is also a flag complex, thanks to the isomorphism given by Theorem 3.2. The weighted version of that result reads:

Theorem 3.7: *Let w be any non degenerate weight function as above:*
i) the natural map $p : \hat{C}(S) \to \overline{C}^w(S)$ is a $\hat{\Gamma}$-equivariant isomorphism of simplicial profinite sets.
ii) Let $\sigma \in C(S)_k$ ($k \geq 0$) be given by a set of commuting twists $\{\tau_0, \ldots, \tau_k\}$. Let G_σ^w be the closed abelian subgroup of $\hat{\Gamma}$ generated by $\tau_0^{w_0}, \ldots, \tau_k^{w_k}$ (for $w = 1$, the constant function, we denote this group simply by G_σ). Then, we have:

$$N_{\hat{\Gamma}}(G_\sigma^w) = Z_{\hat{\Gamma}}(G_\sigma^w) = N_{\hat{\Gamma}}(G_\sigma) = Z_{\hat{\Gamma}}(G_\sigma)$$

and this group coincides with the profinite completion $\hat{\Gamma}_\sigma$ of the centralizer of the τ_i's in Γ.

The proof in large part follows that of Theorem 3.2 (i.e. Theorem 7.3 in [B1]) so that we will sketch the strategy and detail only the necessary modifications. Since point *ii*) immediately follows from point *i*), it is enough to prove *i*). The proof proceeds by induction on the modular dimension $d(S) = d(S_{g,n}) = 3g - 3 + n$, the case of dimension 1 being settled much as in [B1]. So we assume the result has been proved for $d(S) < d$ and prove it for $d(S) = d$ $(d > 1)$. Next, again as in [B1] we can restrict attention to the respective 1-skeleta of $\hat{C}(S)$ and $\overline{C}^w(S)$; here we temporarily denote them $\hat{C}_1(S)$ and $\overline{C}_1^w(S)$ respectively and write p_1 for the restriction of p.

The first point is to show: $p_1 : \hat{C}_1(S) \to \overline{C}_1^w(S)$ *is an étale map*, which means that it is an isomorphism when restricted to the star of a given vertex in $\hat{C}_1(S)$.

Using the action of $\hat{\Gamma}$, we may and will assume that we are dealing with a vertex σ associated to an ordinary loop $\gamma \in \mathcal{S}$. We let S_σ and $\hat{S}_\sigma$ denote the stars of σ in $C_1(S)$ and $\hat{C}_1(S)$ respectively; $\bar{S}_\sigma$ in turn denotes the closure of the image of S_σ in $\overline{C}^w(S)$. Let $S' = S \setminus \gamma$ be the surface S cut along γ, of modular dimension $d(S) - 1$. For definiteness we assume that γ is non separating, so that S' is of type $(g-1, n+2)$. The separating case is treated in a completely analagous way (see [B1]). An element of S_σ is given by a loop γ' which does not intersect γ, so we can clearly identify S_σ with $C(S')_0$, the set of vertices of $C(S')$. From [B1], we find that actually $\hat{S}_\sigma \simeq \hat{C}(S')_0$. Proceeding again as in [B1] we also obtain that the map $\hat{S}_\sigma \simeq \hat{C}(S')_0 \to \overline{C}^w(S')_0$ factors through $\hat{S}_\sigma \to \bar{S}_\sigma$. Since the first map is an isomorphism by the induction hypothesis, the second has to be injective, hence an isomorphism, vindicating the above assertion, i.e. the fact that p_1 is an étale map.

It remains to prove that p_1 is actually an isomorphism, or else that it has degree 1, which can be checked by looking at the preimage of any vertex; again one can choose the image of a discrete non separating loop. The argument repeats again in large part the argument in [B1]; it involves introducing yet another graph such that the stabilizer of a vertex is given by the centralizer, not the normalizer of the corresponding twist or power of twist. We will not repeat the construction here but will concentrate on the key difference, which is actually quite nontrivial and group theoretic. Given any tree $T \subset \hat{C}_1(S)$, it determines a group $G_T^w \subset \mathcal{G}$, namely the closed group generated by the procyclic groups attached to its vertices.

Proceeding as in [B1], one reduces showing that p_1 is an isomorphism to ascertaining that there exists a finite tree T such that the centralizer of G_T^w in $\hat{\Gamma}$ coincides with the center of $\hat{\Gamma}$: $Z_{\hat{\Gamma}}(G_T^w) = Z(\hat{\Gamma})$. The cases of genus 0 and 1 can be settled fairly easily. When the genus of S is zero, the assertion is a consequence of the following lemma:

Lemma 3.8: *For $n \geq 4$, let $\tau_0, \ldots, \tau_k$ be a set of Dehn twist which generate $\Gamma_{0,n}$. Then, for any given weight function w, the centralizer in $\hat{\Gamma}_{0,n}$ of the subgroup $< \tau_0^{w_0}, \ldots, \tau_k^{w_k} >$ is trivial.*

For $n = 4$, the assertion is well-known. For $n \geq 5$, the lemma follows from the short exact sequences:

$$1 \to \hat{\pi}_{0,n} \to \hat{\Gamma}_{0,n+1} \to \hat{\Gamma}_{0,n} \to 1,$$

and a simple induction on n; here $\pi_{0,n} \simeq F_{n-1}$ denotes the topological fundamental group of an n-fold punctured sphere. □

The above lemma thus completes the proof of Theorem 3.7 for $g(S) = 0$. Similarly, for $g(S) = 1$, the theorem follows from:

Lemma 3.9: *For $n \geq 2$, let $\tau_0, \ldots, \tau_k$ be a set of Dehn twist which generate $\Gamma_{1,n}$. Then, for any given weight function w, the centralizer in $\hat{\Gamma}_{1,n}$ of the subgroup $< \tau_0^{w_0}, \ldots, \tau_k^{w_k} >$ is trivial.*

The Teichmüller group $\Gamma_{0,5}$ is naturally identified with the finite index normal subgroup of $\Gamma_{1,2}$ spanned by Dehn twists along separating circles and squares of Dehn twists along non-separating circles. The genus 0 case of Theorem 3.7 then implies the case $g = 1$ and $n = 2$, for weight functions which take an even value on non-separating circles. It is easy to see that this in fact settles the general case. In particular, if an element of $\hat{\Gamma}_{1,2}$ centralizes $\tau_i^{w_i}$, it centralizes τ_i as well, for $i = 0, \ldots, k$, i.e. it is in the center of $\hat{\Gamma}_{1,2}$, which we know is trivial. This completes the proof of the lemma for $n = 2$.

For $n \geq 3$, the lemma then follows as above from the short exact sequences:

$$1 \to \hat{\pi}_{1,n} \to \hat{\Gamma}_{1,n+1} \to \hat{\Gamma}_{1,n} \to 1,$$

and a simple induction on n. $\square$

For $g(S) \geq 2$, the idea is to use the results on centralizers of open subgroups contained in [B1]. It is then enough to find a finite tree T such that G_T^w is open in $\hat{\Gamma}$. For a given finite set of vertices, one can always find a (finite) tree which contains them. Since an open subgroup of $\hat{\Gamma}$ is topologically finitely generated, it is enough to prove that the subgroup of $\hat{\Gamma}$ generated by *all* the vertices (for the given weight function) is open. This however turns out to be quite a nontrivial assertion.

We will first prove it for a particular choice v of weight function, which will finish the proof of the theorem in that case and will then enable us to extend its validity to any w. So we define the weight function v (actually depending on a positive integer m) which assigns the weight m to all non separating loops and the weight 1 to all separating loops. By Corollary 3.11 of [BP] and its generalization, given in Theorem 2.5 of [B2], we have:

Lemma 3.10: *For $g \geq 2$ and any integer $m \geq 1$, the group generated by all m-th powers of twists and all separating twists is open in $\hat{\Gamma}_{g,n}$.*

The above lemma finishes the proof of the theorem in the particular case of a weight function of type v (for any $m \geq 1$). In particular, given a nonseparating loop $\gamma \in \mathcal{S}$, and an integer $m \geq 1$, from point *ii*), we have that $N_{\hat{\Gamma}}(\tau_\gamma^m) = N_{\hat{\Gamma}}(\tau_\gamma)$ and the latter group is identified with $\hat{\Gamma}_\gamma$, the stabilizer of γ as a vertex of $\hat{C}(S)$, whereas the former is the stabilizer of γ as a vertex of $\overline{C}^v(S)$. Note that from [B1] we know that the notation $\hat{\Gamma}_\gamma$ is justified, i.e. it is indeed the profinite completion of the stabilizer Γ_γ of γ in the discrete complex $C(S)$. Moereover we also know (although we will not need it) that $N_{\hat{\Gamma}}(\tau_\gamma) = Z_{\hat{\Gamma}}(\tau_\gamma)$.

Consider now an arbitrary non degenerate weight function w. We know that the map p_1 is étale, and are again trying to determine whether it has degree 1 by looking above a vertex associated to a nonseparating loop γ. This is equivalent to comparing the stabilizers of γ viewed as a vertex of $\hat{C}(S)$ and of $\overline{C}^w(S)$ respectively. But these groups are equal from what we have seen because the stabilizer of $\gamma \in \overline{C}^w(S)_0$ is indeed equal to $N_{\hat{\Gamma}}(\tau_\gamma^m)$ with $m = w(\gamma)$ and so is independent of m. This finishes the proof of the theorem. $\square$

An easy generalization of point *ii*) of Theorem 3.7 is the following:

Corollary 3.11: *Let $\sigma \in C(S)_k$ ($k \geq 0$). Then for any finite index subgroup $U \subset G_\sigma$:*

$$N_{\hat{\Gamma}}(U) = N_{\hat{\Gamma}}(G_\sigma) = Z_{\hat{\Gamma}}(G_\sigma).$$

□

From this corollary and the fact that $Z_{\hat{\Gamma}}(G_\sigma) \simeq \hat{\Gamma}_\sigma$ determines σ we obtain the following useful uniqueness result:

Corollary 3.12: *If $\sigma, \sigma' \in \hat{C}(S)$ are two non degenerate simplices such that the intersection $G_\sigma \cap G_{\sigma'} \subset \hat{\Gamma}(S)$ is open in either of these two groups, then $\sigma = \sigma'$.* □

But we actually expect more, namely a kind of lattice property for the image of $\hat{C}(S)$ in $\mathcal{G}$. It may be useful to state that property explicitly, starting from the discrete case. So we momentarily turn to the discrete setting, with the complex $C(S)$ and $\mathcal{G}^{disc}(\Gamma)$, the set of all subgroups of Γ. All objects now pertain to the discrete topology. In particular to every $\sigma \in C(S)$ we assign the (discrete) free abelian group $G_\sigma^{disc} \in \mathcal{G}(\Gamma)$ spanned by the twists associated to σ. In this context we have the following statement:

Proposition 3.13: *For any two non degenerate simplices $\sigma, \sigma' \in C(S)$ and their associated groups $G_\sigma^{disc}, G_{\sigma'}^{disc} \subset \Gamma(S)$, one has:*

$$G_\sigma^{disc} \cap G_{\sigma'}^{disc} = G_{\sigma\cap\sigma'}^{disc},$$

where $\sigma \cap \sigma' \in C(S)$ is the simplex spanned by the vertices common to σ and σ' (with $G_\emptyset^{disc} = \{1\}$).

In order to prove the proposition, one needs to study identities of the form:

$$\prod_i \tau_i^{m_i} = \prod_j (\tau_j')^{m_j'},$$

in which the τ_i's (resp. τ_j''s) are twists associated with σ (resp. σ'), and $m_i, m_j' \in \mathbb{Z}$. Let us use a multiindex notation and write more compactly $\tau_\sigma^{m_\sigma}$ for the left-hand side, idem for the right-hand side, and call such expressions multitwists. So we write $\tau_\sigma^{m_\sigma} = (\tau_{\sigma'}')^{m_\sigma'} = f$ (this defines f). We can assume that all the m_i's and m_j' are nonzero and we then want to prove that $\sigma = \sigma'$. Now f commutes with τ_j' for all j by looking at its second expression. So τ_j' commutes with the product $\tau_\sigma^{m_\sigma} (= f)$. Assume for a moment that it implies that τ_j' actually commutes with all the τ_i. Then by symmetry we find that all the τ_i and τ_j' form a set of mutually commuting twists. After rearranging we are reduced to showing that a multitwist $\tau_\sigma^{m_\sigma}$ is trivial ($= 1$) if and only if the m_i's vanish, which is obvious (the τ_i are distinct by assumption).

The only serious assertion here is the one we have temporarily assumed, namely that a twist τ commutes with a multitwist $\tau_\sigma^{m_\sigma}$ (if and) only if it commutes with the twists appearing in τ_σ (all the m_i are assumed to be nonzero). This is a consequence of classical results, especially Lemma 4.2 in [I2]. □

One can actually prove more than this last assertion, namely that the centralizer of a multitwist is the intersection of the centralizers of the twists which appear in it. We insist on these statements because their proofs use relatively elementary but typically "archimedean" techniques (e.g. geometric intersection numbers) which are not available in the profinite context so that we suspect but cannot prove that their obvious profinite analogs hold true. For instance, we expect that Proposition 3.13 holds verbatim in the profinite setting but are unable to prove it at the moment. So we stop here this exploration of the profinite curves complex $\hat{C}(S)$ *per se* and will shortly proceed to investigate their automorphisms.

As a final remark in this section however, we note that the above *ipso facto* provides a description of the vertices of the profinite pants graph $\hat{C}_P(S)$, since these coincide with the highest dimensional simplices of $\hat{C}(S)$. However, as mentioned already, we do not have at present a good description of the edges of $\hat{C}_P(S)$, even in the simplest and most important case of modular dimension 1, that is for the Farey prograph $\hat{F}$ (see §2). In other words we are unable to provide an adequate profinite generalization of the "intersection 1 condition" which defines the edges of the discrete pants graph $C_P(S)$.

4. On the automorphism groups of profinite complexes of curves

In this section we investigate the automorphism groups of the complexes $\hat{C}(S)$ and $\hat{C}_P(S)$. It will be plain that concerning the first of these, our present knowledge is fragmentary and that the theme obviously deserves further study. First of all we are refering here of course to continuous automorphisms, which makes sense since for any hyperbolic surface S, $Aut(\hat{C}(S))$ and $Aut(\hat{C}_P(S))$ inherit a natural structure of profinite groups which can be described as follows, working e.g. with $\hat{C}(S)$ for definiteness. This complex is defined as the inverse limit of the finite complexes $C^\lambda(S) = C(S)/\Gamma^\lambda$ for the system of levels $\lambda \in \Lambda$. A continuous automorphism, which is also open, defines for any $\lambda \in \Lambda$ a map $C^\lambda(S) \to C^\mu(S)$ for some $\mu \in \Lambda$ $(\lambda \geq \mu)$. When varying $\lambda \in \Lambda$, a basis of neighborhoods of the identity in $Aut(\hat{C}(S))$ is given by those automorphism which induce the natural projection; note that these open neighborhoods are not subgroups.

We now start investigating the automorphisms of the complex $\hat{C}(S)$ where the surface S is hyperbolic and $d(S) > 1$. The one dimensional case pertains more to the pants graph. In the discrete case, a twist has a topological type and it makes sense to define $Aut^\sharp(C(S)) \subset Aut(C(S))$ as the subgroup of type preserving automorphisms. Note that in group theoretic terms, two twists have the same type if and only if they are conjugate. A basic result (see e.g. [L], §2) states that in fact $Aut^\sharp(C(S))$ coincides with the whole of $Aut(C(S))$ if S is not of type $(1, 2)$. So in the discrete case, all automorphisms of the curves complex are inertia preserving $(Aut^*(C(S) = Aut(C(S)))$ and type preserving $(Aut^\sharp(C(S) = Aut(C(S)))$ with the exception of type $(1, 2)$, which itself is well-understood. Moreover, by Theorem 2.13, automorphisms of $C_P(S)$ are inertia and type preserving too. Actually $(1, 2)$ is no exception then (cf. [M], §5).

In the profinite case we do not know how to compare $Aut^*(\hat{C}(S)$ and $Aut(\hat{C}(S))$ although we suspect these groups coincide. Again this looks like a deep problem, which amounts to set up a "local correspondence". On the other hand, Lemma 3.6 ensures that a protwist has a well-defined topological type: actually two (pro)curves have the same type if and only if the associated (pro)twists are conjugate in $\hat{\Gamma}(S)$. So it makes sense to define $Aut^\sharp(\hat{C}(S)) \subset Aut(\hat{C}(S))$ as the subgroup of type preserving automorphisms. We will now essentially get rid of that decoration, as in the discrete case, and mostly following the discrete proof:

Proposition 4.1: *Assume S is hyperbolic with $d(S) > 1$, then:*

i) $Aut(\hat{C}^{(1)}(S)) = Aut(\hat{C}(S))$*;*

ii) if S *is not of type* $(1,2)$*,* $Aut^\sharp(\hat{C}(S)) = Aut(\hat{C}(S))$*.*

i) is the same as in the discrete case and is proved in the same way. One uses the fact that the orbit of any simplex of $\hat{C}(S)$ contains a discrete simplex (a simplex in $C(S)$) to ascertain that simplices in $\hat{C}(S)$ are exactly the complete subgraphs of $\hat{C}(S)$; see the proof of Theorem 2.1 ii). We included this statement for completeness and because it shows how prographs will tell the whole story.

As mentioned above, the proof of ii) also mimicks that of the corresponding statement in the discrete case. In fact most of §2 in [L] holds essentially *verbatim* in our profinite context so that we only recall the main steps. First, thanks to the fact that we know the homotopy type of $\hat{C}(S)$ ([B1], Theorem 5.2) and that it is again the profinite analog of the homotopy type of $C(S)$, we can distinguish the complexes $\hat{C}(S)$ for various surfaces S in the same way as in the discrete case. Namely there holds (cf. [L], Lemma 2.1):

Lemma 4.2: *Let* S *and* S' *be two connected hyperbolic surfaces of different types. Then* $\hat{C}(S)$ *and* $\hat{C}(S')$ *are not isomorphic, save for the following exceptional cases:* $\hat{C}(S_{2,0}) \simeq \hat{C}(S_{0,6})$*,* $\hat{C}(S_{1,2}) \simeq \hat{C}(S_{0,5})$ *and* $\hat{C}(S_{1,1}) \simeq \hat{C}(S_{0,4}) \simeq \hat{F}$*.*

We mentioned the one dimensional case for the sake of completeness only; the isomorphism is then tautological, provided that $C(S_{1,1})$ and $C(S_{0,4})$ are actually redefined to be both isomorphic to F, which as explained in §2 is not necessarily desirable.

There are only two remarks to be made which make clear that the proof in the discrete case ([L], Lemma 2.1) extends to the profinite situation, modulo again the result on the homotopy type of profinite curve complexes. First the exceptional isomorphisms occuring in the discrete case induce isomorphisms of the corresponding profinite complexes. For instance, one has $C(S_{1,2}) \simeq C(S_{0,5})$. Then $\hat{C}(S_{1,2})$ (resp. $\hat{C}(S_{0,5})$) is the completion of that complex with respect to the action of $\Gamma_{1,[2]}$ (resp. $\Gamma_{0,[5]}$). However, $\Gamma_{1,[2]}$ acts via the quotient by its center $\Gamma_{1,[2]}/Z$ ($Z = Z(\Gamma_{1,[2]}) \simeq \mathbb{Z}/2$) and we have an inclusion $\Gamma_{1,[2]}/Z \subset \Gamma_{0,[5]}$ where $\Gamma_{1,[2]}/Z$ can be identified with the stabilizer of one of the 5 marked points, so is of finite index ($= 5$) inside $\Gamma_{0,[5]}$. This implies that $\hat{C}(S_{1,2}) \simeq \hat{C}(S_{0,5})$.

The second remark is that the topological looking reasoning in the proof of Lemma 4.1 (b) in [L] can be made complex theoretic (see also [I1]). More generally if Σ is a hyperbolic surface, one can recognize on $C(\Sigma)$ (and also on $\hat{C}(\Sigma)$) the connected components of Σ which are not of type $(0,3)$. Let us confine ourselves to the situation which is useful in the proof of the lemma. So consider S a connected hyperbolic surface, $\gamma \in \mathcal{S}(S)$ a loop on S and $S_\gamma = S \setminus \gamma$ the surface cut along γ. First observe that one retrieves $C(S_\gamma)$ as the link of $\gamma \in C(S)^{(0)}$ in $C(S)$. Then let us restrict to the 1-skeleton $C(S)^{(1)}$, so that the link L_γ of γ is (isomorphic) to the 1-skeleton of $C(S_\gamma)$. Let now L_γ^- be the graph defined as follows: the vertices of L_γ^- are the same as that of L_γ, thus corresponding to curves on S disjoint from γ; two vertices of L_γ^- are joined by an edge if and only if this is *not* the case in L_γ. Edges thus correspond to pairs of curves which are disjoint of γ and intersect. Then clearly the connected components of L_γ^- are in one-to-one correspondence with the components of S_γ which are not of type $(0,3)$ ($C(S_{0,3}) = \emptyset$).

Using the above remarks, Lemma 4.2 follows from its discrete version after recalling the oft-mentioned fact that any $\hat{\Gamma}$-orbit of a simplex in $\hat{C}(S)$ contains a discrete representative. $\square$

The next step in proving Proposition 4.1 is the following statement which we record explicitly:

Lemma 4.3: *For any hyperbolic surface S, an automorphism of $\hat{C}(S)$ is type preserving if and only if it preserves the separating classes.*

Only the 'if' part needs proof and it immediately follows from the profinite transposition of the proof of Lemma 2.3 in [L], using Lemma 4.2 above. $\square$

So the only issue for an automorphism to be type preserving is whether it maps a separating (resp. non separating) proloop to a separating (resp. non separating) one. This in turn is the object of Lemma 2.2 in [L] which again, given the above can be readily transposed to the profinite setting, thus completing the proof of Proposition 4.1. $\square$

Our next objective is the profinite analog of Theorem 2.13, which is interesting for its own sake and as a byproduct will ensure, via Proposition 4.1, that automorphisms of the pants graph $\hat{C}_P(S)$ are type preserving. To this end we need to make a detour and prove the analog of Theorem 2.10, which also has independent interest. So we first introduce $\hat{C}_*(S)$, the $\Gamma(S)$-completion of the graph $C_*(S)$ defined in §2. If S is connected and $d(S) > 1$ this is the dual of $\hat{C}^{(1)}(S)$. True, if S is not connected, say $S = \coprod_i S_i$ with connected pieces S_i's, one should define here $\Gamma(S) = \prod_i \Gamma(S_i)$ (a colored modular group) and let each $\Gamma(S_i)$ act naturally on $C_*(S_i)$, so as to extend definitions to non connected situations. Note that possible permutations of the pieces have no effect on completions, since they generate a finite group; in other words, the colored modular group has finite index in the full modular group. Note also that one can proceed in the same way with $\hat{C}_P(S)$ so that this is defined for non connected S as well.

If S is connected with $d(S) = 0$ (a pair of pants), $\hat{C}_*(S) = \hat{C}_P(S)$ is reduced to a point and coincides with its discrete version. If S is connected with $d(S) = 1$, $\hat{C}_*(S) = \hat{G}$ and $\hat{C}_P(S) = \hat{F}$, where the completion can be taken with respect to the natural action of $F_2 \simeq \Gamma_{0,4}$ (which has finite index in $\Gamma_{1,1}$).

If $d(S) > 1$ one can identify $\hat{C}_*(S)$ with the 1-skeleton of the dual of $\hat{C}(S)$ but it is more interesting to give a direct description of $\hat{C}_*(S)$ which is valid in all dimensions $d(S) \geq 0$. From now on we will again skip the prefix "pro", speaking e.g. of graphs rather than prographs, certainly not of prosubgraphs or subprographs; we rather add the adjective "discretei" to point out that certain objects belong to the discrete world. There is a natural action of $\hat{\Gamma}(S)$ on $\hat{C}_*(S)$ and $\hat{C}_P(S)$ and one can describe their common set of vertices $\hat{V}(S)$ as usual, namely as the finite sum $\coprod_{v \in F} \hat{\Gamma} \cdot v$ of $\hat{\Gamma}$-orbits of discrete vertices $v \in V(S)$ representing the types of pants decompositions of S labeled by the finite set F. The set $\hat{E}(S)$ of edges of $\hat{C}_*(S)$ can be described simply as follows:

Lemma 4.4: *The vertices $v, w \in \hat{C}_*(S)$ are joined by an edge if and only the corresponding maximal multicurves differ by exactly one component up to relabeling.*

The satement should be interpreted as follows. Write $v = (\alpha_1, \ldots, \alpha_k)$ (resp. $w = (\beta_1, \ldots, \beta_k)$) where the α_i's and β_j's are (pro)curves. One could assume that v or w corresponds to a discrete pants decomposition but that does not really help. The claim is that the condition for v and w to be joined by an edge in $\hat{C}_*(S)$ is the exact analog of what happens in the discrete case.

The "if" part of the statement is clear and we have to show the "only if" part. In order to do this, let $v = \varprojlim_{\lambda \in \Lambda} v^\lambda$, $w = \varprojlim_{\lambda \in \Lambda} w^\lambda$ where $\lambda \in \Lambda$ runs as usual along the modular levels (we may assume S connected for simplicity) and $v^\lambda, w^\lambda \in C_*^\lambda = C_*(S)/\Gamma^\lambda$. We can write $v^\lambda = (\alpha_i^\lambda)$, $w^\lambda = (\beta_i^\lambda)$ where the α_i^λ and β_i^λ represent Γ^λ-orbits of curves (i.e. they lie in $\mathcal{S}(S)/\Gamma^\lambda$). Moreover since v and w are joined by an edge in $\hat{C}_*$, there exist discrete pants decompositions (A_i^λ), (B_j^λ) in C_* which project to v^λ and w^λ repectively and are joined by an edge in C_*. So (A_i^λ) and (B_j^λ) differ by at most one curve, up to relabeling. For any $\lambda \in \Lambda$ consider the label (in $\{1, \ldots, k\}$) of the curve in the family (A_i^λ) which does not occur in (B_j^λ) (if they coincide pick any label). This may depend also on the chosen lifts of v^λ and w^λ but that does not matter. Now consider a cofinal sequence in Λ and choose a label which occurs infinitely often in the above construction. One finds that v and w can indeed be represented by multicurves (α_i) and (β_i) which coincide save for the entry in v corresponding to that label. □

We can view $\hat{C}_P(S) \subset \hat{C}_*(S)$ as a closed subgraph, with the same set $\hat{V}(S)$ of vertices and a set $\hat{E}_P(S) \subset \hat{E}(S)$ of edges. However, as mentioned several times already, it is not so easy to describe the set $\hat{E}_P(S)$. The notion of algebraic intersection is easily extended to the profinite setting (cap product on homology with $\hat{\mathbb{Z}}$-coefficients) but not so for the geometric intersection; it takes its values in $\mathbb{Z}_+$ which already has no profinite counterpart. We can get a glimpse of the (semi)local structure of $\hat{C}_*$ and $\hat{C}_P$ as follows:

Lemma 4.5: *The analog of Lemma 2.9 holds true for $\hat{C}_*(S)$ and $\hat{C}_P(S)$, replacing G and F by $\hat{G}$ and $\hat{F}$ respectively.*

For $\hat{C}_*(S)$ this is actually a consequence of Lemma 4.4, which enables us repeat the discrete proof. It is possible although not really useful to work with $v \in V(S)$ a discrete vertex. One gets the statement of the lemma for $\hat{C}_*(S)$ except for a possible important difference. Namely one finds that there are copies of some completions $\overline{G}_i$ of G which satisfy the conditions of the statement of Lemma 2.9. That $\overline{G}_i \simeq \hat{G}_i$ is indeed the full profinite completion and not merely a quotient thereof comes from [B1] (Theorem 7.1). Finally one uses that $\hat{C}_P(S)$ is a closed subcomplex of $\hat{C}_*(S)$ in order to conclude the proof in that case. □

From Lemma 4.5, one concludes that the profinite analog of Lemma 2.7 holds true and so does the analog of Lemma 2.8. Actually this last result depends only on the fact that $\hat{C}_*(S)$ and $\hat{C}_P(S)$ share the same set $\hat{V}(S)$ of vertices and that they are flag complexes (an edge is determined by its endpoints) so that any automorphism is determined by its action on the set of vertices. We record it as:

Lemma 4.6: *$Aut(\hat{C}_P(S)) \subset Aut(\hat{C}_*(S))$.* □

Let us now proceed toward the profinite version of Theorem 2.10. Recall from §2 that for $\sigma \in C(S)$ we defined the subsurface $S_\sigma = S \setminus \sigma$ as S slit along σ. The natural action of $\Gamma(S)$ on $C(S)$ translates into an action of $\Gamma(S)$ on the graphs $C_*(S_\sigma)$, viewed as subgraphs of $C_*(S)$. For $g \in \Gamma(S)$, $\sigma \in C(S)$, we have the simple formula: $g \cdot C_*(S_\sigma) = C_*(S_{g\cdot\sigma})$; this also holds with C_* replaced by C_P and for C_*, it says that the reconstruction principle obtained in Corollary 2.11 respects the natural Γ-action.

In the profinite case we first note that for every $\sigma \in C(S)$, $\overline{C}_*(S_\sigma)$ defined as the closure of $C_*(S_\sigma)$ in $\hat{C}_*(S)$ actually coincides with the full completion $\hat{C}_*(S_\sigma)$. This comes from [B1] (§7) as in the proof of Lemma 4.5 and we insist again on the significance of this fact which ultimately represents a geometric result about the abundance of finite étale covers of the moduli stacks. Using the natural action of $\hat{\Gamma}(S)$ on $\hat{C}(S)$, we find that $g \cdot \hat{C}_*(S_\sigma)$ is a well-defined closed subgraph of $\hat{C}_*(S)$ for $g \in \hat{\Gamma}(S)$ and $\sigma \in C(S)$. At this point we can write formally: $g \cdot \hat{C}_*(S_\sigma) = \hat{C}_*(S_{g \cdot \sigma})$ for any $g \in \hat{\Gamma}(S)$, $\sigma \in \hat{C}(S)$. The right-hand side is actually defined by the left-hand side when $\sigma \in C(S)$ is a discrete simplex and the definition extends to any $\sigma \in \hat{C}(S)$ using as usual the fact that the $\hat{\Gamma}(S)$-orbit of any simplex in $\hat{C}(S)$ contains a discrete representative.

One gets a family $(\hat{C}(S_\sigma))$ of closed subgraphs of $C_*(S)$ which is indexed by the profinite simplicial set $\hat{C}(S)$ and is endowed with a natural action of $\hat{\Gamma}(S)$. These subgraphs are distinct, that is $\hat{C}(S_\sigma) = \hat{C}(S_\tau)$ if and only if $\sigma = \tau \in \hat{C}(S)$. In order to show this, it is enough to prove that for any discrete $\sigma \in C(S)$ and any $g \in \hat{\Gamma}(S)$, $g \cdot \hat{C}_*(S_\sigma) = \hat{C}_*(S_\sigma)$ if and only if $g \cdot \sigma = \sigma$, which is not difficult, granted again the results of [B1] (§7). As in the discrete case, reconstructing $\hat{C}(S)$ out of $\hat{C}_*(S)$ consists in graph theoretically detecting or characterizing the family $(\hat{C}(S_\sigma))$, which can be made into a prosimplicial complex using the inclusion as a boundary operator.

In what follows, for $\tau \in \hat{C}(S)$, one can think of $\hat{C}_*(S_\tau)$ via the defining formula $\hat{C}_*(S_\tau) = g \cdot \hat{C}_*(S_\sigma)$ for $\sigma \in C(S)$, $g \in \hat{\Gamma}(S)$, $g \cdot \sigma = \tau$, and avoid making sense directly of the symbol S_τ, that is "S slit along the profinite simplex τ". It may be worth pointing out the possible connection with what Grothendieck calls *discrétifications* in his *Longue Marche à travers la théorie de Galois*. Roughly speaking, and to be specific, given a finitely generated residually finite group G and its profinite completion $\hat{G}$ one can consider the set of its discretifications, that is of the dense injections $G \hookrightarrow \hat{G}$. This is a natural extension of the notion of integral lattice or integral structure in the linear situation. These discretifications will form a torsor under a group which does not seem easy to capture but may be worth keeping in mind. In an analogous way one can view the set of dense embeddings $C(S) \hookrightarrow \hat{C}(S)$ as the set of integral structures on $\hat{C}(S)$ and in that context the above seemingly formal definitions become more natural, since the group $\hat{\Gamma}(S)$ will act naturally on these structures ("discretifications") as well.

We can now state the profinite version of Theorem 2.10 as:

Theorem 4.7: *Let $C \subset \hat{C}_*(S)$ be a subgraph which is (topologically) isomorphic to $\hat{C}_*(\Sigma)$ for a certain surface Σ and is maximal with this property. Then there exists a unique $\sigma \in \hat{C}(S)$ such that $C = \hat{C}_*(S_\sigma)$.*

If one wishes to stick to S_σ for discrete simplices $\sigma \in C(S)$, the assertion can be rephrased as saying that there exist $\sigma \in C(S)$, $g \in \hat{\Gamma}(S)$ such that $C = g \cdot \hat{C}_*(S_\sigma)$. Two solutions (σ, g) and (σ', g') satisfy $g \cdot \sigma = g' \cdot \sigma' \in \hat{C}(S)$. As in the discrete setting, the case $\sigma = \emptyset$ should be included and corresponds to the full complex $\hat{C}_*(S)$.

The proof proceeds along the same line as in the discrete case. Thanks to Lemma 4.5 we may graph theoretically detect the dimension of a surface just as in the discrete case, and we may do so locally around any vertex; see the proof of Lemma 2.9. Turning to the proof of the theorem, we need only show the existence part. The first step consists in showing that a subgraph of the form indicated in the statement is maximal. To this end, one can consider a discrete $\sigma \in C(S)$ and prove that $\hat{C}_*(S_\sigma)$ is maximal in its isomorphy class. Using Lemma 4.5, the proof literally follows the one in the discrete case.

Then one takes care of the low dimensional cases, $d(\Sigma) = 0, 1$. If $d(\Sigma) = 0$, $\hat{C}_*(\Sigma)$ is reduced to a point, and up to the action of $\hat{\Gamma}(S)$ this corresponds to a pants decomposition of S, as in the discrete case. If $d(\Sigma) = 1$, it is still true that any triangle in $\hat{C}_*(S)$ defines a unique subsurface of dimension 1, possibly after twisting by an element of $\hat{\Gamma}(S)$. This was implicitly used in proving Lemma 4.5; it comes directly from Lemma 4.4 and the following fact: given Σ with $d(\Sigma) = 1$, any two curves in $\hat{C}(\Sigma)$ intersect. Otherwise, up to acting via $\hat{\Gamma}(S)$ one would find a pair of non intersecting discrete curves on Σ, which is absurd.

Having disposed of the low-dimensional cases, we argue again by induction on $k = d(\Sigma)$. So we pick $k > 1$, assume the statement is true for $d(\Sigma) < k$ and fix an isomorphism $C \xrightarrow{\sim} \hat{C}_*(\Sigma)$. For $\sigma \in C(S)$ we then define $C_\sigma \simeq \hat{C}_*(\Sigma_\sigma) \subset \hat{C}_*(S)$ as in the discrete case. This time the union of the C_σ's as σ runs over the nonempty simplices of $C(\Sigma)$ form a dense part of C, which is sufficient for the same argument as in the discrete case to go through. Namely in order to conclude the proof, it is enough to show that there exists a k-dimensional subsurface $T \subset S$ and an element $g \in \hat{\Gamma}(S)$ such that for any (nonempty) $\sigma \in C(\Sigma)$, $C_\sigma \subset g \cdot \hat{C}_*(T) \subset \hat{C}_*(S)$.

We may again (as in the discrete case) restrict to $|\sigma| = 1$, i.e. to the discrete loops on Σ. To any such loop $\alpha \in \mathcal{S}(\Sigma)$ we can attach by induction a subsurface $S_{(\alpha)} \subset S$ of dimension $k - 1$ and an element $g_\alpha \in \hat{\Gamma}(S)$ such that $C_\alpha = g_\alpha \cdot \hat{C}_*(S_{(\alpha)}) \subset \hat{C}_*(S)$. As usual, having fixed an isomorphism $C \xrightarrow{\sim} \hat{C}_*(\Sigma)$ we write an equality sign for the sake of simplicity.

Next we use, again as in the discrete case, the connectedness of $C(\Sigma)$ which is ensured by the assumption on the dimension ($k > 1$). So we have to study the following situation. We consider three discrete loops α, β, and γ on Σ such that $\alpha \cap \beta = \beta \cap \gamma = \emptyset$. We attach to them as above pairs $(g_\alpha, S_{(\alpha)} = S_\rho)$, $(g_\beta, S_{(\beta)} = S_\sigma)$ and $(g_\gamma, S_{(\gamma)} = S_\tau)$ for certain simplices $\rho, \sigma, \tau \in C(S)$ with $|\rho| = |\sigma| = |\tau| = d(S) - k + 1$. Moreover ρ and σ (resp. σ and τ) are compatible simplices.

As in the discrete case, the situation should be entirely determined by any two pairs of non intersecting curves on Σ, after which one can worry over a possible overdetermination. The reasoning below may appear more transparent if one recalls that a graph of the form $g \cdot \hat{C}_*(S_\sigma)$ is actually determined by the profinite simplex $g \cdot \sigma$ and so depends on g only up to the subgroup of $\hat{\Gamma}(S)$ fixing σ, which is nothing but the centralizer of the multitwist corresponding to σ. These centralizers are determined in [B1] (§7). So let us first examine what happens when trying to paste the data for α and β. After twisting we may assume that $g_\alpha = 1$ and write $g_\beta = g \in \hat{\Gamma}(S)$. Next we know that the intersection $C_\alpha \cap C_\beta$ has dimension $k - 2$ and indeed is isomorphic to a twist of $\hat{C}_*(\Sigma_{\alpha \cup \beta})$. This implies that $|\rho \cap \sigma| = d(S) - k$ and that g fixes $\varpi = \rho \cap \sigma$, that is $g \in Z_\varpi$. Writing $T = S_\varpi$ we find that $S_{(\alpha)} = S_\rho \subset T$. Moreover, because g fixes ϖ we can find $h \in \hat{\Gamma}(T)$ such that $C_\beta = g \cdot \hat{C}_*(S_\sigma) = h \cdot \hat{C}_*(S_\sigma)$. But then, since $h \in \hat{\Gamma}(T)$, $h \cdot \hat{C}_*(S_\sigma) \subset \hat{C}_*(T)$ and so we get the inclusion $C_\beta \subset \hat{C}_*(T)$. Returning to our original notation, we found a k-dimensional subsurface $T \subset S$ such that $C_\alpha \subset g_\alpha \cdot \hat{C}_*(T)$, $C_\beta \subset g_\beta \cdot \hat{C}_*(T)$ and in fact $g_\beta = g_\alpha = g$. Proceeding in the same way with the pair (β, γ) we get a possibly different pair (g', T'). Now in order to compare T and T', we use again the fact that there is a large intersection, namely that $C_\beta \subset g \cdot \hat{C}_*(T) \cap g' \cdot \hat{C}_*(T')$. This implies that one can modify – say – g' in order to achieve $g = g'$ and then, because T, T' and Σ are all of dimension k, one shows as in the discrete case that $T = T'$. $\square$

We can now draw the consequences of Theorem 4.7 much as in the discrete case. It yields first:

Corollary 4.8: *$\hat{C}(S)$ can be reconstructed from $\hat{C}_*(S)$.*

In fact, as mentioned above, one reconstructs $\hat{C}(S)$ by considering the set of subgraphs of $\hat{C}_*$ satisfying the conditions stated in Theorem 4.7 and making it into a prosimplicial complex using the inclusion as a boundary operator. Theorem 4.7 ensures that the resulting complex is isomorphic to $\hat{C}(S)$. $\square$

From there and Lemma 4.6, we directly get the profinite analogs of Corollaries 2.12 and 2.13, of which we record only the second for the sake of brevity:

Corollary 4.9: *There is a natural injective map: $Aut(\hat{C}_P(S)) \hookrightarrow Aut(\hat{C}(S))$.* $\square$

It will evolve that contrary to what happens in the discrete case, these two groups are far from equal. Indeed the difference (or quotient rather) should essentially consist of the Grohendieck-Teichmüller group. For the time being, Corollary 4.9 together with Proposition 4.1 imply the following useful statement:

Corollary 4.10: *If S is hyperbolic and not of type $(1,2)$, every automorphism of $\hat{C}_P(S)$ is type preserving:*

$$Aut(\hat{C}_P(S)) = Aut^\sharp(\hat{C}_P(S)).$$

$\square$

The statement is actually empty if $d(S) = 1$, so we included that case only formally. More interesting (and unfortunate) is the fact that because of our poor knowledge of $\hat{C}_P(S)$ we were not able to adapt to the profinite setting the specific argument (see [M], §5) which shows that here, contrary to the case of $C(S)$, type $(1,2)$ is not an exception, that is any automorphism of $C_P(S_{1,2})$ is indeed type preserving. We thus have to leave this as a possible exceptional case.

The next statement, which uses Theorem 3.7 in a crucial way should serve to emphasize how computing the automorphisms of curves complexes enables one to study the automorphisms of open subgroups of modular groups in a uniform way. In the discrete setting the (easier) analogous statement leads directly to Theorem 2.4 when combined with Theorem 2.1. Here for the sake of simplicity we will restrict attention to automorphisms rather than general isomorphisms. As usual, the cases with nontrivial center i.e. types $(1,2)$ and $(2,0)$ could be studied specifically much as in the discrete case:

Proposition 4.11: *Assume S is hyperbolic with $d(S) \geq 1$, then for any level $\lambda \in \Lambda$ there is a natural morphism:*

$$n_\lambda : Aut^*(\hat{\Gamma}^\lambda(S)) \to Aut(\hat{C}(S)).$$

This morphism is injective if $\Gamma^\lambda(S)$ has trivial center, thus in particular if $\Gamma(S)$ itself has trivial center.

Recall from §3 that for σ a simplex in $\hat{C}(S)$ we denote by $G_\sigma \subset \Gamma(S)$ the subgroup topologically generated by the (pro)twists attached to the vertices of σ. For any $\sigma \in \hat{C}(S)$, and any level $\lambda \in \Lambda$, $\hat{\Gamma}^\lambda \cap G_\sigma = U_\sigma$ is open in G_σ. Then for $f \in Aut^*(\hat{\Gamma}^\lambda(S))$, we define $\tilde{f} \in Aut(\hat{C}(S))$ through the equality: $f(U_\sigma) = U_{\tilde{f}(\sigma)}$. Here given a simplex τ, U_τ denotes a "generic" open subgroup of G_τ (not a fixed one). The important point is that Corollary 3.12 ensures that $\tilde{f}(\sigma) \in \hat{C}(S)$ is actually well-defined, that is if G_σ and G_τ intersect along a common open subgroup, $\sigma = \tau$.

We now use the natural action of $\hat{\Gamma}(S)$ on $\hat{C}(S)$, namely there is an injective map: $Inn(\hat{\Gamma}(S)) \to Aut(\hat{C}(S))$. For any $\gamma \in \hat{\Gamma}^{\lambda}(S)$ and any $\sigma \in \hat{C}(S)$, we then find that:

$$f(\gamma)(\sigma) = \tilde{f} \circ \gamma \circ \tilde{f}^{-1}(\sigma).$$

Since this is true for any σ, it shows that f is determined by $\tilde{f}$ if $\Gamma(S)$ is centerfree. In fact if $\tilde{f} = 1$ we get $f(\gamma)\gamma^{-1} \in Z(\hat{\Gamma}(S))$ so that $f(\gamma) = \gamma$, provided $\hat{\Gamma}^{\lambda}(S)$ is centerfree, which by [B] is the case if $\Gamma^{\lambda}(S)$ is centerfree (which in turn is implied by the centerfreeness of $\Gamma(S)$). □

Remark: Let $\Gamma^{\lambda'}$ be a finite index characteristic subgroup of Γ^{λ}. Then, there is also a natural representation $r : Aut^*(\hat{\Gamma}^{\lambda}) \to Aut^*(\hat{\Gamma}^{\lambda'})$. It is easy to check that it fits in the commutative diagram:

$$\begin{array}{ccc} Aut^*(\hat{\Gamma}^{\lambda}) & \xrightarrow{n_\lambda} & Aut(\hat{C}(S)). \\ \downarrow r & \nearrow_{n_{\lambda'}} & \\ Aut^*(\hat{\Gamma}^{\lambda'}) & & \end{array}$$

We remark now that we can also define an action of the arithmetic Galois group $G_{\mathbb{Q}}$ on $\hat{C}(S)$. In fact, there is a natural faithful outer action $G_{\mathbb{Q}} \to Out^*(\hat{\Gamma}(S))$ which can be non canonically lifted to a faithful bona fide action $G_{\mathbb{Q}} \to Aut^*(\hat{\Gamma}(S))$ by picking a (possibly tangential) rational basepoint on the moduli space $\mathcal{M}(S)$. By composing with the map constructed in Proposition 4.11 we get:

Proposition 4.12: *Assume S is hyperbolic with $d(S) > 1$; there is a map:*

$$G_{\mathbb{Q}} \hookrightarrow Aut(\hat{C}(S))$$

which is injective and canonical up to composition with the action of $Inn(\hat{\Gamma}(S))$ on $\hat{C}(S)$.

In other words, there is a natural faithful outer action of $G_{\mathbb{Q}}$ on $\hat{C}(S)$. Here the only thing which requires proof is the injectivity in the two cases where $\Gamma(S)$ has nontrivial center. But the kernel of the map which is defined in Proposition 4.11 is then made of involutions because the center has order 2 and it is of course a normal subgroup. Now any involution in $G_{\mathbb{Q}}$ is conjugate to complex conjugacy so that it is enough to check that the image of this latter element is not central; this in turn is clear since it corresponds to a reflection of the surface. We thus find that the image of $G_{\mathbb{Q}}$ in $Aut(\hat{\Gamma}(S))$ does not intersect the kernel of the map in Proposition 4.11 which completes the proof. □

It should be stressed that we get a faithful action of the arithmetic Galois group on a profinite *space*, whereas it is more common to get an action on a profinite *group*, which itself arises as a cohomological or homotopical invariant of an underlying "classical" space. Actually, if X is a (geometrically connected) scheme – say – over $\mathbb{Q}$, one can make its geometric étale covers into a proset by considering a (pro)point in the universal (pro)covering and then let $G_{\mathbb{Q}}$ acts on this proset, much as is done with "dessins d'enfants". The resulting action however, contains no more information and is often no easier to study than the usual action on the geometric fundamental group. Here complexes of curves retain some kind of homotopical information at infinity from the tower of geometric covers of the moduli stacks and they may be more amenable to a direct study.

We can readily extend the above Galois action to an action of the Grothendieck-Teichmüller group $\mathbb{I}\Gamma$ (see in particular [HLS] and [NS] for background material) essentially by the very definition of that group. For the sake of clarity, we record this explicitly in:

Proposition 4.13: *Assume S is hyperbolic with $d(S) > 1$; there is a map:*

$$\mathbb{I}\Gamma \to Aut(\hat{C}(S))$$

which is canonical up to composition with the action of $Inn(\hat{\Gamma}(S))$ on $\hat{C}(S)$ and is injective if $\Gamma(S)$ is centerfree. $\square$

Note that here we cannot *a priori* exclude the existence of a kernel in the two cases when $\Gamma(S)$ has nontrivial center. It may be useful to remind the reader that there is a nested sequence of profinite groups:

$$G_{\mathbb{Q}} \subset \mathbb{I}\Gamma \subset \widehat{GT} \subset Aut^*(\hat{F}_2),$$

where $\widehat{GT}$ is the original "genus 0" Grothendieck-Teichmüller group introduced by V.Drinfeld and $\mathbb{I}\Gamma$ is the version adapted to all genera which is constructed in [HLS] and [NS]. Here F_2 denotes as usual the free group on 2 generators. For any S as in the proposition, there is also an injective map $\mathbb{I}\Gamma \to Aut^*(\hat{\Gamma}(S))$ which gives rise to a canonical injection $\mathbb{I}\Gamma \hookrightarrow Out^*(\hat{\Gamma}(S))$. If S has genus 0, we can enlarge $\mathbb{I}\Gamma$ to $\widehat{GT}$ both here and in Proposition 4.13, that is both in the group and complex theoretic frameworks.

The above perhaps makes the following prediction reasonable. Namely for S hyperbolic with $d(S) > 1$ there could be a split exact sequence:

$$1 \to Inn(\hat{\Gamma}(S)) \to Aut(\hat{C}(S)) \to \mathbb{I}\Gamma \to 1$$

describing the structure of the automorphism group $Aut(\hat{C}(S))$. Clearly there are relatively minor adjustments to be made here. Namely the case $(1,2)$ is exceptional for the usual reasons, so there is a minor (essentially well-understood) correction to be made in that case. More substantial is the fact that here $\mathbb{I}\Gamma$ denotes *a* version of the Grothendieck-Teichmüller group, not necessarily the one appearing in [NS] with that name. There are at present quite a few different versions, some or all of which may or may not turn out to coincide. At any rate, if S has genus 0, one should clearly use the original $\widehat{GT}$. This prediction is vindicated in [Lo], starting with the fact that $Inn(\hat{\Gamma}(S))$ is indeed normal in $Aut(\hat{C}(S))$, so that we can write $Out(\hat{C}(S)) = Aut(\hat{C}(S))/Inn(\hat{\Gamma}(S))$ for the outer automorphism group of the profinite complex.

We refer to [Lo] for detailed statements, proofs and further results. As usual the two dimensional cases of types $(0,5)$ and $(1,2)$ are key. We let $\mathcal{C} = C(S_{0,5}) \simeq C(S_{1,2})$ denote the *graph* we get in these cases and $\hat{\mathcal{C}} = \hat{C}(S_{0,5}) \simeq \hat{C}(S_{1,2})$, noting that these two completions do coincide, as explained in the proof of Lemma 4.2 above. Here we would like to point out the importance of the prograph $\hat{\mathcal{C}}$ and its universality. In particular it follows from Proposition 4.13 in genus 0 that there is a canonical injective map: $\widehat{GT} \hookrightarrow Aut(\hat{\mathcal{C}})/\hat{\Gamma}$, where $\hat{\Gamma} = \hat{\Gamma}_{0,[5]}$. It is shown in [Lo] that this injection is an isomorphism, which provides a rather geometric interpretation of the group $\widehat{GT}$.

We now present the group theoretic counterpart, computing the outer automorphism groups of the profinite modular groups in genus 0. This is slightly off track as it is definitely not complex theoretic but it seems to be the only available statement of the type we are after (see however [Lo]). It can be seen as a fairly direct consequence of previous works but seems to have passed unnoticed:

Proposition 4.14: *For every $n \geq 5$:*

$$Out^*(\hat{\Gamma}_{0,[n]}) = \widehat{GT}.$$

We will derive this result from the main result of [HS] which itself builds on previous works (see references there). In [HS] the authors consider groups they denote $Out^\sharp_n$ and show that $Out^\sharp_n = \widehat{GT}$ for $n \geq 5$. Here $Out^\sharp_n$ is the group of outer automorphisms of the pure group $\hat{\Gamma}_{0,n}$ which are inertia preserving and commute with the natural outer action of the permutation group $\mathcal{S}_n$. One should be cautious about the inertia preserving condition $(*)$ at this point. It means that for any lift F of an element of $Out^\sharp_n$ to an actual automorphism, F permutes the conjugacy classes of the twists in $\hat{\Gamma}_{0,n}$, where of course the classes are intended in that pure group.

Now start from an element in $Out^*(\hat{\Gamma}_{0,[n]})$, where the $*$ refers to conjugacy in the full group $\hat{\Gamma}_{0,[n]}$. Lift that element to some $F \in Aut^*(\hat{\Gamma}_{0,[n]})$. By the inertia preserving condition F induces an element of $Aut(\hat{\Gamma}_{0,n})$, that is an automorphism of the colored group. So it also induces an automorphism of $\mathcal{S}_n$. Now we use that $Out(\mathcal{S}_n) = \{1\}$ for $n \neq 6$ ($Out(\mathcal{S}_6) = \mathbb{Z}/2$) to conclude that for $n \neq 6$, by modifying the lift F by an inner automorphism in $\hat{\Gamma}_{0,[n]}$ we may assume that F preserves the permutations, that is induces the identity on $\mathcal{S}_n$.

So starting from an element of $Out^*(\hat{\Gamma}_{0,[n]})$ with n as in the statement, we get a permutation preserving lift F which induces an element of $Aut(\hat{\Gamma}_{0,n})$ with the same name. It is easy to see that the fact that F preserves the permutations implies that its restriction to the pure group commutes with the action of $\mathcal{S}_n$. There remains to be seen that F is actually inertia preserving in the pure group. To check this, we recall briefly that $\hat{\Gamma}_{0,n}$ is generated by twists x_{ij} corresponding to loops passing through points i and j ($i \neq j$). Moreover there are elements $\sigma_{ij} \in \Gamma_{0,[n]}$ with $\sigma_{ij}^2 = x_{ij}$, where σ_{ij} maps to the transposition $(ij) \in \mathcal{S}_n$ and the σ_{ij} generate $\hat{\Gamma}_{0,[n]}$. We want to check that the procyclic group $\langle F(x_{ij})\rangle$ is conjugate in $\hat{\Gamma}_{0,n}$ with $\langle x_{ij}\rangle$. We may check this by replacing the x_{ij} with the σ_{ij}. Moreover it is readily seen that this condition is invariant under conjugacy in the *full* group $\hat{\Gamma}_{0,[n]}$. Now all the σ_{ij} are actually conjugate in $\Gamma_{0,[n]}$, so it is enough to check it on one of them, say $\sigma_{12} = \sigma$. We know it a priori that $F(\sigma) = f^{-1}\sigma^\lambda f$, where $\lambda \in \hat{\mathbb{Z}}^\times$ does not play any role here and $f \in \hat{\Gamma}_{0,[n]}$. Moreover the perservation of permutation condition says that f induces a permutation which commutes with that of σ, namely the transposition (12). We now leave it to the reader to show that this implies we can find a $g \in \hat{\Gamma}_{0,[n]}$, inducing the same permutation as f and commuting with σ. Then replacing f by $g^{-1}f \in \hat{\Gamma}_{0,n}$, we find that F lies indeed in $Aut^*(\hat{\Gamma}_{0,n})$. We may now apply the main result of [HS] in order to complete the proof of the proposition.

Finally $Out(\mathcal{S}_6) = \mathbb{Z}/2$ so that one has to rule out the possibility that an element of $Aut^*(\hat{\Gamma}_{0,[6]})$ induce the only non inner automorphism of $\mathcal{S}_6$. Here the argument in [DG] (see the proof of Corollary 12) which in a slightly different form goes back to E.Artin carries essentially verbatim to the profinite case. We do not reproduce it here. □

We now turn to the study of the automorphism group of the pants graph $\hat{C}_P(S)$, where S is hyperbolic with this time $d(S) \geq 1$. There is a natural action of $\hat{\Gamma}(S)$ on $\hat{C}_P(S)$ and in fact an injective map: $Inn(\hat{\Gamma}(S)) \hookrightarrow Aut(\hat{C}_P(S))$. In complete parallel with what happens in the discrete setting but in sharp contrast with the case of $\hat{C}(S)$, we will show that this map is almost surjective. Namely we have:

Theorem 4.15: *For any hyperbolic surface S ($d(S) \geq 1$) not of type $(1,2)$, there is a short exact sequence:*

$$1 \to Inn(\hat{\Gamma}(S)) \to Aut(\hat{C}_P(S)) \to \mathbb{Z}/2 \to 1.$$

If S is of type $(1,2)$, one should replace $Aut(\hat{C}_P(S_{1,2}))$ with $Aut^{\sharp}(\hat{C}_P(S_{1,2}))$, the subgroup of type preserving automorphisms.

In other words $Out(\hat{C}_P(S)) \simeq \mathbb{Z}/2$, just as in the discrete case, and the nontrivial outer automorphism comes again from orientation or complex conjugation. Here, as in Corollary 4.10, one actually expects that type $(1,2)$ is no real exception. The overall conclusion is that the pants graph is considerably more rigid than the curves complex. This may sound more natural if one recalls the difference between the graphs $\hat{C}_P$ and $\hat{C}_*$ introduced above and the fact that $Aut(\hat{C}_*(S)) = Aut(\hat{C}(S))$. At any rate, although the above theorem may not look so exciting, we will see in §5 that it does imply a weak anabelian result for the moduli spaces of curves.

Before going into the proof of the theorem we need some geometric remarks. First we relate $Aut(\hat{C}_P(S))$ with an important geometric object which should play a further role by itself. Let $\overline{\mathcal{M}}^{\lambda}$ be the Deligne -Mumford compactification of the level structure of level $\lambda \in \Lambda$ and let $\partial\mathcal{M}^{\lambda} = \overline{\mathcal{M}}^{\lambda} \setminus \mathcal{M}^{\lambda}$ be the divisor at infinity. We now define a curve, or rather a one dimensional D-M stack $\mathcal{F}^{\lambda}$ sitting inside the boundary $\partial\mathcal{M}^{\lambda}$:

Definition 4.16: Let S be hyperbolic with $d(S) > 1$. We define $\mathcal{F}(S) \subset \overline{\mathcal{M}}(S)$ as the complex one dimensional D-M stack whose points represent curves (Riemann surfaces) with $d(S)-1$ singularities (nodes). For an arbitrary level $\lambda \in \Lambda$ we let $\mathcal{F}^{\lambda}(S)$ denote the preimage of $\mathcal{F}(S)$ via the canonical projection $\overline{\mathcal{M}}^{\lambda}(S) \to \overline{\mathcal{M}}(S)$.

In other words $\mathcal{F}$ is the one dimensional stratum in the stable stratification of $\overline{\mathcal{M}}$. A complex point of $\mathcal{F}$ represents an algebraic curve which is a stable graph of copies of $\mathbb{P}^1 \setminus \{0, 1, \infty\}$, save for an irreducible component of type $(0,4)$ or $(1,1)$. If S is of type (g,n), we occasionally write $\mathcal{F}_{g,[n]} = \mathcal{F}(S) \subset \partial\mathcal{M}(S) = \partial\mathcal{M}_{g,[n]}$. The $\mathcal{F}^{\lambda}$'s are stable stack curves, i.e. one dimensional proper D-M stack with nodal singularities and finite groups of automorphisms. It appears with a different purpose in [GKM] which deal with a question first formulated by W.Fulton (hence our $\mathcal{F}$).

Each component of $\mathcal{F}$ is a moduli space of dimension 1 and can be triangulated (it will consist of two triangles) so that by lifting that triangulation to the corresponding Teichmüller space one gets the Farey tesselation F. On the other hand for any level $\lambda \in \Lambda$, one get a cover $\mathcal{F}^{\lambda} \to \mathcal{F}$, which ramifies at most over points representing singular curves (e.g. cusps). Therefore, the triangulation of $\mathcal{F}$ lifts uniquely to $\mathcal{F}^{\lambda}$. Moreover, and this is where the connection between $\mathcal{F}(S)$ and $C_P(S)$ comes in, the dual of that triangulation is naturally isomorphic to the graph $C_P^{\lambda}(S)$ (as usual we identify the Farey graph and the corresponding tesselation). This first enables us to define an orientation on $\hat{C}_P(S)$. Indeed the complex curves $\mathcal{F}^{\lambda}$ are oriented, and this defines an orientation on $C_P^{\lambda}(S)$; the natural projections $\mathcal{F}^{\lambda} \to \mathcal{F}^{\mu}$ for $\lambda \geq \mu$ are complex maps and thus preserve the orientation. So in turn $\hat{C}_P(S)$ inherit a natural orientation. Theorem 4.15 states that any continuous orientation preserving automorphism of $\hat{C}_P(S)$ is induced by an element of $\hat{\Gamma}(S)$.

Going to the proof of Theorem 4.15, we start with $\phi \in Aut(\hat{C}_P(S))$ and assume it is orientation preserving. If S is of type $(1,2)$ we assume it is also type preserving; this condition will be usually implicit in the sequel. For $\lambda \in \Lambda$ and for some $\mu \leq \lambda$ we get a surjective map $C_P^\lambda(S) \to C_P^\mu(S)$ of finite complexes. Since ϕ is type preserving it commutes with the projection $C^\lambda(S) \to C^0(S)$ (resp. μ) where 0 refers to the trivial level. It is now easy to show:

Lemma 4.17: *Given an automorphism ϕ as above, for any level μ, there is a Galois level λ dominating μ, such that ϕ induces a surjective map (which we call ϕ again for simplicity) $\phi : \mathcal{F}^\lambda \to \mathcal{F}^\mu$ of D-M stack curves commuting with the natural projections $p_\lambda : \mathcal{F}^\lambda \to \mathcal{F}$ (resp. $p_\mu : \mathcal{F}^\mu \to \mathcal{F}$).*

In fact, as noticed above $C_P^\lambda(S)$ determines an ideal triangulation of $\mathcal{F}^\lambda$, so that a morphism $C_P^\lambda(S) \to C_P^\mu(S)$ extends uniquely to an analytic, and in fact algebraic morphism: $\mathcal{F}^\lambda \to \mathcal{F}^\mu$; the commutation with the respective natural projections also follows. $\square$

Starting from $\phi \in Aut(\hat{\Gamma}(S))$, we thus get a coherent system of surjective analytic maps $\mathcal{F}^\lambda \to \mathcal{F}^\mu$ commuting with the natural projections to $\mathcal{F} \subset \partial\mathcal{M}(S)$. The theorem now immediately follows from:

Proposition 4.18: *Let S be hyperbolic ($d(S) \geq 1$) and ϕ, λ, μ be as in Lemma 4.17. Then the map $\phi : \mathcal{F}^\lambda \to \mathcal{F}^\mu$ is induced by an automorphism of the Galois cover $\mathcal{F}^\lambda \to \mathcal{F}$; in other words $\phi_\mu = \gamma \circ \pi$ where $\gamma \in \Gamma/\Gamma^\lambda$ and $\pi : \mathcal{F}^\lambda \to \mathcal{F}^\mu$ is the natural projection.*

Before proving the proposition, and hence the theorem, we need to introduce yet some more geometry especially objects which appear already in [B1] (among other places) and play the role of real algebraic tubular neighborhoods of the boundary for the level structures $\mathcal{M}^\lambda$ and Teichmüller space $\mathcal{T}$. We define $\widetilde{\mathcal{M}}$ as the oriented real blowup of $\overline{\mathcal{M}}$ along $\partial\mathcal{M}$; more generally $\widetilde{\mathcal{M}}^\lambda$ is the blowup of $\overline{\mathcal{M}}^\lambda$ along $\partial\mathcal{M}^\lambda$. Finally $\widetilde{\mathcal{T}}$ is the universal cover of $\widetilde{\mathcal{M}}$ (we use tildes in the present paper; the corresponding objects are denoted with hats in [B1]). We let $\partial\widetilde{\mathcal{M}}^\lambda = \widetilde{\mathcal{M}}^\lambda \setminus \mathcal{M}^\lambda$ and $\partial\widetilde{\mathcal{T}} = \widetilde{\mathcal{T}} \setminus \mathcal{T}$ and remark that the description of the space $\widetilde{\mathcal{T}}$ is actually at the origin of the introduction of the complex $C(S)$. In fact $C(S)$ is nothing but the nerve of the covering of $\partial\widetilde{\mathcal{T}}$ by its irreducible components. Since moreover these components and their mutual intersections are contractible, $\partial\widetilde{\mathcal{T}}$ is actually homotopically equivalent to the geometric realization of $C(S)$. Turning to a finite Galois level structure $\mathcal{M}^\lambda$ we find that, since $\partial\widetilde{\mathcal{T}} \to \partial\widetilde{\mathcal{M}}^\lambda$ is a Galois étale cover with group Γ^λ, $C^\lambda(S) = C(S)/\Gamma^\lambda$ is the nerve of the covering of $\partial\widetilde{\mathcal{M}}^\lambda$ (or equivalently of $\partial\mathcal{M}^\lambda$) by its irreducible components.

Now we let $\widetilde{\mathcal{F}}^\lambda \subset \partial\widetilde{\mathcal{M}}^\lambda$ denote the preimage of $\mathcal{F}^\lambda$ via the natural retraction map $q : \widetilde{\mathcal{M}}^\lambda \to \mathcal{M}^\lambda$. We also let $\widetilde{\mathcal{H}} \subset \partial\widetilde{\mathcal{T}}$ denote the preimage of $\widetilde{\mathcal{F}} \subset \widetilde{\mathcal{M}}$ via the natural projection $p : \widetilde{\mathcal{T}} \to \widetilde{\mathcal{M}}$. We observe that the (analytically) irreducible components of $\widetilde{\mathcal{H}}$ are contractible and likewise their mutual intersections. Therefore, $\widetilde{\mathcal{H}}$ is homotopically equivalent to the geometric realization of the nerve $N(\widetilde{\mathcal{H}})$ of the covering of $\widetilde{\mathcal{H}}$ by its irreducible components.

Let us then proceed with the proof of Proposition 4.18. For any level $\lambda \in \Lambda$ there is a canonical isomorphism $\widetilde{\mathcal{F}}^\lambda \simeq \mathcal{F}^\lambda \times_{\mathcal{F}} \widetilde{\mathcal{F}}$. It follows that there is a canonical lift of ϕ to a real analytic morphism $\tilde{\phi} : \widetilde{\mathcal{F}}^\lambda \to \widetilde{\mathcal{F}}^\mu$. Let us observe that for any level λ, the natural morphism $\tilde{p}_\lambda : \widetilde{\mathcal{F}}^\lambda \to \widetilde{\mathcal{F}}$ is étale.

Therefore, the same holds for the morphism $\tilde{\phi}$.

For any given Galois level λ, there is a short exact sequence:

$$1 \to \pi_1(\widetilde{\mathcal{H}}) \to \pi_1(\widetilde{\mathcal{F}}^\lambda) \to \Gamma^\lambda \to 1.$$

The central point in the proof of the proposition is to show that $\tilde{\phi}$ lifts to an analytic automorphism $h : \widetilde{\mathcal{H}} \to \widetilde{\mathcal{H}}$ which is determined up to the action of Γ^λ. This is equivalent to showing that:

$$\tilde{\phi}_*(\pi_1(\widetilde{\mathcal{H}})) \subseteq \pi_1(\widetilde{\mathcal{H}}). \tag{$*$}$$

The set of Farey graphs $\{F_\alpha\}_{\alpha\in A}$ covers the pants complex $C_P(S)$. Therefore, also the profinite pants graph $\hat{C}_P(S)$ is covered by a profinite set $\{\hat{F}_\alpha\}_{\alpha\in\overline{A}}$ of profinite Farey graphs. Let us denote by $N(\hat{C}_P(S))$ the nerve of this cover. It is a profinite simplicial set isomorphic to the inverse limit of the finite nerves of the covers of the curves $\mathcal{F}^\lambda$ by their irreducible components. The first of the lemmas we need in order to prove the inclusion $(*)$ reads as follows, where as usual in the profinite case all automorphisms are assumed to be continuous:

Lemma 4.19: *An automorphism of $\hat{C}_P(S)$ induces an automorphism of $N(\hat{C}_P(S))$.*

Indeed an automorphism of $\hat{C}_P(S)$ preserves the set of its Farey subgraphs. $\square$

From here to the end of the present section we will write π_1 for the topological fundamental group and $\hat{\pi}_1$ for the profinite (or geometric in the sense of algebraic geometry) fundamental group. Then we have:

Lemma 4.20: *There is a natural continuous isomorphism:*

$$\hat{\pi}_1(N(\hat{C}_P(S))) \simeq \varprojlim_{\lambda\in\Lambda} \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda).$$

By definition $N(\hat{C}_P(S)) = \varprojlim_\lambda N(\widetilde{\mathcal{F}}^\lambda)$; now for every λ, there is an exact sequence:

$$\hat{*}_{\alpha\in A}\, \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda_\alpha) \to \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda) \to \hat{\pi}_1(N(\widetilde{\mathcal{F}}^\lambda)) \to 1.$$

Since $\varprojlim_\lambda \hat{\pi}_1(\widetilde{\mathcal{F}}_\alpha) = \cap_\lambda \hat{\pi}_1(\widetilde{\mathcal{F}}_\alpha) = \{1\}$, taking the inverse limit of the above sequences one gets:

$$\varprojlim_\lambda \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda) \simeq \varprojlim_\lambda \hat{\pi}_1(N(\widetilde{\mathcal{F}}^\lambda)) = \hat{\pi}_1(N(\hat{C}_P(S))).$$

$\square$

Lemma 4.21: *i) For every level λ, the group $\pi_1(\widetilde{\mathcal{F}}^\lambda)$ is residually finite.*
ii) Let us denote by $\overline{\pi_1(\widetilde{\mathcal{H}})}$ the closure of the group $\pi_1(\widetilde{\mathcal{H}})$ inside the profinite group $\hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda)$. Then:

$$\pi_1(\widetilde{\mathcal{H}}) = \pi_1(\widetilde{\mathcal{F}}^\lambda) \cap \overline{\pi_1(\widetilde{\mathcal{H}})}.$$

For every Galois level $\mu \leq \lambda$, the finite group $\Gamma^\lambda/\Gamma^\mu$ acts on the finite graph $N(\widetilde{\mathcal{H}})/\Gamma^\mu = N(\mathcal{F}^\mu)$ with quotient $N(\mathcal{F}^\lambda)$. Therefore one can build a graph of groups $\mathcal{G}^\mu$ associated to this action with nerve $N(\mathcal{F}^\lambda)$ and whose fundamental group is described by the short exact sequence:

$$1 \to \pi_1(N(\mathcal{F}^\mu)) \to \pi_1(\mathcal{G}^\mu) \to \Gamma^\lambda/\Gamma^\mu \to 1.$$

Let us observe that $\pi_1(\widetilde{\mathcal{F}}^\lambda)$ is the fundamental group of the graph of groups determined by the action of Γ^λ on the graph $N(\widetilde{\mathcal{H}})$. Therefore, there is a natural homomorphism $\pi_1(\widetilde{\mathcal{F}}^\lambda) \to \pi_1(\mathcal{G}^\mu)$ which fits into a commutative diagram with exact rows:

$$\begin{array}{ccccccccc} 1 & \to & \pi_1(N(\widetilde{\mathcal{H}})) & \to & \pi_1(\widetilde{\mathcal{F}}^\lambda) & \to & \Gamma^\lambda & \to & 1 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 1 & \to & \pi_1(N(\mathcal{F}^\mu)) & \to & \pi_1(\mathcal{G}^\mu) & \to & \Gamma^\lambda/\Gamma^\mu & \to & 1. \end{array}$$

Taking first profinite completions and then inverse limits on the bottom line, one gets a commutative diagram with exact rows:

$$\begin{array}{ccccccccc} 1 & \to & \pi_1(N(\widetilde{\mathcal{H}})) & \to & \pi_1(\widetilde{\mathcal{F}}^\lambda) & \to & \Gamma^\lambda & \to & 1 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 1 & \to & \hat{\pi}_1(N(\hat{C}_P(S))) & \to & \varprojlim_\mu \pi_1(\mathcal{G}^\mu) & \to & \hat{\Gamma}^\lambda & \to & 1. \end{array}$$

By Proposition 5.1 in [B1], the natural map $C(S) \to \hat{C}(S)$ is injective. In particular the map $C_P(S)_0 \to \hat{C}_P(S)_0$ is also injective. Since an edge of $C_P(S)$ is determined by its vertices, it follows that the natural map $C_P(S) \to \hat{C}_P(S)$ is injective and then the same holds for the natural map $N(\widetilde{\mathcal{H}}) \to N(\hat{C}_P(S))$. From this, it easily follows that the induced homomorphism $\pi_1(N(\widetilde{\mathcal{H}})) \to \hat{\pi}_1(N(\hat{C}_P(S)))$ is injective as well. The above commutative diagram and the fact that Γ^λ is residually finite then imply that the homomorphism $\pi_1(\widetilde{\mathcal{F}}^\lambda) \to \varprojlim_\mu \hat{\pi}_1(\mathcal{G}^\mu)$ is injective. Therefore, since it admits an injective homomorphism into a profinite group, the group $\pi_1(\widetilde{\mathcal{F}}^\lambda)$ is residually finite.

In order to check the identity claimed in ii) it is now enough to note the existence of the following commutative diagram with exact rows and injective vertical maps:

$$\begin{array}{ccccccccc} 1 & \to & \pi_1(\widetilde{\mathcal{H}}) & \to & \pi_1(\widetilde{\mathcal{F}}^\lambda) & \to & \Gamma^\lambda & \to & 1 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 1 & \to & \overline{\pi_1(\widetilde{\mathcal{H}})} & \to & \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda) & \to & \hat{\Gamma}^\lambda & \to & 1. \end{array}$$

□

Lemma 4.22: *Let $\overline{\pi_1(\widetilde{\mathcal{H}})}$ denote as above the closure of the group $\pi_1(\widetilde{\mathcal{H}})$ inside any of the profinite groups $\hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda)$. Then there is a natural isomorphism:*

$$\varprojlim_{\lambda\in\Lambda} \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda) \simeq \overline{\pi_1(\widetilde{\mathcal{H}})}.$$

Indeed, consider the inverse limit of the short exact sequences:

$$1 \to \overline{\pi_1(\widetilde{\mathcal{H}})} \to \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda) \to \hat{\Gamma}^\lambda \to 1$$

and observe that $\varprojlim_\lambda \hat{\Gamma}^\lambda = \cap_\lambda \hat{\Gamma}^\lambda = \{1\}$. $\square$

Now in order to prove the inclusion $(*)$ let us take a look at the commutative diagram:

$$\begin{array}{ccc} \pi_1(\widetilde{\mathcal{F}}^\lambda) & \xrightarrow{\phi_*} & \pi_1(\widetilde{\mathcal{F}}^\mu) \\ \downarrow & & \downarrow \\ \hat{\pi}_1(\widetilde{\mathcal{F}}^\lambda) & \xrightarrow{\hat{\phi}_*} & \hat{\pi}_1(\widetilde{\mathcal{F}}^\mu) \end{array}$$

in which we write ϕ for $\tilde{\phi}$ in order to make room for the completed map in the bottom row. By Lemma 4.21 the vertical maps are injective. By Lemmas 4.19, 4.20 and 4.22 the map $\hat{\phi}_*$ preserves the image of the closure of the subgroup $\pi_1(\widetilde{\mathcal{H}})$, that is:

$$\hat{\phi}_*(\overline{\pi_1(\widetilde{\mathcal{H}})}) \subseteq \overline{\pi_1(\widetilde{\mathcal{H}})}.$$

In particular:

$$\phi_*(\pi_1(\widetilde{\mathcal{H}})) = \hat{\phi}_*(\pi_1(\widetilde{\mathcal{H}})) \subseteq \overline{\pi_1(\widetilde{\mathcal{H}})} \cap \pi_1(\widetilde{\mathcal{F}}^\mu).$$

The inclusion $(*)$ then follows immediately from ii) in Lemma 4.14.

Let now $\overline{\mathcal{T}}$ be the Bers bordification of the Teichmüller space $\mathcal{T}$. Let then $\mathcal{H}$ be the image of $\widetilde{\mathcal{H}}$ in the Bers boundary $\partial\overline{\mathcal{T}}$ of $\overline{\mathcal{T}}$, which is actually a retract of $\partial\widetilde{\mathcal{T}}$ along the semiaxis $\mathbb{R}^+$. The real analytic automorphism $h : \widetilde{\mathcal{H}} \to \widetilde{\mathcal{H}}$ induces an automorphism $h : \mathcal{H} \to \mathcal{H}$, which is complex analytic because it lifts the holomorphic epimorphism $\mathcal{F}^\lambda \to \mathcal{F}^\mu$. Moreover h preserves the natural triangulation of $\mathcal{H}$ inherited from the $\mathcal{F}^\lambda$, the dual of which is nothing but the pants graph $C_P(S)$ (the vertices correspond to the cusps of $\mathcal{H}$). Therefore h determines an automorphism of the pants graph $C_P(S)$. By Margalit's result (Theorem 2.5 above) such an automorphism is induced by an element of $Mod(S)$ and since ϕ preserves the orientation, that element γ actually lies in $\Gamma(S)$. Finally it is determined only modulo Γ^λ. This completes the proof of Proposition 4.18, thus also of Theorem 4.15. $\square$

5. Anabelian properties of moduli stacks of curves

Let k be a field over which the anabelian conjecture for hyperbolic curves is valid. As we already mentioned, according to [Mo], we can take k to be a sub-p-adic field, that is a subfield of a finitetely generated extension of $\mathbb{Q}_p$ for some prime p. Especially noteworthy are the cases where k is a finite extension of $\mathbb{Q}$ or $\mathbb{Q}_p$. The latter is particularly interesting here since it corresponds so to speak to particularly "small" Galois groups G_k.

In order to prove the anabelian statement exposed in Section 1, we need first to restrict to a peculiar situation. So let $\mathcal{M}^\lambda \to \mathcal{M}_{g,[n]}$ be a Galois level structure defined over a sub-p-adic field k such that all geometric (i.e. defined over $\overline{k}$) automorphism of the cover are already defined over k. Note that this property is satisfied by all level structures after a finite extension of the base field. By the theory of the algebraic fundamental group, we know that the algebraic fundamental group of $\mathcal{M}_{g,[n]} \otimes \overline{k}$ can be identified with the profinite completion $\hat{\Gamma}_{g,[n]}$ of the mapping class group $\Gamma_{g,[n]}$ of the Riemann surface $S_{g,n}$ (see Section 1) and that of $\mathcal{M}^\lambda \otimes \overline{k}$ with the profinite completion $\hat{\Gamma}^\lambda$ of a finite index normal subgroup Γ^λ of $\Gamma_{g,[n]}$. From Teichmüller theory (Royden Theorem, to be precise), we know that the group of all geometric automorphisms of $\mathcal{M}^\lambda$ equals the group of geometric automorphisms of the cover $\mathcal{M}^\lambda \to \mathcal{M}_{g,[n]}$ and the latter is then identified with the quotient group $G^\lambda := \Gamma_{g,[n]}/\Gamma^\lambda$.

The boundary of $\overline{\mathcal{M}}_{g,[n]}$ contains $\mathbb{Q}$-rational points. In particular, it is possible to find k-rational points on the boundary of $\overline{\mathcal{M}}^\lambda$. Let us then pick a k-rational tangential base point at infinity on $\mathcal{M}^\lambda$ and get an action $G_k \to Aut(\hat{\Gamma}^\lambda)$. In this situation we define Galois invariant automorphisms as:

Definition 5.1: $Aut_{G_k}(\hat{\Gamma}^\lambda)$ is the subgroup of automorphisms of $\hat{\Gamma}^\lambda$ which, modulo inner automorphisms, commute with the Galois action, i.e. those $f \in Aut(\hat{\Gamma}^\lambda)$ such that for any $\sigma \in G_k$, the commutator $[\sigma, f] \in Inn(\hat{\Gamma}^\lambda)$.

We can now state:

Theorem 5.2: *Let $\mathcal{M}^\lambda$, for $2g-2+n>0$, be a Galois level structure over $\mathcal{M}_{g,[n]}$ and k a sub-p-adic field of definition for both $\mathcal{M}^\lambda$ and all its geometric automorphisms; then:*

$$Aut^*_{G_k}(\hat{\Gamma}^\lambda) = Inn(\hat{\Gamma}_{g,[n]}).$$

In order to prove the result, we first note that $\hat{\Gamma}$ acts on $\hat{\Gamma}^\lambda$ by conjugation; so there is a natural inclusion $Inn(\hat{\Gamma}) \subset Aut^*(\hat{\Gamma}^\lambda)$. The induced representation $\hat{\Gamma} \to Out(\hat{\Gamma}^\lambda)$ factors through the natural outer action of the geometric Galois group G^λ of the cover $\mathcal{M}^\lambda \to \mathcal{M}$ on the fundamental group of the level structure $\mathcal{M}^\lambda$. The automorphisms of the cover commute with the action of G_k by the very definition of the field k. We thus get the easy inclusion $Inn(\hat{\Gamma}) \subset Aut^*_{G_k}(\hat{\Gamma}^\lambda)$ and have to show that conversely any $f \in Aut^*_{G_k}(\hat{\Gamma}^\lambda)$ comes from an inner automorphism of $\hat{\Gamma}$.

So let $f \in Aut^*_{G_k}(\hat{\Gamma}^\lambda)$ be a Galois invariant automorphism. Like any inertia preserving automorphism, it defines a continuous automorphism of $\hat{C}(S)$ (see Proposition 4.11 above) and in particular of the profinite set of the 0-simplices of $\hat{C}_P(S)$. For simplicity, we denote these automorphisms again by f. *If* we can prove that actually $f \in Aut(\hat{C}_P(S))$ and it is orientation preserving, applying Theorem 4.15, that is using the rigidity of the pants complex (or graph) will complete the proof. To be complete, one should note that type $(1,2)$ is no exception here. Indeed in order to show that f is type preserving in that case, it is enough to show that it preserves the set of non separating (pro)curves. But f comes from a *group* homomorphism and in $\hat{\Gamma}_{1,[2]}$ the centralizers of separating and non separating twists are not isomorphic, which implies that indeed the induced automorphism on $\hat{C}(S)$ is type preserving. So there remains to show that a Galois invariant automorphism of $\hat{\Gamma}^\lambda$ induces an automorphism of $\hat{C}(S)$ which stabilizes the pants graph $\hat{C}_P(S)$.

By the remark to Proposition 4.11, for any finite index characteristic subgroup $\Gamma^{\lambda'}$ of Γ^λ, the representation n_λ factors through $n_{\lambda'}$ and the natural representation $Aut^*(\hat{\Gamma}^\lambda) \to Aut^*(\hat{\Gamma}^{\lambda'})$. So, it is actually enough to prove the above statement for any Galois level $\Gamma^{\lambda'}$ contained in Γ^λ. Let us then assume that $\mathcal{M}^\lambda$ is a representable level structure lying over $\mathcal{M}_{g,n}$.

Let us consider first the one dimensional case. If (g,n) is either $(0,4)$ or $(1,1)$, we have to show that (say for type $(0,4)$):

$$Out_{G_k}(\hat{\Gamma}^\lambda_{0,4}) = G^\lambda.$$

This however is precisely (part of) the content of the anabelian result for hyperbolic curves (in [Mo] as far as general sub-p-adic fields are concerned), applied here to the finite Galois cover of $\mathcal{M}_{0,4} \cong \mathbb{P}^1$ corresponding to the given level λ. Similarly, we do have that for $(g,n) = (1,1)$:

$$Aut_{G_k}(\hat{\Gamma}^{\lambda}_{1,1}) = Inn(\hat{\Gamma}_{1,1}).$$

We remark that in terms of complexes, $C_P(S_{1,1}) = C_P(S_{0,4}) = F$ and from Theorem 4.15 or by a more direct proof we know that $Aut(\hat{F}) = Inn(\hat{\Gamma}_{1,1})$. From that perspective, in order to treat the one dimensional case one needs to prove directly that $Aut_{G_k}(\hat{\Gamma}^{\lambda}) \subset Aut(\hat{F})$, that is again, a Galois invariant automorphism does induces an automorphism of the profinite Farey tesselation. It would of course be extremely interesting to get an alternative proof of that deep fact.

We now proceed to reduce the general case to the one dimensional case. We have to show that for any given edge $e \in \hat{C}_P(S)_1$ the two vertices $\{f(\partial_i e)\}_{i=0,1}$ are connected by an edge $e' = f(e) \in \hat{C}_P(S)$. The pants graph $C_P(S)$ is a union of Farey graphs $\{F_\alpha\}_{\alpha \in A}$. Thererefore, also the profinite pants graph $\hat{C}_P(S)$ is covered by a family $\{\hat{F}_\alpha\}_{\alpha \in \overline{A}}$ of profinite Farey graphs. We may and will assume that $e \in \hat{F}_\alpha$, where $\hat{F}_\alpha$ is the closure of a discrete Farey graph F_α contained in $C_P(S)$ corresponding to a set $\{\gamma_1, \ldots, \gamma_{d-1}\}$ of nontrivial disjoint circles on S ($d = 3g-3+n$). We now make the simple but essential remark that the set $\{f(\partial_i e)\}_{i=0,1}$ satisfies the required property if and only if the set $\{a^{-1} f a(\partial_i e)\}_{i=0,1}$ does, for some and then in fact for any $a \in \hat{\Gamma}$. This amounts to saying that the required property is $\hat{\Gamma}$-invariant for the natural action of $\hat{\Gamma}$ on $\hat{C}_P(S)$. Moreover for any $a \in \hat{\Gamma}$, $a^{-1} f a$ is G_k-equivariant since $Inn(\hat{\Gamma}) \subset Aut^*_{G_k}(\hat{\Gamma}^{\lambda})$.

Thus, after twisting f by a suitable $a \in \hat{\Gamma}$, we can assume that $f(\vec{\gamma}_i) = \vec{\gamma}_i$ for $i = 1, ..., d-1$ where $\vec{\gamma}_i$ denotes the oriented loop γ_i. In particular, f then stabilizes the profinite set $(\hat{F}_\alpha)_0$ of the 0-simplices of $\hat{F}_\alpha \subset \hat{C}_P(S)$ and, with the notation of [B1], induces an automorphism of $\hat{\Gamma}_{\vec{\sigma}}$, the stabilizer of $\sigma = \{\gamma_1, \ldots, \gamma_{d-1}\}$ which is a $d-2$ oriented simplex of $C(S)$. By Theorem 7.1 in [B1] there is a short exact sequence:

$$1 \to \bigoplus_{i=1}^{d} \hat{\mathbb{Z}} \tau_{\gamma_i} \to \hat{\Gamma}_{\vec{\sigma}} \to \hat{\Gamma}_{h,k} \to 1,$$

where (h,k) is equal to $(1,1)$ or $(0,4)$.

Let us assume we are in the first case for definiteness; the second one can be treated in the same way. By the above assumptions, $f \in Aut^*_{G_k}(\hat{\Gamma}^{\lambda})$ acts on the subgroup $\hat{\Gamma}_{\vec{\sigma}} \cap \hat{\Gamma}^{\lambda}$ and the action descends to the image $\hat{\Gamma}^{\lambda}_{1,1}$ of the subgroup $\hat{\Gamma}_{\vec{\sigma}} \cap \hat{\Gamma}^{\lambda}$ in $\hat{\Gamma}_{1,1}$. The profinite pants graph $\hat{C}_P(S_{1,1})$ associated with the latter group is naturally isomorphic to the profinite Farey graph $\hat{F}_\alpha$ via an isomorphism which is compatible with the action of f. From the fact that $f \in Aut^*_{G_k}(\hat{\Gamma}^{\lambda})$, it follows that $f \in Aut^*_{G_k}(\hat{\Gamma}^{\lambda}_{1,1})$ as well. Moreover, by the assumptions we made on λ, the corresponding level structure $\mathcal{M}^{\lambda}_{1,1}$ is a representable one dimensional D-M stack, i.e. a hyperbolic curve. Therefore, by Mochizuki's Theorem, we get that f acts on $\hat{\Gamma}^{\lambda}_{1,1}$ via an inner automorphism. Thus, the action of f on the 0-simplices of $\hat{F}_\alpha$ extends to an action on the entire profinite Farey graph which, moreover, preserves its orientation. This concludes the proof of the theorem. $\square$

Theorem 5.2 immediately implies the following anabelianity result:

Theorem 5.3: *Let $\mathcal{M}^\lambda$ be a level structure over $\mathcal{M}_{g,[n]}$, with $2g-2+n>0$, defined over a sub-p-adic field k. Let A_λ be the automorphisms group of the generic point of $\mathcal{M}^\lambda$ and denote by $\mathcal{M}^\lambda /\!\!/ A_\lambda$ the stack obtained rigidifying $\mathcal{M}^\lambda$ with respect to A_λ. The fundamental group functor then induces an isomorphism:*

$$Aut_k(\mathcal{M}^\lambda /\!\!/ A_\lambda) \cong Out^*_{G_k}(\pi_1(\mathcal{M}^\lambda \otimes \overline{k})).$$

Remark: The automorphisms group of the generic point of $\mathcal{M}^\lambda$ is trivial unless the couple (g,n) belongs to the set $\{(0,4),(1,1),(1,2),(2,0)\}$ and the corresponding level Γ^λ of $\Gamma_{g,[n]}$, in the first case, intersects non-trivially the Klein subgroup of $\Gamma_{0,[4]}$ and, in all other cases, contains the hyperelliptic involution.

Like in the proof of the anabelianity conjecture for curves (see [Mo]), it is enough to prove the statement of the theorem modulo a finite extension of the base field k and modulo a connected étale cover of $\mathcal{M}^\lambda$ and then apply geometric and Galois descent. Thus, we can assume that the hypothesis of Theorem 5.2 are all satisfied.

Let us then identify the geometric fundamental groups of $\mathcal{M}^\lambda$ and $\mathcal{M}_{g,[n]}$ with $\hat{\Gamma}^\lambda$ and $\hat{\Gamma}_{g,[n]}$, respectively, as done in the statement of Theorem 5.2. By Theorem 7.4 in [B1], the only possible non-trivial element in the center of $\hat{\Gamma}^\lambda$ is the hyperelliptic involution ι which then is also a generic automorphism of the stack $\mathcal{M}^\lambda$. Therefore, if $\mathcal{M}^\lambda$ has at most a central generic automorphism, it holds:

$$Out^*_{G_k}(\hat{\Gamma}^\lambda) \cong Inn\hat{\Gamma}_{g,[n]}/Inn\hat{\Gamma}^\lambda \cong \hat{\Gamma}_{g,[n]}/(\hat{\Gamma}^\lambda \cdot Z),$$

where Z denotes the center of $\Gamma_{g,[n]}$, i.e. either $Z=\{1\}$ or $Z=<\iota>$. Let us observe, in case Z is not trivial, that ι determines a non-trivial automorphism of the cover $\mathcal{M}^\lambda \to \mathcal{M}_{g,[n]}$ but a trivial one of the stack $\mathcal{M}^\lambda /\!\!/ A_\lambda$, where $A_\lambda = \Gamma^\lambda \cap <\iota>$. Therefore, in any case, it holds:

$$\hat{\Gamma}_{g,[n]}/(\hat{\Gamma}^\lambda \cdot Z) \cong Aut_k(\mathcal{M}^\lambda /\!\!/ A_\lambda)$$

and the claim of the theorem follows.

The only case which remains to consider is $(g,n)=(0,4)$. Since the center of $\hat{\Gamma}_{0,[4]}$ and of any of its open subgroups is trivial, we have:

$$Out^*_{G_k}(\hat{\Gamma}^\lambda) \cong Inn\hat{\Gamma}_{0,[4]}/Inn\hat{\Gamma}^\lambda \cong \hat{\Gamma}_{0,[4]}/\hat{\Gamma}^\lambda \cong Aut_k(\mathcal{M}^\lambda /\!\!/ A_\lambda).$$

□

In closing we stress again the fact amply illustrated in §§3, 4, 5 above that complexes of curves make it possible to get results pertaining uniformly to open subgroups of the modular groups. In Theorem 5.2 for instance, the right-hand side is independent of λ; Γ^λ can be taken arbitrarily small. Also, and more classically, the same phenomenon occurs on the arithmetic side: in Theorem 5.3, F can be arbitrarily large and one still gets the finite group G^λ. In other words the results are quite robust under passing to geometric or arithmetic covers.

References

[B1] M.Boggi. Profinite Teichmüller theory. Math. Nachr. **279** (2006), 953-987.

[B2] M.Boggi. Fundamental groups of moduli stacks of stable curves of compact type. ArXiv: math.AG/0604271 v2 (2006).

[BP] M.Boggi, M.Pikaart. Galois covers of moduli of curves. Compositio Math. **120** (2000), 171–191.

[DG] J.L.Dyer, E.K.Grossman. The automorphism groups of the braid groups. American Journal of Mathematics **103** (1981), 1151-1169.

[EK] C.J.Earle, I.Kra. On isometries between Teichmüller spaces. Duke Math. J. **41** (1974), 583-591.

[FTP] *Travaux de Thurston sur les surfaces* (seconde édition), A.Fathi, F.Laudenbach, V.Poénaru eds., Astérisque **66-67**, SMF Publ., 1991.

[G] A. Grothendieck. Brief an G. Faltings. In *Geometric Galois Actions I*, edited by L. Schneps and P. Lochak, Lon. Math. Soc. LNS **242** (1997), 49-58.

[GKM] A.Gibney, S.Keel, I.Morrison. Towards the ample cone of $\overline{\mathcal{M}}_{g,n}$. J. Amer. Math. Soc. **15** (2002), 273–294.

[HLS] A.Hatcher, P.Lochak, L.Schneps. On the Teichmüller tower of mapping class groups. J. reine und angew. Math. **521** (2000), 1–24.

[HS] D.Harbater, L.Schneps. Fundamental groups of moduli and the Grothendieck-Teichmüller group. Transactions AMS **352** (2000), 3117-3148.

[IN] Y.Ihara, H.Nakamura. Some illustrative examples for anabelian geometry in high dimensions. In *Geometric Galois Actions I*, edited by L. Schneps and P. Lochak, Lon. Math. Soc. LNS **242** (1997), 127–138.

[I1] N.V.Ivanov. Automorphisms of complexes of curves and of Teichmüller spaces. Internat. Math. Res. Notices **14** (1997), 651–666.

[I2] N.V.Ivanov. *Subgroups of Teichmüller modular groups*. Translations of Mathematical Monographs **115**, AMS Publ., 1992.

[L] F.Luo. Automorphisms of the complex of curves. Topology **39** (2000), 283–298.

[Lo] P.Lochak. Automorphism groups of profinite complexes of curves and the Grothendieck-Teichmüller group. Preprint, 2007.

[M] D.Margalit. Automorphisms of the pants complex. Duke Math. J. **121** (2004), 457–479.

[Mo] S.Mochizuki. The local pro-p anabelian geometry of curves. Inventiones math. **138** (1999), 319-423.

[McC] J.D.McCarthy. Automorphisms of surface mapping class groups: A recent theorem of N.Ivanov. Inventiones math. **84** (1986), 49-71.

[MM] H.A.Masur, Y.N.Minsky. Geometry of the complex of curves I: Hyperbolicity. Inventiones math. **138** (1999), 103-149.

[MT] S.Mochizuki, A.Tamagawa. The algebraic and anabelian geometry of configuration spaces. Preprint, 2006.

[MNT] S.Mochizuki, H.Nakamura, A.Tamagawa. The Grothendieck conjecture on the fundamental group of algebraic curves. Sugaku Expositions **14** (2001), 31-53.

[N] H.Nakamura. Galois rigidity of pure sphere braid groups and profinite calculus. J. Math. Sci. Univ. Tokyo **1** (1994), 72-136.

[No] B.Noohi. Fundamental groups of algebraic stacks. J. Inst. Math. Jussieu **3** (2004), 69–103.

[NS] H.Nakamura, L.Schneps. On a subgroup of the Grothendieck-Teichmüller group acting on the tower of profinite Teichmüller modular groups. Inventiones math. **141** (2000), 503-560.

[P] F.Pop. Glimpses of Grothendieck's Anabelian Geometry. In *Geometric Galois Actions I*, edited by L. Schneps and P. Lochak, Lon. Math. Soc. LNS **242** (1997), 113–126.

[PdJ] M.Pikaart, A.J. de Jong. Moduli of curves with non abelian level structure. In *The moduli space of curves*, 483–509, Progr. Math. **129**, Birkhäuser, 1995.

[R] M.Romagny. Group actions on stacks and applications. Michigan Math. J. **53** (2005), 209–236.

[S] C.Scheiderer. Farell cohomology and Brown theorems for profinite groups. Manuscripta Math. **91** (1996), 247-281.

[Sx] J.Stix. Maps to anabelian varieties and extending curves. Preprint, 2004.

[Sz] T.Szamuely. Groupes de Galois de corps de type fini (d'après Pop). Astérisque **294** (2004), SMF Publ., 403–431.

Institut de Mathématiques de Jussieu
175 rue du Chevaleret, F-75013 Paris
boggi@math.jussieu.fr, lochak@math.jussieu.fr