GOMPF GLUING AND NON-COMPUTABILITY OF CLASSIFICATION

ARTHUR LEI QIU

1. Introduction

A fundamental question in the study of manifolds is that of classification up to isomorphism (in some appropriate category). Closely linked to this is the question of tractability: how *difficult* is it to determine when two manifolds are isomorphic? Ideally, one would like not only to have such a classification, but also to find explicit recipes or decision processes which determine when two manifolds are isomorphic. The answers to these questions depend highly on the dimension of the manifolds in consideration.

(0) Zero-dimensional manifolds are trivial: they are completely classified by cardinality.
(1) One-dimensional manifolds are not much more complicated: up to isomorphism, the only connected 1-manifolds without boundary are the line $\mathbb{R}$ and the circle S^1, and these are distinguished by simple criteria such as compactness or homotopy type.
(2) The classification of surfaces famously states that a connected closed (i.e., compact and without boundary) 2-manifold is homeomorphic to either the sphere S^2, a connected sum of tori $\mathbb{T}^2 = S^1 \times S^1$, or a connected sum of real projective planes $\mathbb{R}P^2$. These are distinguished by Euler characteristic and orientability; thus, by using triangulations, we can come up with algorithms which distinguish surfaces from each other.
(3) In dimension 3, the situation is considerably more subtle, but there is still a semi-affirmative answer: oriented closed 3-manifolds are classified by Thurston's geometrization, and from this it follows that there is an algorithm for classifying them up to orientation-preserving homeomorphism. However, the computational complexity of this algorithm is, as a function of the number of simplices in a triangulation, a tower of exponentials [4].

Our luck begins to run out in dimension 4, however. This is due to the following result, which implies that classifying manifolds in dimensions 4 and up is at least as hard as classifying finitely presentable groups—and the latter problem is quite difficult, to say the least (Section 3).

Theorem 1.1. *For any finitely presentable group G and fixed $n \geq 4$, there exists a closed smooth n-manifold M such that $\pi_1(M) \cong G$.*

Proof sketch. Given a finite presentation $G \cong \langle g_1, \ldots, g_k \mid r_1, \ldots, r_\ell \rangle$, consider the k-fold connected sum

$$M_0 = \underbrace{(S^1 \times S^{n-1}) \# \cdots \# (S^1 \times S^{n-1})}_{k \text{ times}}.$$

Repeated application of van Kampen's theorem shows that $\pi_1(M_0) \cong \langle g_1, \ldots, g_k \rangle$ is isomorphic to the free group of rank k. Represent the word r_1 in $\pi_1(M_0)$ by a simple closed curve γ in M_0 (which can be assumed to be smoothly embedded by the Whitney approximation theorem and a general position argument). Since M_0 is orientable, a small tubular neighbourhood of γ is diffeomorphic to $S^1 \times D^{n-1}$; perform surgery on this neighbourhood to replace it with $D^2 \times S^{n-2}$ and call the resulting smooth manifold M_1. By van Kampen's theorem, M_1 has fundamental group given by the amalgamated free product

$$\pi_1(M_1) \cong \pi_1(M_0) *_{\pi_1(S^1 \times S^{n-2})} \pi_1(D^2 \times S^{n-2}).$$

Since $n \geq 4$, we have $\pi_1(D^2 \times S^{n-2}) = 0$; on the other hand, $\pi_1(S^1 \times S^{n-2})$ is the cyclic subgroup generated by γ. It follows that $\pi_1(M_1) \cong \langle g_1, \ldots, g_k \mid r_1 \rangle$. Repeat this process with M_1 by finding a representative curve of r_2 and performing surgery on a tubular neighbourhood to obtain a manifold M_2 with $\pi_1(M_2) \cong \langle g_1, \ldots, g_k \mid r_1, r_2 \rangle$. By further iterating this process, we obtain a sequence of manifolds with an additional relator killed off at each step; after ℓ iterations, we obtain the desired manifold $M = M_\ell$. ■

One might hope that imposing requirements such as the existence of certain geometric structures might make the classification problem more tractable. For instance, recall that a *symplectic form* on a (necessarily even-dimensional) smooth manifold M is a closed, non-degenerate 2-form ω on M; the pair (M, ω) is called a *symplectic manifold*. Is it possible to classify symplectic manifolds up to isomorphism? In this essay we will show that, in fact, all finitely presentable groups are also realized as fundamental groups of symplectic manifolds; thus, the classification problem is no easier for symplectic manifolds than it is for general smooth manifolds.

Theorem 1.2 (Gompf [3]). *For any finitely presentable group G, there exists a closed symplectic 4-manifold (M, ω) such that $\pi_1(M) \cong G$. (The same is true in all even dimensions greater than 4 by taking products with S^2 suitably many times.)*

There is little hope in adapting the proof sketch of Theorem 1.1 as written. For one, the starting manifold M_0 has vanishing second de Rham cohomology and therefore never admits a symplectic form. Even if it did, one cannot generally expect the usual notion of surgery to preserve additional geometric structures on a manifold. To remedy this, Gompf introduced a new kind of surgery operation which allows one to form connected sums of symplectic manifolds along certain symplectic submanifolds of codimension 2; the resulting manifold has a canonical isotopy class of symplectic forms, and therefore this surgery operation yields a new construction of symplectic manifolds. In particular, this operation allows us to explicitly cook up symplectic manifolds whose fundamental groups realize any finitely presentable group.

2. Gompf Gluing

We begin by stating a special case of Weinstein's tubular neighbourhood theorem which says, in essence, that symplectic submanifolds admit symplectic tubular neighbourhoods. Denote by $B^2(\varepsilon) = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < \varepsilon^2\}$ the open disk in $\mathbb{R}^2$ centred at the origin with radius $\varepsilon > 0$, taken with the symplectic form $\omega_{\text{std}} = dx \wedge dy$ induced from $\mathbb{R}^2$.

Theorem 2.1 (Weinstein [9]). *Let (M^{2n}, ω) and (Q^{2n-2}, τ) be symplectic manifolds, where Q is compact. For any symplectic embedding $\iota \colon Q \to M$ with trivial normal bundle, there exists a symplectic embedding $f \colon Q \times B^2(\varepsilon) \to M$ such that $f(q, 0) = \iota(q)$ for all $q \in Q$, where $Q \times B^2(\varepsilon)$ has the product symplectic form $\tau \times \omega_{\text{std}}$.*

Now let (M_1, ω_1) and (M_2, ω_2) be symplectic manifolds of equal dimension $2n$, and let (Q, τ) be a compact symplectic manifold of dimension $2n - 2$. Suppose we are given symplectic embeddings $\iota_1 \colon Q \to M_1$ and $\iota_2 \colon Q \to M_2$ with trivial normal bundles, and let $f_j \colon Q \times B^2(\varepsilon) \to M_j$ be the symplectic embeddings given by Theorem 2.1. From each fibre $B^2(\varepsilon)$, remove a smaller disk of radius $\delta < \varepsilon$. The resulting annulus $A(\delta, \varepsilon)$ admits a symplectomorphism $\phi \colon A(\delta, \varepsilon) \to A(\delta, \varepsilon)$ which interchanges the two boundary components (for example, the map given in polar coordinates by $(r, \theta) \mapsto (\sqrt{\delta^2 + \varepsilon^2 - r^2}, -\theta))$. This allows us to glue $M_1 \setminus \iota_1(Q)$ and $M_2 \setminus \iota_2(Q)$ along their punctured neighbourhoods.

Definition 2.2. *The **fibre connected sum** of (M_1, ω_1) and (M_2, ω_2) along Q is given by*

$$M_1 \#_Q M_2 = \left(M_1 \setminus f_1(Q \times B^2(\delta))\right) \cup_\phi \left(M_2 \setminus f_2(Q \times B^2(\delta))\right),$$

where we identify $f_2(q, z)$ with $f_1(q, \phi(z))$ for $q \in Q$ and $\delta < |z| < \varepsilon$.

The symplectic forms ω_1 and ω_2 agree on the overlap $Q \times A(\delta, \varepsilon)$, and therefore yield a symplectic form on $M_1 \#_Q M_2$. While this construction depends on many choices (the constant δ and the framings f_j, for instance), we shall not concern ourselves with questions of uniqueness. The curious reader is referred to [3, Theorem 1.3] for details.

Remark 2.3. Why is the codimension 2 condition necessary for symplectic summing? This is because a symplectomorphism of a punctured d-disk which swaps the boundary components exists only when $d = 2$; otherwise, one could construct a symplectic structure on a homotopy d-sphere, which is impossible as the second Betti number of S^d vanishes for $d \neq 2$.

Remark 2.4. More generally, the fibre connected sum can be performed in the case where the normal bundles ν_1 and ν_2 of the embeddings are non-trivial, provided that their Euler classes satisfy $e(\nu_1) = -e(\nu_2)$. However, we will not need this.

Note that if $\dim M_1 = \dim M_2 = 2$ and Q is a single point, then the fibre connected sum $M_1 \#_Q M_2$ is the usual connected sum. Besides this special case, our main application of the fibre connected sum will be to kill elements of a symplectic manifold's fundamental group, as in the following example.

EXAMPLE 2.5. We shall use the following fact: there exists a symplectic 4-manifold V and a symplectically embedded 2-torus $T \subset V$ with trivial normal bundle such that $V \setminus T$ is simply connected. (Such a V can be found by considering the elliptic surface obtained by blowing up $\mathbb{C}P^2$ at nine distinct points lying on the intersection of two transverse nonsingular cubics—see [5, Example 7.1.7] for details.) Thus, given a symplectic 4-manifold X and a symplectic 2-torus $T' \subset X$ with trivial normal bundle, we may form the fibre connected sum

$$X' = X \#_{T'} V,$$

perhaps after rescaling the symplectic form on V so that T and T' have equal area. (In dimension 2, the notions of symplectomorphism and area-preserving diffeomorphism coincide.) From van Kampen's theorem, it follows that this kills $\pi_1(T')$; that is, if $\iota \colon T' \to X$ denotes the inclusion, then

$$\pi_1(X') \cong \pi_1(X) / \langle \iota_* \pi_1(T') \rangle,$$

where $\langle \iota_* \pi_1(T') \rangle$ denotes the normal closure of $\iota_* \pi_1(T')$.

Proof of Theorem 1.2. Fix a finite presentation $G \cong \langle g_1, \ldots, g_k \mid r_1, \ldots, r_\ell \rangle$. Let Σ_k be a compact, oriented Riemann surface of genus k; this is, in particular, a symplectic manifold. Choose a standard homology basis of $H_1(\Sigma_k)$ represented by oriented simple closed curves $\alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_k$ such that the α_i's are pairwise disjoint, as are the β_j's, and the intersection numbers satisfy $\alpha_i \cdot \beta_j = \delta_{ij}$.[1] The genus $k = 3$ case is depicted in Figure 2.1.

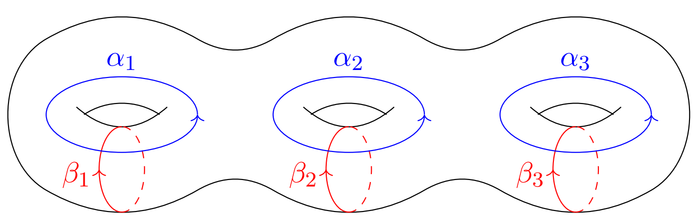

FIGURE 2.1. A standard homology basis of $H_1(\Sigma_3)$.

Upon attaching all curves to a common base point, we have

$$\pi_1(\Sigma_k) \cong \langle \alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_k \mid [\alpha_1, \beta_1] \cdots [\alpha_k, \beta_k] \rangle,$$

[1] c.f. the notion of a Darboux basis of a symplectic vector space.

so the quotient $\pi_1(\Sigma_k)/\langle \beta_1, \ldots, \beta_k \rangle$ is the free group generated by the α_i's. For $i \in \{1, \ldots, \ell\}$, let γ_i be an immersed oriented closed curve in Σ_k which represents the word r_i in $\pi_1(\Sigma_k)$. Set $\gamma_{\ell+i} = \beta_i$ for $i \in \{1, \ldots, k\}$. Then

$$G \cong \pi_1(\Sigma_k)/\langle \gamma_1, \ldots, \gamma_{\ell+k} \rangle.$$

For the moment, suppose that there exists a closed 1-form ρ on Σ_k whose restrictions to each γ_i is a volume form compatible with the orientation on γ_i. (Such a 1-form may not always exist for homological reasons, but we will show later how to modify the proof for such cases.) Denoting by $\mathbb{T}^2 = S^1 \times S^1$ the standard symplectic 2-torus, we consider the 4-manifold $X = \Sigma_k \times \mathbb{T}^2$ with the product symplectic form ω. Let $\alpha \subset \mathbb{T}^2$ be an oriented simple closed curve which is nontrivial in $H_1(\mathbb{T}^2)$, and choose a 1-form θ on $\mathbb{T}^2$ that restricts to a volume form on α which is compatible with its orientation. Set $\eta = \rho \wedge \theta$.

For $i \in \{1, \ldots, \ell+k\}$, let $T_i = \gamma_i \times \alpha$; then T_i is an immersed torus in X which is Lagrangian with respect to ω (being a product of Lagrangian submanifolds), and η restricts to a symplectic form on T_i. If ρ and θ are chosen sufficiently small (as measured by some Riemannian metric, say), then $\omega' := \omega + \eta$ is symplectic, and (X, ω') contains the symplectically immersed tori T_i and $\{z\} \times \mathbb{T}^2$, where $z \in \Sigma_k$ is disjoint from the γ_i's. Since each T_i lies in the 3-dimensional submanifold $\Sigma_k \times \alpha$, we may perturb the T_i's and $\{z\} \times \mathbb{T}^2$ to make them disjoint and symplectically embedded, with trivial normal bundle. As in Example 2.5, we attach $\ell + k + 1$ copies of the elliptic surface V to X along these tori; this has the effect of killing the homotopy classes represented by the γ_i's and by $\pi_1(\mathbb{T}^2)$. The resulting symplectic manifold M therefore has $\pi_1(M) \cong G$.

We are left with explaining what to do when the desired closed 1-form ρ on Σ_k does not exist. Let $x, y \in S^1$ be distinct points and consider the curves $\alpha = S^1 \times \{x\}$, $\beta = \{x\} \times S^1$, and $\gamma = \{y\} \times S^1$ in $\mathbb{T}^2$, where γ is oriented to be parallel to β. Choose a disk $D \subset \mathbb{T}^2$ that is disjoint from α and β, and intersects γ in an arc. Then there exists a closed 1-form ρ^* on $\mathbb{T}^2$ that vanishes in a neighbourhood of D such that $\int_\alpha \rho^*$, $\int_\beta \rho^*$, and $\int_{\gamma \setminus D} \rho^*$ are all strictly positive; this can be seen by collapsing a neighbourhood of D to a point and pulling back a volume form on the diagonal copy of $S^1 \subset \mathbb{T}^2$ along a projection $\mathbb{T}^2 \to S^1$ which maps α, β, and γ with degree 1.

By a general position argument, we may assume that the γ_i's have pairwise transversal intersections; their union therefore can be viewed as an oriented graph Γ on Σ_k. Let E denote the set of oriented edges and isolated circles in Γ. For each $e \in E$, choose a disk in Σ_k which intersects the interior of e in an arc and is disjoint from all other edges; take the connected sum with a copy of $\mathbb{T}^2$ by gluing in a copy of $\mathbb{T}^2 \setminus D$ so that $\gamma \setminus D$ "matches up" with e, including orientations, and the curves $\alpha, \beta \subset \mathbb{T}^2$ are added to the graph Γ (Figure 2.2).

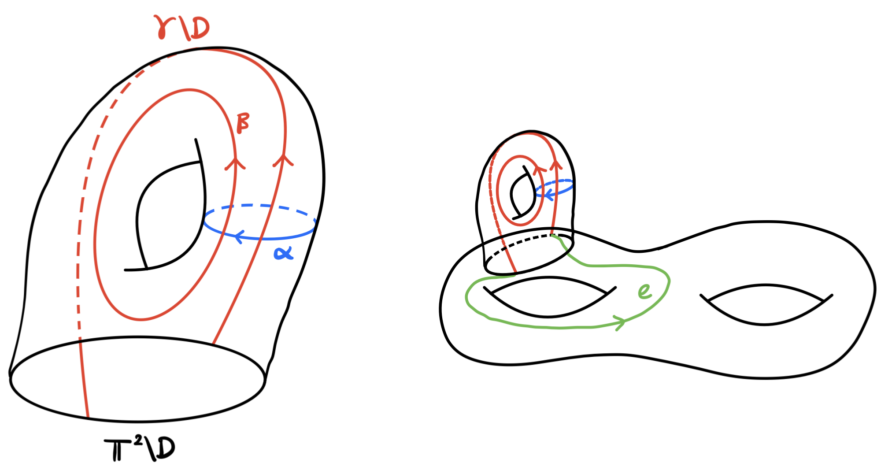

FIGURE 2.2. Attaching a copy of $\mathbb{T}^2 \setminus D$ so that $\gamma \setminus D$ matches up with e.

Call the resulting surface F. Since ρ^* vanishes near each copy of $D \subset \mathbb{T}^2$, we can extend by zero to obtain a closed 1-form, which we continue to denote by ρ^*, defined on all of F such that $\int_e \rho^* > 0$ for all edges $e \in \Gamma$. Let $\{\gamma_1, \ldots, \gamma_m\}$ be the set of curves containing the original γ_i's after summing with copies of γ, as well as each added copy of α and β in F. For each $i \in \{1, \ldots, m\}$, there exists a volume form θ_i on γ_i such that $\int_e \theta_i = \int_e \rho^*$ for every $e \in \Gamma$ which lies in γ_i. Then $(\theta_i - \rho^*)|_{\gamma_i} = df_i$ for some $f_i \in C^\infty(\gamma_i)$ which vanishes at the vertices of Γ lying on γ_i. From the f_i's, we obtain a smooth function $f \in C^\infty(F)$ such that $\rho := \rho^* + df$ is a closed 1-form on F which restricts to compatible volume forms on each γ_i. This is the desired ρ; since

$$G \cong \pi_1(F)/\langle \gamma_1, \ldots, \gamma_m \rangle,$$

the argument from before works with F replacing Σ_k and $\gamma_1, \ldots, \gamma_m$ replacing the original γ_i's. ■

Gompf used Theorem 1.2 to give an alternative proof of the following.

Theorem 2.6 (A'campo–Kotschick [1]). *Every finitely presentable group G is realized as the fundamental group of a closed contact 5-manifold.*

Recall that a *contact structure* on a $(2n+1)$-dimensional smooth manifold M is a $2n$-dimensional distribution $\xi \subset TM$ which is "maximally non-integrable" in the following sense: for any open set $U \subseteq M$ such that $\xi|_U = \ker(\alpha)$ for a 1-form α defined on U (such α's always exist locally), we have $\alpha \wedge (d\alpha)^n \neq 0$ everywhere on U. The pair (M, ξ) is called a *contact manifold*. Contact manifolds are, in many ways, the odd-dimensional analogues of symplectic manifolds.

We shall need the following result to prove Theorem 2.6.

Lemma 2.7. *If a closed manifold M admits a symplectic form ω, then it admits an integral symplectic form, i.e., a symplectic form ω' whose cohomology class $[\omega']$ lies in $H^2(M; \mathbb{Z})$.*

Proof. Fix a Riemannian metric on M and let B_ε be the open ball of radius $\varepsilon > 0$ in the space of all harmonic 2-forms with respect to this metric. Since non-degeneracy is an open condition, $\omega + B_\varepsilon$ consists entirely of symplectic forms for sufficiently small ε, and it contains an open set in $H^2_{\mathrm{dR}}(M)$. In particular, $\omega + B_\varepsilon$ contains some symplectic form ω'' whose cohomology class $[\omega'']$ lies in $H^2(M; \mathbb{Q})$. Scaling ω'' by a suitable integer, we obtain the desired symplectic form ω'. ■

Proof of Theorem 2.6. By Theorem 1.2 and Lemma 2.7, there exists a closed symplectic 4-manifold (M, ω) such that $\pi_1(M) \cong G$ and $[\omega] \in H^2(M; \mathbb{Z})$. Up to rescaling ω by an integer, we may symplectically blow up (M, ω) to obtain another symplectic 4-manifold (M', ω') with $\pi_1(M') \cong G$, $[\omega'] \in H^2(M', \mathbb{Z})$, and an embedded 2-sphere $S \subset M'$ such that $\int_S \omega' = 1$.

Let $P \to M'$ denote the principal circle bundle over M' with first Chern class $[\omega']$. Then $P|_S \to S$ is the Hopf fibration, and so inclusion of a fibre into P yields the zero map on π_1. From the long exact sequence of homotopy groups, we deduce that $\pi_1(P) \cong \pi_1(M') \cong G$. A result of Boothby and Wang [2, Theorem 3] furnishes a contact structure on P; it is essentially determined by the connection 1-form on P whose curvature is given by ω' up to a constant. ■

Let us summarize the main results we have seen so far and note a small generalization to higher dimensions, thus fully subsuming Theorem 1.1.

Corollary 2.8. *For any finitely presentable group G and fixed $n \geq 4$, there exists a closed smooth n-manifold M such that $\pi_1(M) \cong G$, which can be taken to be symplectic (if n is even) or contact (if n is odd).*

Proof. For even $n \geq 4$, this is the content of Theorem 1.2; the $n = 5$ case is covered by Theorem 2.6. For odd $n = 2k - 1$ with $k \geq 4$, there exists a closed smooth k-manifold N such that $\pi_1(N) \cong G$ by Theorem 1.1. The unit cotangent sphere bundle $M = ST^*N$ of N with respect to any Riemannian metric then yields the desired contact n-manifold, as $\pi_1(M) \cong \pi_1(N) \cong G$ by the long exact sequence of homotopy groups. ■

3. Non-Computability of Classification

We conclude by briefly discussing some consequences of the fact that all finitely presentable groups are realized as fundamental groups of n-manifolds for any $n \geq 4$. A famous problem in group theory is the *word problem* for a finite presentation of a group $G \cong \langle S \mid R \rangle$: given a word in the alphabet determined by elements of S and their formal inverses, determine whether this word represents the identity element in G.

Theorem 3.1 (Novikov–Boone [7, Chapter 12]). *There exists a finitely presented group $G \cong \langle S \mid R \rangle$ for which there does not exist an algorithm that takes as input a word in the alphabet $S \cup S^{-1}$ and outputs whether this word represents the identity element in G or not.*

Closely related to the word problem is the *isomorphism problem*: given two finite presentations, determine whether they represent isomorphic groups. In fact, it can be deduced from Theorem 3.1 that the isomorphism problem is also undecidable [8, Theorem 5]: there is no algorithm to determine whether or not two finite presentations represent the same group. As a consequence, we see that there can also be no algorithm which determines whether or not two n-manifolds are homeomorphic (or even homotopy equivalent), for any $n \geq 4$.

One might ask what questions about groups *do* admit solvable decision problems. A reasonably large class of properties that includes many elementary properties of groups is the following:

Definition 3.2. *A property P of finitely presentable groups is called a* ***Markov property*** *if the following conditions hold:*

- *P is preserved by group isomorphism.*
- *There exists a finitely presentable group satisfying P.*
- *There exists a finitely presentable group that does not embed into any finitely presentable group satisfying P.*

For instance, the following are Markov properties:

- Being abelian.
- Being finite.
- Being trivial.
- Being free.
- Being torsion-free.
- Being cyclic.

Unfortunately, the following "no-go" theorem shows that decision problems for Markov properties are undecidable.

Theorem 3.3 (Adian–Rabin). *For any Markov property P, there does not exist an algorithm which takes as input a finite presentation for a group G and outputs whether G satisfies P or not.*

The proof of this theorem can be reduced to the Novikov–Boone theorem. An English translation of the original papers proving this theorem is recently available on arXiv [6].

In light of the results we have seen in Sections 1 and 2, we see that many decision problems about the topology of n-manifolds are undecidable for $n \geq 4$. For example, there is no algorithm to decide whether a manifold is simply connected: if there were, then there would be an algorithm to determine whether a finitely presented group is trivial, which is impossible by Theorem 3.3. Moreover, Corollary 2.8 shows that the situation does not improve even if we restrict attention to symplectic or contact manifolds.

References

1. N. A'campo and D. Kotschick, *Contact structures, foliations, and the fundamental group*, Bulletin of the London Mathematical Society **26** (1994), no. 1, 102–106.
2. W. M. Boothby and H. C. Wang, *On contact manifolds*, Annals of Mathematics **68** (1958), no. 3, 721.
3. Robert E. Gompf, *A new construction of symplectic manifolds*, Annals of Mathematics **142** (1995), no. 3, 527–595.
4. Greg Kuperberg, *Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization*, Pacific Journal of Mathematics **301** (2019), no. 1, 189–241.
5. Dusa McDuff and Dietmar Salamon, *Introduction to symplectic topology*, 3rd ed., Oxford graduate texts in mathematics, no. 27, Oxford University Press, Oxford New York, NY, 2017.
6. Carl-Fredrik Nyberg-Brodda, *The Adian–Rabin theorem – an English translation*, August 2022, arXiv:2208.08560 [math].
7. Joseph J. Rotman, *An introduction to the theory of groups*, 4th ed., Graduate texts in mathematics, no. 148, Springer-Verlag, New York, 1995.
8. John Stillwell, *The word problem and the isomorphism problem for groups*, Bulletin of the American Mathematical Society **61** (1982), no. 1, 33–56.
9. Alan Weinstein, *Symplectic manifolds and their Lagrangian submanifolds*, Advances in Mathematics **6** (1971), no. 3, 329–346.