

From Invariants to Canonization in Parallel

Johannes Köbler¹ and Oleg Verbitsky^{2,*}

¹ Institut für Informatik, Humboldt Universität zu Berlin, D-10099 Berlin, Germany

koebler@informatik.hu-berlin.de

² Institute for Applied Problems of Mechanics and Mathematics,

Naukova 3b, 79060 Lviv, Ukraine

verbitsky@informatik.hu-berlin.de

Abstract. A function f of a graph is called a *complete graph invariant* if two given graphs G and H are isomorphic exactly when $f(G) = f(H)$. If additionally, $f(G)$ is a graph isomorphic to G , then f is called a *canonical form* for graphs. Gurevich [9] proves that any polynomial-time computable complete invariant can be transformed into a polynomial-time computable canonical form. We extend this equivalence to the polylogarithmic-time model of parallel computation for classes of graphs having either bounded rigidity index or small separators. In particular, our results apply to three representative classes of graphs embeddable into a fixed surface, namely, to 3-connected graphs admitting either a polyhedral or a large-edge-width embedding as well as to all embeddable 5-connected graphs. Another application covers graphs with treewidth bounded by a constant k . Since for the latter class of graphs a complete invariant is computable in NC, it follows that graphs of bounded treewidth have a canonical form (and even a canonical labeling) computable in NC.

1 Introduction

We write $G \cong H$ to indicate that G and H are isomorphic graphs. A *complete invariant* is a function f on graphs such that $f(G) = f(H)$ if and only if $G \cong H$. If, in addition, $f(G)$ is a graph isomorphic to G , then f is called a *canonical form* for graphs. For a given graph G and a one-to-one map σ on the vertices of G , we use G^σ to denote the isomorphic image of G under σ . A *canonical labeling* assigns to each graph G a map σ so that the function f defined as $f(G) = G^\sigma$ is a complete invariant. Note that f is even a canonical form. Thus, the notion of a canonical labeling is formally stronger than that of a canonical form which in turn is formally stronger than that of a complete invariant.

Obviously, a polynomial-time computable complete invariant can be used to decide in polynomial time whether two given graphs are isomorphic. Conversely, it is not known whether a polynomial-time decision algorithm for graph isomorphism implies the existence of a polynomial-time complete invariant (cf. the discussion in [2, Sect. 5]). However, for many classes of graphs for which we

* Supported by an Alexander von Humboldt fellowship.

have an efficient isomorphism test, we also have a canonical labeling algorithm of comparable complexity (see, e.g., [17,14]); but this often requires substantial additional efforts (cf., e.g., [3,1]).

Gurevich [9] proves that a polynomial-time computable complete graph invariant can be used to compute a canonical labeling in polynomial time. This result is really enlightening because there are approaches to the graph isomorphism problem which are based on computing a graph invariant and, without an extra work, do not provide us with a canonical form. An important example is the k -dimensional Weisfeiler-Lehman algorithm WL^k . Given an input graph G , the algorithm outputs a coloring of its vertices in polynomial time (where the degree of the polynomial bounding the running time depends on k). WL^k always produces the same output for isomorphic input graphs. Whether the algorithm is able to distinguish G from every non-isomorphic input graph H depends on whether the dimension k is chosen large enough for G . In particular, $k = 1$ suffices for all trees T . However, notice that the coloring computed by WL^1 on input T partitions the vertex set of T into the orbits of the automorphism group of T and hence WL^1 does not provide a canonical labeling unless T is rigid (i.e., T has only the trivial automorphism). We mention that an appropriate modification of the 1-dimensional Weisfeiler-Lehman algorithm to a canonical labeling algorithm is suggested in [11].

The reduction of a canonical labeling to a complete invariant presented in [9] (as well as in [11]) is inherently sequential and thus leaves open the following question.

Question 1. Suppose that for the graphs in a certain class C we are able to compute a complete invariant in NC. Is it then possible to compute also a canonical labeling for these graphs in NC?

For several classes of graphs, NC algorithms for computing a complete invariant are known (see, e.g., [4,17,14,10]). For example, in [10] it is shown that a k -dimensional Weisfeiler-Lehman algorithm making logarithmically many rounds can be implemented in $TC^1 \subseteq NC^2$ and that such an algorithm succeeds for graphs of bounded treewidth. Similar techniques apply also to planar graphs but for this class a canonical labeling algorithm in AC^1 is known from an earlier work [17]. Nevertheless, also in this case it is an interesting question whether the approach to the planar graph isomorphism problem suggested in [10], which is different from the approach of [17], can be adapted for finding a canonical labeling. Finally, Question 1 even makes sense for classes C for which we don't know of any NC-computable complete invariant since such an invariant may be found in the future.

We notice that a positive answer to Question 1 also implies that the search problem of computing an isomorphism between two given graphs in C , if it exists, is solvable in NC whenever for C we have a complete invariant in NC (notice that the known polynomial-time reduction of this search problem to the decision version of the graph isomorphism problem is very sequential in nature, see [12]).

As our main result we give an affirmative answer to Question 1 for any class of graphs having either small *separators* (Theorem 2) or bounded *rigidity index* (Theorem 5). A quite general example for a class of graphs having small separators is the class of graphs whose treewidth is bounded by a constant. Since, as mentioned above, a complete invariant for these graphs is computable in TC^1 [10], Theorem 2 immediately provides us with an NC (in fact TC^2) canonical labeling algorithm for such graphs (Corollary 3). As a further application we also get a TC^2 algorithm for solving the search problem for pairs of graphs in this class (Corollary 4).

Regarding the second condition we mention the following representative classes of graphs with bounded rigidity index:

- 3-connected graphs having a *large-edge-width* embedding into a fixed surface S (Corollary 7).¹
- 3-connected graphs having a *polyhedral* embedding into a fixed surface S (Corollary 9).
- 5-connected graphs embeddable into a fixed surface S (Corollary 10).

As shown by Miller and Reif [17], the canonization problem for any hereditary class of graphs C (meaning that C is closed under induced subgraphs) AC^1 reduces to the canonization problem for the class of all 3-connected graphs in C . Thus, with respect to the canonization problem, the 3-connected case is of major interest.

The rest of the paper is organized as follows. In Sect. 2 we provide the necessary notions and fix notation. Graphs with small separators are considered in Sect. 3 and graphs with bounded rigidity index are considered in Sect. 4. Section 5 summarizes our results and discusses remaining open problems.

2 Preliminaries

The concept of polylogarithmic parallel time is captured by the hierarchy of complexity classes $\text{NC} = \bigcup_{i \geq 1} \text{NC}^i$, where NC^i consists of functions computable by DLOGTIME-constructible boolean circuits of polynomial size and depth $O(\log^i n)$. The class AC^i is the extension of NC^i to circuits with unbounded fan-in and TC^i is a further extension allowing threshold gates as well. Recall also that $\text{AC}^0 \subseteq \text{TC}^0 \subseteq \text{NC}^1 \subseteq \text{L} \subseteq \text{NL} \subseteq \text{AC}^1$ and $\text{NC}^i \subseteq \text{AC}^i \subseteq \text{TC}^i \subseteq \text{NC}^{i+1}$, where L (resp. NL) is the set of languages accepted by (non)deterministic Turing machines using logarithmic space. Alternatively, the AC^i level of the NC hierarchy can be characterized as the class of all functions computable by a CRCW PRAM with polynomially many processors in time $O(\log^i n)$.

The vertex set of a graph G is denoted by $V(G)$. The set of all vertices adjacent to a vertex $v \in V(G)$ is called the *neighborhood* of v and denoted by $\Gamma(v)$.

A *colored graph* G , besides the binary adjacency relation, has unary relations U_1, \dots, U_n defined on $V(G)$. If a vertex v satisfies U_i , we say that v has color i . A vertex is allowed to have more than one color or none. It is supposed that the number of colors is equal to the number of vertices in a graph, though some

¹ This result is actually stated in a stronger form, without referring to a parameter S .

of the color relations may be empty. A colored graph $\langle G, U_1, \dots, U_n \rangle$ will be called a *coloring* of the underlying graph G . An *isomorphism* between colored graphs must preserve the adjacency relation as well as the color relations. Thus, different colorings of the same underlying graph need not be isomorphic.

We consider only classes of graphs that are closed under isomorphism. For a given class of graphs C we use C^* to denote the class containing all colorings of any graph in C .

Let C be a class of graphs and let f be a function mapping graphs to strings over a finite alphabet. We call f a *complete invariant for C* if for any pair of graphs G and H in C we have $G \cong H$ exactly when $f(G) = f(H)$. A *canonical labeling for C* assigns to each graph G on n vertices a one-to-one map $\sigma : V(G) \rightarrow \{1, \dots, n\}$ such that $f(G) = G^\sigma$ is a complete invariant for C . Note that a complete invariant f originating from a canonical labeling has an advantageous additional property: $f(G) \neq f(H)$ whenever G is in C and H is not. Moreover, it provides us with an isomorphism between G and H whenever $f(G) = f(H)$.

The notions of a complete invariant and of a canonical labeling are easily extensible to colored graphs. In our proofs, extending these notions to colored graphs will be technically beneficial and, at the same time, will not restrict the applicability of our results. In fact, any available complete-invariant algorithm for some class of graphs C can be easily extended to C^* without increasing the required computational resources. In particular, this is true for the parallelized version of the multi-dimensional Weisfeiler-Lehman algorithm suggested in [10].

3 Small Separators

For a given graph G and a set X of vertices in G , let $G - X$ denote the graph obtained by removing all vertices in X from G . A set X is called a *separator* if every connected component of $G - X$ has at most $n/2$ vertices, where n denotes the number of vertices of G . A class of graphs C is called *hereditary* if for every $G \in C$, every induced subgraph of G also belongs to C .

Theorem 2. *Let C be a hereditary class of graphs such that for a constant r , every graph $G \in C$ has an r -vertex separator. Suppose that C^* has a complete invariant f computable in TC^k (resp. AC^k) for some $k \geq 1$. Then C has a canonical labeling in TC^{k+1} (resp. AC^{k+1}).*

Proof. Having f in our disposal, we design a canonical labeling algorithm for C . Let G be an input graph with vertex set $V(G) = \{1, \dots, n\}$ and assume that G has an r -vertex separator. We describe a recursive algorithm for finding a canonical renumbering $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. In the following, the parameter d refers to the recursion depth. Initially $d = 1$. Further, set $R = 2^r + r$.

For a given sequence $s = (v_1, \dots, v_r)$ of vertices, let G_s denote the coloring of G in which v_i receives color $(d-1)R + i$.

For each sequence $s = (v_1, \dots, v_r)$ in parallel we do the following. First of all, we check if the set $\{v_1, \dots, v_r\}$ is a separator. We are able to do this in AC^1 since checking if two vertices are in the same connected component reduces to

the s - t -connectivity problem, which is easily solvable in NL, and the remaining job can be easily organized in TC^0 . If the verification is positive, we mark the sequence s as *separating*. If no such sequence s is separating, i.e., G has no r -vertex separator, we terminate and output the identity permutation. Otherwise, for each separating sequence s in parallel we compute $f(G_s)$. Then in AC^1 we find a sequence $s = (v_1, \dots, v_r)$ for which the value $f(G_s)$ is lexicographically minimum. For this purpose we use the fact that lexicographic comparison can be done in AC^0 and employ a known TC^0 sorting algorithm.

At this stage we are able to determine the renumbering σ only in a few points. Namely, we set $\sigma(v_i) = (d-1)R + i$ for each $i \leq r$.

To proceed further, let F_1, \dots, F_m be the connected components of $G - X$ where $X = \{v_1, \dots, v_r\}$. We color each $v \notin X$ by its adjacency pattern to X , that is, by the set of all neighbors of v in X , encoding this set by a number in the range between $(d-1)R + r + 1$ and dR . Each F_j , regarded as a colored graph, will be called an X -flap. For each X -flap F_j in parallel, we now compute $f(F_j)$ and establish the lexicographic order between these values. At this stage we fix the following partial information about the renumbering σ under construction: $\sigma(u) < \sigma(v)$ whenever we have $f(F_i) < f(F_j)$ for the two flaps F_i and F_j containing u and v , respectively. Thus, we split $V(G) \setminus X$ into blocks $V(F_1), \dots, V(F_m)$ and determine the renumbering σ first between the blocks. It may happen that for some flaps we have $f(F_i) = f(F_j)$. We fix the σ -order between the corresponding blocks arbitrarily. Note that the output will not depend on a particular choice made at this point.

It remains to determine σ inside each block $V(F_j)$. We do this in parallel. For F_j with more than r vertices we repeat the same procedure as above with the value of d increased by 1. If $F = F_j$ has $t \leq r$ vertices, we proceed as follows. Let a be the largest color present in F . We choose a bijection $\tau : V(F) \rightarrow \{1, \dots, t\}$ and define σ on $V(F)$ by $\sigma(u) < \sigma(v)$ if and only if $\tau(u) < \tau(v)$. To make the choice, with each such τ we associate the colored graph F_τ obtained from F by adding new colors, namely, by coloring each $v \in V(F)$ with color $a + \tau(v)$. For each τ we compute $f(F_\tau)$ and finally choose the τ minimizing $f(F_\tau)$ in the lexicographic order. Note that, if the minimum is attained by more than one τ , the output will not depend on a particular choice.

Finally, we have to estimate the depth of the TC (resp. AC) circuit implementing the described algorithm. At the recursive step of depth d we deal with graphs having at most $n/2^{d-1}$ vertices. It follows that the circuit depth does not exceed $\log_2^k n + \log_2^k(n/2) + \log_2^k(n/4) + \dots + \log_2^k(r) \leq \log_2^{k+1} n$. \square

It is well known that all graphs of treewidth t have a $(t+1)$ -vertex separator [21]. By [10], this class has a complete invariant computable in TC^1 and therefore is in the scope of Theorem 2.

Corollary 3. *For each constant t , a canonical labeling for graphs of treewidth at most t can be computed in TC^2 .*

Theorem 2 also has relevance to the complexity-theoretic *decision-versus-search* paradigm. Let C be a class of graphs. It is well known (see, e.g., [12]) that, if we

are able to test isomorphism of graphs in C^* in polynomial time, we are also able to find an isomorphism between two given isomorphic graphs in C in polynomial time. As the standard reduction is very sequential in nature, it is questionable if this implication stays true in the model of parallel computation. Nevertheless, a canonical labeling immediately provides us with an isomorphism between two isomorphic graphs.

Corollary 4. *For each constant t , an isomorphism between isomorphic graphs of treewidth at most t can be computed in TC^2 .*

4 Bounded Rigidity Index

In this section we show that the canonization problem for any class of graphs with bounded rigidity index NC reduces to the corresponding complete invariant problem. Further we show that certain embeddability properties of a given class of graphs C imply a bound on the rigidity index of the graphs in C .

4.1 Canonizing Rigid Graphs

A set $S \subseteq V(G)$ of vertices is called *fixing* if every non-trivial automorphism of G moves at least one vertex in S . The *rigidity index* of a graph G is defined to be the minimum cardinality of a fixing set in G and denoted by $\text{rig}(G)$.

Theorem 5. *Let C be a class of graphs such that for a constant r , we have $\text{rig}(G) \leq r$ for all $G \in C$. Suppose that C^* has a complete invariant f computable in AC^k , for some $k \geq 1$. Then C has a canonical labeling also in AC^k .*

Proof. Let an input graph G with vertex set $V(G) = \{1, \dots, n\}$ be given. We describe an algorithm that uses f as a subroutine in order to find a canonical renumbering $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ for G , provided that $G \in C$.

For a given sequence $s = (v_1, \dots, v_r)$ of vertices, let G_s denote the coloring of G in which v_i receives color i . If v is another vertex, $G_{s,v}$ denotes the coloring where vertex v additionally gets color $r + 1$.

For each such sequence s in parallel we do the following. For each v in parallel we compute $f(G_{s,v})$. If all the values $f(G_{s,v})$, $v \in V(G)$, are pairwise distinct, which is decidable in AC^0 , mark s as *fixing*. If no fixing sequence s of length r exists, which implies $G \notin C$, we terminate and output the identity permutation. Otherwise, for each fixing sequence s in parallel, we compute $f(G_s)$ and determine a sequence $s = (v_1, \dots, v_r)$ for which $f(G_s)$ is lexicographically minimum (as was already mentioned in the proof of Theorem 2, this can be done in TC^0). The output permutation σ is now computed as follows. For each $i \leq r$, we set $\sigma(v_i) = i$. To determine σ everywhere else, we sort the values $f(G_{s,v})$ for all $v \in V(G) \setminus \{v_1, \dots, v_r\}$ lexicographically and set $\sigma(v)$ to be the number of v in this order increased by r . \square

4.2 Basics of Topological Graph Theory

A detailed exposition of the concepts discussed in this section can be found in [19]. We are interested in embeddability of an abstract graph G into a surface S . All surfaces are supposed to be 2-dimensional, connected, and closed.

In an *embedding* Π of G into S , each vertex v of G is represented by a point on S (labeled by v and called *vertex of the Π -embedded graph G*) and each edge uv of G is drawn on S as a continuous curve with endpoints u and v . The curves are supposed to be non-self-crossing and any two such curves either have no common point or share a common endpoint. A *face* of Π is a connected component of the space obtained from S by removing the curves. We consider only *cellular* embeddings meaning that every face is homeomorphic to an open disc. A *closed walk* in a graph is a sequence of vertices $v_1 v_2 \cdots v_k$ such that v_i and v_{i+1} are adjacent for any $i < k$, and v_1 and v_k are adjacent as well. Notice that some of the vertices may coincide. We will not distinguish between a closed walk $v_1 v_2 \cdots v_k$ and any cyclic shift of it or of its reversal $v_k v_{k-1} \cdots v_1$. A closed walk $v_1 v_2 \cdots v_k$ is called *Π -facial*, if there exists a face F of Π , such that the vertices v_1, v_2, \dots, v_k occur in this order as labels along the boundary of F .

Two embeddings Π and Π' of G into S are called *equivalent* if they can be obtained from each other by a homeomorphism of S onto itself (respecting vertex labels). Since such a homeomorphism takes faces of one embedding to faces of the other embedding, we see that equivalent embeddings have equal sets of facial walks. In fact, the converse is also true: if the set of the Π -facial walks is equal to the set of the Π' -facial walks, then Π and Π' are equivalent. This follows from the fact that up to homeomorphism, the surface S is reconstructible from the set of facial walks by attaching an open disc along each facial walk.

A closed walk $v_1 v_2 \cdots v_k$ can be alternatively thought of as the sequence of edges $e_1 e_2 \cdots e_k$ where $e_i = v_i v_{i+1}$ ($i < k$) and $e_k = v_1 v_k$. Every edge either appears in two Π -facial walks (exactly once in each) or has exactly two occurrences in a single Π -facial walk. An embedding Π is called *polyhedral* if every Π -facial walk is a cycle (i.e., contains at most one occurrence of any vertex) and every two Π -facial walks either have at most one vertex in common or share exactly one edge (and no other vertex).

Let $\text{Aut}(G)$ denote the automorphism group of G . For a given automorphism $\alpha \in \text{Aut}(G)$, let Π^α denote the embedding of G obtained from Π by relabeling the vertices according to α . Note that Π^α and Π are not necessarily equivalent (they are *topologically isomorphic*, that is, obtainable from one another by a surface homeomorphism which is allowed to ignore the vertex labeling). An embedding Π is called *faithful* if Π^α is equivalent to Π for every automorphism $\alpha \in \text{Aut}(G)$.

Recall that a graph G is *k -connected* if it has at least $k + 1$ vertices and stays connected after removing any set of at most $k - 1$ vertices. We now summarize known results showing that, for $k \geq 3$, the flexibility of embedding a k -connected graph into certain surfaces is fairly restricted.

The Whitney Theorem. [23] *Up to equivalence, every 3-connected planar graph has a unique embedding into the sphere.*

The Mohar-Robertson Theorem. [18] *Up to equivalence, every connected² graph has at most c_S polyhedral embeddings into a surface S , where c_S is a constant depending only on S .*

A closed curve in a surface is *contractible* if it is homotopic to a point. The *edge-width* of an embedding Π is the minimum length of a non-contractible cycle in the Π -embedded graph. Π is called a *large-edge-width embedding* (abbreviated as *LEW embedding*) if its edge-width is larger than the maximum length of a Π -facial walk.

The Thomassen Theorem. [22] (see also [19, Corollary 5.1.6]) *Every 3-connected graph having a LEW embedding into a surface S has, up to equivalence, a unique embedding into S . Moreover, such a surface S is unique.*

Note that if a graph has a unique embedding into a surface (as in the Whitney Theorem or the Thomassen Theorem), then this embedding is faithful.

As we have seen, an embedding is determined by its set of facial walks (up to equivalence). We will need yet another combinatorial specification of an embedding. To simplify the current exposition, we restrict ourselves to the case of orientable surfaces.

Let G be a graph and let T be a ternary relation on the vertex set $V(G)$ of G . We call $R = \langle G, T \rangle$ a *rotation system* of G if T fulfills the following two conditions:

- (1) If $T(a, b, c)$ holds, then b and c are in $\Gamma(a)$, the neighborhood of a in G .
- (2) For every vertex a , the binary relation $T(a, \cdot, \cdot)$ is a directed cycle on $\Gamma(a)$ (i.e., for every b there is exactly one c such that $T(a, b, c)$, for every c there is exactly one b such that $T(a, b, c)$, and the digraph $T(a, \cdot, \cdot)$ is connected on $\Gamma(a)$).

An embedding Π of a graph G into an orientable surface S determines a rotation system $R_\Pi = \langle G, T_\Pi \rangle$ in a natural geometric way. Namely, for $a \in V(G)$ and $b, c \in \Gamma(a)$ we set $T_\Pi(a, b, c) = 1$ if, looking at the neighborhood of a in the Π -embedded graph G from the outside of S , b is followed by c in the clockwise order.

The *conjugate* of a rotation system $R = \langle G, T \rangle$, denoted by R^* , is the rotation system $\langle G, T^* \rangle$, where T^* is defined as $T^*(a, b, c) = T(a, c, b)$. This notion has two geometric interpretations. First, $(R_\Pi)^*$ is a variant of R_Π where we look at the Π -embedded graph from the inside rather than from the outside of the surface (or, staying outside, just change the clockwise order to the counter-clockwise order). Second, $(R_\Pi)^* = R_{\Pi^*}$ where Π^* is a mirror image of Π .

It can be shown that two embeddings Π and Π' of G into S are equivalent if and only if $R_\Pi = R_{\Pi'}$ or $R_\Pi = R_{\Pi'}^*$ (see [19, Corollary 3.2.5]).

Further, for a given rotation system $R = \langle G, T \rangle$ and automorphism $\alpha \in \text{Aut}(G)$, we define another rotation system $R^\alpha = \langle G, T^\alpha \rangle$ by $T^\alpha(a, b, c) = T(\alpha^{-1}(a), \alpha^{-1}(b), \alpha^{-1}(c))$. It is not hard to see that $R_\Pi^\alpha = R_{\Pi^\alpha}$. If $R = \langle G, T \rangle$ and $R' = \langle G, T' \rangle$ are two rotation systems of the same graph G and $R' = R^\alpha$ for some $\alpha \in \text{Aut}(G)$, then this equality means that α is an isomorphism from R' onto R (respecting not only the binary adjacency relation but also the ternary relations of these structures).

² It is known that only 3-connected graphs have polyhedral embeddings.

4.3 Rigidity from Non-flexible Embeddability

Let α be a mapping defined on a set V . We say that α *fixes an element* $x \in V$ if $\alpha(x) = x$. Furthermore, we say that α *fixes a set* $X \subseteq V$ if α fixes every element of X .

Lemma 6. *If a graph G has a faithful embedding Π into some surface S , then $\text{rig}(G) \leq 3$.*

Proof. Clearly, G is connected as disconnected graphs don't have a cellular embedding. If G is a path or a cycle, then $\text{rig}(G) \leq 2$. Otherwise, G contains some vertex v with at least 3 neighbors. Notice that a facial walk cannot contain a segment of the form uvu . Therefore, some facial walk W contains a segment uvw , where u and w are two different neighbors of v . As v has at least one further neighbor that is distinct from u and w , uvw cannot be a segment of any other facial walk than W .

We now show that $\{u, v, w\}$ is a fixing set. Assume that α is an automorphism of G that fixes the vertices u, v and w . We have to prove that α is the identity.

Note that $v_1v_2 \cdots v_k$ is a Π -facial walk if and only if $\alpha(v_1)\alpha(v_2) \cdots \alpha(v_k)$ is a Π^α -facial walk. Since Π and Π^α are equivalent and hence, have the same facial walks, α takes each Π -facial walk to a Π -facial walk. It follows that α takes W onto itself. Since α fixes two consecutive vertices of W , it actually fixes W .

Call two Π -facial walks W_1 and W_2 *adjacent* if they share an edge. Suppose that adjacent facial walks W_1 and W_2 share an edge u_1u_2 and that α fixes W_1 . Since u_1u_2 cannot participate in any third facial walk, α takes W_2 onto itself. Since u_1 and u_2 are fixed, α fixes W_2 , too.

Now consider the graph whose vertices are the Π -facial walks with the adjacency relation defined as above. It is not hard to see that this graph is connected, implying that α is the identity on the whole vertex set $V(G)$. □

By the Thomassen Theorem and by Lemma 6 it follows that every 3-connected LEW embeddable graph has rigidity index at most 3. Hence we can apply Theorem 5 to obtain the following result.

Corollary 7. *Let C be any class consisting only of 3-connected LEW embeddable graphs. If C^* has a complete invariant computable in AC^k , $k \geq 1$, then C has a canonical labeling in AC^k .*

Noteworthy, the class of all 3-connected LEW embeddable graphs is recognizable in polynomial time [19, Theorem 5.1.8].

Lemma 8. *If a connected graph G has a polyhedral embedding into a surface S , then we have $\text{rig}(G) \leq 4c$, where c is the total number of non-equivalent polyhedral embeddings of G into S .*

Proof. To simplify the current exposition, we prove the lemma only for the case that S is orientable. Let $a \in V(G)$ be a vertex in G . We call two rotation systems $R = \langle G, T \rangle$ and $R' = \langle G, T' \rangle$ of G *a-coherent* if the binary relations $T(a, \cdot, \cdot)$ and $T'(a, \cdot, \cdot)$ coincide.

Claim 1. Let ab be an edge in G . Then any isomorphism α between two a -coherent rotation systems $R = \langle G, T \rangle$ and $R' = \langle G, T' \rangle$ of G that fixes both a and b , fixes also $\Gamma(a)$.

Proof of Claim. Since α fixes a , it takes $\Gamma(a)$ onto itself. Since $T(a, \cdot, \cdot) = T'(a, \cdot, \cdot)$, α is an automorphism of this binary relation. The latter is a directed cycle and α must be a shift thereof. Since α fixes b , it has to fix the whole cycle. \triangleleft

Let R_1, \dots, R_{2c} (where $R_i = \langle G, T_i \rangle$) be the rotation systems representing all polyhedral embeddings of G into S (i.e., each of the c embeddings is represented by two mutually conjugated rotation systems). Pick an arbitrary edge xy in G . For each i , $1 < i \leq 2c$, select a vertex x_i so that R_i and R_1 are not x_i -coherent and the distance between x and x_i is minimum (it may happen that $x_i = x$). Furthermore, select y_i and z_i in $\Gamma(x_i)$ so that $T_1(x_i, y_i, z_i) \neq T_i(x_i, y_i, z_i)$. We will show that $\{x, y, y_2, z_2, \dots, y_{2c}, z_{2c}\}$ is a fixing set. Assume that $\alpha \in \text{Aut}(G)$ fixes all these vertices. We have to show that α is the identity.

Notice that R_1^α is a polyhedral embedding of G into S because so is R_1 . Therefore $R_1^\alpha = R_k$ for some $k \leq 2c$. Suppose first that R_k and R_1 are x -coherent. We will apply Claim 1 repeatedly to $R = R_k$ and $R' = R_1$. We first put $a = x$ and $b = y$ and see that α fixes $\Gamma(x)$. If the distance between x and x_k is more than 1, we apply Claim 1 once again for xx' being the first edge of a shortest path P from x to x_k (now $a = x'$ and $b = x$; we have $\alpha(x') = x'$ as $x' \in \Gamma(x)$, and R_k and R_1 are x' -coherent by our choice of x_k). Applying Claim 1 successively for all edges along P except the last one, we arrive at the conclusion that $\alpha(x_k) = x_k$. This also applies for the case that R_k and R_1 are not x -coherent, when we have $x_k = x$ by definition.

It follows that α is an isomorphism between the cycles $T_k(x_k, \cdot, \cdot)$ and $T_1(x_k, \cdot, \cdot)$. Our choice of y_k and z_k rules out the possibility that $k \geq 2$ and we conclude that $k = 1$. In other words, R and R' are coherent everywhere. Therefore, we are able to apply Claim 1 along any path starting from the edge $ab = xy$. Since G is connected, we see that α is the identity permutation on $V(G)$. \square

By the Mohar-Robertson Theorem and Lemma 8 it follows that every connected graph having a polyhedral embedding into a surface S has rigidity index bounded by a constant depending only on S .³ Applying Theorem 5, we obtain the following result.

Corollary 9. *Let C be any class containing only graphs having a polyhedral embedding into a fixed surface S . If C^* has a complete invariant computable in AC^k , $k \geq 1$, then C has a canonical labeling in AC^k .*

We conclude this section by applying a ready-to-use result on the rigidity index of 5-connected graphs that are embeddable into a fixed surface S .

The Fijavž-Mohar Theorem. [7] *The rigidity index of 5-connected graphs embeddable into a surface S is bounded by a constant depending only on S .*

³ As we recently learned, this result has been independently obtained in [7] by using a different argument.

Corollary 10. *Let C be the class of 5-connected graphs embeddable into a fixed surface S . If C^* has a complete invariant computable in AC^k , $k \geq 1$, then C has a canonical labeling in AC^k .*

5 Conclusion and Open Problems

For several important classes of graphs, we provide NC Turing-reductions of canonical labeling to computing a complete invariant. As a consequence, we get a canonical labeling NC algorithm for graphs with bounded treewidth by using a known [10] NC-computable complete invariant for such graphs.

We also consider classes of graphs embeddable into a fixed surface. Though we currently cannot cover this case in full extent, we provide NC reductions between the canonical labeling and complete invariant problems for some representative subclasses (namely, 3-connected graphs with either a polyhedral or an LEW embedding as well as all embeddable 5-connected graphs).

To the best of our knowledge, complete invariants (even isomorphism tests) in NC are only known for the sphere but not for any other surface. The known isomorphism tests and complete-invariant algorithms designed in [5,13,15,16,8] run in sequential polynomial time. Nevertheless, the hypothesis that the complexity of some of these algorithms can be improved from P to NC seems rather plausible. By this reason it would be desirable to extend the reductions proved in the present paper to the whole class of graphs embeddable into S , for any fixed surface S . As a first step in this direction one could consider the class of 4-connected toroidal graphs.⁴

A more ambitious research project is to find an NC-reduction of the canonical labeling problem to computing a complete invariant for classes of graphs that are defined by excluding certain graphs as minors or, equivalently, for classes of graphs closed under minors. A polynomial-time canonization algorithm for such classes has been worked out by Ponomarenko [20]. Note that any class of graphs with bounded treewidth as well as any class consisting of all graphs embeddable into a fixed surface is closed under minors.

Acknowledgement. We thank Gašper Fijavž and Bojan Mohar for sending us their manuscript [7].

References

1. Arvind, V., Das, B., Mukhopadhyay, P.: On isomorphism and canonization of tournaments and hypertournaments. In: Asano, T. (ed.) ISAAC 2006. LNCS, vol. 4288, pp. 449–459. Springer, Heidelberg (2006)
2. Arvind, V., Torán, J.: Isomorphism testing: Perspective and open problems. Bulletin of the European Association of Theoretical Computer Science 86, 66–84 (2005)
3. Babai, L., Luks, E.: Canonical labeling of graphs. In: Proc. 15th ACM Symposium on Theory of Computing, pp. 171–183. ACM Press, New York (1983)

⁴ As shown in [7], graphs in this class can have arbitrarily large rigidity index.

4. Chlebus, B.S., Diks, K., Radzik, T.: Testing isomorphism of outerplanar graphs in parallel. In: Koubek, V., Janiga, L., Chytil, M.P. (eds.) MFCS 1988. LNCS, vol. 324, pp. 220–230. Springer, Heidelberg (1988)
5. Filotti, I.S., Mayer, J.N.: A polynomial-time algorithm for determining the isomorphism of graphs of fixed genus. In: Proc. 12th ACM Symposium on Theory of Computing, pp. 236–243. ACM Press, New York (1980)
6. Fijavž, G., Mohar, B.: Rigidity and separation indices of Paley graphs. *Discrete Mathematics* 289, 157–161 (2004)
7. Fijavž, G., Mohar, B.: Rigidity and separation indices of graphs in surfaces. A manuscript in preparation, cited in [6]
8. Grohe, M.: Isomorphism testing for embeddable graphs through definability. In: Proc. 32th ACM Symposium on Theory of Computing, pp. 63–72. ACM Press, New York (2000)
9. Gurevich, Y.: From invariants to canonization. *Bulletin of the European Association of Theoretical Computer Science* 63, 115–119 (1997)
10. Grohe, M., Verbitsky, O.: Testing graph isomorphism in parallel by playing a game. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 3–14. Springer, Heidelberg (2006)
11. Immerman, N., Lander, E.: Describing graphs: a first order approach to graph canonization. In: Selman, A.L. (ed.) *Complexity Theory Retrospective*, pp. 59–81. Springer, Heidelberg (1990)
12. Köbler, J., Schöning, U., Torán, J.: *The Graph Isomorphism Problem: Its Structural Complexity*. Birkhäuser, Boston (1993)
13. Lichtenstein, D.: Isomorphism for graphs embeddable on the projective plane. In: Proc. 12th ACM Symposium on Theory of Computing, pp. 218–224. ACM Press, New York (1980)
14. Lindell, S.: A logspace algorithm for tree canonization. In: Proc. 24th ACM Symposium on Theory of Computing, pp. 400–404. ACM Press, New York (1992)
15. Miller, G.L.: Isomorphism testing for graphs of bounded genus. In: Proc. 12th ACM Symposium on Theory of Computing, pp. 225–235. ACM Press, New York (1980)
16. Miller, G.L.: Isomorphism of k -contractible graphs. A generalization of bounded valence and bounded genus. *Information and Computation* 56, 1–20 (1983)
17. Miller, G.L., Reif, J.H.: Parallel tree contraction. Part 2: Further applications. *SIAM Journal on Computing* 20, 1128–1147 (1991)
18. Mohar, B., Robertson, N.: Flexibility of polyhedral embeddings of graphs in surfaces. *J. Combin. Theory, Ser. B* 83, 38–57 (2001)
19. Mohar, B., Thomassen, C.: *Graphs on surfaces*. The John Hopkins University Press (2001)
20. Ponomarenko, I.: The isomorphism problem for classes of graphs closed under contraction. In: *Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov (in Russian)*, vol. 174, pp. 147–177 (1988); English translation in: *Journal of Mathematical Sciences* 55, 1621–1643 (1991)
21. Robertson, N., Seymour, P.D.: Graph minors II. Algorithmic aspects of tree-width. *J. Algorithms* 7, 309–322 (1986)
22. Thomassen, C.: Embeddings of graphs with no short noncontractible cycles. *J. Combin. Theory, Ser. B* 48, 155–177 (1990)
23. Whitney, H.: 2-isomorphic graphs. *Amer. Math. J.* 55, 245–254 (1933)