On the Complexity and Combinatorics of Covering Finite Complexes

James Abello
Department of Computer Science
Texas A & M University
College Station, Texas 77843, U.S.A.

Michael R. Fellows
Department of Computer Science
University of Victoria
Victoria, British Columbia V8W 3P6, Canada

John C. Stillwell
Department of Mathematics
Monash University
Clayton, Victoria 3168, Australia

Abstract

Some aspects of the theory and computational complexity of covering projections of finite complexes are considered, from both the combinatorial and topological perspectives. The relationship between these two perspectives is explored. It is shown that there are 1-complexes (graphs) Y for which the computational decision problem which takes as input a finite 1-complex X and determines if X covers Y is NP-complete for both simplicial (combinatorial) and topological covering projections. A theorem of Leighton concerning finite common covers of 1-complexes, which holds both combinatorially and topologically, is shown to fail topologically for 2-complexes. Some results concerning 1-complexes which are mutual covers are also presented. The discussion is intended to be accessible to both combinatorialists and topologists.

1. Introduction.

There are two principal motivations for the present study. The first of these is the following elegant theorem of Leighton [Le].

Theorem. If G and H are finite graphs having the same universal cover, then there is a finite graph K that covers both G and H.

The theorem has an interesting history. It was first conjectured in the context of a study of complexity classes of distributed algorithms by Angluin [An]. A proof of the theorem was first offered from the combinatorial perspective in [Le] (and later, independently [Mo]). Combinatorial covering projections for graphs are equivalent to the notion of a *simplicial covering projection* as it is defined below for complexes of arbitrary dimension. Leighton's theorem holds for both simplicial and topological covering projections (the relationship between these is studied in the next section of this paper).

It is natural to wonder: (1) whether Leighton's theorem generalizes to higher dimensions, and (2) whether there might be an efficient algebraic proof of Leighton's theorem. We offer some limited results on these questions, showing that there is no higher dimensional generalization of Leighton's theorem in the topological category (we conjecture that there *is* such a generalization in the simplicial category). We also give a relatively short proof based on quasigroups (or Latin squares) of Leighton's theorem for regular graphs.

A second motivation for our work is provided by the reflection that covering projections for 1-complexes (graphs) are a kind of *local isomorphism*. The computational complexity of graph *isomorphism* appears to occupy a position of interesting intermediate difficulty between P and NP (see [GJ] for a discussion), assuming $P \neq NP$. What then of the complexity of local isomorphism? We show that it is apparently more difficult, exhibiting a fixed graph H for which it is NP-complete to determine if an input graph G covers H.

Generalizing the notion of common covering in a different direction, we consider pairs of spaces which are *mutual covers*. Finite graphs that are (combinatorially) mutual covers are necessarily isomorphic. We show that this statement is false for infinite 1-complexes. A number of interesting questions left open by our limited results are discussed in the concluding section.

2. Preliminaries.

A graph may be viewed as the combinatorial data for a 1-dimensional CW-complex [Ma]. One may think of this quite simply as the topological space corresponding to a "string model" of the graph. For example, the two graphs depicted in figure 1 (in the latter part of this section) are nonisomorphic combinatorially as graphs (one has order two and the other order ten), yet their string models are topologically equivalent. More precisely, the string model of a graph is formed by identifying, according to the incidence data, the endpoints of copies of $[0,1]$. The topology (the open sets) may be taken to be given by the natural distance metric on the model.

If G is a graph then we write $space(G)$ to denote the 1-CW-complex of G. If $space(G)$ and $space(H)$ are homeomorphic, we write $space(G) \approx space(H)$. Two spaces have the same *topological type* if they are homeomorphic. A space has *finite topological type* if it is homeomorphic to a finite CW-complex. It is evident that vertices of degree two make essentially no difference topologically. Thus for a graph G, $space(G)$ is canonically represented by the graph obtained from G by "suppressing" all irrelevant vertices of degree two. We make this precise in the following definition.

Definition. A vertex v of degree two in a graph G is *irrelevant* if there are two distinct edges of G incident on v. (In other words, v is not the only vertex of a component that is just v and a loop.) If v is an irrelevant vertex of a graph G with incident edges e and f, where the endpoints of e are u, v and the endpoints of f are v, w, then the graph G' obtained from G by *suppressing* v has vertex set $V(G') = V(G) - \{v\}$ and edge set $E(G') = E(G) - \{e, f\} \cup \{g\}$ where g is a new edge incident on u and w.

It should be clear that for each graph G there is a unique (up to isomorphism) *reduced graph*, $red(G)$ that canonically represents $space(G)$ in the sense that $red(G)$ has no irrelevant vertices and $space(G) \approx space(red(G))$. The following definition is standard (see, for example, [Ma]).

Definition. For topological spaces $\tilde{X}$ and X, a continuous map $p : \tilde{X} \to X$ is a *covering projection* if for each $x \in X$ there is an arcwise connected open neighborhood O of x such that each arcwise connected component of $p^{-1}(O)$ is mapped homeomorphically onto O by p. We say that $\tilde{X}$ *covers* X if there is a covering projection $p : \tilde{X} \to X$.

Definition. Topological spaces X and Y are *mutual covers* if X covers Y and Y covers X. A *common cover* of X and Y is a space Z that covers both X and Y.

A graph G is a *topological cover* of a graph H if there is a covering projection from $space(G)$ to $space(H)$. For the spaces of graphs it is generally easier to work combinatorially. The following lemma is easily verified.

Lemma 1. If G, H are connected graphs with no irrelevant vertices then G is a topological cover of H if and only if there is a surjective coloring from the vertices of G to the vertices of H and from the edges of G to the edges of H such that the incidence relation of edge and vertex colors in G is exactly as in H. That is, there are surjective maps $p_1 : V(G) \to V(H)$ and $p_2 : E(G) \to E(H)$ which satisfy

1. If $e \in E(G)$ has endpoints u and v (not necessarily distinct) then $p_2(e)$ has endpoints $p_1(u)$ and $p_1(v)$.
2. If $p_1(u) = x$ and $f \in E(H)$ has endpoints x, y with $x \neq y$, then there is a unique edge $e \in E(G)$ such that $p_2(e) = f$ and the endpoints of e are u, v with $p_1(v) = y$.
3. If $p_1(v) = x$ and f is a loop at x in H then either:

(*i*) there is a unique loop e at v in G with $p_2(e) = f$ and there are no edges g incident once on v with $p_2(g) = f$, or

(*ii*) there are exactly two edges e_1 and e_2 incident once each on v in G with $p_2(e_1) = p_2(e_2) = f$. □

A general combinatorial point of view is provided by the (standard) notion of a simplicial complex.

Definition. A *simplicial complex* (K, Σ) is a set K together with a family Σ of finite subsets (*simplices*) of K that satisfies: (1) $\emptyset \notin \Sigma$, (2) if $v \in K$ (v is a *vertex*) then $\{v\} \in \Sigma$, and (3) if $s \in \Sigma$ and $\emptyset \neq s' \subseteq s$ then $s' \in \Sigma$. The *dimension* of a simplex $s \in \Sigma$ is $dim(s) = |s| - 1$ and the *dimension* of (K, Σ) is $max\ dim(s)$.

Definition. A *simplicial map* $\phi : (K, \Sigma) \to (K', \Sigma')$ is a function $\phi : K \to K'$ such that if $\{v_o, \ldots, v_n\} = s \in \Sigma$ then $\phi(s) = \{\phi(v_0), \ldots, \phi(v_n)\} \in \Sigma'$.

A simplicial complex of dimension 1 is just a *simple graph* (a graph without loops or multiple edges) and a simplicial map between simple graphs is a graph homomorphism. A combinatorial notion of covering projection is described as follows.

Definition. A simplicial map ϕ from (K, Σ) to (K', Σ') is a *simplicial covering projection* if and only if: (1) $\phi : K \to K'$ is surjective, and (2) if $\phi(s) = s'$, $s \in \Sigma$, $s' \in \Sigma'$ and $\exists v' \in K'$ such that $s' \cup \{v'\} \in \Sigma'$ then there is a unique vertex $v \in K$ with $\phi(v) = v'$ and $s \cup \{v\} \in \Sigma$.

Associated to each simplicial complex (K, Σ) there is a topological space $|(K, \Sigma)|$, the *polyhedron* of (K, Σ). For a one-dimensional complex (simple graph) G the polyhedron of G is homeomorphic to *space* (G). If there is simplicial covering projection $\phi : (K, \Sigma) \to (K', \Sigma')$ then there is a topological covering projection from $|(K, \Sigma)|$ to $|(K', \Sigma')|$. The nice thing about dimension one is that, conversely, topological covering projections can be uniformly represented simplicially. This is not true for higher dimensions.

Definition. If G is a (general) graph then *simp*(G) is the simple graph obtained from *red*(G) by introducing two vertices of degree two into each edge and each loop of *red*(G). (See figure 1.)

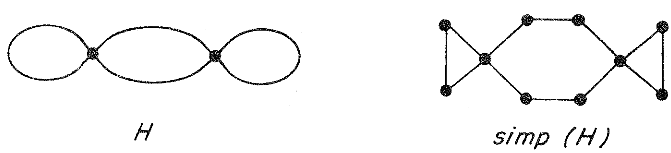

Figure 1.

Lemma 2. For graphs G, H there is a (topological) covering projection from *space*(G) to *space*(H) if and only if there is a simplicial covering projection from *simp*(G) to *simp*(H). $\square$

The above lemma, which the reader may routinely verify using Lemma 1, has the consequences: (1) Leighton's theorem, proved for the simplicial category, holds as well topologically in dimension one, and (2) the complexity result proved in the next section holds in both the simplicial and topological categories.

3. The complexity of covering in dimension one.

We consider the computational complexity of the following decision problem.

H-COVER
Instance: A graph G.
Question:] Is G a cover of H?

We may simplify our considerations by using Lemma 1. Given G, the canonical representation $red(G)$ of $space(G)$ can be easily computed, so we henceforth assume that G and H are reduced.

We note in passing that for some graphs H, H-COVER can be solved in polynomial time; for example, if H is the graph with one vertex and two loops or if H is the graph with two vertices and three edges between them.

The following theorem shows that for some graphs H, H-COVER may be computationally intractable.

Theorem 1. For the graph H of figure 1, H-COVER is NP-complete.

Proof. The problem is clearly in NP. To show that it is NP-hard we reduce from the problem NOT-ALL-EQUAL 3SAT [GJ].

An instance of this decision problem is a Boolean expression B in conjunctive normal form with exactly 3 literals of distinct variables per clause, and the question is whether there exists a truth assignment τ to the variables of B that makes at least one literal in each clause true and at least one literal in each clause false. We show how to construct, in polynomial time, a graph G_B such that G_B covers H if and only if B is not-all-equally satisfiable. A small example illustrating the main ideas of the of the construction is shown in figure 2. (In this example, the variable components have been made smaller – just big enough to supply the necessary connections to the clause components – than in the general construction of G_B which we next describe.)

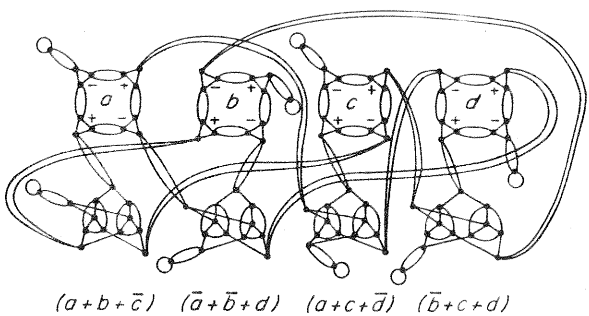

Figure 2.

Let $B = \prod_{i=0}^{m-1}(l_{i1} + l_{i2} + l_{i3})$ where l_{ij} is the j^{th} literal in the i^{th} clause and the set of variables of B is $\{x_v : 1 \le v \le n\}$. For $v = 1, \dots, n$ let $S_v^+ = \{i :$ the variable x_v occurs in the i^{th} clause of $B\}$ and let $S_v^- = \{i :$ the negation of x_v occurs in the i^{th} clause of $B\}$.

The set of vertices of the variable components is $V_1 = \{u(v,k,t) : 1 \le v \le n, 0 \le k \le 2m-1, 0 \le t \le 3\}$. The set of vertices of the clause components is $V_2 = \{w(i,j,r) : 0 \le i \le m-1, 0 \le j \le 3, 0 \le r \le 2\} \cup \{z(i) : 0 \le i \le m-1\}$.

The edge set of G_B is $E(G_B) = E_1 \cup E_2$ where E_1 is the set of edges of the variable components and E_2 is the set of edges of the clause components. These are described as follows.

E_1: For $1 \le v \le n$ and for $0 \le k \le 2m-1$ the set of vertices $\{u(v,k,0), u(v,k,1), u(v,k,2)\}$ induces a subgraph isomorphic to K_3, and there are two edges joining $u(v,k,2)$ to $u(v,k,3)$. Additionally, there are two edges joining $u(v,k,0)$ to $u(v,k+1 \pmod{2m},1)$. If $q \notin S_v^+$ then there is a loop on the vertex $u(v,2q \pmod{2m},3)$. If $q \notin S_v^-$ then there is a loop on the vertex $u(v,2q+1 \pmod{2m},3)$.

E_2: For $1 \le i \le m$ and for $r = 0,1$ the set of vertices $\{w(i,0,r), w(i,1,r), w(i,2,r), w(i,3,r)\}$ induces a subgraph isomorphic to K_4. For $1 \le i \le m$ and for $0 \le j \le 3$ the vertex $w(i,j,2)$ is joined by a single edge to each of the vertices $w(i,j,0)$ and $w(i,j,1)$. For $1 \le i \le m$ the vertex $z(i)$ is joined by two edges to the vertex $w(i,0,2)$ and there is a loop on the vertex $z(i)$.

The graph G_B is obtained from the disjoint union of the sets of variable and clause components described above by identifying certain pairs of vertices. (Note that this does not alter the edge set.) These vertex identifications are described:

1. For $1 \le v \le n$, if $q \in S_v^+$ and $l_{qh} = x_v$ then identify the vertex $u(v,2q \pmod{2m},3)$ with the vertex $w(q,h,2)$.

2. For $1 \le v \le n$, if $q \in S_v^-$ and $l_{qh} = \overline{x_v}$ then identify the vertex $u(v,2q+1 \pmod{2m},3)$ with the vertex $w(q,h,2)$.

Note that, counting incidence twice for each loop, there are four edges incident on each vertex of G_B and H. If the vertices and edges of G_B are labeled according to a covering projection from G_B to H then each vertex u of G_B must be adjacent twice to a vertex labeled x and twice to a vertex labeled y, where we consider that a loop makes a vertex adjacent twice to itself. An immediate consequence of this observation is the following.

Claim 1. If there is a covering projection from G_B to H then for each variable component, the vertices of each subgraph isomorphic to K_3 must all be colored the same.

The next claim establishes that "signals" are sent consistently from each variable component to the clause components for the clauses in which the variable occurs as a literal. The verification is left to the reader.

Claim 2. Let $T_0 \ldots, T_{2m-1}$ denote the sequence of subgraphs isomorphic to K_3 of a variable component. If the vertices of T_0 are all colored x according to a covering projection (as per Claim 1) then for each even index i all the vertices of T_i must be colored x, and for each odd index j all the vertices of T_j must be colored y.

Thus the sample lines (see figure 2) emerging from a variable component can be grouped into two equivalence classes according to the parity of the indices of the K_3's to which they are attached. Denote by S_0 and S_1 the two equivalence classes of sample lines for a given variable component.

If τ is a not-all-equally satisfying truth assignment then a coloring that is valid on the variable components can be described so that for each clause component, two of the three vertices $w(i,1,2), \ldots, w(i,3,2)$ are colored x and one is colored y (or *vice versa*). In this case $w(i,0,2)$ should be colored y and $z(i)$ should be colored x. Note that there is a symmetry of the clause component (ignoring $z(i)$) taking any two vertices of $\{w(i,0,2), \ldots, w(i,3,2)\}$ to any other two. This makes it easy to check that the coloring can be extended to all of the clause components.

Conversely, one easily checks that no coloring of $w(i,0,2), \ldots, w(i,3,2)$ that colors 3 vertices the same can be extended to a valid coloring of the entire clause component. Thus any valid coloring, by Claims 1 and 2, corresponds to a not-all-equally satisfying truth assignment (for example, the one that assigns the variable x_v the value "true" if and only if the vertices of the odd-indexed triangle of the corresponding variable component are colored with the vertex y of H). □

4. Nonisomorphic mutual covers.

It is easy to see that if G and H are simple finite graphs for which there are simplicial covering projections $p : G \to H$ and $q : H \to G$ then G and H must be isomorphic. By Lemmas 1 and 2 it follows that if G and H are finite graphs such that *space*(G) and *space*(H) are mutual covers then *space*$(G) \approx$ *space*(H). This implication fails for infinite graphs (and therefore, if you like, it fails for mutual simplicial covers).

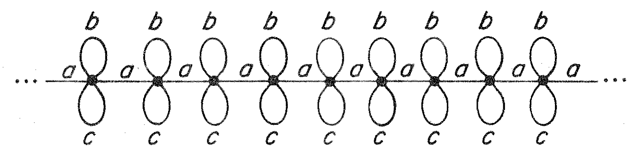

The graph S

Figure 3.

Let T denote the graph that has a single vertex and 3 loops denoted a, b, c. Let U be the universal cover of T. There are covering projections from U to S and from S to T where S is the string of loops shown in figure 3. The edges of S in figure 3 are labeled in accordance with the projection onto T. Similarly, there are covering projections from U to G, and from G to S where G is depicted in figure 4 and the projection respects the edge labeling.

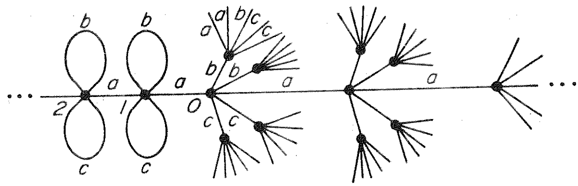

The graph G

Figure 4.

The graph G is a cover of itself by a projection $p : G \to G$ that maps vertex i to vertex $i+1$. The projection p can be factored as $p = p_2 \circ p_1$ where $p_1 : G \to H$ maps vertex i of G to vertex i of H and respects the edge labeling shown for H in figure 5.

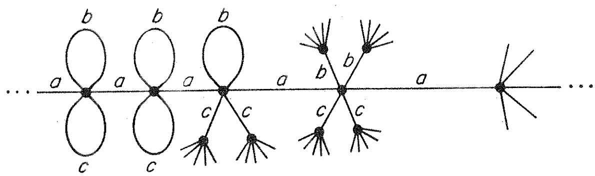

The graph H

Figure 5.

The projection $p_2 : H \to G$ sends vertex i of H to vertex $i+1$ of G. Thus G and H are mutual covers, but they are not isomorphic graphs.

5. Common covers of regular graphs.

The following pretty theorem was originally conjectured by Angluin in the context of a study of complexity classes of distributed algorithms [An]. The special case of regular graphs was proven by Angluin and Gardiner [AG] and the general case was first established by Leighton [Le] and later, independently, by Mohar [Mo].

Theorem 2. (Leighton) For finite graphs G and H, $space(G)$ and $space(H)$ have homeomorphic universal covers if and only if there is a finite graph K such that $space(K)$ is a common cover of $space(G)$ and $space(H)$.

We offer here a short proof for the special case of simple regular graphs by an argument employing Latin squares.

Proof. Assume that $G = (V_1, E_1) and H = (V_2, E_2)$ with $V_1 \cap V_2 = \emptyset$. We argue that there is a common finite simplicial cover $K = (V, E)$ defined as follows. For each pair of vertices (c, z) in $V_1 \times V_2$ choose

(1) a Latin square $L_{c,z}$ with entries in the set $\{1, \ldots \delta\}$ where δ is the degree of the graphs G and H, and

(2) bijections between the rows of $L_{c,z}$ and the set of neighbors of c in G, and between the columns of $L_{c,z}$ and the set of neighbors of z in H.

The vertex and edge sets of K are

$$V = \{(a, x, i) | a \in V_1, x \in V_2, i \in \{1, \ldots, \delta\}\}$$

$$E = \{(a, x, i)(b, y, j) | L_{b,y}(a, x) = j, L_{a,x}(b, y) = i, ab \in E_1, xy \in E_2\}$$

where $L_{b,y}(a, x)$ denotes the entry of $L_{b,y}$ in the row corresponding to a and in the column corresponding to x according to the bijections chosen in (2).

The projections are defined by $p : V \to V_1 : (a, x, i) \mapsto a$ and $q : V \to V_2 : (a, x, i) \mapsto x$. We verify that p is a simplicial covering projection (the verification for q is just the same). Suppose $p((a, x, i)) = a$ and $ab \in E_1$. We must argue, according to Lemma 2, that there is a unique vertex (a', x', i') adjacent to (a, x, i) in K with $p((a', x', i')) = b$. By the definition of p the only possibility for a' is $a' = b$, and x' must be adjacent to x in H. By the properties of a Latin square, there is a unique neighbor y of x such that $L_{a,x}(b, y) = i$. Given that $a' = b$ and $x' = y$, the value of i' is uniquely determined, $i' = L_{b,y}(a, x)$. □

Theorem 2 fails to generalize to spaces of 2-dimensional topological type as shown by the following.

Counterexample. Recall that if X is a cover of Y then the Euler characteristic $\chi(Y)$ of Y divides the Euler characteristic $\chi(X)$ of X. Let S_1 and S_2 denote, respectively, the orientable surfaces of genus 1 and genus 2. The universal cover of each is homeomorphic to the plane. Suppose S is a common finite cover. (That is, suppose $S \approx space(K)$ for a finite simplicial complex K). Since S must be a compact orientable surface with $\chi(S) = 0$, we must have $S \approx S_1$. But S_1 does not cover S_2. □

Note, however, that Theorem 2 may still be true in the simplicial category, since S_1 and S_2 do not have a common universal simplicial cover. We conjecture that Leighton's theorem generalizes in the simplicial category.

6. Open Problems.

There are a number of open problems connected with this work that we have so far been unable to settle. Among these is the problem of classifying the graphs H for which H-COVER can be solved in polynomial time. Are there graphs H for which the problem of determining if an input graph G is a *regular* cover of H is NP-hard?

Our construction showing that for 1-complexes mutual covers need not be isomorphic involves covering projections that are not finite-fold. Are mutual finite-fold covers isomorphic? Are there nonisomorphic mutual regular covers?

References

[An] D. Angluin, "Local and Global Properties in Networks of Processors," *in* Proceedings of the Twelfth ACM Symposium on the Theory of Computing, Los Angeles, 1980.

[AG] D. Angluin and A. Gardiner, "Finite Common Coverings of Pairs of Regular Graphs," *J. Comb. Th. B* 30 (1981), 184–187.

[GJ] M.R. Garey and D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W.H. Freeman, San Francisco, 1979.

[Le] F.T. Leighton, "Finite Common Coverings of Graphs," *J. Comb. Th. B* 33 (1982), 231–238.

[Ma] W.S. Massey, Algebraic Topology: An Introduction, Springer-Verlag, New York, 1967.

[Mo] B. Mohar, " A Common Cover of Graphs and 2-Cell Embeddings," *J. Comb. Th. B* 40 (1986), 94–106.