Figure 3.41. Two imbeddings of a graph.

Theorem 3.5.1. *Let $G \to S$ be an imbedding such that every boundary walk is a cycle, and such that the only boundary walks passing through both endpoints of an edge are the two boundary walks containing that edge. Then every adjacent imbedding has larger genus.*

Proof. Let e be any edge of the graph G, and let $G' \to S'$ be the imbedding obtained by deletion surgery on edge e. This imbedding has one less face than the original imbedding, since the two faces containing edge e have been joined into one. Moreover, no face of the imbedding $G' \to S'$ passes through both endpoints of edge e, except the joined face. Therefore, when the edge e is reattached in a different way, it must form a bridge joining two more faces. Hence, any adjacent imbedding has fewer faces and larger genus. □

Example 3.5.1 (Gross and Tucker, 1979b). *It is easy to verify that the toroidal imbedding on the left of Figure* 3.41 *satisfies the hypothesis of Theorem* 3.5.1. *Thus, every adjacent imbedding has larger genus. However, the given graph has the planar imbedding given on the right. It follows that the imbedding at the left is a local minimum that is not a global minimum.*

One might try to salvage an algorithm by defining a different kind of adjacency for graph imbeddings. One obvious candidate is " vertex-adjacency": two imbeddings are vertex-adjacent if their rotation systems agree at every vertex except one. However, a regular graph of valence 3, such as that in Example 3.5.1, has only two rotations at any vertex, so that vertex-adjacent imbeddings of these graphs would also be adjacent, presenting us with the same problem of local genus minima. In general, algorithms based on rotation systems have not been sufficiently fast.

3.5.2. Genus of an Amalgamation

Suppose that H and K are graphs imbeddable in surfaces S_h and S_k, respectively, and that G is a graph obtained by amalgamating H and K at a vertex. Then there is no difficulty in constructing an imbedding of G into a surface of genus $h + k$. First, one chooses a disk on S_h, whose interior is disjoint from the imbedded graph H and whose boundary intersects H only at

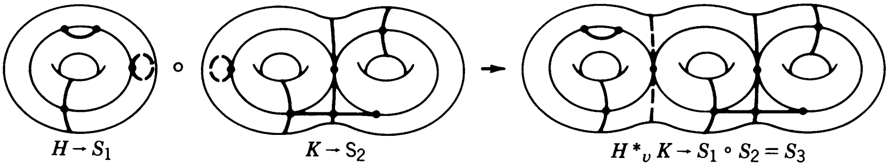

Figure 3.42. Any amalgamation of two graphs is imbeddable into a disk sum of any two of their respective imbedding surfaces.

the amalgamation vertex. Next, one chooses an analogous disk on the surface S_k. Then, one deletes the interiors of both disks and identifies the resulting surface boundaries so that the graphs H and K are amalgamated, as illustrated in Figure 3.42. As we have mentioned in Exercise 3.3.11, this composition of surfaces is called a "disk sum".

For this simplified type of graph amalgamation, we often use the notation $H *_v K$, where v is the vertex of amalgamation, to represent the resulting graph. From the above discussion, it is obvious that

$$\gamma(H *_v K) \leq \gamma(H) + \gamma(K)$$

The so-called *BHKY* theorem (for Battle, Harary, Kodama, and Youngs) asserts the nonobvious fact that

$$\gamma(H *_v K) = \gamma(H) + \gamma(K)$$

By way of contrast, the analogous assertion, is not true for maximum genus (see Exercise 5), and it is not true for crosscap number (see Theorem 3.5.5).

The straightforward way to try to prove the *BHKY* theorem is to attempt to show that an amalgamated graph $H *_v K$ must have a minimum imbedding that permits reversal of the combined disk sum and graph amalgamation operation. That is, one hopes for an imbedded circle in the surface that intersects $H *_v K$ only at the amalgamation vertex v and that separates the imbedding surface into two parts, one of which contains H and the other of which contains K. Necessarily, all the H-edges would have to be grouped contiguously in the rotation at v (as would the K-edges). If one started with a minimum imbedding of $H *_v K$ in which the H-edges were not continuously grouped, then one would wish to regroup, without increasing the genus of the imbedding surface. The following lemma, whose application is not restricted to imbeddings of amalgamated graphs, is what facilitates regrouping.

Lemma 3.5.2. *Let*

$$v. \quad ab \cdots cd$$

be the rotation at a vertex v of an orientable imbedding $G \to S$ such that the corners ab and cd lie on the same face. Let $G \to S'$ be the adjacent imbedding *such that the edge a is inserted between edges c and d, so that vertex v has the resulting rotation*

$$v.\quad b \ldots c a d$$

and all other rotations remain the same. Then $\gamma(S') \le \gamma(S)$.

Proof. Let $G \to S''$ be the imbedding obtained from $G \to S$ by deletion surgery on edge a. We shall use F, F', and F'' to denote the face sets of the respective imbeddings $G \to S$, $G \to S'$, and $G \to S''$. If the edge a appears twice on the same face boundary of the imbedding $G \to S$, then $\#F'' = \#F + 1$, from which it follows that $\#F' \ge \#F'' - 1 = \#F$. Otherwise, the deletion surgery on edge a in the imbedding $G \to S$ unites the face having the corners ab and cd with another face, and the resulting "new" face in $G \to S''$ contains the corner cd plus another occurrence of v across the "new" face to the corner cd, thereby resplitting the "new" face. Thus, $\#F' = \#F$. □

Theorem 3.5.3 [The *BHKY* Theorem (Battle, Harary, Kodama, and Youngs, 1963)]. $\gamma(H *_v K) = \gamma(H) + \gamma(K)$.

Proof. In any imbedding of $H *_v K$, the rotation at vertex v contains sequences of edges from the graph H interspersed with sequences from K. Maximal sequences from H and K are called "H-segments" and "K-segments", respectively, relative to that imbedding. Let $H *_v K \to S$ be a minimum-genus imbedding that maximizes the sum of the squares of the segment lengths at v, taken over all imbeddings of $H *_v K$ into the surface S. We assert that for this imbedding, there is but one H-segment and one K-segment.

In order to verify the assertion, we observe that in any imbedding of $H *_v K$, there exists an edge d from H and an edge e from K such that d immediately precedes e in the rotation at v. Moreover, whatever face boundary contains the corner de must also contain a corner $e'd'$ where e' is from K and d' from H. Since edges of H and K meet only at v, it follows that the rotation v has the form

$$v.\quad d e \ldots e' d' \ldots$$

We claim that the edges d and d' must lie in the same H-segment. If not, then suppose that the lengths of the H-segments containing d and d' are n and n', respectively. Without loss of generality, we shall assume that $n \le n'$. Apply Lemma 3.5.2, with edges d, e, e', and d' filling the roles of a, b, c, and d, respectively. Evidently, the adjacent imbedding with the modified rotation

$$v.\quad e \ldots e' d d'$$

(and no other changes) would have the same imbedding surface—lower genus is impossible, since S is a minimum-genus imbedding surface—but a larger sum of squares of segment lengths, since

$$(n-1)^2 + (n'+1)^2 > n^2 + (n')^2$$

because $n \leq n'$. Thus, the edges d and d' are in the same H-segment. Similarly, the edges e and e' are in the same K-segment. Therefore, we have established an assertion that there is only one H-segment and only one K-segment.

We must still prove that a closed arc C through vertex v from corner de to corner $e'd'$ in the interior of their common face must separate the imbedding surface. However, if not, then closed arc C would lie on a handle of surface S. If the surface S were cut open along C and reclosed with two disks, the result would be an imbedding of the disjoint union of the graphs H and K on a surface S' of genus $\gamma(S) - 1$. Such an imbedding would be noncellular, since any region that contained both a closed H-walk and a closed K-walk in its boundary would fail to be simply connected. By the construction of S', no such region could arise. It follows that the closed arc C separates the surface S. Since graph H lies on one side of the separation, that side must have genus at least $\gamma(H)$. Similarly, the other side must have genus at least $\gamma(K)$. It follows that

$$\gamma(H *_v K) \geq \gamma(H) + \gamma(K)$$

Since we have previously established the opposite direction of this inequality, the $BHKY$ theorem is proved. □

A maximal 2-connected subgraph of a graph G is called a "block" of the graph G. For example, the graph in Figure 3.43 has four blocks. A slight generalization of Theorem 3.5.3 is the following.

Theorem 3.5.4. *If the blocks of graph G are $G_1, \ldots, G_n$, then*

$$\gamma(G) = \gamma(G_1) + \cdots + \gamma(G_n).$$

Proof. One uses a simple induction on the number of vertices of the graph G, while applying the fact that a subgraph H is a block graph G if and only if $G = H *_v K$ for some vertex v and some subgraph K (of course, K itself may have several blocks). □

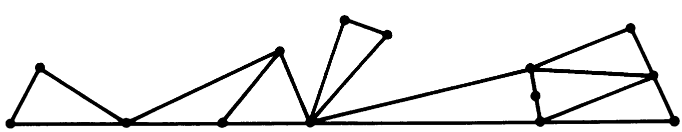

Figure 3.43. A graph having four blocks.