

THREE-DIMENSIONAL MANIFOLDS AND THEIR HEEGAARD DIAGRAMS*

BY
JAMES SINGER

INTRODUCTION

One of the outstanding problems in topology today is the classification of n -dimensional manifolds, $n \geq 3$. Poincaré, the founder of modern analysis situs, devoted several papers to it and allied problems. † Heegaard ‡, in a paper concerned primarily with another aspect of the subject, found it convenient to construct a pseudo-normal form for a 3-dimensional manifold, a form which we now call the *Heegaard diagram*. Dehn § and Veblen || gave modifications of his construction.

The Heegaard diagram of a 3-dimensional manifold consists of a closed 2-dimensional manifold upon which are drawn a certain number of non-intersecting simple closed curves. Any diagram is an adequate representation of a 3-dimensional manifold in the sense that it completely determines such a manifold, but, unfortunately, a 3-dimensional manifold gives rise to an infinity of diagrams. The problem of classifying manifolds is thus transferred to the problem of classifying diagrams.

Heegaard, in the paper cited above, studied (although not completely) the modifications that can be made on the curves and surface of a diagram which transformed it into another diagram but yet did not change the manifold which it represented. In this paper we extend Heegaard's results and study more completely the relationships between manifolds and their diagrams.

We begin then (Part I) by introducing the notions of a *canonical region*, *canonical surface* and *canonical curve* of a manifold. The *Heegaard diagram* is then constructed from a canonical surface and curves. Then, before proceeding to a discussion of manifolds and their diagrams, we show (Part II) how to read off the usual invariants of a manifold from any one of its representative diagrams.

* Presented to the Society, October 29, 1932; received by the editors June 10, 1932.

† Poincaré, H. For references see his paper in the Rendiconti del Circolo Matematico di Palermo, vol. 18 (1904).

‡ Heegaard, P. A translation into French is given in the Bulletin de la Société Mathématique de France, vol. 44 (1916).

§ Dehn, M., Mathematische Annalen, vol. 69 (1910).

|| Veblen, O., The Cambridge Colloquium, *Analysis Situs*, 2d edition, p. 155.

We then define (Part III) a set of *moves* which operate on the curves and surface of a diagram and transform it into another. Two diagrams are called *equivalent* if one can be obtained from the other by a finite number of these moves. We then prove by a sequence of theorems (Part IV) in which we make use of a specially constructed canonical surface and diagram that any two representative diagrams of a manifold are equivalent. Several other theorems (Part V) lead up to the general theorem of equivalence to the effect that equivalent manifolds (equivalent in the sense of semi-linear analysis situs) arise from and give rise to equivalent Heegaard diagrams (equivalent in our sense) and vice versa.

I wish to thank Professor J. W. Alexander for his many suggestions in the preparation of this paper.

I. PRELIMINARY DEFINITIONS

1. We will assume that the reader is familiar with the simplex, complex, manifold, incidence matrix, etc., as used in combinatorial analysis situs as, for example, in the book by Veblen, loc. cit. Throughout this paper, simplexes, cells, etc., will be at most 3-dimensional.

2. We derive here several elementary properties of the simplex which we will need later. Let $\bar{\sigma}_3 = A^0 A^1 A^2 A^3$ be a closed 3-simplex, σ_k, σ_{2-k} ($k=0$ or 1) two opposite faces of σ_3 (e.g., $\sigma_k = A^0, \sigma_{2-k} = A^1 A^2 A^3$); let $\bar{\sigma}'_3$ be the derived complex (i.e., a regular subdivision) of $\bar{\sigma}_3$. Every 3-simplex of $\bar{\sigma}'_3$ contains a vertex of σ_k or of σ_{2-k} , but not of both. It follows at once that all the 3-simplexes of $\bar{\sigma}'_3$ fall into two groups, R_1 and R_2 , where R_1 contains all the simplexes incident with σ_k and R_2 contains all those incident with σ_{2-k} . Moreover, R_1 and R_2 have as boundaries $B_1 + B$ and $B_2 + B$, respectively, where B_1 and B_2 consist of 2-simplexes of $\bar{\sigma}'_3$ on the boundary of $\bar{\sigma}_3$, and the common portion B is a 2-cell whose boundary is a circuit on the boundary of $\bar{\sigma}_3$ which "separates" σ_k and σ_{2-k} .

3. **The canonical region, cell and curve.** Given a 3-cell with its boundary sphere S_2 , we may decompose S_2 into a sum $E'_2 + E''_2 + \bar{B}$ where E'_2 and E''_2 are two cells such that \bar{E}'_2 and \bar{E}''_2 do not meet and B is a spherical band.* We can think of this system as a handle with the E 's as its bases.

We will now take a euclidean 3-sphere and attach p handles to it thus obtaining a 3-dimensional region R which we call *canonical*. Explicitly, we have $R = E^p_3 + \sum_{i=1}^p (E^i_3 + E^{i'}_2 + E^{i''}_2)$ where E^p_3 is the spherical region and the elements in the sum represent the handles and their bases. The (point set)-boundary of R will be designated by L . † Incidentally, L is not to be construed

* The bar over a symbol for a set denotes the closure of the set.

† We shall hereafter omit the words *point set*, understanding that whenever we speak of the boundary of a canonical region, we mean the (point set)-boundary.

to lie in euclidean 3-space, for it may very well happen that L is non-orientable.

We recognize in R the following property: if we sever each handle by a cross cut in the form of a 2-cell E_2^i , what is left is a 3-cell. In other words we can think of a canonical region as a region R within which there are p 2-cells E_2^i , with boundaries e^i on the boundary L of R (no two \bar{E} 's intersecting) such that $R - \sum E_2^i$ is a 3-cell. The cells E_2^i will be called *canonical cells* and their boundaries *canonical curves*.

4. **The canonical surface and the Heegaard diagram.** A surface L is said to be a *canonical surface* of a manifold M if it satisfies these conditions:

(a) L is a subcomplex of M and is a closed, connected 2-dimensional manifold;

(b) M can be decomposed into $R_1 + L + R_2$, where R_1 and R_2 are canonical regions with the common boundary L .

We note four properties of the canonical surface and regions which follow directly from their definitions:

A. If M' is a subdivision of M , and L' , R_1' and R_2' the induced subdivisions of L , R_1 and R_2 , then L' is a canonical surface of M' dividing it into the canonical regions R_1' and R_2' . We will also call L' a canonical surface of M .

B. The number of 2-cells that must be removed from R_1 to reduce it to a 3-cell is the same as the number that must be removed from R_2 , and each is precisely the maximum number of non-intersecting circuits that can be drawn on L without disconnecting it (Heegaard).

C. R_1 and R_2 are homeomorphic, since they have the same boundary.

D. L and M are both orientable or non-orientable (Heegaard).

5. A problem as yet unsolved is the determination of the minimum genus (or connectivity) of all the canonical surfaces of a given manifold M . This number is clearly a topological invariant. A simpler question is to ask: under what conditions can the genus of a canonical surface of a manifold be lowered or raised? Lemmas 1 and 2 below do not state the most general circumstances but they are sufficient for our needs.

6. In this paragraph we shall use the following notation: Let R be a canonical region, L its boundary; E_3 , E_3' and \mathcal{E}_3 3-cells with boundary spheres S_2 , S_2' and \mathcal{S}_2 . Let \mathcal{S}_2 be separated by a circuit into the two 2-cells \mathcal{E}_2 and \mathcal{E}_2' , and let $E_3[E_3']$ be separated by two non-intersecting circuits into two 2-cells E_2^1, E_2^2 [$E_2'^1, E_2'^2$] and a band B [B']. Finally, let B [B'] be separated into two 2-cells E_2^{*1}, E_2^{*2} [$E_2'^{*1}, E_2'^{*2}$] by two arcs, each running from one of its bounding circuits to the other.

LEMMA 1a. *If $\bar{\mathcal{E}}_3$ and $R+L$ have in common only $\bar{\mathcal{E}}_2$ on their boundaries, then $R + \mathcal{E}_3 + \mathcal{E}_2$ is a canonical region whose boundary is $L + \mathcal{E}_2' - \mathcal{E}_2$.*

LEMMA 2a. *If E_3 is a subcomplex of R and if S_2 and L have \bar{E}'_2 in common, then $R - E_3 - E_2$ is a canonical region whose boundary is $L + E_2 - E'_2$.*

These lemmas need no proofs. Obviously, the genus of L is neither raised nor lowered by the operations of the lemmas.

We can, however, change the genus of L by removing or adding a handle, E_3 .

LEMMA 1b. *If E_3 is a subcomplex of R and if S_2 and L have in common \bar{B} , then $R - (E_3 + E_2^1 + E_2^2)$ is a canonical region whose boundary is $L + (E_2^1 + E_2^2 - B)$.*

LEMMA 2b. *If \bar{E}_3 and $R + L$ have in common only \bar{E}_2^1 and \bar{E}_2^2 on their boundaries, then $R + (E_3 + E_2^1 + E_2^2)$ is a canonical region whose boundary is $L + (B - E_2^1 - E_2^2)$.*

These lemmas, too, need no proof. In the first case, the genus of L is lowered by the removal of a handle, in the second, raised by the addition of a handle.

We can lower the genus of L by attaching to it a 3-cell E'_3 along a band and we can raise the genus by removing such a 3-cell, i.e. by "boring" a hole through the region. However, we can add or remove a 3-cell only under certain conditions which are stated in the lemmas below.

LEMMA 1c. *Let E'_3 and $R + L$ have no points in common and S'_2 and L have \bar{B}' in common, and let E_3 be as in Lemma 1b; then, if S_2 and S'_2 have only $\bar{E}_2^{*1} \equiv \bar{E}_2'^{*1}$ in common, $R + E'_3 + B'$ is a canonical region whose boundary is $L + (E_2^1 + E_2^2 - B')$.*

For since, by Lemma 1b, $R - (E_3 + E_2^1 + E_2^2)$ is a canonical region whose boundary has the closure of the 2-cell $E_2^1 + E_2^2 + E_2^{*2}$ in common with the boundary of the 3-cell $E_3 + E'_3 + E_2^{*1}$, it follows from Lemma 1a that

$$[R - (E_3 + E_2^1 + E_2^2)] + [E_3 + E'_3 + E_2^{*1}] + [E_2^1 + E_2^2 + E_2'^{*2}] \equiv R + E'_3 + B'$$

is a canonical region. It is clear that the boundary of $R + E'_3 + B'$ is $L + E_2^1 + E_2^2 - B'$.

LEMMA 2c. *If E'_3 is a subcomplex of R such that S'_2 and L have $\bar{E}_2'^{*1}$ and $\bar{E}_2'^{*2}$ in common, and if there exists a 3-cell E_3 , subcomplex of R , such that $E_2^{*1} \equiv E_2'^{*1}$ and E_2^{*2} is a 2-cell on L , then $R - E'_3 - B'$ is a canonical region whose boundary is $L - E_2'^{*1} - E_2'^{*2} + B'$.*

For, since $E_3 + E'_3 + E_2'^{*1}$ is a 3-cell, it follows from Lemma 2a that $R - (E_3 + E'_3 + E_2'^{*1}) - (E_2^1 + E_2^2 + E_2'^{*2})$ is a canonical region. Hence, by

Lemma 2b, $R - (E_3 + E'_3 + E_2'^{*1}) - (E_2^1 + E_2^2 + E_2'^{*2}) + E_3 + E_2^1 + E_2^2 \equiv R - E'_3 - B'$ is a canonical region, whose boundary, as can be readily seen, is $L - E_2^1 - E_2^2 + B'$.

7. Let the canonical surface L of a manifold M divide it into the two regions R_1 and R_2 ; let E and F be canonical 2-cells of R_1 and R_2 , respectively, e and f , their boundaries. Let E_3 be a 3-cell as in Lemma 1b, where we put $E \equiv E_2^1$ and $R = R_1$. Let A be a 1-cell interior to R_1 with end points on L . Let E'_3 be a 3-cell as in Lemma 2c which is a neighborhood of A in R_1 . The 2-cells E_2^1 and E_2^2 will be neighborhoods on L of the end points of A . We now have

LEMMA 1. *If the canonical curves e and f meet once and only once, then $L + (E_2^1 + E_2^2 - B)$ is a canonical surface dividing M into the two canonical regions $R_1 - (E_3 + E_2^1 + E_2^2)$ and $R_2 + (E_3 + B)$.*

LEMMA 2. *If there exists a 1-cell A' on L such that $A + A'$ bounds a 2-cell of R_1 , then $L - (E_2^1 + E_2^2 - B')$ is a canonical surface dividing M into the canonical regions $R_1 - (E'_3 + B')$ and $R_2 + (E'_3 + E_2^1 + E_2^2)$.*

The proofs of the two lemmas follow at once from Lemmas 1 abc, 2 abc. We note that the effect of the first lemma is to remove a handle from R_1 (as in Lemma 1b) and to add a 3-cell to R_2 (as in Lemma 1c) and the effect of the second lemma is to remove a 3-cell from R_1 (as in Lemma 2c) and to add a handle to R_2 (as in Lemma 2b). The genus of the canonical surface is lowered in the first lemma, raised in the second.

8. **Construction of the Heegaard diagram.** The utility of the canonical surface and curves lies in the fact that they give us an adequate representation of the manifold. Indeed, let L be a canonical surface of a manifold M , and let $e^1, \dots, e^p, f^1, \dots, f^p$ be two sets of canonical curves, boundaries of canonical sets of 2-cells of R_1 and R_2 , respectively. Then, having $L, e^1, \dots, e^p, f^1, \dots, f^p$, we can dispense with the rest of M entirely, for any information that can be derived from M can be derived from them. To reconstruct a 3-dimensional manifold, we attach 2-cells to each of the canonical curves, and then add two 3-cells in the obvious way. We thus obtain a manifold N which in general will not be identical to M , but which must always be homeomorphic to it because of the construction.

Of great use in the study of manifolds is the fact that a model of the canonical surface and curves of any manifold can be constructed in ordinary spherical 3-space, S_3 . If the canonical surface were orientable, we could immerse it in S_3 immediately; this is impossible if it is non-orientable. To treat both cases at the same time, we adopt one of the normal forms for a 2-dimensional manifold which can be immersed in S_3 , i.e., the plane (plus a

point at infinity) from which the interiors of $2p$ circles have been removed.* We call the surface Λ , its $2p$ -bounding circles, $\epsilon_1^i, \epsilon_2^i, i=1, \dots, p$. For the sake of definiteness later on, let us assume that the circles are of equal radii with centers equally spaced along a straight line.

We now establish a continuous correspondence between the points of L and Λ which is $(1, 1)$ everywhere except that a point on e^i has for image a point on ϵ_1^i and a point on ϵ_2^i . Thus, the image of an e curve will be a pair of the ϵ circles on Λ . If a particular f^i does not meet any of the e curves, then its image on Λ will be a circuit; if it does meet the e 's, then its image on Λ will consist of a set of arcs, each joining two points on the circles. We call the circuit or aggregate of arcs corresponding to f^i, ϕ^i .

From what has been said above, it is clear that $\Lambda, \epsilon_1^i, \epsilon_2^i$, and $\phi^i (i=1, \dots, p)$ also serve as an adequate representation of the manifold M . We call this representation a *Heegaard diagram* of M .

For some purposes it is convenient to have a Heegaard diagram in which the f 's are mapped on the pairs of circles, and the e 's become aggregates of arcs. In such cases, we shall use another plane Λ' , and introduce notation as needed.

To reconstruct a manifold N from a Heegaard diagram Δ , we first subdivide Λ , if necessary, and then construct a closed 2-dimensional manifold L equivalent to Λ where, however, a single circuit e^i corresponds to the pair of circles ϵ_1^i and ϵ_2^i . We must take care that ϵ_1^i and ϵ_2^i are matched upon e^i with the proper orientations. We then proceed as before, successively adding the 2-cells and finally the two 3-cells.

If a Heegaard diagram is constructed from a manifold, we shall say that the manifold *gives rise* to the diagram and that the diagram *arises from* the manifold; similarly we shall say that a diagram *gives rise* to a manifold and that the manifold *arises from* the diagram when the manifold is constructed from the diagram.

9. Since a Heegaard diagram is an adequate representation of its manifold, all invariants of the latter should be obtainable from the former. We show how to get the Poincaré group, homology characters, and some intersection invariants in the next section.

Since any one manifold can give rise to a great variety of diagrams, we do not seem to be any nearer the solution of the problem of the classification of manifolds. We reserve all such questions for Parts III, IV, and V; at the present we merely note that the form of a Heegaard diagram arising from a manifold is by no means unique.

* See Alexander, J. W., *Normal forms for one- and two-sided surfaces*, Annals of Mathematics, (2), vol. 16, No. 4, June, 1915.

II. THE KNOWN INVARIANTS

1. **The Poincaré group.** Let Δ be a Heegaard diagram of a manifold M , where Δ consists of Λ , ϵ_1^i , ϵ_2^i and ϕ^i , $i = 1, 2, \dots, p$, all defined as in Part I.

It is well known that there exist, on Λ , $2p$ circuits, a_i and b_i , $i = 1, \dots, p$, all passing through a fixed point O but having no other points in common and such that every other closed curve of Λ is deformable into a sum of the a 's and b 's. The a 's and b 's can then be taken as the generators of the Poincaré group of Λ . Moreover, we can always choose the curves in such a manner that a_k is isotopic to ϵ_1^k on Λ and b_k consists of two arcs joining O to congruent points of ϵ_1^k and ϵ_2^k . We can choose a similar base on Λ' (another representation of Δ , in which a pair of circles ϕ_1^i and ϕ_2^i represents a canonical curve f^i), namely c_i and d_i , $i = 1, \dots, p$, where c_k is isotopic to ϕ_1^k on Λ' and d_k consists of two arcs joining O' (image of O) to congruent points of ϕ_1^k and ϕ_2^k .

Since the surfaces Λ and Λ' are representations of the same surface L of M , we can express every curve of one base as a product of the generators of the other base, i.e., $c_k = a_i^{\epsilon_i} b_i^{\epsilon_i} a_i^{\epsilon_i} b_i^{\epsilon_i} \dots$, where i_j is one of the integers $1, 2, \dots, p$ and ϵ_j is 0, 1 or -1 . Symbolically we can write

$$(1) \quad c_i = \Pi_i^1 ab, \quad d_i = \Pi_i^2 ab,$$

$$(1') \quad a_i = \Pi_i^3 cd, \quad b_i = \Pi_i^4 cd.$$

The only identity relation among the generators is

$$(2) \quad a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1} \dots a_p b_p a_p^{-1} b_p^{-1} = 1$$

in case A is orientable, or

$$(2') \quad a_1 b_1 a_1 b_1^{-1} a_2 b_2 a_2 b_2^{-1} \dots a_p b_p a_p b_p^{-1} = 1$$

in case A is non-orientable.* There are, of course, similar expressions in terms of the c 's and d 's.

Let us now adjoin to Λ the pairs of congruent 2-cells E_1^i and E_2^i , interiors of the circles ϵ_1^i and ϵ_2^i , where E_1^i and E_2^i correspond to the same canonical 2-cell E^i of R_1 . To note what modification the group of Λ undergoes, we shall make use of a theorem concerning the Poincaré group of an arbitrary 2-dimensional complex K_2 : If the Poincaré group of K_2 is generated by the elements

$$g_1, g_2, \dots, g_p$$

with the identity relations

$$r_i = g_{i_1}^{\epsilon_1} g_{i_2}^{\epsilon_2} g_{i_3}^{\epsilon_3} \dots = 1 \quad (i = 1, 2, \dots, q)$$

* The identity relation in the non-orientable case is not one of the usual forms; it is, however, equivalent to them in the sense that we can so choose the generators of the Poincaré group that (2') holds.

where g_{ij}^k is either a generating element or its inverse and we adjoin to K_2 a 2-cell E having no points in common with K_2 but whose boundary is h , where h is a circuit of K_2 , then the group of $K_2 + E$ is the same as the group of K_2 plus the additional relation

$$r = g_{i_1}^{\epsilon_1} g_{i_2}^{\epsilon_2} g_{i_3}^{\epsilon_3} \cdots = 1,$$

where r represents a curve isotopic to h on K_2 .* In particular, if h is isotopic to a generator, say g_1 , of the group of K_2 , the group of $K_2 + E$ can be obtained by replacing g_1 by 1 wherever it occurs in the identity relations of the group of K_2 .

It follows at once that the group of $\Lambda + \sum_{i=1}^p (E_1^i + E_2^i)$ is generated by the p elements

$$(3) \quad b_1, b_2, \cdots, b_p$$

with no generating relations, since (2) or (2') reduces to the unit element identically. The products (1) and (1') take on new forms which we write symbolically as

$$(4) \quad c_i = \Pi_i'^1 b, \quad d_i = \Pi_i'^2 b,$$

$$(4') \quad 1 = \Pi_i'^3 c d, \quad b_i = \Pi_i'^4 c d.$$

We now add the p 2-cells F^i to $\Lambda + \sum (E_1^i + E_2^i) \equiv \Lambda + \sum E^i$ obtaining the 2-dimensional complex $L + \sum E^i + \sum F^i$. The products (4) and (4') take on the new forms

$$(5) \quad c_i = \Pi_i''^1 b, \quad d_i = \Pi_i''^2 b,$$

$$(5') \quad 1 = \Pi_i''^3 d, \quad b_i = \Pi_i''^4 d.$$

From what has just been said, it is clear that the group of $\Lambda + \sum E^i + \sum F^i$ is generated by the p elements (3) connected by the p relations $\Pi_i''^1 b = 1$ of (5). We obtain an equivalent Poincaré group if we take for basis the p elements

$$(3') \quad c_1, c_2, \cdots, c_p$$

connected by the p relations $\Pi_i''^3 d = 1$ of (5').

The addition of the remainder of the manifold M (that is, the two 3-cells) can have no effect on the Poincaré group since a curve homotopic to a point

* If r' is another product of the g 's also representing a curve isotopic to h on K_2 , then the group obtained by adding r' is equivalent to the group obtained by adding r , since $1 = r' = sr s^{-1}$, where s is a product of g 's.

and the homologies

$$(7) \quad \phi^i \sim \sum_j \alpha_j^i b^j \quad (\text{on } L),$$

where the matrix $\|\alpha_j^i\|$ is the same as the matrix of exponents in the identity relations of G^{M_0} . It is well known that there exist unimodular transformations $F^i = \sum_j \mu_j^i F'^j$, $\phi^i = \sum_j \mu_j^i \phi'^j$ and $b^i = \sum_j \nu_j^i b'^j$ such that the bounding relations (6) and the homologies (7) become

$$(8) \quad F'^i \rightarrow \phi'^i \quad (\text{on } M)$$

and

$$(9) \quad \phi'^i \sim \sum_j \alpha'_j{}^i b'^j \quad (\text{on } L)$$

respectively, where $\|\alpha'_j{}^i\|$ is a diagonal matrix. The first $\rho - \tau$ terms of the main diagonal of $\|\alpha'_j{}^i\|$ are equal to 1, the next τ terms are the invariant factors (1-dimensional coefficients of torsion) and all the other terms are zero.

Since every 1-cycle of M is homologous to a linear combination of the b 's, and hence to the b 's, and since the 1-dimensional Betti number of M is $p - \rho$, it follows that the curves

$$(10) \quad b'^i \quad (i = \rho + 1, \rho + 2, \dots, p)$$

form a basis for the non-bounding 1-cycles of M in this sense: if b is any 1-cycle of M , then

$$b \sim \sum_{i=\rho+1}^p a_i b'^i + \sum_{i=\rho-\tau+1}^{\rho} \bar{a}_i b'^i.$$

But since some multiple of b'^i ($\rho - \tau + 1 \leq i \leq \rho$) bounds, it will meet any 2-cycle of M zero times algebraically. Hence as far as the intersection of b with a 2-cycle is concerned, only the terms of the first sum in b need to be considered.

It can be shown that the 2-dimensional complexes

$$(11) \quad F'^i \quad (i = \rho + 1, \rho + 2, \dots, p)$$

can serve as a base for the non-bounding 2-cycles as far as intersections are concerned. To be sure, none of the F 's is a cycle, nor is any combination of them a cycle since every sum $\sum \lambda_i F'^i$ has for boundary $\sum \lambda_i \phi'^i$, but it can be shown that if the complexes F_0^i , $i = \rho + 1, \dots, p$, do form a base for the non-bounding 2-cycles of M in the same sense as the b 's just above, then the intersection of b^i and F_0^j on M is the same as the intersection of b^i and the

boundary of F^i on L . Hence the intersection matrix of the non-bounding 1- and 2-cycles of M is given by

$$(12) \quad (b'^i \cdot F'^j) \quad (i, j = \rho + 1, \dots, p)$$

where the element in the i th row and j th column is the Kronecker index of b'^i and ϕ'^j on Λ . In case M is orientable, the matrix (12) is unimodular and can be transformed into a diagonal matrix by further unimodular transformations of the b'' 's and F'' 's.

The 2-complexes (11) also yield the intersections of the non-bounding 2-cycles among themselves. Indeed, we can compute from them a cubic matrix

$$(13) \quad \|\sigma_{ijk}\| \quad (i, j = \rho + 1, \rho + 2, \dots, p; k = 1, 2, \dots, p),$$

where we can show that $\sum_k \sigma_{ijk} b^k$ is a cycle homologous to the intersection cycle of two non-bounding 2-cycles of M , as defined, for example, by Lefschetz in his Colloquium Lectures, Chapter IV. The σ 's are precisely the same as those defined by Alexander in Proceedings of the National Academy of Sciences, 1924, p. 99.

They are obtained in this fashion. If $i = j$, then $\sigma_{ijk} = 0$ for all k 's. If $i \neq j$, we proceed as follows.

We consider a particular pair of (11), say F'^1 and F'^2 , and a particular canonical circle, say ϵ_1^3 . Since F'^1 can be considered a two-cycle, its boundary on Λ , namely ϕ'^1 , meets ϵ_1^3 zero times algebraically, if we count each intersection with its proper orientation and multiplicity. Now imagine that the curve ϵ_1^3 is shrunk to its center A . The point A is met zero times algebraically by the segments of ϕ'^1 abutting on it.

Similarly, the point A will be met zero times, algebraically, by the segments of ϕ'^2 incident with it.

If no segment is common to ϕ'^1 and ϕ'^2 , we define the number σ_{123} as the Kronecker index of ϕ'^1 and ϕ'^2 at the point A , where we take into account the multiplicity of the various branches.

Suppose now that the segment a belongs to ϕ'^1 with the multiplicity m_1 and to ϕ'^2 with the multiplicity m_2 . Replace a by two distinct segments a^1 and a^2 lying very close to it, and let a^1 belong to ϕ'^1 with the multiplicity m_1 and a^2 to ϕ'^2 with the multiplicity m_2 . If we do this for all segments common to ϕ'^1 and ϕ'^2 we are led to the former case and can define σ_{123} as there. However, it is apparent that σ_{123} as so defined is not unique, for the Kronecker index will depend on the method of replacing a by a^1 and a^2 , that is, on whether we take a^1 to the right or left of a^2 . We therefore define as the σ_{123}

the number of smallest numerical value so obtained. As a matter of fact, the difference between any two values represents a cycle which is homologous to zero; we choose the smallest numerical value purely for convenience.

We can now define the σ 's for all i, j , and k in the same way. It is clear that we need not consider the canonical circles ϵ_2^i , since the intersections are the same as on ϵ_1^i .

III. DEFINITIONS OF THE MOVES AND EQUIVALENT HEEGAARD DIAGRAMS

1. Two Heegaard diagrams Δ and Δ' with canonical curves $\epsilon^1, \dots, \epsilon^p, \phi^1, \dots, \phi^p$ and $\epsilon'^1, \dots, \epsilon'^q, \phi'^1, \dots, \phi'^q$, respectively, are said to be *identical* if

(a) $p = q$, and

(b) Δ and Δ' are mapped on the same euclidean plane Λ in such a way that each pair of circles ϵ_1^k and ϵ_2^k coincides in position and orientation with a pair ϵ_1^i and ϵ_2^i and each ϕ^k is isotopic to a ϕ^i on Λ .

2. We now define a set of transformations, called *moves*, which operate on and modify the curves and surface of a Heegaard diagram. The moves will fall into three classes or types; the first will modify the canonical curves, the second will modify the plane Λ by rotating a portion of it through a multiple of π , the third will modify the canonical curves and the plane Λ by the addition (or subtraction) of canonical curves.

Type I. A. A move of Type IA simply changes the orientation of a given canonical curve.

Let g be a 1-sphere on Λ not meeting any of the circles (images of the canonical e curves) and containing in its interior at least one circle, say ϵ_1^k , and at most two such circles. In the latter case, the two circles must be adjacent but not the two circles of a pair, i.e., images of the same e . Sever Λ along g and call the inner and outer lips of the cut g and g' , respectively. Remove the interior of g from Λ and attach it once more, this time, however, by matching corresponding points of ϵ_1^k and ϵ_2^k . We have lost the canonical pair ϵ_1^k and ϵ_2^k and have gained a new pair of circuits, g and g' . The two cases give us the moves of Types IB and IC, i.e.,

B. g contains only one circle,

C. g contains two circles.

A move of Type IA, B or C is its own inverse. The effect of a move of Type IB is to replace a canonical curve by one isotopic to it. The effect of a move of Type IC is to replace a canonical curve by the "sum" of it and another canonical curve.

D. A move of Type ID is the replacement of the image on Λ of a canonical f curve, say ϕ^i , by the sum of the images of f^i and some f^j , i.e., by $\phi^i + \phi^j$.

The move is effected by deforming ϕ^i on Λ until it has a 1-simplex in common with ϕ^j . Then by removing this 1-simplex we obtain a circuit $\phi'^i = \phi^i + \phi^j$, which after a slight deformation will have no points in common with ϕ^j . The replacement of ϕ^i by ϕ'^i is the move of Type ID.

When a diagram has been modified by a move of Type IB or C, its circles will no longer be in standard form, i.e., of equal radii and equally spaced along a straight line. However, an isotopic deformation of Λ in S_3 will bring them into standard form. We shall always suppose this done when we operate on Λ by a move of Type IB or C.

Type II. Let g be a circuit on Λ having no points in common with the canonical set $\{e\}$ and let g' be its position after a small isotopic deformation such that g and g' have no points in common, i.e., they have the appearance of two concentric circles, where g' is interior to g , say. Let B be the band-shaped region bounded by g and g' . Now rotate g' and that part of Λ interior to it through a positive angle of π or 2π , keeping g and that part of Λ exterior to it fixed. We note particularly that if one of a pair of canonical circles lies within g' and the other without, only one of them is rotated. The result is a distortion in the band B ; however, if B is suitably subdivided into simplexes of fine enough mesh, its structure will remain unaltered.

However, we do not wish to employ this type of transformation in its most general form but only in two certain cases when the region interior to g' contains either one or two canonical curves.

A. Rotation through a positive angle π when the 1-sphere g' contains two consecutive circles.

B. Rotation through a positive angle 2π when the 1-sphere g' contains only a single circle.

Inverses of moves of Types IIA and B will be rotations through negative angles.

Type III. A move of this type adds a pair of canonical circles and a new canonical curve to the plane Λ . It is effected as follows: Let ϵ_1^{p+1} and ϵ_2^{p+1} be a pair of circuits on Λ each bounding a 2-cell, and let ϕ^{p+1} be an arc joining the point P_1 of ϵ_1^{p+1} and the point P_2 of ϵ_2^{p+1} . The arc ϕ^{p+1} must have no other points in common with ϵ_1^{p+1} or ϵ_2^{p+1} and no points in common with the two 2-cells bounded by ϵ_1^{p+1} and ϵ_2^{p+1} . Also ϵ_1^{p+1} , ϵ_2^{p+1} and ϕ^{p+1} must not meet any of the other canonical curves or arcs. Then if we remove the interiors of ϵ_1^{p+1} and ϵ_2^{p+1} from Λ and identify the points of ϵ_1^{p+1} and ϵ_2^{p+1} so that P_1 and P_2 are matched, we can add the pair ϵ_1^{p+1} and ϵ_2^{p+1} to our canonical circles and ϕ^{p+1} to our canonical arcs. Its inverse, the removal of a pair of canonical curves ϵ^{p+1} and ϕ^{p+1} under the conditions just described, will be denoted by III'.

The surface Λ seems to play a preferred rôle in the description and definition of these moves. However, all our moves are also applicable to the surface Λ' , where the ϕ curves are represented by pairs of circles. It is not necessary to note the effect on Λ' of one of the moves operating on Λ ; we only note that when we modify Λ by a move of Type IC, Λ' is modified by a move of Type ID, and vice versa.

3. Let now Δ be a Heegaard diagram with the two canonical sets of curves $\{\epsilon\}$ and $\{\phi\}$. We state six lemmas whose proofs follow immediately from the definitions of the moves.

LEMMA 3. *If g is a circuit on Λ not meeting any of the canonical circles and containing in its interior ϵ_1^k and any number of other circles but not ϵ_2^k , then the result of severing Λ along g and reattaching the piece so removed to Λ along ϵ_1^k and ϵ_2^k may be obtained by successively applying moves of Type IC.*

LEMMA 4. *If the set of curves $\{\epsilon'\}$ is derived from the canonical set $\{\epsilon\}$ by a finite number of moves of Type I, then the set $\{\epsilon'\}$ is canonical.*

LEMMA 5. *If the canonical set $\{\epsilon'\}$ is derived from the canonical set $\{\epsilon\}$ by a finite number of moves of Type I, then $\{\epsilon\}$ can be derived from $\{\epsilon'\}$ by the same number of moves.*

LEMMA 6. *If the surface Λ and the canonical sets of curves $\{\epsilon\}$ and $\{\phi\}$ of a Heegaard diagram Δ are transformed into a surface Λ' with the sets of curves $\{\epsilon'\}$ and $\{\phi'\}$ by a finite number of moves, then Λ' , $\{\epsilon'\}$ and $\{\phi'\}$ form a Heegaard diagram.*

LEMMA 7. *If the Heegaard diagram Δ' is derived from the Heegaard diagram Δ by a finite number of moves, then Δ can be derived from Δ' also by a finite number of moves.*

LEMMA 8. *If the Heegaard diagram Δ gives rise to the manifolds N and N' , then N and N' are homeomorphic.*

COROLLARY. *If the manifold M gives rise to the diagram Δ , which in turn gives rise to the manifold N , then M and N are homeomorphic.*

The proofs of this lemma and corollary follow at once from the construction of N and N' as outlined in §8, Part I.

4. We are now in a position to give a formal definition of equivalence: two Heegaard diagrams are said to be *equivalent* if one can be transformed into a diagram identical to the other by a finite number of moves of Types I, II and III.

5. We prove the following theorem:

THEOREM 1. *If the equivalent Heegaard diagrams Δ and Δ' give rise to the manifolds N and N' , then N and N' are homeomorphic.*

Since Δ and Δ' are equivalent, Δ can be obtained from Δ' by a finite sequence of moves of Types I, II, and III which transform Δ' successively into the diagrams $\Delta^1, \Delta^2, \dots, \Delta^m \equiv \Delta$. It is therefore necessary and sufficient to prove that at each stage the transformed diagram Δ^{k+1} gives rise to a manifold N^{k+1} homeomorphic to the one, N^k , arising from the preceding diagram Δ^k .

This is obviously true when we modify Δ^k by a move of Type IA, IB, or a move of Type II, hence we need only prove that the theorem holds when we employ a move of Type IC, ID, or a move of Type III.

Let, then, Δ^k be transformed into Δ^{k+1} by a move of Type IC and let Δ^k and Δ^{k+1} give rise to the manifolds N^k and N^{k+1} . We can suppose that the move replaces the canonical curve ϵ_k^1 of Δ^k by the canonical curve $\epsilon_{k+1}^1 = \epsilon_k^1 + \epsilon_k^2$. Let us retain the circuit ϵ_k^1 on the diagram Δ^{k+1} and slightly deform it so that it does not meet any of the canonical curves of Δ^{k+1} . We can then find in the manifold N^{k+1} a 2-cell E_k^1 lying in the region containing the canonical 2-cells E_{k+1}^1 and whose boundary is ϵ_k^1 . The 2-cell E_k^1 will divide this region, a 3-cell, into two parts. If we unite these two parts along E_{k+1}^1 and consider E_k^1 as part of the boundary, we have again a 3-cell. We have not changed N^{k+1} at all, but obviously from this point of view it may be considered a manifold arising from Δ^k ; hence by Lemma 8, N^{k+1} is homeomorphic to N^k .

A similar argument holds when Δ^k is transformed by a move of Type ID.

Suppose now that Δ^k is transformed into Δ^{k+1} by a move of Type III which adds the canonical curves ϵ^{p+1} and ϕ^{p+1} to the canonical curves of Δ^k . In the manifold N^{k+1} the curves ϵ^{p+1} and ϕ^{p+1} are such that Lemma 1 is applicable where the 2-cell E^{p+1} plays the rôle of E_2^1 of the lemma. If we make the necessary modifications, we transform the canonical surface and regions of N^{k+1} into new ones. But with these latter canonical surface and regions N^{k+1} may be considered as a manifold arising from Δ^k ; hence, once more N^k and N^{k+1} are homeomorphic and the theorem is proved.

6. If, then, the Heegaard diagrams Δ and Δ' are equivalent, i.e., if one can be transformed into the other by means of the moves, then any two manifolds N and N' that arise from them are homeomorphic. To give greater justification to the definition and notion of equivalence, we must prove conversely that if two manifolds M and M' are homeomorphic and they give rise to the Heegaard diagrams Δ and Δ' then the two diagrams are equivalent. In other words, we must prove that the moves are indeed sufficient to transform Δ into Δ' .

IV. HEEGAARD DIAGRAMS ARISING FROM A MANIFOLD

1. Variations in a Heegaard diagram may arise in

- (a) the choice of the canonical surface,
- (b) the choice of the canonical 2-cells of one of the canonical regions,
- (c) the choice of the canonical 2-cells of the other canonical region,
- (d) the method of immersion in S_3 , i.e., the method of mapping L on Λ .

In this section we study the relationship between any two diagrams arising from a manifold and prove (Theorems 2-8) that any two such diagrams are equivalent.

We shall use consistently the following notation: M, M', N , etc., shall denote a 3-dimensional manifold; L , a canonical surface, R_1 and R_2 the two canonical regions. Canonical sets of 2-cells of R_1 shall be denoted by $\{E\}$, $\{E'\}$, etc., of R_2 by $\{F\}$, $\{F'\}$, etc. The canonical sets of curves (boundaries of the canonical 2-cells) will be denoted by $\{e\}$, $\{f\}$, etc. We shall denote a Heegaard diagram by Δ, Δ' , etc., and we shall use the notation $\Delta \equiv (\Lambda, \epsilon, \phi)$ to signify that the Heegaard diagram Δ consists of the plane Λ (image of a canonical surface L), the pairs of circles, $\epsilon_1^i, \epsilon_2^i$ (images of the canonical e curves), and the arcs ϕ^i (images of the canonical f curves).

2. In Theorem 2 below we prove that two diagrams arising from the same canonical surface and curves are equivalent, i.e. two methods of immersing L , $\{e\}$ and $\{f\}$ in S_3 (see (d) of §1) yield equivalent diagrams.

THEOREM 2. *Let L be a canonical surface of a manifold M , $\{e\}$ and $\{f\}$ canonical sets of curves of the two regions; then if $L, \{e\}$ and $\{f\}$ give rise to the Heegaard diagrams $\Delta \equiv (\Lambda, \epsilon, \phi)$ and $\Delta' \equiv (\Lambda', \epsilon', \phi')$, Δ and Δ' are equivalent.*

In this theorem Λ and Λ' , ϵ^i and ϵ'^i , ϕ^i and ϕ'^i correspond to the same L , e^i and f^i , respectively, of M . Let us superimpose the two planes Λ and Λ' ; we can assume, without any loss of generality, that the canonical circles coincide. We choose our notation so that $\epsilon_1^k, \epsilon_2^k$ and $\epsilon_1'^k, \epsilon_2'^k$ (ϕ^k and ϕ'^k) are images of the same e^k (f^k) of L .

The circles $\epsilon_1^k, \epsilon_2^k$ and $\epsilon_1'^k, \epsilon_2'^k$ will not coincide, in general, but we shall show how to modify Λ' by means of our moves so that the two images of each e do coincide and the two images of each f are isotopic.

We regard the two planes Λ and Λ' as infinitely close but distinct, as, for example, two sheets of a Riemann surface. There will be no misunderstanding then when we speak of the intersection of a curve on Λ with a curve on Λ' , and at the same time insist that when we modify Λ' by a move, no change occurs on Λ .

Let us denote by $\lambda^i, i=1, 2, \dots, 2p-1$, the $2p-1$ segments of the straight line in Λ through the centers of the canonical circles. The segment

λ^k then joins the k th and the $(k+1)$ st circles. By subdivision of Λ we can make these segments 1-cells of Λ . Let $\lambda'^i, i=1, 2, \dots, 2p-1$, be the images of λ^i in Λ' . Then if λ^1 joins ϵ_1^5 to ϵ_2^2 , λ'^1 joins $\epsilon_1'^5$ to $\epsilon_2'^2$.

Since $\Lambda - \sum_{i=1}^p (\epsilon_1^i + \epsilon_2^i) - \sum_{j=1}^{2p-1} \lambda^j$ is a 2-cell, to prove Δ is equivalent to Δ' it is sufficient to prove that we can modify Λ' by our moves so that ϵ_1^k and ϵ_2^k coincide with $\epsilon_1'^k$ and $\epsilon_2'^k$, respectively, and λ^k is isotopic to λ'^k . But this is quite obviously possible by means of moves IIA and IIB. Indeed, let ϵ_1^1 be the first circle on Λ (reading from left to right, say). The 1-cell λ^1 joins ϵ_1^1 to ϵ^a , say. By a finite number of moves of Type IIA (acting on ϵ_1^1 of Λ') we can make $\epsilon_1'^1$ coincide with ϵ_1^1 . Then by another finite sequence of moves we can make ϵ'^a coincide with ϵ^a . Moreover, we can choose this sequence of moves in such a fashion that λ'^1 loses all its intersections with all $\lambda^i, i \neq 1$. Then by a finite number of moves of Type IIB acting on ϵ'^a , we can make λ'^1 lose all its intersections with λ^1 . The 1-cells λ^1 and λ'^1 are now isotopic. By continuing this process, we can make each λ^k isotopic to its corresponding λ'^k and each ϵ^k (ϵ_2^k) coincident with $\epsilon_1'^k$ ($\epsilon_2'^k$). The theorem is therefore proved.

COROLLARY 1. *Let $\Delta \equiv (\Lambda, \epsilon, \phi)$ be a given Heegaard diagram, and let $\bar{\Delta} \equiv (\bar{\Delta}, \bar{\phi}, \bar{\epsilon})$ be a diagram obtained from Δ by mapping the arcs ϕ^i of Λ as pairs of circles on $\bar{\Delta}$ and the pairs of circles ϵ_1^i and ϵ_2^i of Λ on arcs $\bar{\phi}^i$ of $\bar{\Delta}$. Then if $\bar{\Delta}' \equiv (\bar{\Delta}', \bar{\phi}', \bar{\epsilon}')$ is another diagram similarly obtained, $\bar{\Delta}$ and $\bar{\Delta}'$ are equivalent.**

COROLLARY 2. *Let the diagram Δ be equivalent to the diagram Δ' ; then if the diagrams $\bar{\Delta}$ and $\bar{\Delta}'$ are obtained from Δ and Δ' , respectively, as above, they are equivalent.*

The proofs of these corollaries follow directly from Theorem 2. We only need to point out that a move of Type IC acting on Δ is equivalent to a move of Type ID acting on $\bar{\Delta}$.

3. Theorem 3 below states that two Heegaard diagrams arising from two choices of the canonical 2-cells of R_1 (see (b) §1) are equivalent.

THEOREM 3. *If $\{e\}$ and $\{e'\}$ are two sets of canonical curves for the same region R_1 of M , then any diagram $\Delta \equiv (\Lambda, \epsilon, \phi)$ arising from $L, \{e\}$ and $\{f\}$ is equivalent to any diagram $\Delta' \equiv (\Lambda', \epsilon', \phi')$ arising from $L, \{e\}$ and $\{f\}$.*

In this theorem, Λ and Λ', ϕ^i and ϕ'^i correspond to the same L and f^i , respectively, of M , but ϵ corresponds to an unprimed e whereas ϵ' corresponds to a primed e .

Let ϵ''^i be the image of e^i on Λ' . We will prove the theorem by showing how to modify Λ' so that the ϵ'' 's are replaced, one by one, by the ϵ'' 's. The

* From here on, $\bar{\Delta}$, etc., will indicate not the closure of Δ , etc., but a new Δ , etc.

diagram Δ' will be then transformed into the diagram Δ'' . By definition, Δ' and Δ'' are equivalent, by Theorem 2, Δ and Δ'' are equivalent, hence Δ and Δ' are equivalent. The process by which we modify Λ' will be broken up into three cases.

Case 1. We suppose first that no ϵ'' meets an ϵ_1' or an ϵ_2' and that no ϵ'' is interior to another.

The ϵ'' curves will then be circuits on Λ' not meeting any of the circles, ϵ_1^i or ϵ_2^i . Each ϵ'' must contain at least one ϵ' circle; for any i , either ϵ_1^i or ϵ_2^i or both are contained in an ϵ'' ; and there is at least one ϵ' circle not contained in any ϵ'' .

Let ϵ_1^k be a circle not contained in any ϵ'' . Its partner ϵ_2^k must be contained in at least one ϵ'' . Several subcases now arise.

Suppose first only one ϵ'' circle, say ϵ''^a , contains ϵ_2^k and that ϵ''^a contains no other ϵ' . Then by a move of Type IB, in which we sever Λ' along ϵ''^a and heal it up again by matching ϵ_1^k and ϵ_2^k , we transform one of the ϵ'' 's into one of the ϵ'' 's.

Suppose secondly that only one ϵ'' circle, say ϵ''^b , contains ϵ_2^k , but that ϵ''^b contains other ϵ' circles. If ϵ''^b contains only one other such circle, say ϵ_1^m , then by a move of Type IC by which we sever Λ' along ϵ''^b and heal it up again by matching ϵ_1^k and ϵ_2^k , we again transform one of the ϵ'' 's into one of the ϵ'' 's. The circle ϵ_1^m is now not contained in any ϵ'' curve.

If ϵ''^b contains more than one circle besides ϵ_2^k , we modify Λ' by a move of Type IC by which we sever Λ' along a circuit not meeting ϵ''^b and containing ϵ_2^k and one other circle and then patch it up by matching ϵ_1^k and ϵ_2^k . We thus obtain a new configuration in which there are one fewer circles interior to ϵ''^b . One of the new canonical circles is interior and the other is exterior to ϵ''^b . By repeating this process we can remove the circles interior to ϵ''^b one by one until only two remain, and then we are back to the preceding case.

By a finite number of steps we can modify Λ' so that all the ϵ'' 's become ϵ'' 's. Hence, Δ' is transformed into a diagram Δ'' representing L , $\{e\}$ and $\{f\}$, and is therefore equivalent to Δ by Theorem 2, as we wished to prove.

Case 2. Let no ϵ'' meet an ϵ' circle, but suppose that some ϵ'' 's are interior to others.

Suppose first that the ϵ'' 's form "nests" of circuits, i.e. if ϵ''^a contains ϵ''^b and ϵ''^c , then either ϵ''^b contains ϵ''^c or ϵ''^c contains ϵ''^b . By considerations entirely analogous to those of Case 1 we can modify Λ' so that the innermost ϵ'' of a nest is changed into an ϵ' and hence the nest has one fewer ϵ'' curves. After a finite number of steps we are led back to Case 1.

Now suppose that the ϵ'' 's are contained in other ϵ'' 's in a general fashion.

Several cases arise, but it is not necessary to go into details. We can always transform Λ' so that the ϵ'' 's are transformed into nests of curves and we are back to the former case.

Case 3. We suppose that the ϵ'' 's do meet the ϵ' 's.

Let us examine, first of all, the nature of the intersection of an unprimed canonical 2-cell, say E^i , with a primed canonical 2-cell, say E'^k , in the manifold M . The intersection will consist of a number of 1-cells and circuits. By a slight deformation of E'^k and its boundary e'^k we can arrange so that e'^k meets e^i , the boundary of E^i , only in a finite number of points, such that the end points of any 1-cell common to E^i and E'^k are distinct.

Keep j fixed and let k run from 1 to p . We obtain on e^j a certain number of points, grouped into pairs, where each pair is the boundary of a 1-cell common to E^j and some E' and no pair "separates" another on e^j . Hence there is at least one pair, say P_1 and P_2 , such that one of the two arcs into which P_1 and P_2 divide e^j contains no other intersection point. Let us call this arc a . The points P_1 and P_2 are on some e' , say e'^k .

We now return to the Heegaard diagram Δ' on Λ' . The canonical curve e'^k is represented by the pair of circles $\epsilon_1'^k$ and $\epsilon_2'^k$; each circle has on it the images of P_1 and P_2 , which we continue to call P_1 and P_2 , and a is now mapped on a 1-cell α of Λ' joining P_1 and P_2 of $\epsilon_2'^k$, say. The arc α meets no other circle.

We now show that we can always modify Λ' so as to lose the two intersections P_1 and P_2 . The circle $\epsilon_2'^k$ and the arc α divide Λ' into two regions. Let A be the one which does not contain $\epsilon_1'^k$. Several cases arise. Suppose first that A contains no canonical circle at all. Then, obviously, α can be deformed so that the intersections P_1 and P_2 are lost.

Suppose next that A contains only one canonical circle. Let g be a circuit in $\Lambda' - A$ which lies very close to α and that part of $\epsilon_2'^k$ bounding $\Lambda' - A$. Then if we modify Λ' by a move of Type IC in which we sever Λ' along g and patch it up again by matching $\epsilon_1'^k$ and $\epsilon_2'^k$, we find that the two intersections P_1 and P_2 have disappeared.

Suppose, finally, that A contains several canonical circles. Choose g as before and again sever Λ' along it and patch it up again along corresponding points of $\epsilon_1'^k$ and $\epsilon_2'^k$. Again, we lose the two intersections, P_1 and P_2 , and, by Lemma 3, this operation is the product of moves of Type IC. Hence since we can remove all the intersections two by two this case is reduced to the former and the theorem is completely proved.

4. Theorem 4 states that two Heegaard diagrams arising from two choices of the canonical 2-cells of R_2 (see (c), §1) are equivalent. In the previous theorem, the circles on Λ and Λ' were images of different e 's; in this theorem, the ϕ 's are images of different f 's.

THEOREM 4. *If $\{f\}$ and $\{f'\}$ are two sets of canonical curves for the same region R_2 of M , then any diagram $\Delta \equiv (\Lambda, \epsilon, \phi)$ arising from $L, \{e\}$ and $\{f\}$ is equivalent to any diagram $\Delta' \equiv (\Lambda', \epsilon', \phi')$ arising from $L, \{e\}$ and $\{f'\}$.*

This theorem follows at once from Corollaries 1 and 2 of Theorem 2 and Theorem 3 by mapping Δ and Δ' on $\bar{\Delta}$ and $\bar{\Delta}'$ as in Corollary 1.

5. We have proved, Theorems 2, 3, and 4, that any two Heegaard diagrams arising from a manifold are equivalent, provided that we chose the same canonical surface in each case. We have left to prove that any two diagrams whatsoever arising from a manifold are equivalent (see (a), §1). To prove this we make use of a special canonical surface whose construction is given below.

6. **Construction of the special canonical surface.** Let M be a manifold which we assume simplicial, G the linear graph consisting of all the 0- and 1-simplexes of M , H a subcomplex of G . Let M' be the first derived complex of M , and M'' the first derived complex of M' ; and let G', G'', H', H'' be the induced subdivisions on G and H , respectively. It may happen that the boundary† of the M'' -neighborhood of H is a canonical surface, in which case we call it a *special canonical surface* and any diagram arising from it a *special Heegaard diagram*. In particular, if H is G itself, we prove

LEMMA 9. *The boundary, L , of the M'' -neighborhood of G is a special canonical surface.*

Let M^* be the dual of M , G^* the linear graph consisting of all the 0- and 1-cells of M^* , and $G^{*'}$ the subdivision of G^* induced by M' . Further, let R_1 and R_2 be the M'' -neighborhoods of G and G^* , respectively. We prove that L is a canonical surface by showing

- (a) that every 3-simplex of M'' belongs either to R_1 or R_2 ;
- (b) L is the common boundary of R_1 and R_2 , and is a closed and connected 2-dimensional manifold;
- (c) R_1 and R_2 are canonical regions.

It then follows from our definitions (Part I, §4) that L is a canonical surface dividing M into the canonical regions R_1 and R_2 .

(a) Let A_0^i be the vertices of M ; B_1^i, B_2^i, B_3^i the vertices of M' on the 1-, 2- and 3-cells of M , respectively, and C_1^i, C_2^i, C_3^i the vertices of M'' on the 1-, 2- and 3-cells of M' , respectively. Every 3-simplex of M' is of the form $A_0 B_1 B_2 B_3$, every 3-simplex of M'' is of the form $A_0 C_1 C_2 C_3$ or $B_a C_1 C_2 C_3$, where $a = 1, 2$, or 3 . Hence every 3-simplex of M'' is incident with an A_0 or a B_1 , i.e. a vertex of G' or else with a B_2 or a B_3 , i.e. a vertex of $G^{*'}$. That is, every 3-simplex of M'' belongs to R_1 or to R_2 .

† See second footnote on page 89.

(b) In any 3-simplex, σ_3^i , of M' , the component 3-simplexes of R_1 and R_2 are grouped as the simplexes of §2 of Part I, hence in any 3-simplex of M' , the boundaries of R_1 and R_2 have in common a 2-cell which we call E_2^i . The boundary of E_2^i is a circuit which may be thought of as composed of four 1-cells, each on one of the faces of σ_3^i . Let us call the 1-cell on the 2-simplex σ_2^i , E_1^i . It is quite clear that the incidence matrix of the 2- and 3-simplexes of M' is the same as the incidence matrix of all the 1- and 2-cells E_1^i and E_2^i . It follows at once from the fact that M is a manifold that $L = \sum E_2^i$ (the sum being taken over all i 's) is a closed, connected 2-dimensional cellular manifold, from which it can be deduced that, upon subdivision, L will become a closed, connected 2-dimensional simplicial manifold. Furthermore, it follows that L is orientable or non-orientable according as M is orientable or non-orientable, and conversely.

(c) Since the M'' -neighborhood of a 1-simplex of M is a 3-cell, the M'' -neighborhood of any tree of 1-simplexes of M is also a 3-cell. Let the linear graph G contain p independent circuits; then the removal of p properly chosen 1-simplexes, say $\sigma_1^1, \sigma_1^2, \dots, \sigma_1^p$, from G will reduce it to a tree T . Let B_i^1 be the vertex of M' on σ_1^i ($i=1, \dots, p$) and let E^i be the aggregate of 2-simplexes of M'' that lie in the 2-cell of M^* dual to σ_1^i and are incident with B_i^1 ($i=1, \dots, p$). The M'' -neighborhood of T is a 3-cell and if we add to it the remainder of the M'' -neighborhood of G excepting the 2-cells E^1, \dots, E^p , we obtain a 3-cell. Hence R_1 is a canonical region.

By exactly the same argument it can be shown that R_2 is also a canonical region. It therefore follows that L is a canonical surface dividing the manifold M into the two canonical regions, R_1 and R_2 , as we wished to prove.

7. We have proved incidentally that L and M are together orientable or non-orientable. Also, since the number of 2-cells which when removed reduce R_1 to a 3-cell is equal to the maximum number of non-intersecting circuits that can be drawn on L without disconnecting it, it follows that the number of 2-cells removed from R_1 is equal to the number removed from R_2 . In other words, the cyclomatic number of G is equal to the cyclomatic number of $G^*\dagger$.

8. We prove the following lemma:

LEMMA 10. *Let H and J be subcomplexes of the linear graph G of a manifold M such that it is possible to build up H from J by successively adding to J closed 1-simplexes of G in such fashion that at every step the 1-simplex which is being added either has one and only one end point in common with the subgraph already built up or else is the third side of a 2-simplex of which the other two sides*

\dagger The cyclomatic number of a linear graph is the number of independent 1-circuits on it, i.e. the minimum number of 1-simplexes that can be removed which reduce the graph to a tree.

already belong to the subgraph. Then if the boundary of the M'' -neighborhood of J is a canonical surface, so is the boundary of the M'' -neighborhood of H , and conversely.

The proof of this lemma follows by induction. Suppose that at the n th step we have built up the subgraph J_n from J and that the boundary of the M'' -neighborhood of J_n is a canonical surface of M . We add the closed 1-simplex σ_1^{n+1} to J_n , obtaining the subgraph J_{n+1} . By hypothesis, either σ_1^{n+1} has only one end point in common with J_n , or else it completes a triangle, of which the other two sides are already in J_n . It is obvious in the first case that the boundary of the M'' -neighborhood of J_{n+1} is a canonical surface of M .

Let then σ_1^{n+1} complete a triangle of which the other two sides belong to J_n . Then if we call R_2 the M'' -neighborhood of J_n , L its boundary and R_1 the remainder of M , that part of σ_1^{n+1} lying in R_1 can play the rôle of A of Lemma 2. Hence, applying Lemma 2, we obtain a new canonical surface of M which is precisely the boundary of the M'' -neighborhood of J_{n+1} . It follows that the boundary of the M'' -neighborhood of H is a canonical surface of M .

Conversely, let us suppose that the boundary of the M'' -neighborhood of J_{n+1} is a canonical surface. Again, if the 1-simplex σ_1^{n+1} which was added to J_n to form J_{n+1} has only one end point in common with J_n , then the boundary of the M'' -neighborhood of J_n is a canonical surface of M .

In the second case, let us call that part of the 2-simplex whose boundary σ_1^{n+1} completes lying in R_1 , E , and F that part of the 2-cell dual to σ_1^{n+1} lying in R_2 . We can choose E and F canonical 2-cells of R_1 and R_2 , respectively, and so arrange that their boundaries (which meet once and only once) do not meet any of the other canonical curves. Lemma 1 is then applicable, and therefore the boundary of the M'' -neighborhood of J_n is a canonical surface of M . It follows that the boundary of the M'' -neighborhood of J is a canonical surface of M and the lemma is proved.

9. Let L and L' be the boundaries of the M'' -neighborhoods of H and J , respectively, as above, and suppose that one (and hence the other) is a canonical surface. As an immediate consequence of Lemma 10 we have

COROLLARY 1. *Any special Heegaard diagram arising from L is equivalent to any special Heegaard diagram arising from L' .*

The proof follows at once from the fact that when we pass from J_n to J_{n+1} , we either do not change the Heegaard diagram at all, or else we modify it by a move of Type III.

10. At this point it becomes necessary to restrict the notion of homeomorphism. Let M and M_0 be two manifolds which are not only homeomorphic

but also equivalent in the sense of semi-linear analysis situs, i.e., it is possible to subdivide each of them so that the subdivisions have identical structures (cf. Alexander, J. W., *The combinatorial theory of complexes*, Annals of Mathematics, (2), vol. 31 (1930)). M and M_0 will then be called *equivalent manifolds*.

From this definition and from Lemma 10 and its corollary follows at once

THEOREM 5. *If M and M_0 are equivalent manifolds, G and G_0 their linear graphs, L and L_0 the boundaries of the M'' - and M_0'' -neighborhoods of G and G_0 , respectively, then any special Heegaard diagram of M arising from L is equivalent to any special Heegaard diagram of M_0 arising from L_0 .*

11. We have one more theorem to prove before we can prove our objective, Theorem 7.

THEOREM 6. *Let G be the linear graph of a manifold M , L_0 the boundary of the M'' -neighborhood of G , and L any canonical surface whatsoever of M . Then any Heegaard diagram Δ of M arising from L is equivalent to any special Heegaard diagram Δ_0 arising from L_0 .*

Let us note that if the canonical surface L is isotopic to a surface L' , then L' is also a canonical surface of M .

Let L divide the manifold M into the canonical regions R_1 and R_2 and let E_2^1, \dots, E_2^p be a set of canonical 2-cells for R_1 . After subdivision of M , if necessary, we can group the simplexes of M into cells, so that R_1 is composed of $p+1$ 3-cells, E_3^1, \dots, E_3^{p+1} , where E_3^i is bounded by the 2-cell E_2^i and $E_2^{i'}$ (a 2-cell isotopic to E_2^i) and 2-simplexes on L , for $i=1, 2, \dots, p$; and E_3^{p+1} is a 3-cell bounded by $\sum(E_2^i + E_2^{i'})$ and simplexes on L . By elementary operations we can further change the structure of M so that all of the 3-cells E_3^1, \dots, E_3^{p+1} and all of the 2-cells $E_2^1, \dots, E_2^p, E_2^{1'}, \dots, E_2^{p'}$, become stars of simplexes having for centers the 0-cells $E_0^1, \dots, E_0^{p+1}, P_0^1, \dots, P_0^p$, and $P_0^{1'}, \dots, P_0^{p'}$, respectively. The manifold M is thus transformed into an equivalent manifold N . Let L_n be the image of L in N .

Consider the linear graph H of N , where

$$H = \sum_{i=1}^p (E_0^{p+1}P_0^i + P_0^iE_0^i + E_0^iP_0^{i'} + P_0^{i'}E_0^{p+1})$$

in which $E_0^{p+1}P_0^i$, etc., stands for the 1-simplex joining E_0^{p+1} and P_0^i . By construction, G_n , the linear graph of N can be built up from H as is required in Lemma 10. Hence any special Heegaard diagram of N , say Δ_n , that arises from the boundary of the N'' -neighborhood of G_n is equivalent to any special Heegaard diagram that arises from the boundary of the N'' -neighborhood of

H . But this latter surface is isotopic to L_n . Hence Δ_N is equivalent to any Heegaard diagram arising from L_n . But since L_n is the image of L , Δ is such a diagram, hence Δ and Δ_N are equivalent. But by Theorem 5, Δ_N and Δ_0 are equivalent, therefore Δ and Δ_0 are equivalent as we wished to prove.

12. From Theorems 2, 3, 4, 5, and 6 follows the all inclusive theorem

THEOREM 7. *If Δ and Δ' are any two Heegaard diagrams whatsoever arising from a manifold M , then Δ and Δ' are equivalent.*

V. THE GENERAL THEOREM

1. Several theorems follow immediately from those of Part IV.

THEOREM 8. *If the equivalent manifolds M and N give rise to the Heegaard diagrams Δ and Δ' , then Δ and Δ' are equivalent.*

THEOREM 9. *If the diagram Δ gives rise to the manifold M , which in turn gives rise to the diagram Δ' , then Δ and Δ' are equivalent.*

From Theorem 9 and the preceding theorems follows

THEOREM 10. *If the Heegaard diagrams Δ and Δ' give rise to equivalent manifolds, then Δ and Δ' are equivalent.*

THEOREM 11. *If the manifolds M and N give rise to equivalent diagrams, then M and N are equivalent.*

From Theorems 1, 8, 10, and 11 follows the general theorem

THEOREM 12. *Equivalent manifolds give rise to and arise from equivalent Heegaard diagrams and equivalent Heegaard diagrams give rise to and arise from equivalent manifolds.*

The problem of determining when two given 3-dimensional manifolds are equivalent (in the sense that we can pass from one to the other by elementary operations) is thus reduced to the problem of determining when two given Heegaard diagrams are equivalent (in the sense that we can pass from one to the other by means of the moves defined in Part III).

PRINCETON UNIVERSITY,
PRINCETON, N. J.