EVERY PLANAR MAP IS FOUR COLORABLE PART II: REDUCIBILITY[1]

BY

K. APPEL, W. HAKEN, AND J. KOCH

1. Introduction

In Part I of this paper, a discharging procedure is defined which yields the unavoidability (in planar triangulations) of a set $\mathcal{U}$ of configurations of ring size fourteen or less. In this part, $\mathcal{U}$ is presented (as Table $\mathcal{U}$ consisting of Figures 1-63) together with a discussion of the reducibility proofs of its members.

When the term reducible is used above it is used in the following formal sense. Every configuration in $\mathcal{U}$ has the property that it is not only C- or D-reducible in the sense of [16], [27] (references are to the bibliography of Part I), but also if it is arbitrarily immersed in a planar map (i.e., not necessarily "properly embedded") then that planar map cannot be a minimal five chromatic map. A rather detailed study of such "immersion reducibility" is included in this paper.

Every configuration in $\mathcal{U}$ of ring size eleven or greater has been checked by our computer programs, with one exception.[2] For the reducibility of configurations of smaller ring size we rely on the tables in [2]. We do not claim to have been first to reduce all of these configurations. In particular we understand that F. Allaire has made a complete list of reducible eleven-rings and that H. Heesch has a large list of reducible configurations which has not been published. Furthermore, since we did not apply splicing arguments, there are C-reducible configurations, some of which appear in [25] and [1], for which we were not able to find reducers. But, since it meant only a small enlargement of our set $\mathcal{U}$ we preferred to include in $\mathcal{U}$ only such configurations as we could verify with our programs.[2] (See the note at the bottom of page 490.)

Received July 23, 1976.

[1] We should like to express our appreciation to the Research Board of the University of Illinois for supporting the computing effort. We have received tremendous help from the Computer Services Office (C.S.O.) at University of Illinois in using not only the IBM 360-75 computer at Urbana but also the IBM 370-158 computer at Chicago Circle and the 370-168 computer of the University Administrative Data Processing Unit. We should like to especially thank the consultants and systems programmers at C.S.O. for their excellent help and advice and the operations staff for their superb cooperation. We should also like to thank Laurel, Peter, and Andrew Appel for careful checking of diagrams and verifying the occurrence of configurations in the results of the discharging procedure.

In particular, we want to thank Michael Rolle, Charles Mills, and William Mills for pointing out copying errors in the earlier preprints of this paper.

[2] There is one major exception to our policy of reducing all required configurations of ring size greater than ten. Early in our work we realized that Configuration 16-14, which we could not reduce, would, if reducible, enable us to simplify our argument. We asked Frank

It is clear that our set $\mathcal{U}$ is not nearly best possible in the sense of being the shortest list of smallest configurations which might yield a proof of the conjecture (see also Section 5 of Part I). One might shorten the list by using splicing techniques for finding reducers for some configurations we were not able to reduce by our programs or by admitting a few fifteen-ring configurations to avoid a much larger number of smaller configurations, or by taking a few more steps in the discharging discussion to avoid certain configurations, or by choosing different sub-configurations of some of the configurations yielded by the discharging procedure. We chose not to do this because there seems neither to be a shortening which would significantly change the complexity of the proof nor a natural place to stop attempting such "simplifications." However, in working out the microfiche supplement to Part I of this paper we have found that 352 configurations may be omitted from $\mathcal{U}$ so that the remaining set $\mathcal{U}'$ of 1482 configurations is still unavoidable. The 352 configurations in $\mathcal{U} - \mathcal{U}'$ are listed in the microfiche supplement (see back cover).

The reducibility computations took place simultaneously with the final development of the procedure described in Part I, enabling us to modify the procedure where necessary to circumvent reduction failures. (Here, the term reduction failure will mean a configuration which was not reduced by our programs within the computer time allotted. We know that a significant number of configurations which we call reduction failures are actually C-reducible.)

2. The computer programs

D-reduction was done dynamically (see [2]) in the sense that once a coloration was proved good, its goodness was immediately available for use in the testing of other colorations. The programs for D-reduction are extensions and modifications of those in [20].

While the task of D-reduction is rather straightforward, C-reduction presents certain choices. When dealing with a good-sized configuration it is virtually impossible to try all possible reducers. In general it seems reasonable to make a limited effort to find C-reducers and, if necessary, accept a certain number of configurations as reduction failures which, with some more effort, could be shown reducible. While, in retrospect, our choices are probably not best possible, especially in light of the theory of splicing, they did easily C-reduce a large number of configurations.

For each ring size, eleven through fourteen, we tried a "best guess" reducer, then exhausted a certain class of potential reducers. If either no reducer was found in this manner or if, in the case of fourteen-rings, our time limit (90 minutes on an IBM 370-158 or 30 minutes on a 370-168) was exceeded, the con-

---

Allaire if he could reduce it. His elegant splicing argument provided a reducer for the C-reduction. We later discovered that this configuration eliminated at least eight configurations from $\mathcal{U}$. While we realized that Allaire's methods could greatly reduce the size of $\mathcal{U}$, we felt it unreasonable to ask him about all of the other configurations which were of interest to us. We certainly appreciate his help in this instance, however.

figuration was called a reduction failure. There was one class of exceptions to this procedure. For eleven-rings, the programs developed for [20] enabled us to test all reducers of up to four interior vertices.

We will illustrate the reducer-choosing algorithms in the case of twelve-rings. First, the best guess reducer was chosen as follows. For each pair of ring positions, the number of good colorations (from D-reduction) which gave both the same color was tabulated. The reducer was formed by choosing the largest possible number of high-ranking compatible identifications, choosing identifications in order of rank of the pairs in the number of such agreements (discarding identifications which caused conflicts with previously chosen identifications).

The second class of potential reducers was chosen as follows. First, all nonconflicting triples of identifications, each member of which identified vertices at distance two on the ring, were tried. If none of these reduced the configuration, the twenty with the smallest number of bad colorations were further processed. First, note that three such identifications convert a twelve-ring into a six-ring. Thus, we may consider reducers for this six-ring along with the three identifications to form a reducer for the twelve-ring. This was done to the twenty best triples using each six-ring reducer without interior vertices (see [16]) with each six-ring reducer in every possible position. Eleven-, thirteen-, and fourteen-rings were similarly treated.

The computer programs were greatly influenced by the facilities available. We had access to IBM computers (a 360-75 at Urbana-Champaign, a 370-158 at the University's Chicago Circle Campus, and later a 370-168 of the University of Illinois administrative data processing unit). For this reason the programs were written in IBM assembler language to attempt to maximize efficiency. When we inquired, the operations staff suggested that we use less computer time at the expense of larger amounts of core storage. Therefore, to save steps we chose to use large tables. The core storage requirements were as follows: for twelve-rings, 220,000 bytes; for thirteen-rings, 600,000 bytes; for fourteen-rings, 1,700,000 bytes.

Samples of reduction time for D-reducible configurations using four passes were roughly as follows. Eleven-rings took about 40 seconds on the 360-75; twelve-rings took about one minute on the 370-158; thirteen-rings took about 15 minutes on the 360-75 (slow core) and about 5 minutes on the 370-158; fourteen-rings required about 25 minutes on the 370-158 or 6 minutes on the 370-168. In general, the C-reduction, if it were to succeed, would succeed rather quickly so that a majority of the time was usually spent on D-reduction. Reduction failures, however, often took a great many passes in D-reduction and extra time attempting C-reduction and hence often took between four and eight times as long as the D-reductions mentioned above.

3. Immersion reducibility

The unavoidability of $\mathscr{U}$ in planar triangulations, as established in Part I, means that every planar triangulation Δ (without vertices of degree smaller than five) contains at least one member, say C, of $\mathcal{U}$ as an image of an immersion $f: C \to \Delta$ which respects the degree-specifications. (For definitions see Section 2 of Part I.) It is not sufficient to show that every configuration C in $\mathcal{U}$ is C- or D-reducible since C might be immersed in a planar triangulation Δ in such a way as to make untrue some of the hypotheses of independence which are basic to C- and D-reducibility. Thus we must discuss such immersion reducibility in some detail.

A configuration is sometimes thought of as an n-ring plus its interior, but we shall use the following terminology. We use the term *configuration* (as we did in Part I) for the vertices interior to the ring and the edges joining pairs of such vertices and the triangles bounded by the 3-circuits which are formed by these edges. For every vertex V of a configuration C a specification of its degree, deg (V), is given; in particular, we assume from now on that in every instant deg (V) is specified to be precisely one of the numbers $5, \ldots, 11$. (All configurations in $\mathcal{U}$ have this property; in the drawings we indicate the degree-specifications by coding according to the left column of Figure 1 in Part I.) Now every configuration C can be extended to a *ringed configuration* $\bar{C}$, see Figure A, so that $\bar{C}$ is a triangulation of a disk (in the plane) in which every vertex of C is an interior vertex and every triangle is incident to at least one vertex of C; in particular, each vertex V of C is incident to precisely as many edges of $\bar{C}$ as is specified for deg (V). Note that $\bar{C}$ is uniquely (up to isomorphism) determined by C. We call the boundary circuit R of $\bar{C}$ the *ring of* $\bar{C}$ and also the *ring of* C and we denote the number of vertices in R by n. (Every vertex of $\bar{C}$ is either a ring vertex or a vertex of C.) The edges between C and R (dashed in Figure A) are called the *legs* of C. A ringed configuration $\bar{C}$ is said to be *properly imbedded* in a triangulation Δ if the images of any pair of vertices are adjacent in Δ only if the vertices are adjacent in $\bar{C}$.

The definitions of C- and D-reducibility guarantee that configurations which satisfy them cannot have their ringed configurations properly imbedded in minimal five-chromatic planar triangulations. However our argument concerns immersions rather than imbeddings and we must establish that no member of $\mathcal{U}$ can be immersed in a minimal five-chromatic planar triangulation.

In order to formulate our theorem on immersion-reducibility we need several definitions. Let C be a configuration (with fully specified vertex degrees), let $\bar{C}$ be the corresponding ringed configuration, and let R be the ring. Suppose that V is a vertex of R which has precisely k neighbors in C and that $k \geq 3$. This means that precisely $k - 2$ neighbors of V lie on "1-legger outer sectors of C" and that V has precisely $k + 2$ neighbors in $\bar{C}$. For example in Figure A, the configuration C is 21-34 from our set $\mathcal{U}$; the "legs" are drawn as dashed lines, and for the vertex V we have $k = 5$. (k-values greater than five do not occur in any configuration of $\mathcal{U}$.) Now we may derive a configuration C' from C by "adding" vertex V to C and giving it the degree-specification d ($d \geq 5$). The ring size n' of C' will be smaller than the ring size n of C if (and only if) $d \leq k + 2$. In this case we call C' an *n-decreased extension of* C; in order to make this concept transitive we call any n-decreased extension of C' also an n-decreased extension of C. Obviously the total number of n-decreased extensions of C is finite. In our example Figure A, three n-decreased extensions $C'_{(1)}$, $C'_{(2)}$, $C'_{(3)}$ of C are obtained by adding V with $d = 7$, 6, or 5, respectively. In [26] Stromquist considers extensions with $d \geq k + 2$ and calls them *r-extensions* where $r = d - (k + 2)$; i.e., $C'_{(1)}$ in Figure A is a 0-extension of C. Generalizing this terminology we should call $C'_{(2)}$ a (-1)-extension and $C'_{(3)}$ a (-2)-extension.

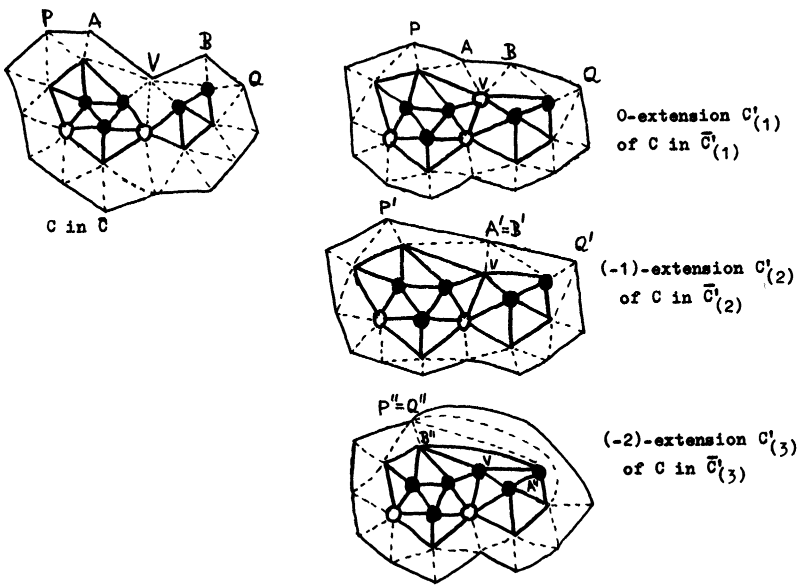

Figure A

Now let f: $C \to \Delta$ be a simplicial immersion. Certainly, f can be extended to a simplicial and dimension-preserving mapping $\bar{f}$: $\bar{C} \to \Delta$ of the ringed configuration $\bar{C}$; (for definition see Section 2 of Part I). However, $\bar{f}$ need not be an immersion since it may fail to be locally one-to-one in the neighborhood of some ring vertices. For instance the image $\bar{f}(\bar{C})$ of $\bar{C}$ in Figure A may be (isomorphic to) one of the ringed configurations $\bar{C}'_{(2)}$ or $\bar{C}'_{(3)}$, in case the ring vertex V is mapped to a vertex of degree six or five in Δ; then in the image, the edges VA and VB of $\bar{C}$ would be identified to each other or to legs from V to C. We say that the immersion f causes an *interior overlap* if the restriction of $\bar{f}$ to the interior of $\bar{C}$ is not one-to-one, i.e., if $\bar{f}$ maps some two different triangles of $\bar{C}$ to the same triangle of Δ. Suppose that C is C-reducible and that S is a reducer. Then f is called *compatible with S* if any two vertices of R which are identified by S are mapped into nonadjacent vertices of Δ and any two vertices of R which are joined by a diagonal of S are mapped into nonidentical vertices of Δ. Then we have the following.

LEMMA I. *Let f: $C \to \Delta$ be an immersion of a configuration C into a planar triangulation Δ which does not cause an interior overlap. Suppose that C is D-reducible, or that C is C-reducible and has a reducer S which is compatible with f. Then Δ cannot be minimal five-chromatic.*

Proof. Let R be the ring of C, $\bar{C}$ the ringed configuration, and $\bar{f}$: $\bar{C} \to \Delta$ the extension of f to $\bar{C}$. We make the following observations.

(i) Every coloration ϕ' of $\bar{f}(R)$ has a unique pre-image coloration ϕ of R (so that $\bar{f}$ carries ϕ into ϕ').

(ii) If the pre-image ϕ of ϕ' extends to a coloration ψ of $\bar{C}$ then $\bar{f}$ carries ψ into a coloration ψ' which extends ϕ' over $\bar{f}(\bar{C})$.

(iii) Now let ϕ' be a coloration of $\bar{f}(R)$ which extends over $\Delta - f(C)$. If π is a partition of the four colors into two pairs and K' is the corresponding Kempe chain disposition in $\Delta - f(C)$ then there exists an (abstract) Kempe chain disposition K in the exterior of $\bar{C}$, corresponding to π and the pre-image ϕ of ϕ', so that K is the pre-image of K' in the following sense. (iii)(a) If V' is a vertex of $\bar{f}(R)$ then all pre-images of V' are joined by Kempe chains of K; (iii)(b) if two vertices of $\bar{f}(R)$ are joined by a Kempe chain of K' then their pre-images are joined by a chain of K.

(iv) If a coloration $\tilde{\phi}$ of R is obtained from ϕ by a Kempe interchange according to π and K (as considered in (iii)) then $\tilde{\phi}$ is carried by $\bar{f}$ into a coloration $\tilde{\phi}'$ of $\bar{f}(R)$ which can be obtained from ϕ' by a Kempe interchange (according to π and K) in $\Delta - f(C)$.

(v) If a coloration ϕ' of $\bar{f}(R)$ is bad in the sense that it cannot be converted by iterated Kempe interchanges in $\Delta - f(C)$ into any coloration which also extends over $f(C)$, then, by (i), . . . , (iv), the pre-image ϕ of ϕ' is also a bad coloration and in particular, C is not D-reducible.

This implies the lemma in the case that C is D-reducible. If C is not D-reducible then, by hypothesis, the reducer S is incompatible with all bad colorations on R, and thus the image of S under $\bar{f}$ is, by (v), incompatible with all bad colorations of $\bar{f}(R)$. This implies the lemma in the case that C is not D-reducible. ■

COROLLARY. *If C is D-reducible and if C' is a* 0- *or* (-1)-*extension of C then C' is also D-reducible.*

Proof. We consider the ringed configurations $\bar{C}$ and $\bar{C}'$ and the "natural" immersion f: $C \to C'$ and its extension $\bar{f}$: $\bar{C} \to \bar{C}'$. Then f does not cause an interior overlap and we conclude as in the proof of the lemma above (where, in (iii), K' means an abstract Kempe chain disposition in the exterior of $\bar{f}(C)$). Note that for 0-extensions the corollary follows also from much stronger theorems of Stromquist (the 2-extension theorem and the gradient theorem in [26]). ■

Remark. It is conceivable (although unlikely) that there exists a configuration C and a 0- or (-1)-extension C' of C so that C is C-reducible but C' is not. The fact that C' "contains" C cannot be used for concluding reducibility of C' since $\bar{C}$ is not properly imbedded in $\bar{C}'$. It might be that all possible reducers of C are incompatible with the immersion f: $C \to \bar{C}'$.

Now we consider a simple edge path P in a ringed configuration $\bar{C}$ which joins two different vertices A and B of the ring R; (for examples see Figures B

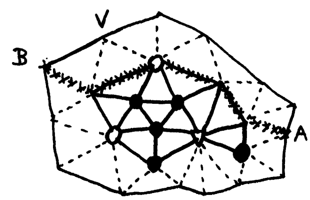

I-path from A to B

Figure B

and C). We say that *P fulfills a bend condition* if there is a vertex V in $\bar{C}$ which is not on P but is adjacent to (at least) three consecutive vertices on P. (The path marked in Figure B fulfills a bend condition, but the path marked in Figure C

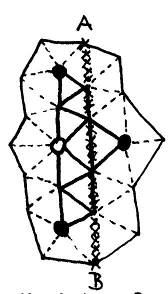

D-path, but no I-path from A to B

Figure C

does not.) A path P in $\bar{C}$ is called an *I-path* if it joins two vertices A and B of R and if it has the following properties.

(I.1) The length of P (i.e., the number of edges on P) is at most five.
(I.2) If the length of P is five then P fulfills a bend condition.
(I.3) The distance of A and B *on* R is at least four.
(I.4) P does not contain any vertices of R besides A and B.

A path P in $\bar{C}$ is called a *D-path* if it joins two vertices A and B of R and if it has the following properties.

(D.1) The length of P is at most six.
(D.2) If the length of P is six then P fulfills a bend condition.
(D.3) The distance of A and B on R is at least five.
(D.4) P does not contain any vertices of R besides A and B.

Now assume that C is C-reducible and that S is a reducer of C. Then S is called *fine* if it has the following properties.

(F.1) If two vertices A and B of the ring R are identified by S and if the distance of A and B on R is at least four then there exists an I-path from A to B in $\bar{C}$.
(F.2) If two vertices A and B of R are joined by a diagonal of S and if the distance of A and B on R is at least five then there exists a D-path from A to B in $\bar{C}$.

Now we have the following.

THEOREM. *Let C be a configuration which contains precisely m vertices and is of ring size n so that $n \leq 14$ and $n + m \leq 28$, and let $C'_{(1)}, \ldots, C'_{(p)}$ be all n-decreased extensions of C. Suppose that each one of $C, C'_{(1)}, \ldots, C'_{(p)}$ is D-reducible or is C-reducible with a fine reducer. Then C cannot be immersed into a minimal five-chromatic planar triangulation.*

Proof. Let $f\colon C \to \Delta$ be an immersion of C into a planar triangulation Δ and let $\bar{f}\colon \bar{C} \to \Delta$ be the extension of f to the ringed configuration $\bar{C}$ with ring R. We have to prove that Δ cannot be minimal five-chromatic.

(1) If $n \leq 5$ then the proof is easy since then the image $\bar{f}(R)$ of R contains a reducible ring by Birkhoff [10] or Δ contains at most 29 vertices.
(2) From now on we assume that $n > 5$ and we assume by induction that the theorem is proved for all ring-sizes smaller than n.
(3) If $\bar{f}$ is not an immersion then $\bar{f}$ fails to be one-to-one in the neighborhood of some vertex, say V, of R (compare Figure A). Then Δ contains an n-decreased extension C' of C which is obtained by adding V to C. Thus Δ cannot be minimal five-chromatic by induction hypothesis (2).
(4) From now on we assume that $\bar{f}$ is an immersion. Then in particular, $\bar{f}$ cannot identify any two vertices whose distance in $\bar{C}$ is smaller than three.
(5) If $\bar{f}$ identifies two vertices A and B of R which are at distance three or four *on* R then a path Q from A to B on R maps onto a 3- or 4-circuit Q' in Δ. If there is any vertex of Δ on that side of Q' which is opposite to $f(C)$ then Q' is a reducible ring by Birkhoff [10] and thus Δ is not minimal five-chromatic. If there is no vertex of Δ on that side of Q' then Δ contains an n-decreased extension C' of C (since at least one of the vertices between A and B on Q can

be added to C as a 2-legger) and we conclude again by induction that Δ is not minimal five-chromatic.

(6) From now on we assume that $\bar{f}$ does not identify any two vertices of R which are at distance less than five on R.

(7) If $\bar{f}$ causes an interior overlap then we claim that the following holds (compare Figure D).

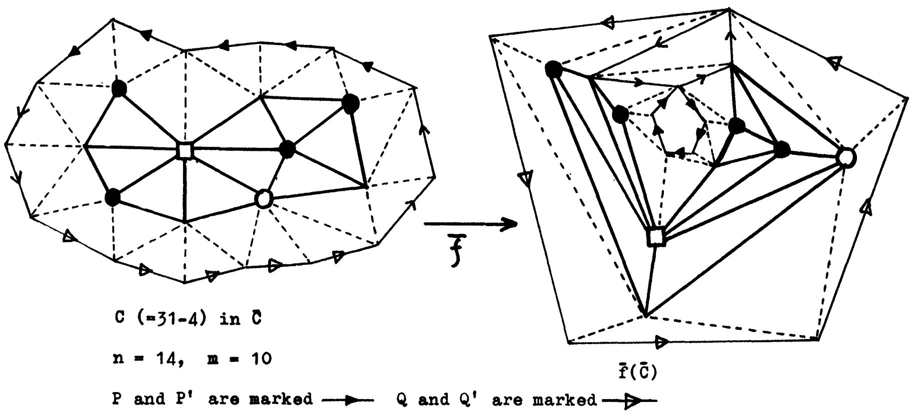

Figure D

(7.1) $\bar{f}(\bar{C})$ is an annulus the boundary of which consists of two disjoint 5-circuits P' and Q' in $\bar{f}(R)$; moreover, $n = 14$.

Proof of claim (7.1).[3] $\bar{C}$ is a triangulated disk; let t be the number of triangles in $\bar{C}$ (it follows from Euler's formula that $t = 2m + n - 2$ but we shall not use this fact). If D_i is a triangulated disk which contains precisely i triangles and if $i > 1$ then at least one of the triangles, which we denote by T_i, fulfills one of the two following conditions (see Figure E).

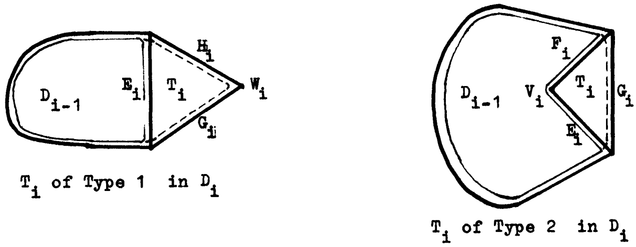

Figure E

[3] We want to thank George Francis for supplying us with a much shorter version of the main part of this proof which, however, uses more advanced topological concepts.

T_i is of *Type* 1 (in D_i) if precisely one of its edges, say E_i, is an interior edge of D_i and its two other edges, say G_i and H_i, are boundary edges of D_i.

T_i is of *Type* 2 (in D_i) if precisely two of its edges, say E_i and F_i, are interior edges of D_i and its third edge, say G_i, is a boundary edge of D_i so that the vertex V_i which is incident to E_i and F_i is an interior vertex of D_i.

If T_i is of Type 1 then removal of $T_i - E_i$ from D_i (i.e., removal of the interiors of T_i, G_i, and H_i and of the vertex, say W_i, which is incident to G_i and H_i) yields a triangulated disk D_{i-1}. If T_i is of Type 2 then removal of the interiors of T_i and of G_i, yields a triangulated disk D_{i-1}. We say that D_i can be obtained from D_{i-1} by an *expansion* of Type 1 or 2, respectively, according to whether T_i is of Type 1 or 2 in D_i. Consequently we may build up $\bar{C}$ in t steps by a sequence $D_1, D_2, \ldots, D_t = \bar{C}$ of triangulated disks so that D_1 is one triangle and D_i is an expansion of D_{i-1} (for $i = 2, \ldots, t$).

We consider the sequence $\bar{f}(D_1), \bar{f}(D_2), \ldots, \bar{f}(D_t) = \bar{f}(\bar{C})$ of images of the disks D_i under the immersion $\bar{f}$. Certainly, $\bar{f} \mid D_1$ is an imbedding. Let u be the smallest index so that $\bar{f} \mid D_u$ is not an imbedding; then we have the situation of Figure F, i.e., D_u is an expansion of Type 1 of D_{u-1}, and the image W'_u of the

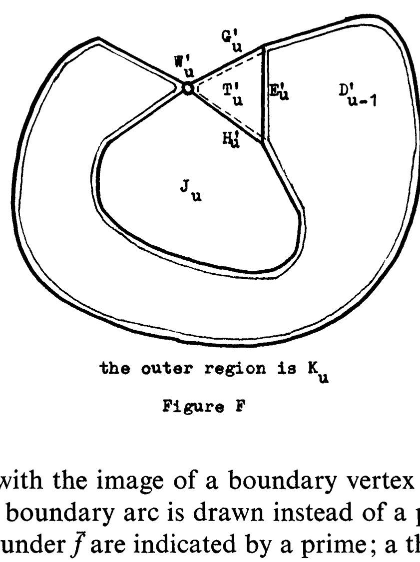

the outer region is K_u

Figure F

vertex W_u coincides with the image of a boundary vertex of D_{u-1}. (In Figures E, F, G, H a curved boundary arc is drawn instead of a polygonal arc of arbitrary length; images under $\bar{f}$ are indicated by a prime; a thin line parallel to the boundary of D_{i-1} or D_i indicates how the interior of D_{i-1} or D_i is mapped under $\bar{f}$.) In particular, the exterior of $\bar{f}(D_u)$ consists of two disjoint open regions which we denote by J_u and K_u (see Figure F). The boundaries of J_u and K_u are circuits which we denote by P'_u and Q'_u, respectively; there are corresponding arcs P_u and Q_u in the boundary of D_u the interiors of which are mapped one-to-one into P'_u and Q'_u; (the end points of P_u and Q_u are mapped to W'_u).

Now we conclude by induction (see Figure G) that for each $i = u, u + 1, \ldots, t$, there are arcs P_i and Q_i (with disjoint interiors) in the boundary of D_i so that

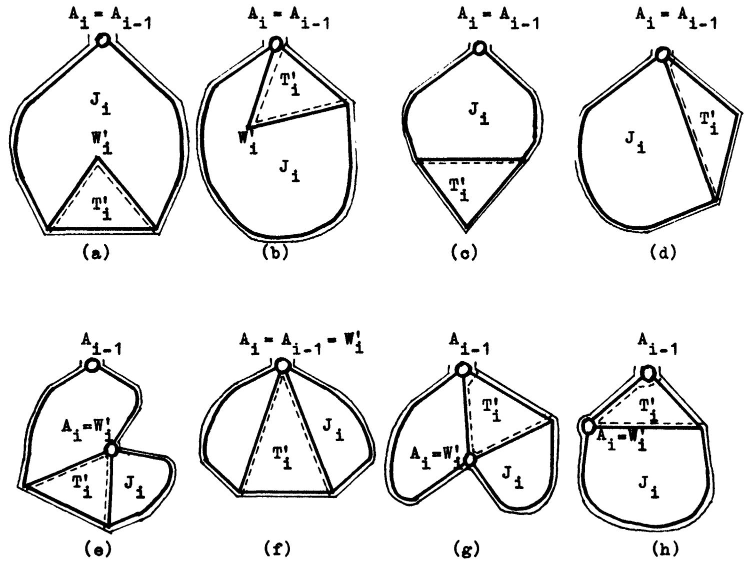

Figure G. J_i in J_{i-1}

the interiors of P_i and Q_i are mapped one-to-one into circuits P'_i and Q'_i and so that the end points of P_i are mapped to a vertex A_i in P'_i and the end points of Q_i are mapped to a vertex B_i in Q'_i (A_i and B_i may be distinct or may be the same vertex of Δ). In particular, P'_i is the boundary of an open region J_i and Q'_i is the boundary of an open region K_i so that $J_i \subseteq J_{i-1} \subseteq \cdots \subseteq J_u$ and $K_i \subseteq K_{i-1} \subseteq \cdots \subseteq K_u$. For if T'_i is disjoint from J_{i-1} (from K_{i-1}) then we may choose $J_i = J_{i-1}$ and $A_i = A_{i-1}$ ($K_i = K_{i-1}$ and $B_i = B_{i-1}$); if the interior of T'_i lies in J_i then we have one of the cases (a), . . . , (h) indicated in Figure G. If T_i is of Type 1 and W'_i is in J_{i-1} (cases (a) and (b)) or if T_i is of Type 2 (cases (c) and (d)) then we choose $J_i = J_{i-1} - T'_i$ and $A_i = A_{i-1}$; if T_i is of Type 1 and W'_i is in the boundary P'_{i-1} of J_{i-1} (cases (e), . . . , (h)) then at least one of the connected components of $J_{i-1} - T'_i$ can be chosen for J_i and W'_i for A_i.

Now we consider the last member of the sequence, $i = t$ (i.e., $D_i = D_t = \bar{C}$) and we omit the indices. Then, by hypothesis (6), each one of the arcs P and Q in R has length at least five. Moreover, it is not possible that the end points of P are identical with the end points of Q for otherwise $\bar{f}(\bar{C})$ would be homeomorphic to D'_u (see Figure F) and $\bar{f}$ would not cause an interior overlap. Thus $R - (P \cup Q)$ is not empty; on the other hand, $R - (P \cup Q)$ contains (the interiors of) at most four edges (since by hypothesis, $n \leq 14$). Thus, again by hypothesis (6), the images A and B of the end points of P and of Q must be distinct. Now it is not possible that $R - (P \cup Q)$ contains (the interiors of) fewer than four edges for otherwise one of the cases (α), (β), (γ) indicated in Figure H would apply. If $R - (P \cup Q)$ contains only two edges (case (α)) then

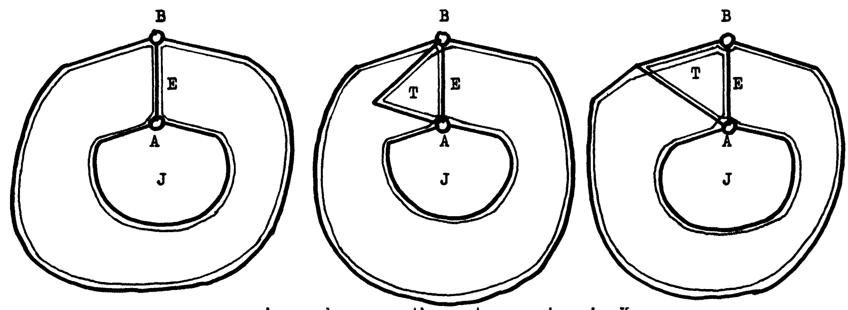

in each case, the outer region is K

(α) (no interior overlap) (β) (γ)

Figure H

these map to the same edge E (since Δ does not contain any 2-circuits) and no interior overlap is possible. If $R - (P \cup Q)$ contains precisely three edges (cases (β) and (γ)) then these map to the boundary of a triangle T so that one side, say E, of T joins A and B (since Δ does not contain any 3-circuits other than boundaries of triangles); but then no interior overlap is possible for otherwise one of the two pre-images of E (under $\bar{f}$) would have to be a diagonal in $\bar{C}$ (but a ringed configuration cannot contain any diagonal edges).

Thus the only remaining case is that $R - (P \cup Q)$ contains precisely four edges and that each of P and Q contains precisely five edges and $n = 14$. This proves our claim (7.1).

(7.2) By (7.1), Δ cannot be minimal five-chromatic since at least one of the two 5-circuits P', Q' is reducible or Δ contains at most 30 vertices.

(8) From now on we assume that f does not cause an interior overlap.

(9) If C is D-reducible then it follows directly from Lemma I that Δ cannot be minimal five-chromatic.

(10) From now on we assume that C is not D-reducible but is C-reducible and has a fine reducer S.

(11) If $\bar{f}$ maps two vertices A and B of R which are of distance two or three on R to vertices A' and B' which are joined by an edge E' of Δ then there is a path Q from A to B on R so that $\bar{f}(Q) \cup E'$ is a 3- or 4-circuit Q' in Δ. Then we can conclude (as in (5)) that Q' is a reducible circuit in Δ or that Δ contains an n-decreased extension C' of C and is not minimal five-chromatic by induction hypothesis (2).

(12) From now on we assume that $\bar{f}$ does not map any two vertices of R which are of distance two or three on R to adjacent vertices of Δ.

(13) If f is not compatible with S then we can conclude that Δ contains a circuit Q' of length at most six which is the image of an I-path in $\bar{C}$ plus an edge of Δ which joins its end points, or which is the image of a D-path with identified

end points. On either side of Q' there are at least three distinct vertices of Δ which belong to $\bar{f}(R)$. Now we distinguish three cases.

(13.1) If the length of Q' is less than six then Q' is reducible by Birkhoff [10] and Δ is not minimal five-chromatic.

(13.2) If on some side of Q' there are precisely three vertices of Δ then these are the images of three consecutive vertices, say U, V, W, of R. U, V, W form a 5-5-5 triangle (since otherwise the length of Q' would be greater than six.) But this means that an edge of Δ joins U and W which contradicts (12) and thus this case is ruled out.

(13.3) If on either side of Q' there are at least four vertices of Δ and if the length of Q' is six, then Q' fulfills a bend condition and is therefore reducible by Arthur Bernhart [6] and Δ is not minimal five-chromatic.

(14) Finally we assume that f is compatible with S. But then Lemma I implies that Δ cannot be minimal five-chromatic. This finishes the proof of the theorem. ■

Now it is not difficult to verify that all configurations of our set $\mathscr{U}$ fulfill the hypotheses of the theorem. For the D-reducible members of $\mathscr{U}$ the D-reducibility of all 0- and (-1)-extensions follows immediately from the Corollary to Lemma I and we have to individually check only the very few cases in which a (-2)-extension is possible. For those configurations which are not D-reducible we have to check the 0- and (-1)-extensions also. But in most cases these extensions contain, properly embedded, other configurations of $\mathscr{U}$ which are of smaller ring size. We also have to check the fineness of the reducers. But most of the configurations are so small that every reducer is fine. (See the supplement.)

4. The unavoidable set $\mathscr{U}$ of reducible configurations

The members of $\mathscr{U}$ are displayed in Table $\mathscr{U}$ which consists of Figures 1-63. The order was chosen in an attempt to aid the reader in finding configurations and to display similar configurations in the same part of the table. While it is easiest to understand the ordering of the table by just glancing through it, a few remarks seem in order.

The table is organized into primary parts determined by the number of major vertices of each degree in the configuration. For example, the configurations in Figure 1 have no major vertices, while those in Figure 16 have a pair of V_7's. Except in the case of rather small such classes, for example configurations with two vertices of degree 9 (see Figure 45), classes span a consecutive sequence of full figures. Within such classes the ordering is more arbitrary and not wholly consistent. Similar looking configurations are usually on the same page, for example all of the configurations in Figure 23 have a V_7 adjacent to a V_8 with a single V_5 below and adjacent to both and a V_6 above and adjacent to both. Occasionally a configuration cannot fit into the appropriate figure and will usually then be found within two pages of its natural position, towards the bottom of the figure.

The attempt to obtain a reasonable visual display has caused some figures to have fewer than the 35 possible entries. In addition, when a redundancy was discovered in the final examination of the table the redundant configuration was deleted without the potentially error-introducing procedure of moving up the remaining configurations.

Section (β) of the microfiche supplement provides more information on the C-reducible (but not D-reducible) configurations of Table $\mathscr{U}$. Above each diagram in the supplement is a pair *f-p* which give the figure number and position in the figure of the configuration which is discussed. The order of the supplement is the same as that of the table except that certain configurations were added after the figures in the supplement were drawn. This is indicated by a heavy carat in the figure. For example, the carat between 5-32 and 6-13 indicates that a configuration has been added at the end of the table which should fit in this position (the configuration is 5-35).

Since almost all of our C-reducers consist of diagonals and identifications, the diagram is drawn as follows. First, the configuration of Table $\mathscr{U}$ is copied. The n-ring (which is not shown) can easily be drawn by the reader using the information given by the degrees of the boundary vertices of the configuration. The vertices on the n-ring are numbered from 1 through n in clockwise order. All edges leading from the configuration to vertex 1 on the n-ring are added to the figure as dashed lines. This enables the reader to appropriately number all of the vertices on the ring. Below the figure all identifications and diagonals are listed. For example, Configuration 2-28 has identifications which identify vertices 3 and 5, and vertices 7, 9, and 11, and diagonals connecting vertices 1 and 6, vertices 2 and 6, and vertices 6 and 10. Further m (the number of vertices) and n (the ring size) are supplied (in the case of 2-28, $m = 8$, $n = 12$). The last number (6655 in Configuration 2-28) gives the number of bad colorations after the algorithm for testing D-reducibility has been applied.

The following numbers shed some additional light on $\mathscr{U}$. Of the 1834 configurations in $\mathscr{U}$, 7 are of ring size eight or less; 8 are of ring size nine; 35 are of ring size ten; 89 are of ring size eleven; 334 are of ring size twelve; 701 are of ring size thirteen; and 660 are of ring size fourteen. There are 504 configurations in $\mathscr{U}$ which are C-reducible but not D-reducible. For more details see Section (ε) of the microfiche supplement.

There are 13 C-reducible configurations of ring size less than eleven which are not included in the supplement. These are considered in the tables in [2].

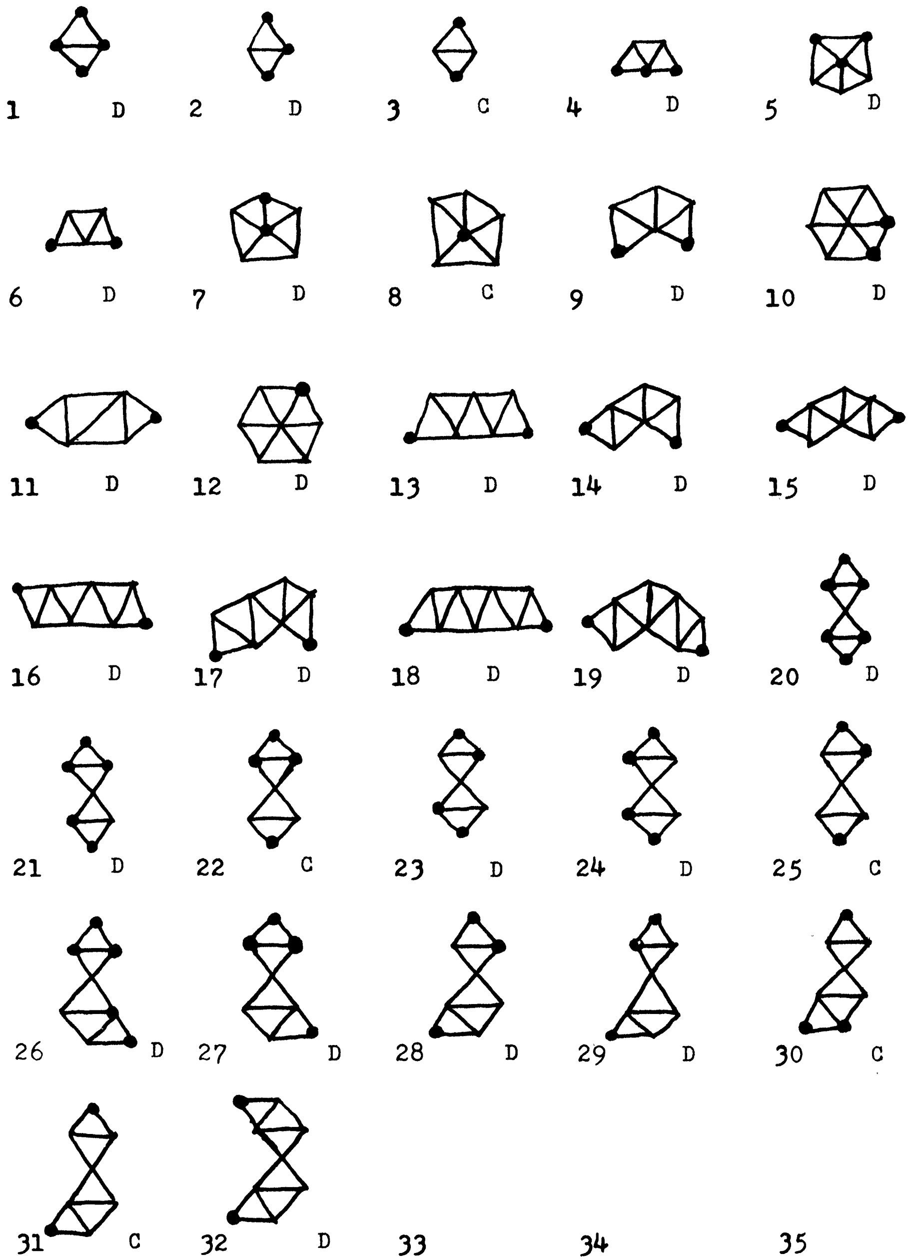
1 D
2 D
3 C
4 D
5 D
6 D
7 D
8 C
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 C
23 D
24 D
25 C
26 D
27 D
28 D
29 D
30 C
31 C
32 D
33
34
35

Figure 1

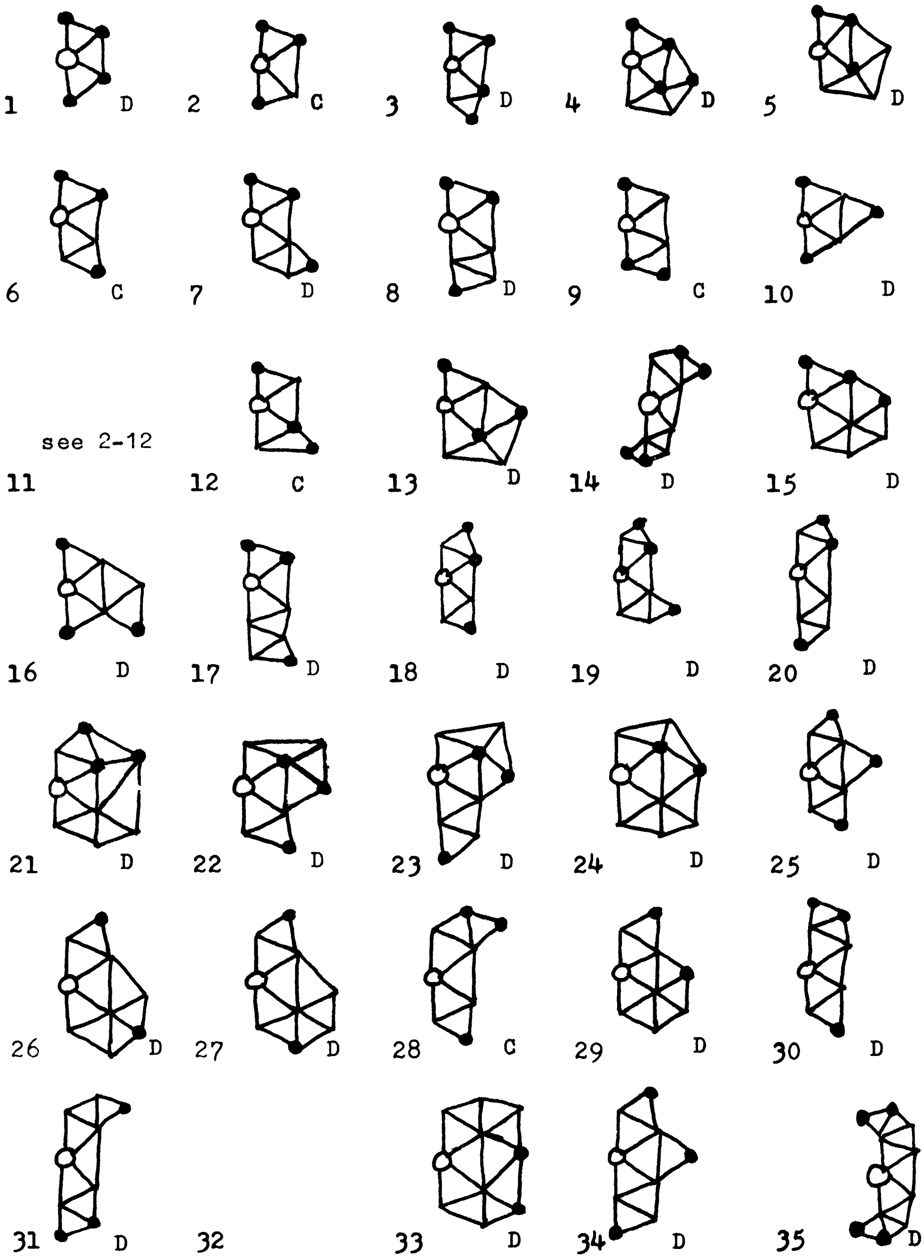
1 D
2 C
3 D
4 D
5 D
6 C
7 D
8 D
9 C
10 D
see 2-12
11
12 C
13 D
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 C
29 D
30 D
31 D
32
33 D
34 D
35 D

Figure 2

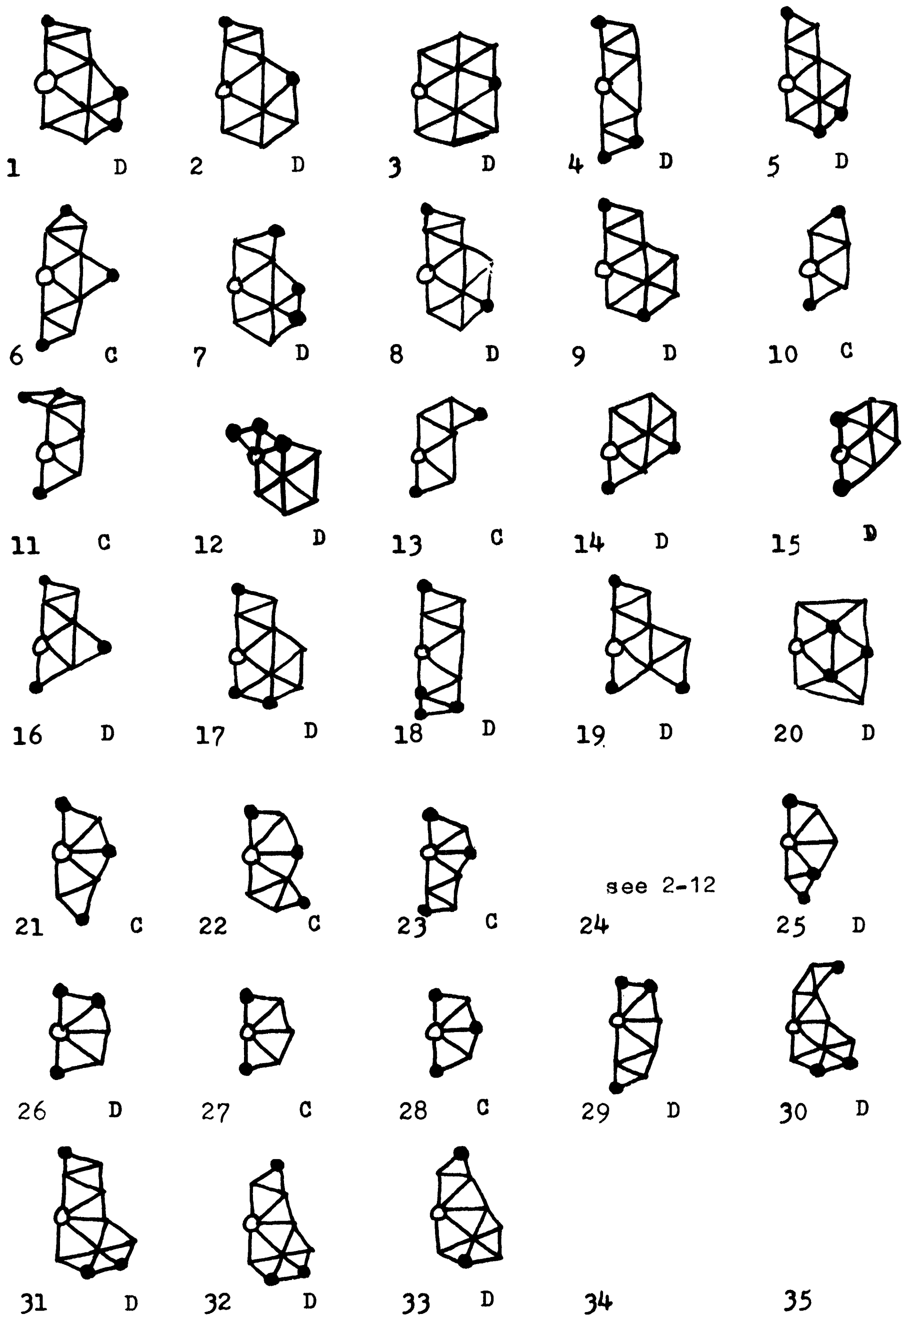
1 D
2 D
3 D
4 D
5 D
6 C
7 D
8 D
9 D
10 C
11 C
12 D
13 C
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21 C
22 C
23 C
see 2-12
24
25 D
26 D
27 C
28 C
29 D
30 D
31 D
32 D
33 D
34
35

Figure 3

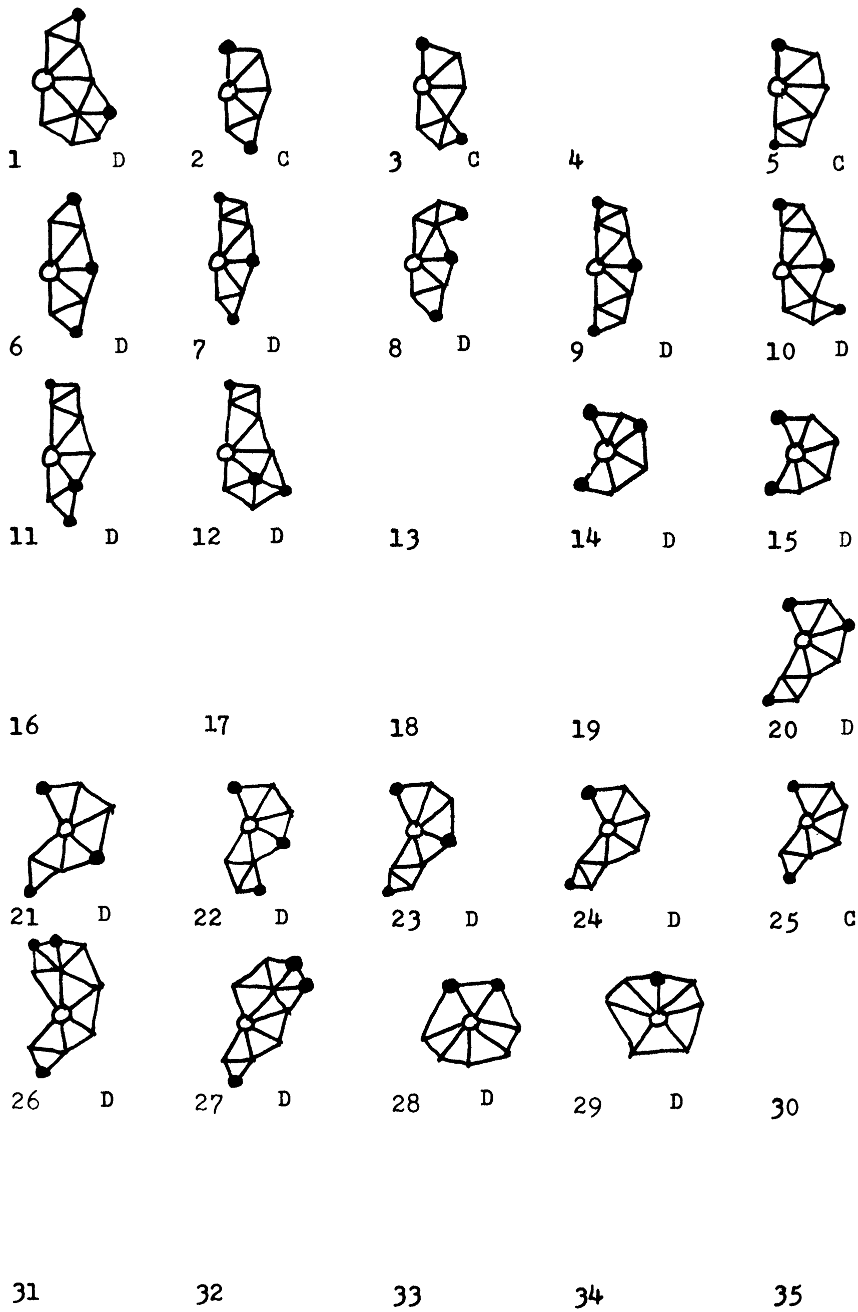
1 D
2 C
3 C
4
5 C
6 D
7 D
8 D
9 D
10 D
11 D
12 D
13
14 D
15 D
16
17
18
19
20 D
21 D
22 D
23 D
24 D
25 C
26 D
27 D
28 D
29 D
30
31
32
33
34
35

Figure 4

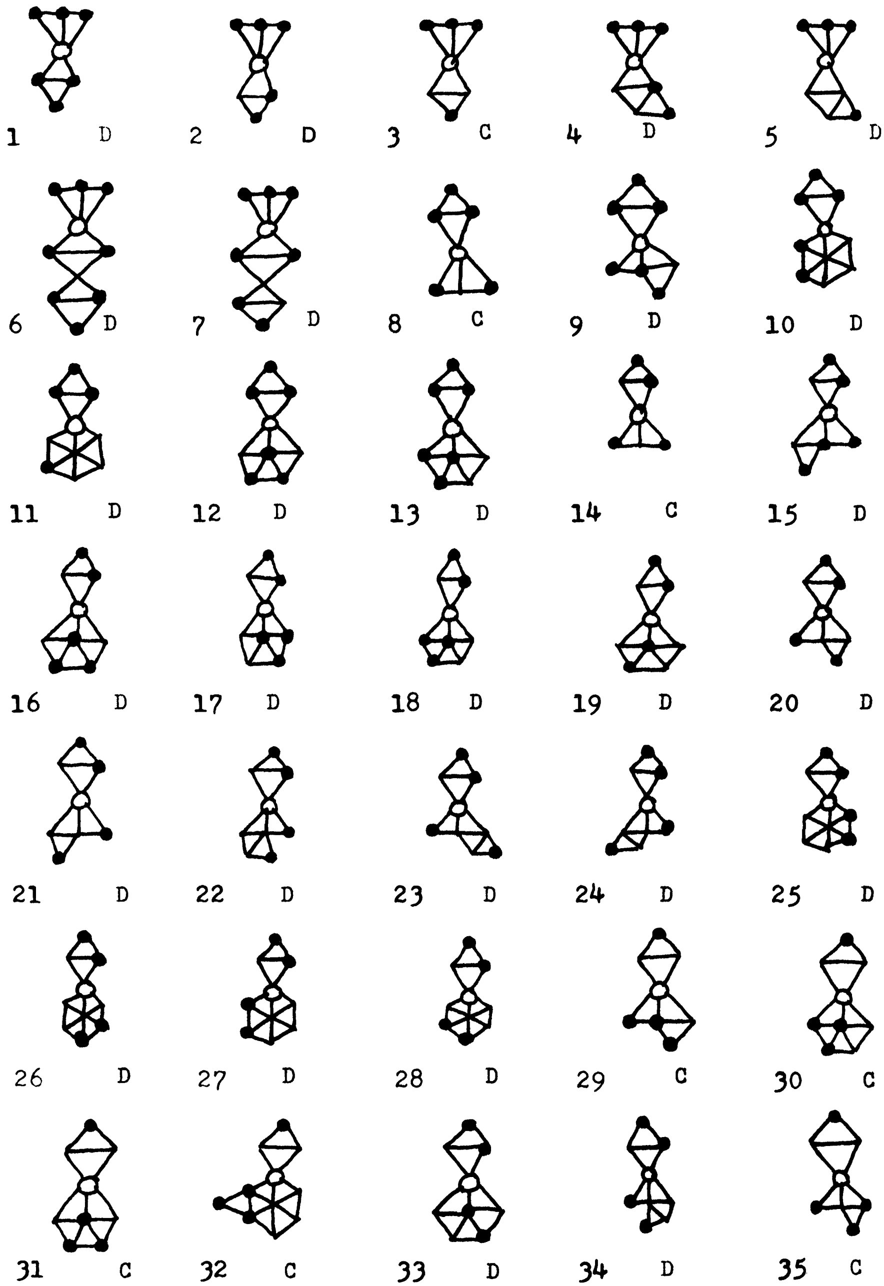
1 D
2 D
3 C
4 D
5 D
6 D
7 D
8 C
9 D
10 D
11 D
12 D
13 D
14 C
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 C
30 C
31 C
32 C
33 D
34 D
35 C

Figure 5

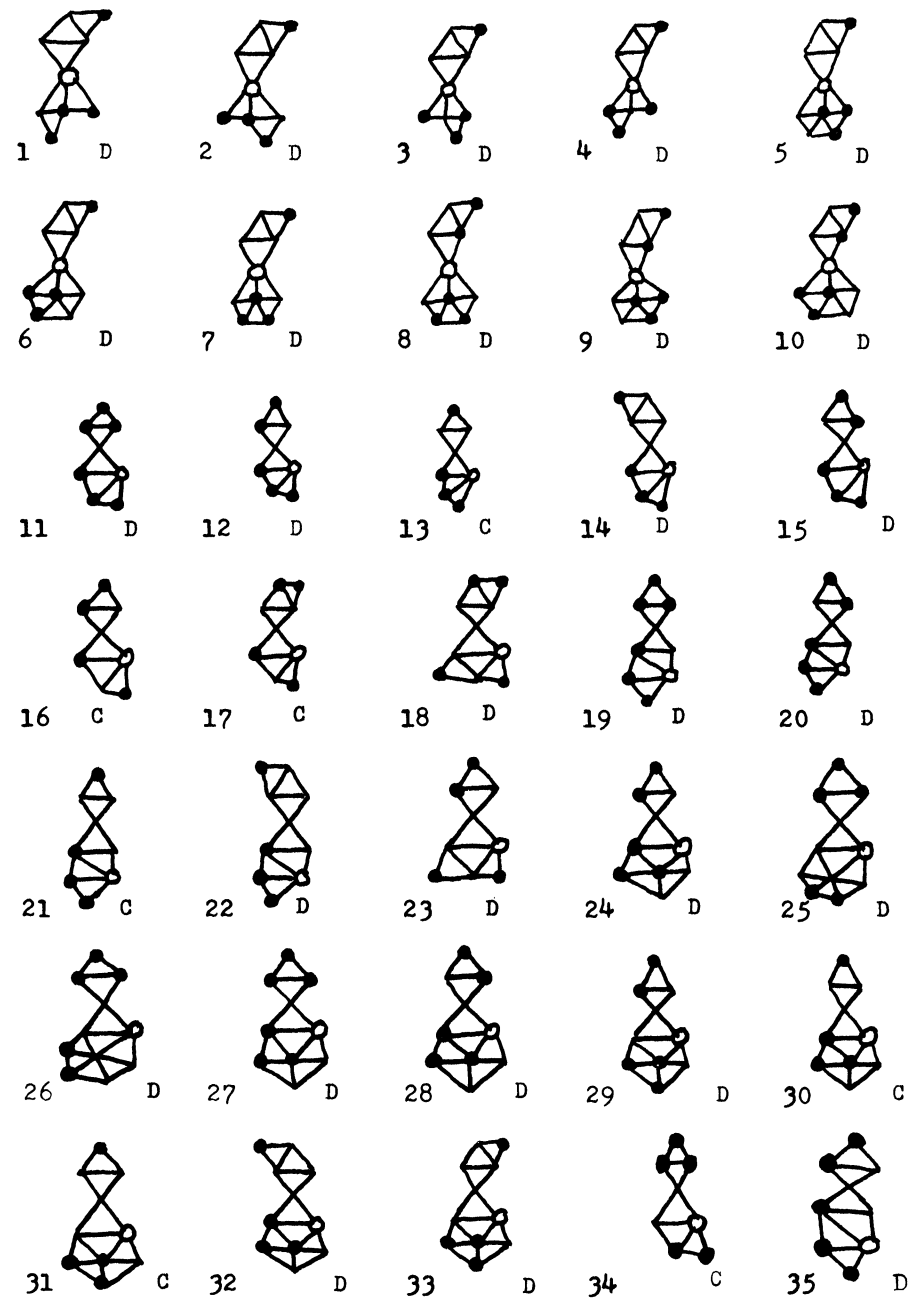
1 D
2 D
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 D
13 C
14 D
15 D
16 C
17 C
18 D
19 D
20 D
21 C
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 C
31 C
32 D
33 D
34 C
35 D

Figure 6

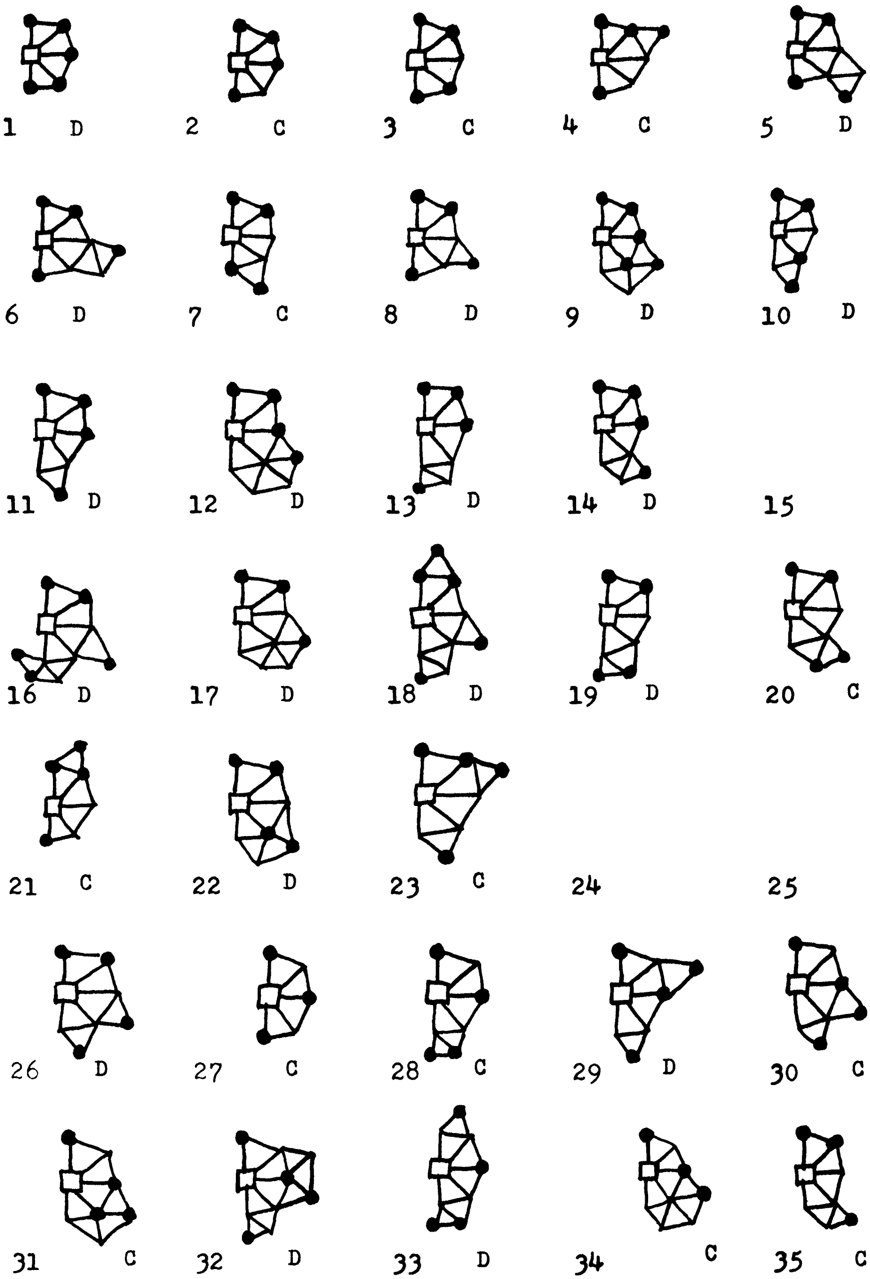
1 D
2 C
3 C
4 C
5 D
6 D
7 C
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15
16 D
17 D
18 D
19 D
20 C
21 C
22 D
23 C
24
25
26 D
27 C
28 C
29 D
30 C
31 C
32 D
33 D
34 C
35 C

Figure 7

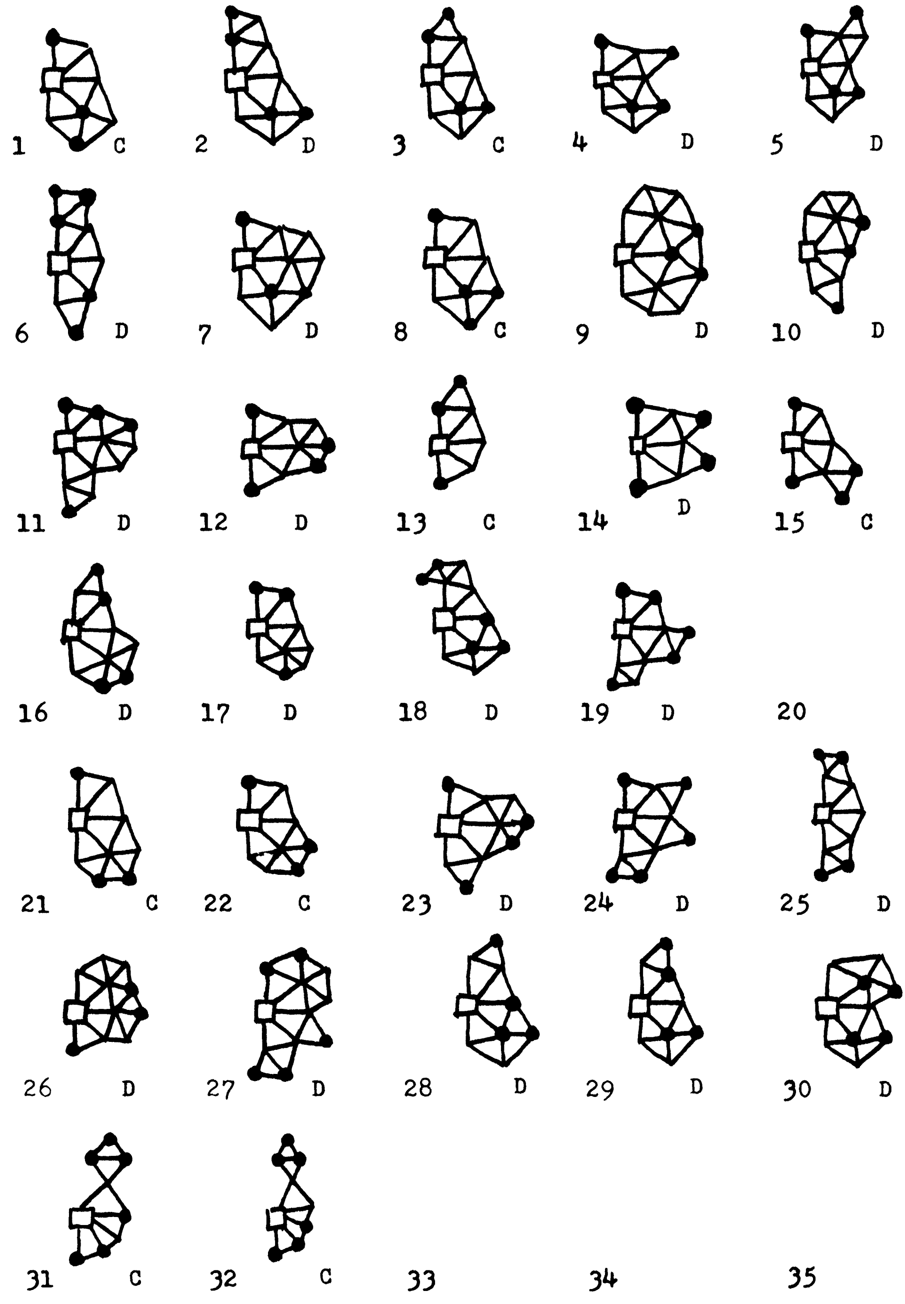
1 C
2 D
3 C
4 D
5 D
6 D
7 D
8 C
9 D
10 D
11 D
12 D
13 C
14 D
15 C
16 D
17 D
18 D
19 D
20
21 C
22 C
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 C
32 C
33
34
35

Figure 8

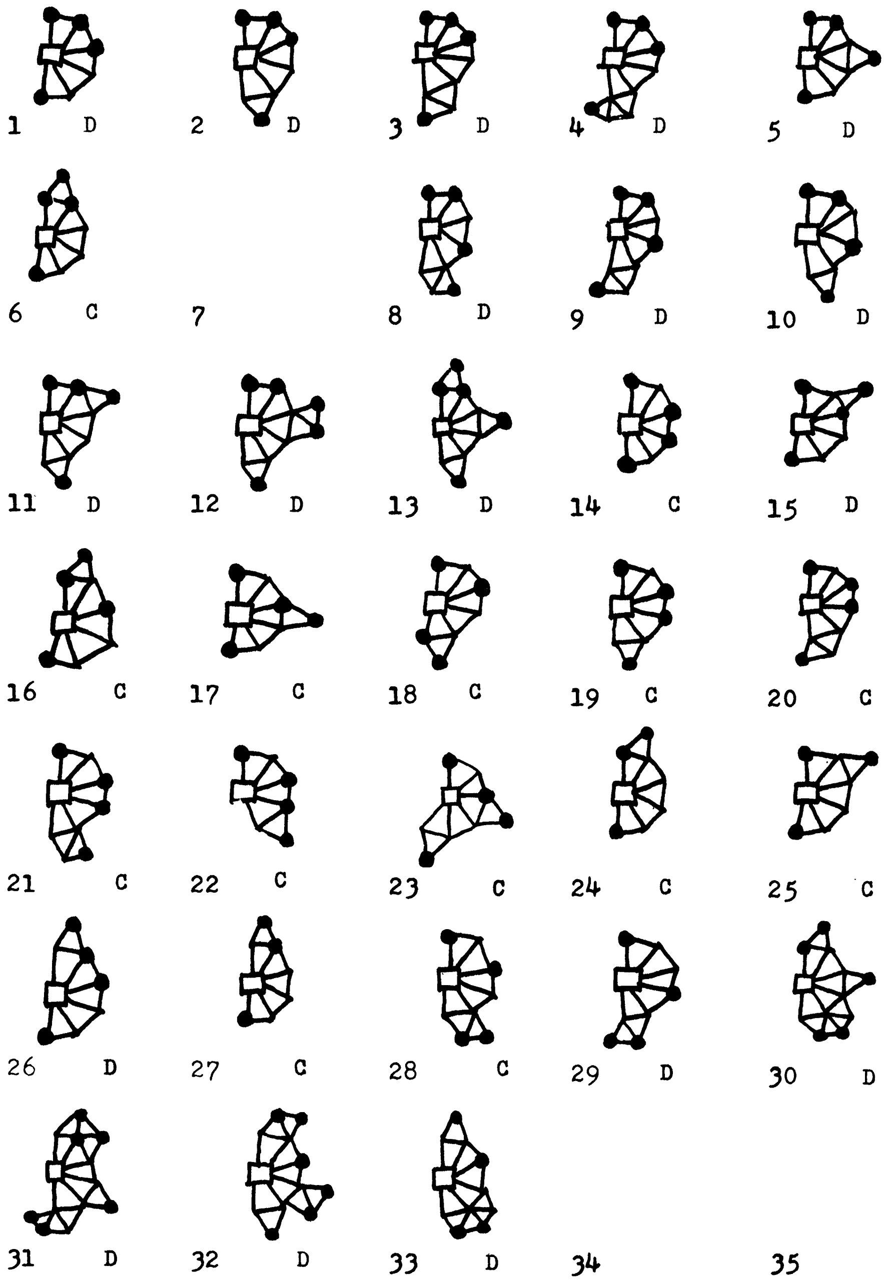
1 D
2 D
3 D
4 D
5 D
6 C
7
8 D
9 D
10 D
11 D
12 D
13 D
14 C
15 D
16 C
17 C
18 C
19 C
20 C
21 C
22 C
23 C
24 C
25 C
26 D
27 C
28 C
29 D
30 D
31 D
32 D
33 D
34
35

Figure 9

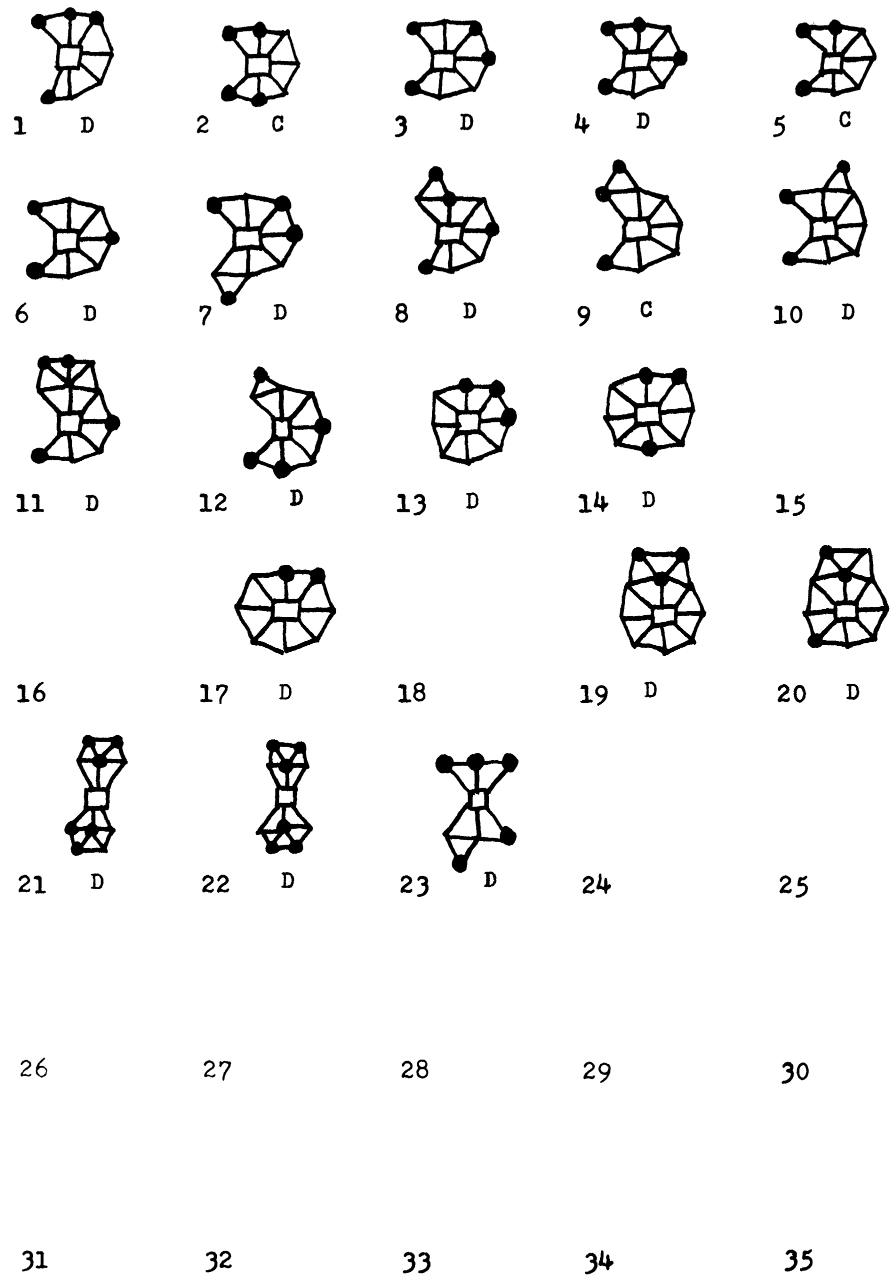
1 D
2 C
3 D
4 D
5 C
6 D
7 D
8 D
9 C
10 D
11 D
12 D
13 D
14 D
15
16
17 D
18
19 D
20 D
21 D
22 D
23 D
24
25
26
27
28
29
30
31
32
33
34
35

Figure 10

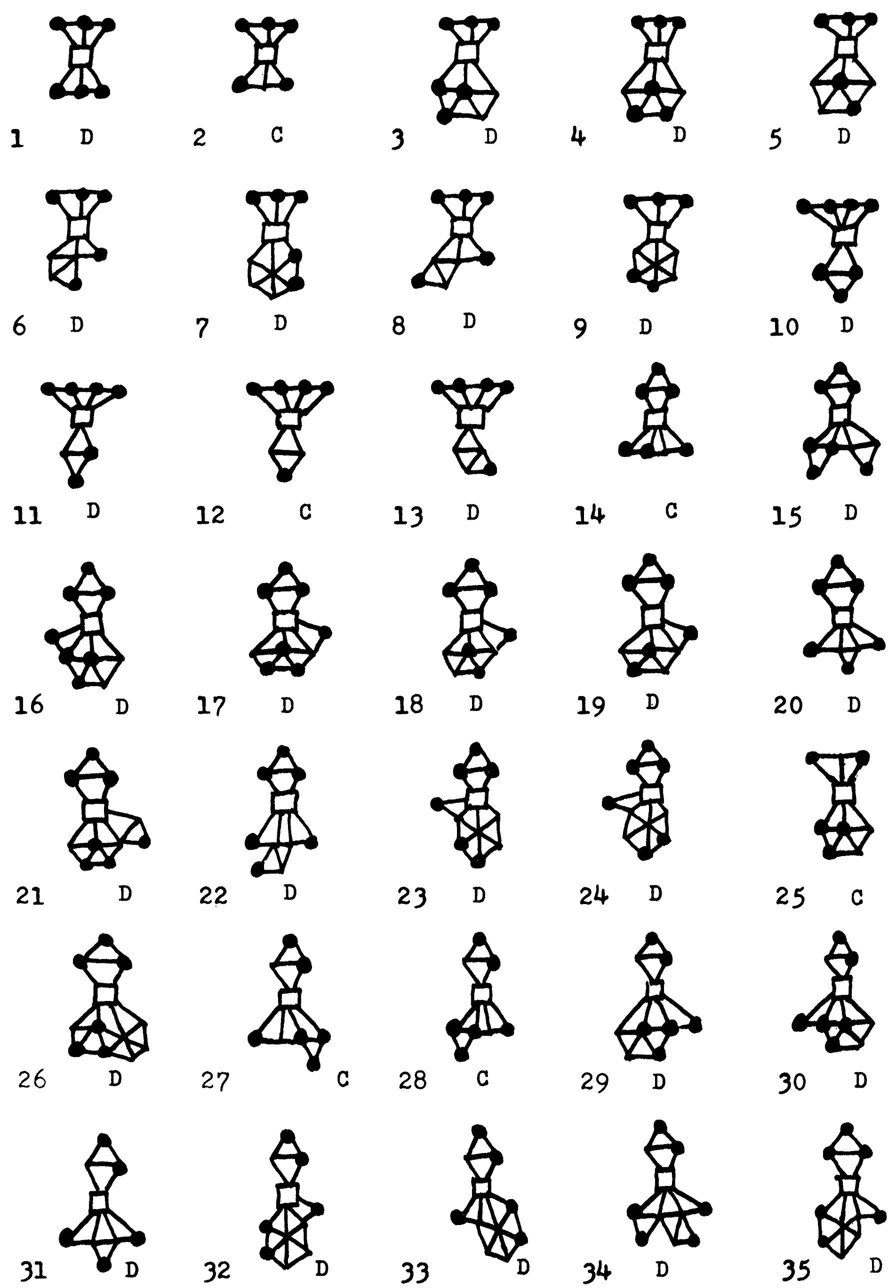
1 D
2 C
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 C
13 D
14 C
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 C
26 D
27 C
28 C
29 D
30 D
31 D
32 D
33 D
34 D
35 D

Figure 11

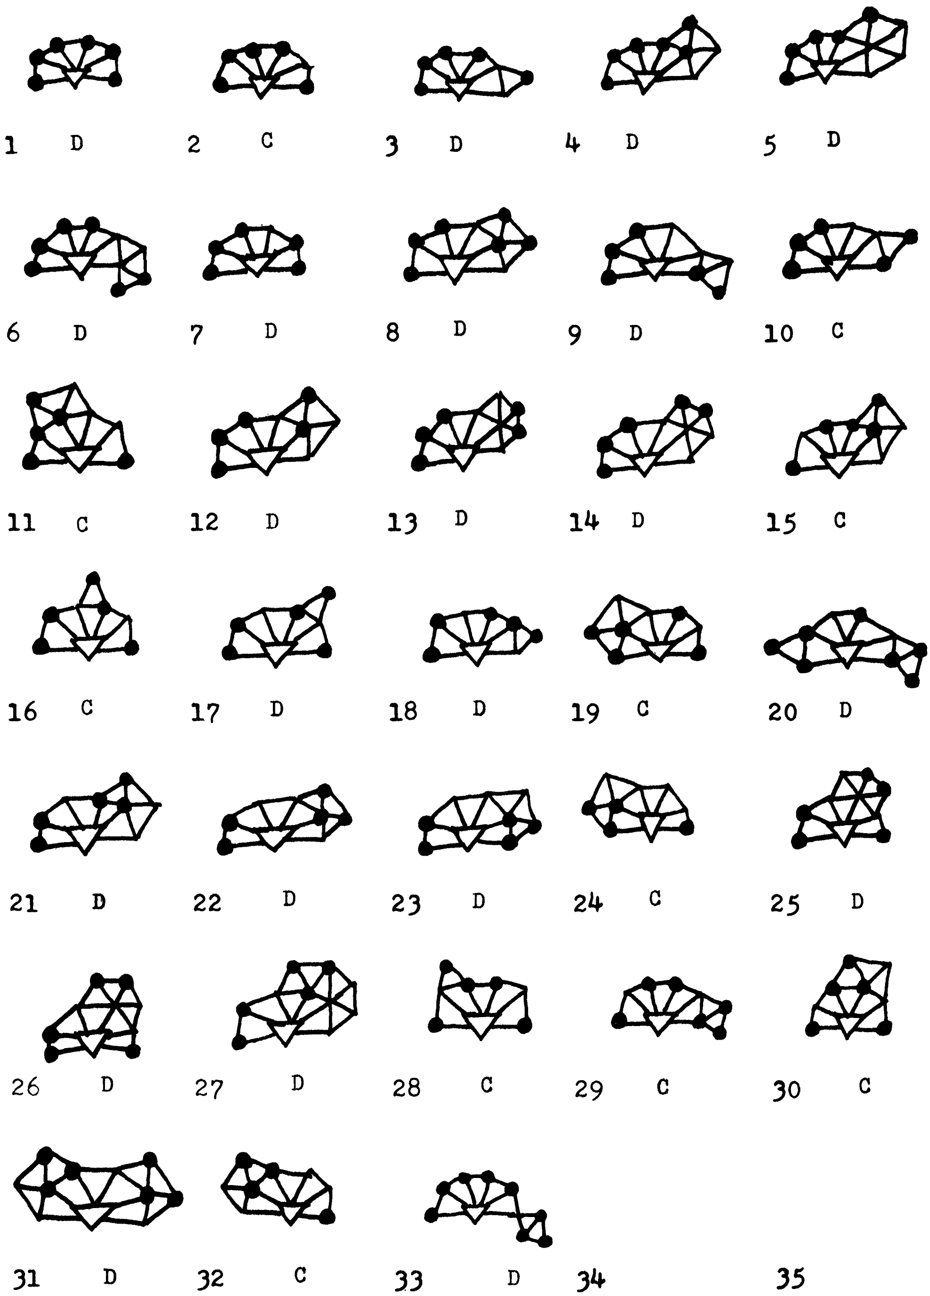
1 D
2 C
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 C
11 C
12 D
13 D
14 D
15 C
16 C
17 D
18 D
19 C
20 D
21 D
22 D
23 D
24 C
25 D
26 D
27 D
28 C
29 C
30 C
31 D
32 C
33 D
34
35

Figure 12

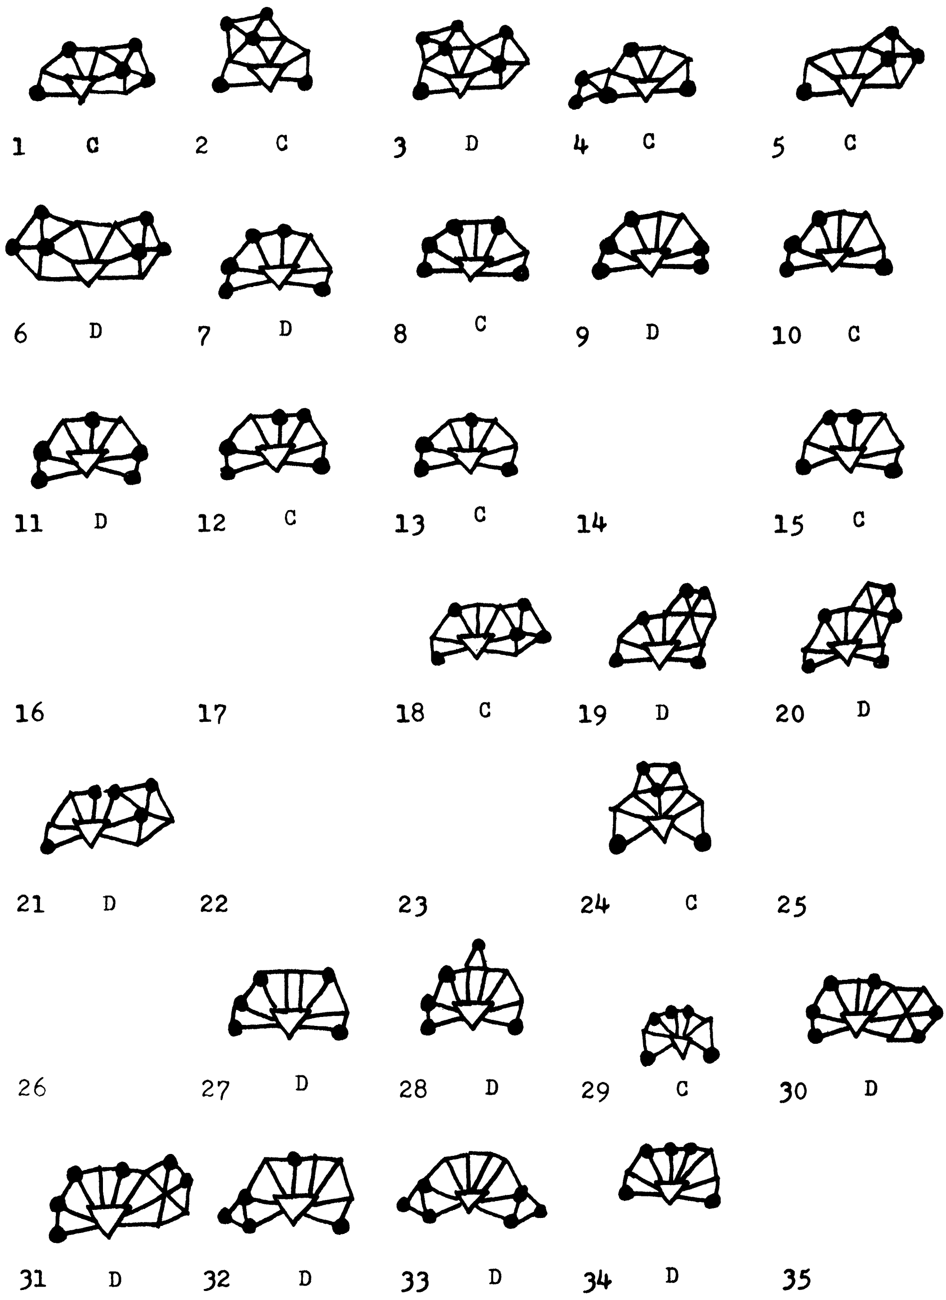
1 C
2 C
3 D
4 C
5 C
6 D
7 D
8 C
9 D
10 C
11 D
12 C
13 C
14
15 C
16
17
18 C
19 D
20 D
21 D
22
23
24 C
25
26
27 D
28 D
29 C
30 D
31 D
32 D
33 D
34 D
35

Figure 13

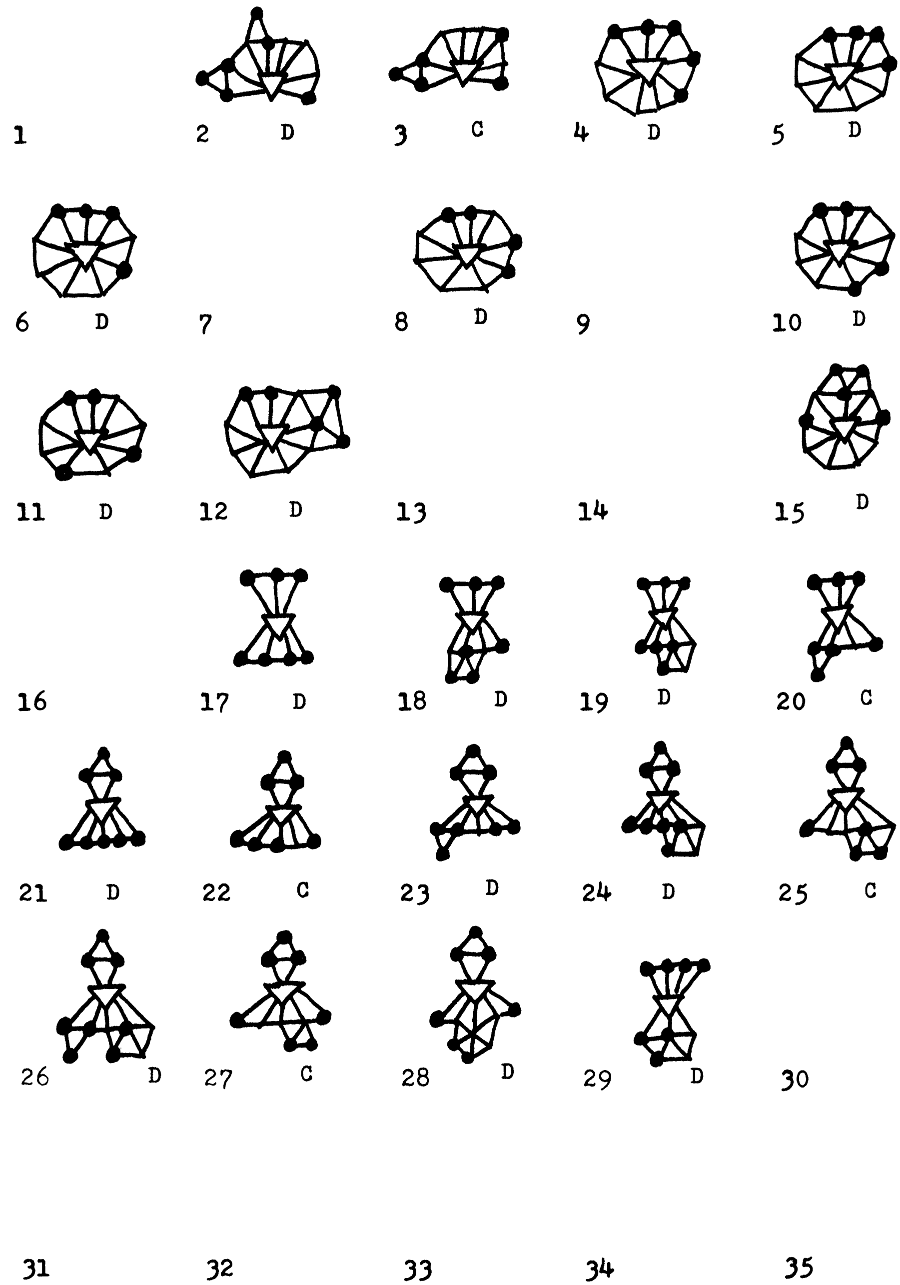
1
2 D
3 C
4 D
5 D
6 D
7
8 D
9
10 D
11 D
12 D
13
14
15 D
16
17 D
18 D
19 D
20 C
21 D
22 C
23 D
24 D
25 C
26 D
27 C
28 D
29 D
30
31
32
33
34
35

Figure 14

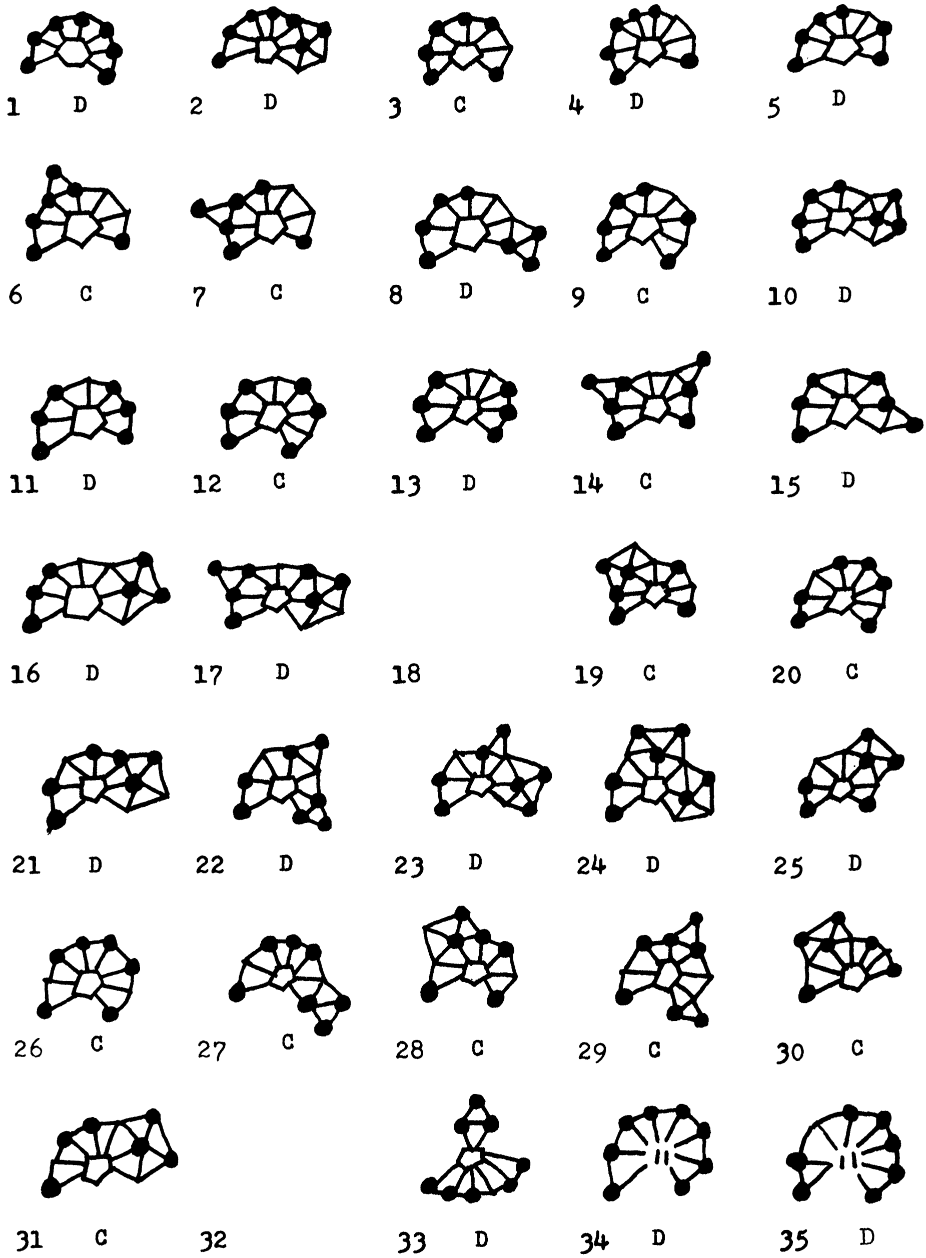
1 D
2 D
3 C
4 D
5 D
6 C
7 C
8 D
9 C
10 D
11 D
12 C
13 D
14 C
15 D
16 D
17 D
18
19 C
20 C
21 D
22 D
23 D
24 D
25 D
26 C
27 C
28 C
29 C
30 C
31 C
32
33 D
34 D
35 D

Figure 15

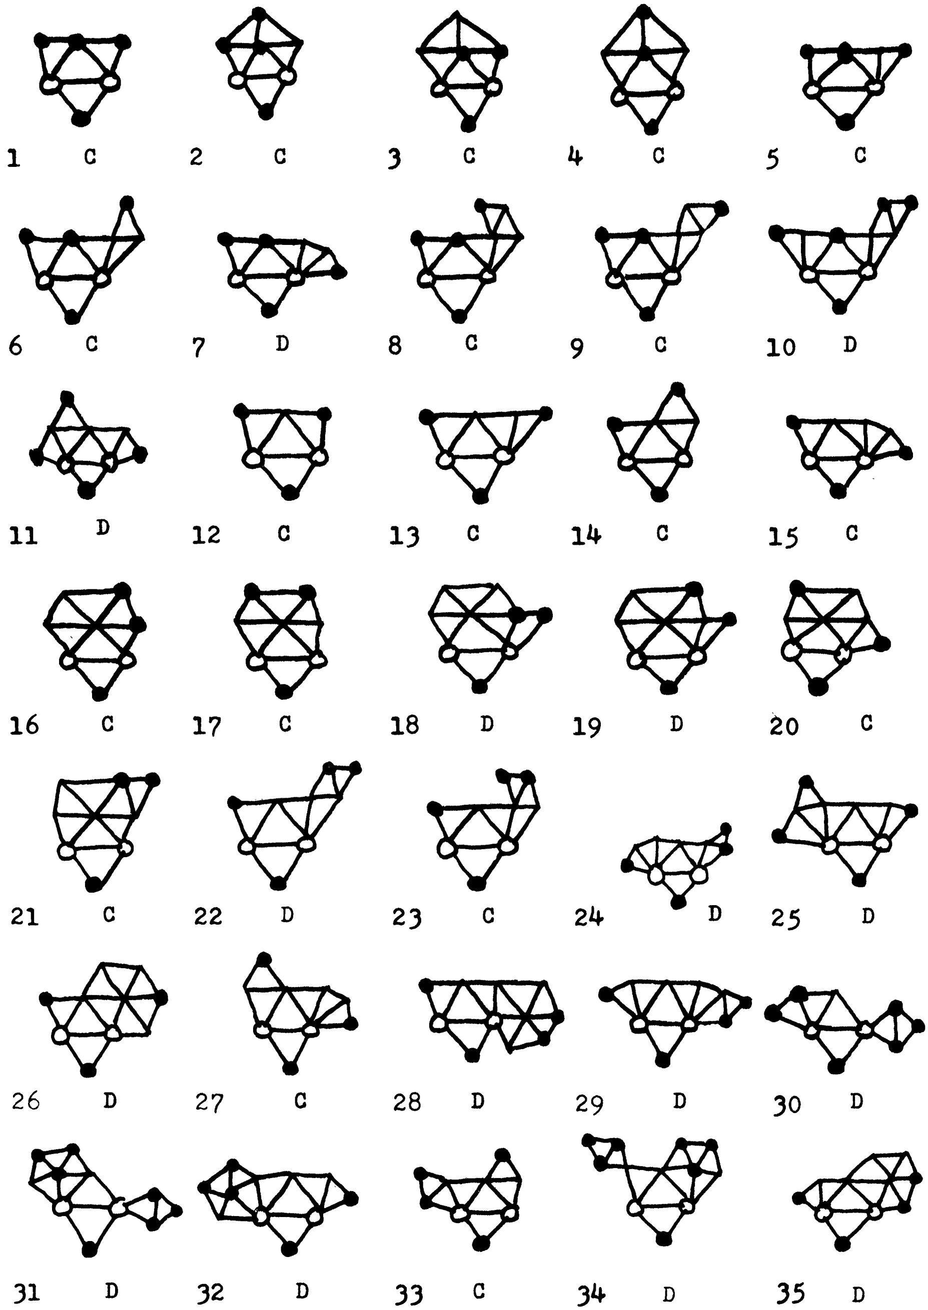
1 C
2 C
3 C
4 C
5 C
6 C
7 D
8 C
9 C
10 D
11 D
12 C
13 C
14 C
15 C
16 C
17 C
18 D
19 D
20 C
21 C
22 D
23 C
24 D
25 D
26 D
27 C
28 D
29 D
30 D
31 D
32 D
33 C
34 D
35 D

Figure 16

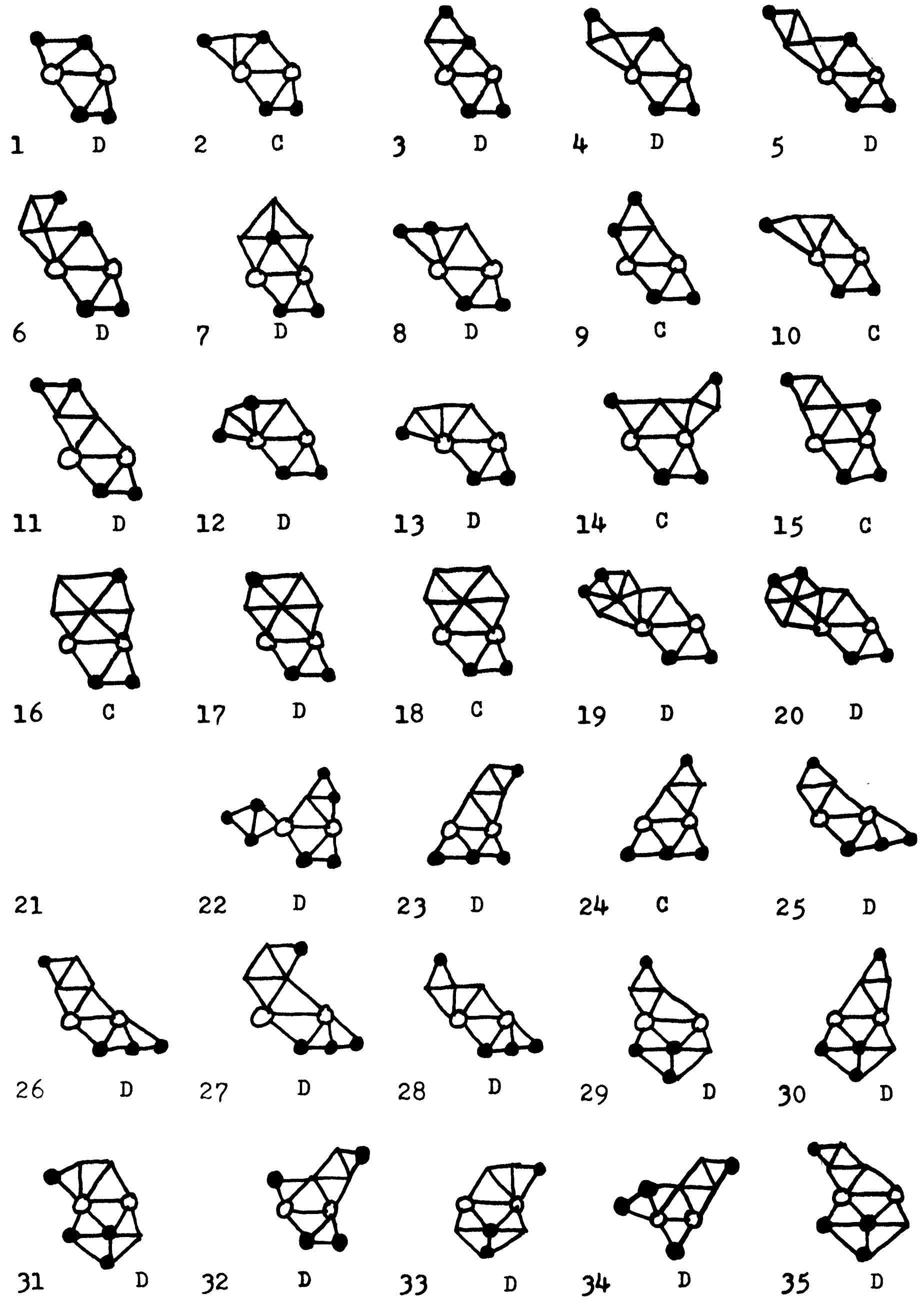
1 D
2 C
3 D
4 D
5 D
6 D
7 D
8 D
9 C
10 C
11 D
12 D
13 D
14 C
15 C
16 C
17 D
18 C
19 D
20 D
21
22 D
23 D
24 C
25 D
26 D
27 D
28 D
29 D
30 D
31 D
32 D
33 D
34 D
35 D

Figure 17

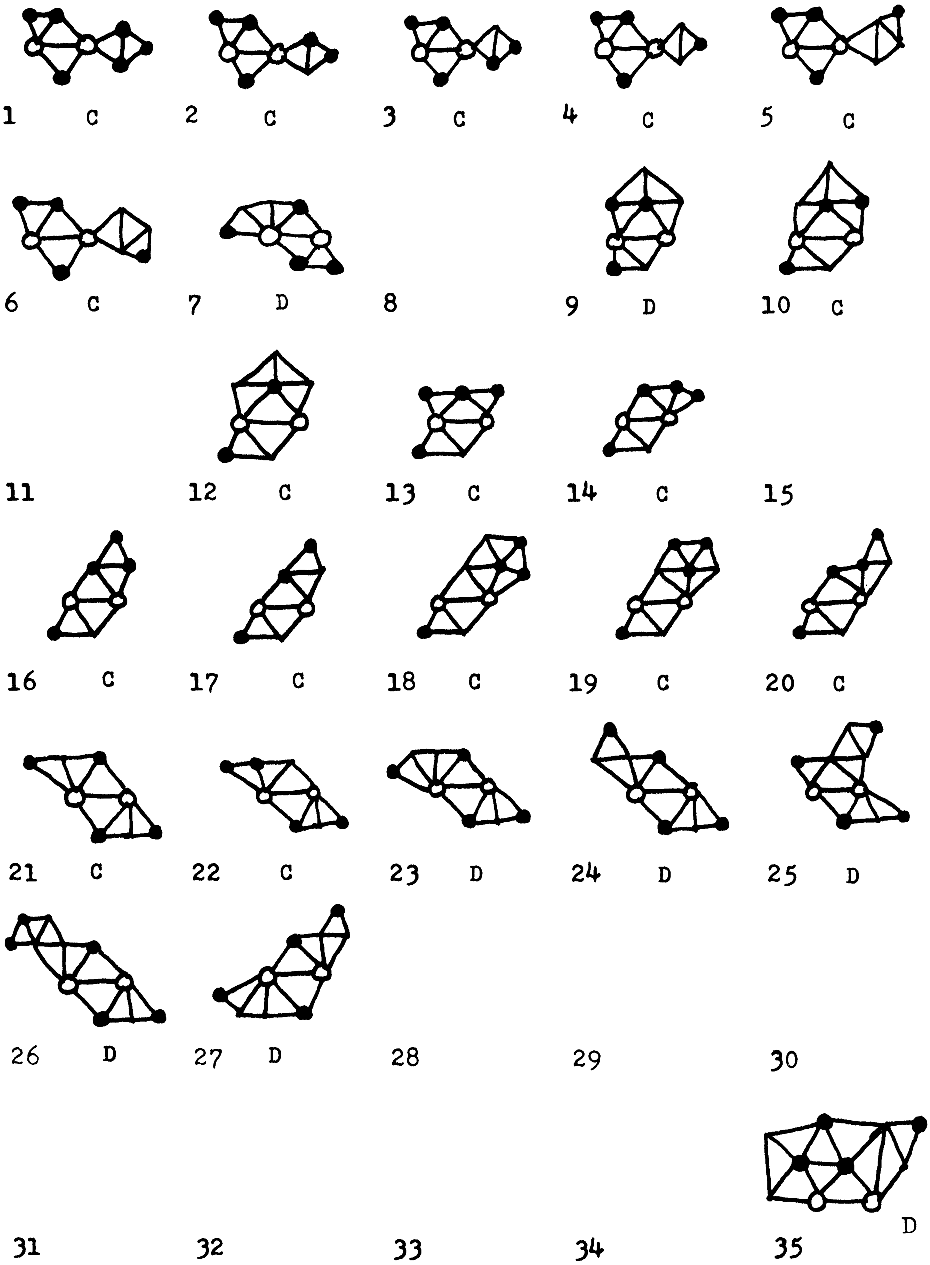
1 C
2 C
3 C
4 C
5 C
6 C
7 D
8
9 D
10 C
11
12 C
13 C
14 C
15
16 C
17 C
18 C
19 C
20 C
21 C
22 C
23 D
24 D
25 D
26 D
27 D
28
29
30
31
32
33
34
35 D

Figure 18

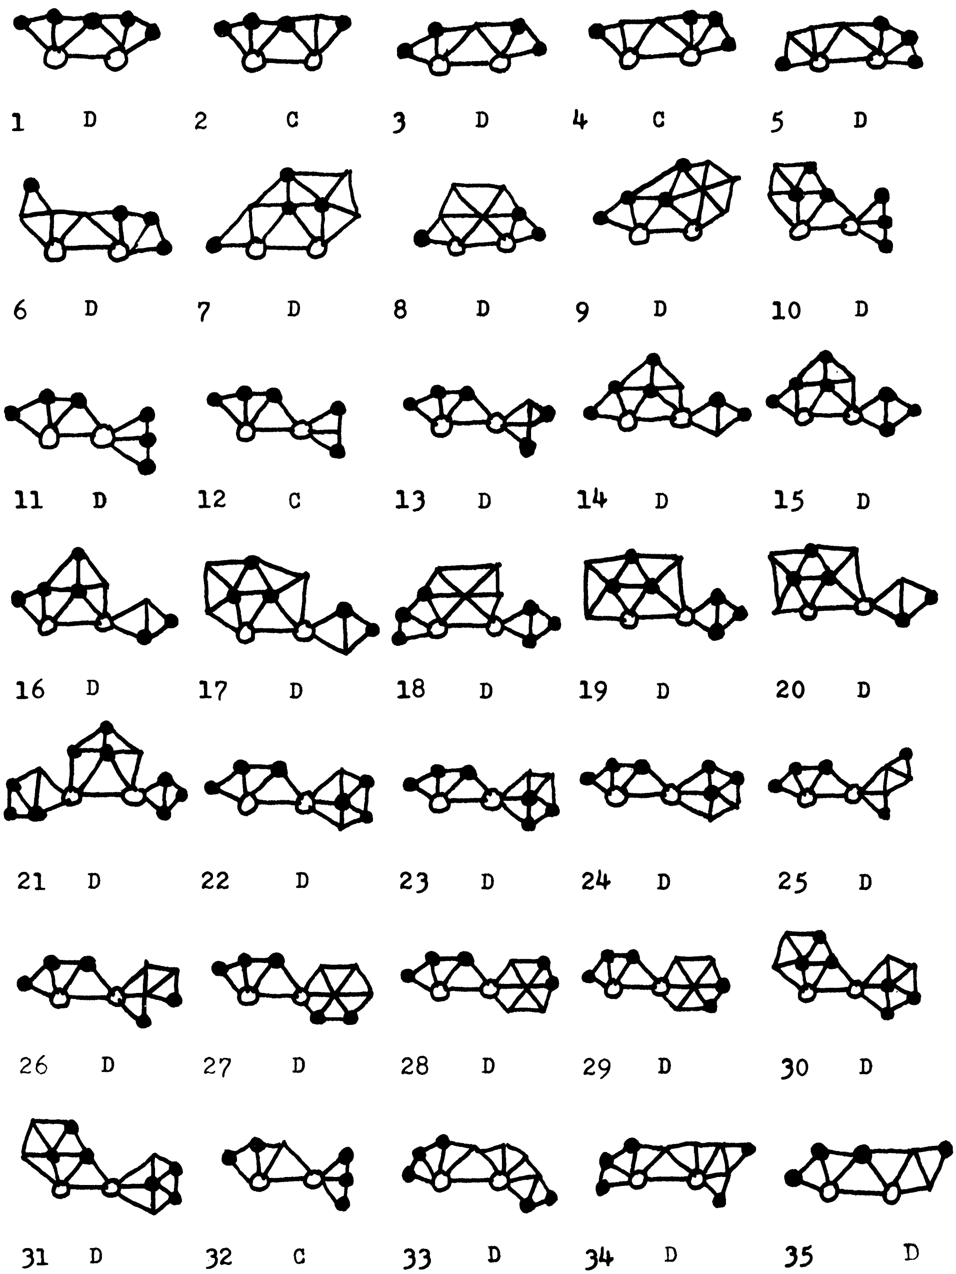
1 D
2 C
3 D
4 C
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 C
13 D
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 D
32 C
33 D
34 D
35 D

Figure 19

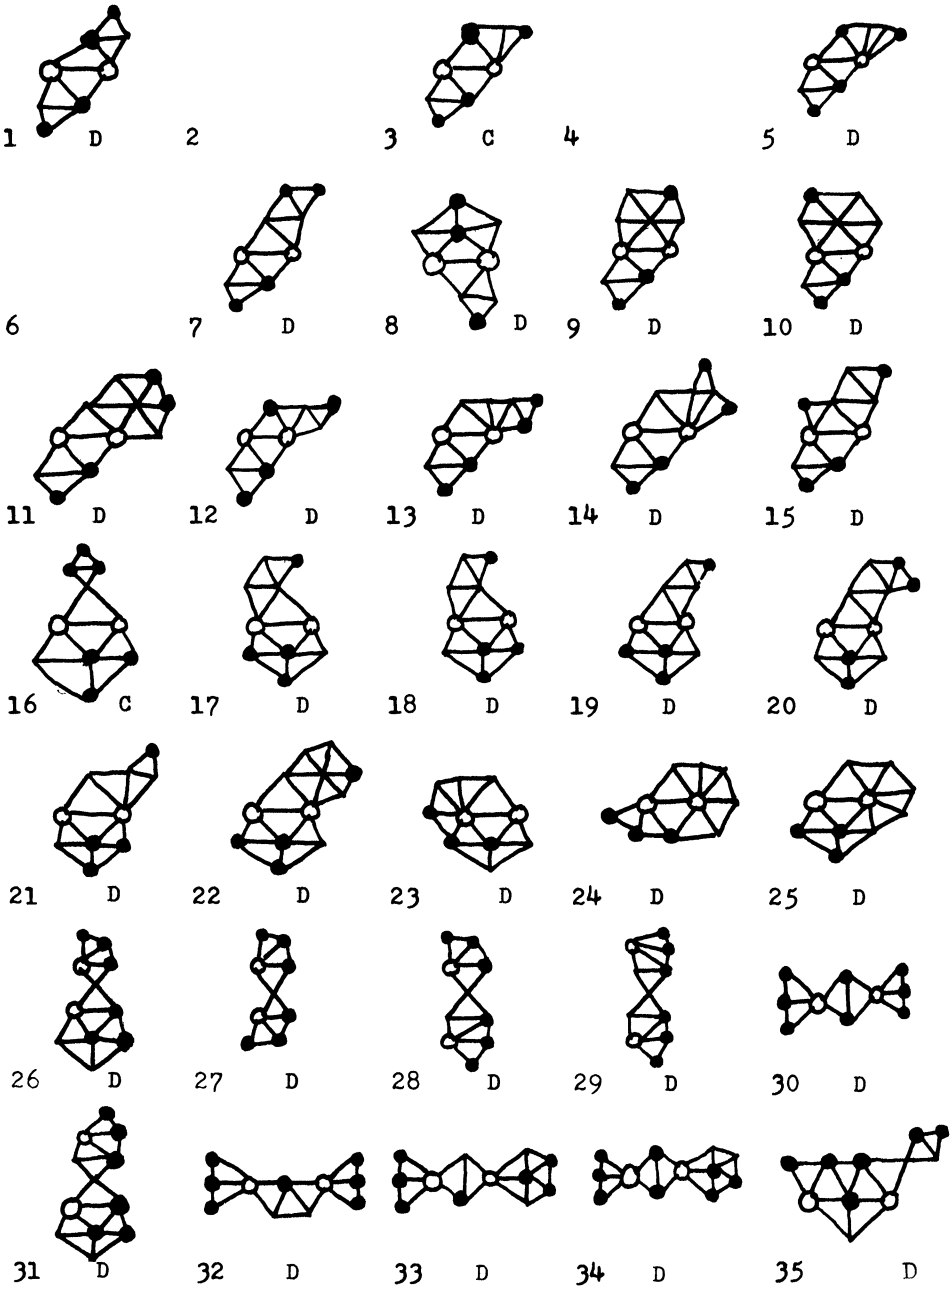
1 D
2
3 C
4
5 D
6
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 C
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 D
32 D
33 D
34 D
35 D

Figure 20

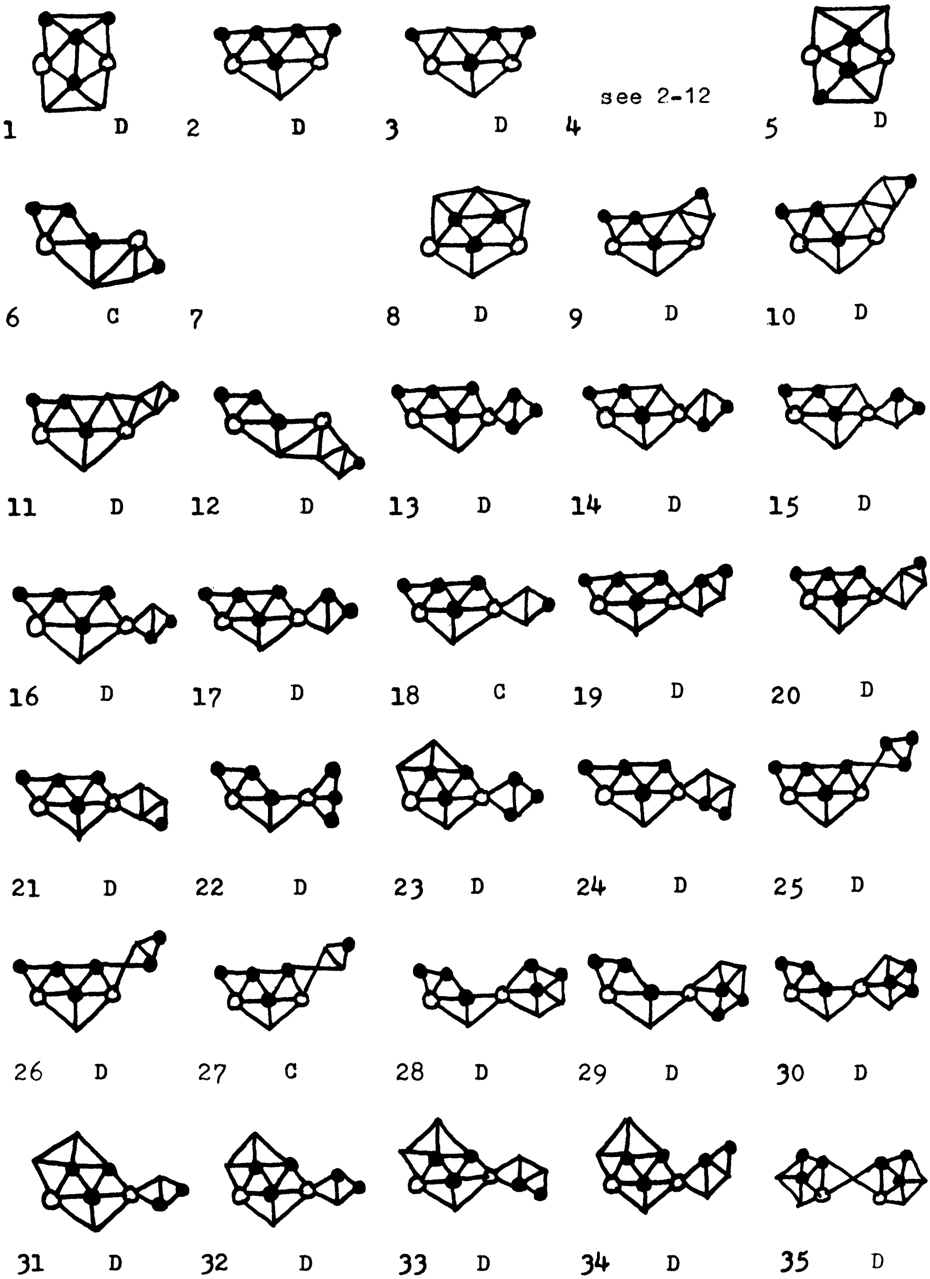
1 D
2 D
3 D
4 see 2-12
5 D
6 C
7
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17 D
18 C
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 C
28 D
29 D
30 D
31 D
32 D
33 D
34 D
35 D

Figure 21

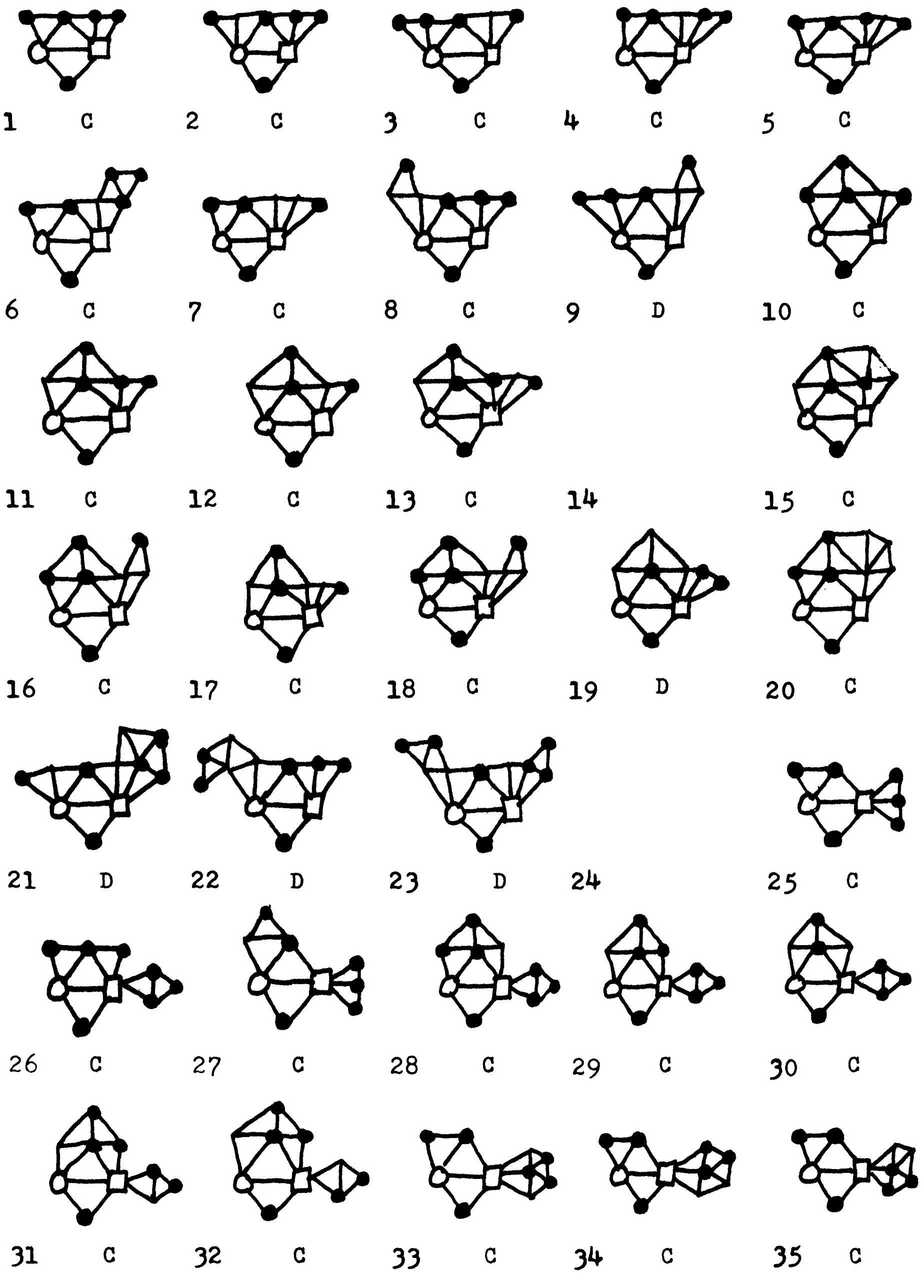
1 C
2 C
3 C
4 C
5 C
6 C
7 C
8 C
9 D
10 C
11 C
12 C
13 C
14
15 C
16 C
17 C
18 C
19 D
20 C
21 D
22 D
23 D
24
25 C
26 C
27 C
28 C
29 C
30 C
31 C
32 C
33 C
34 C
35 C

Figure 22

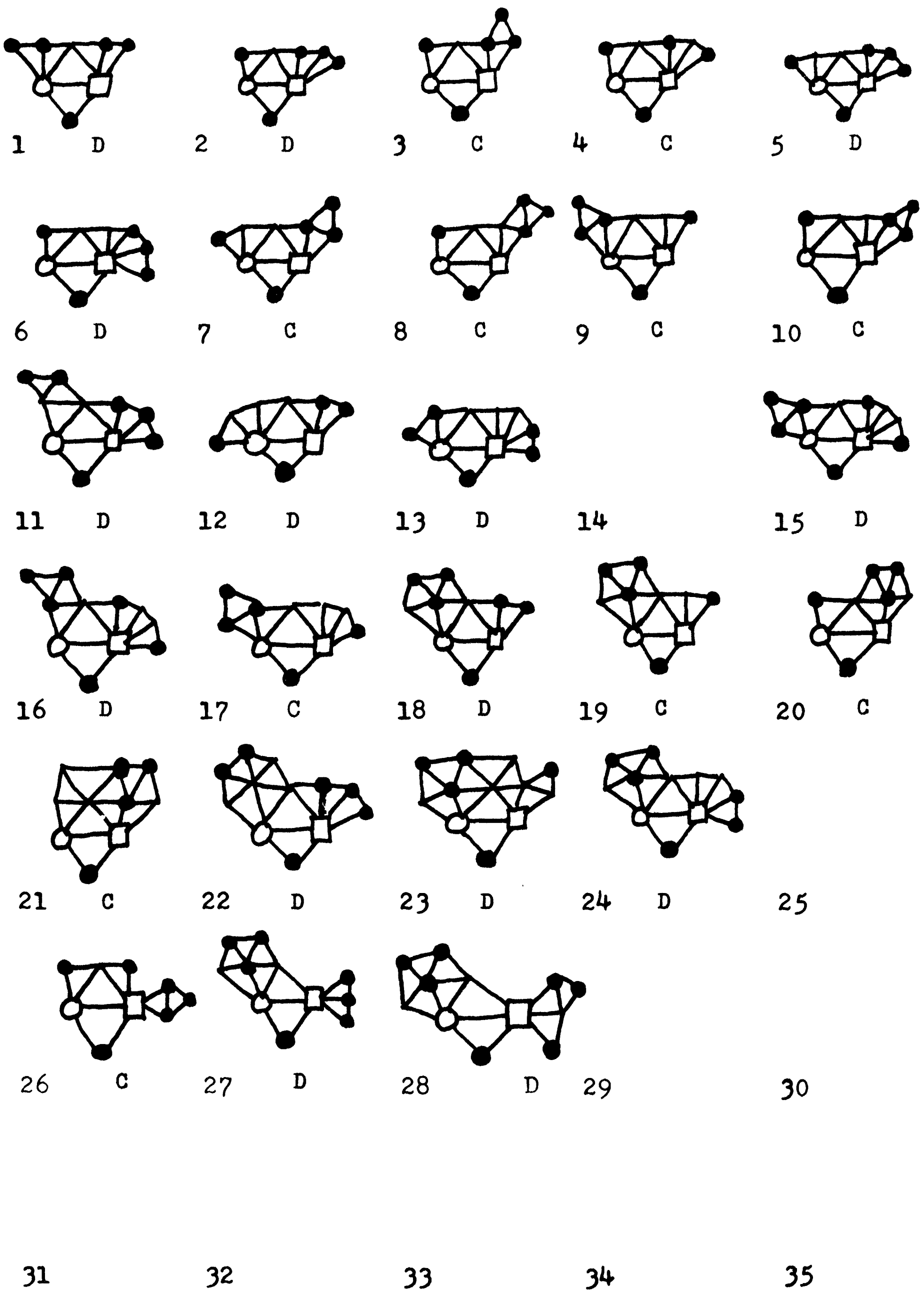
1 D
2 D
3 C
4 C
5 D
6 D
7 C
8 C
9 C
10 C
11 D
12 D
13 D
14
15 D
16 D
17 C
18 D
19 C
20 C
21 C
22 D
23 D
24 D
25
26 C
27 D
28 D
29
30
31
32
33
34
35

Figure 23

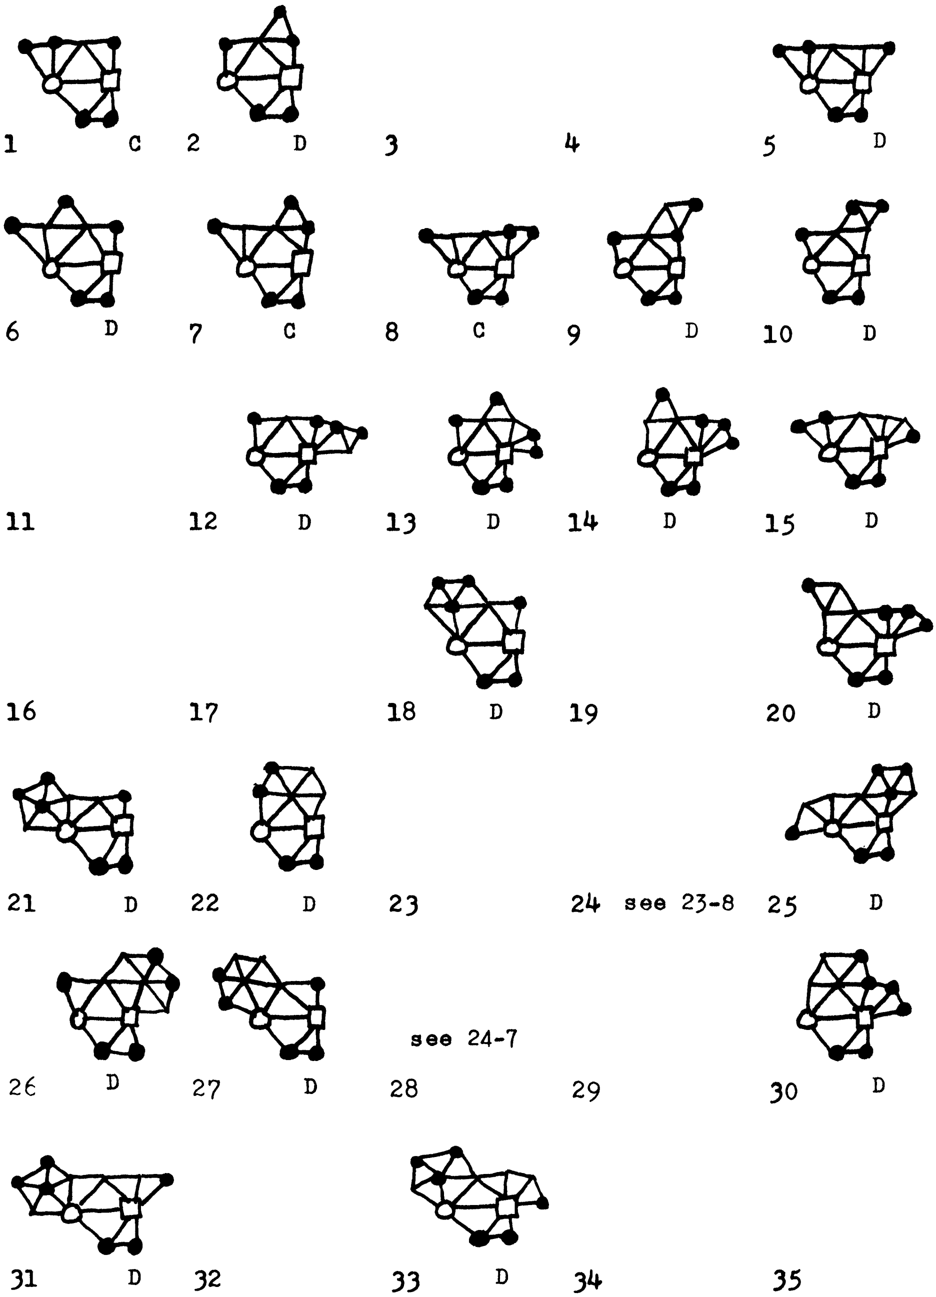
1 C
2 D
3
4
5 D
6 D
7 C
8 C
9 D
10 D
11
12 D
13 D
14 D
15 D
16
17
18 D
19
20 D
21 D
22 D
23
24 see 23-8
25 D
26 D
27 D
see 24-7
28
29
30 D
31 D
32
33 D
34
35

Figure 24

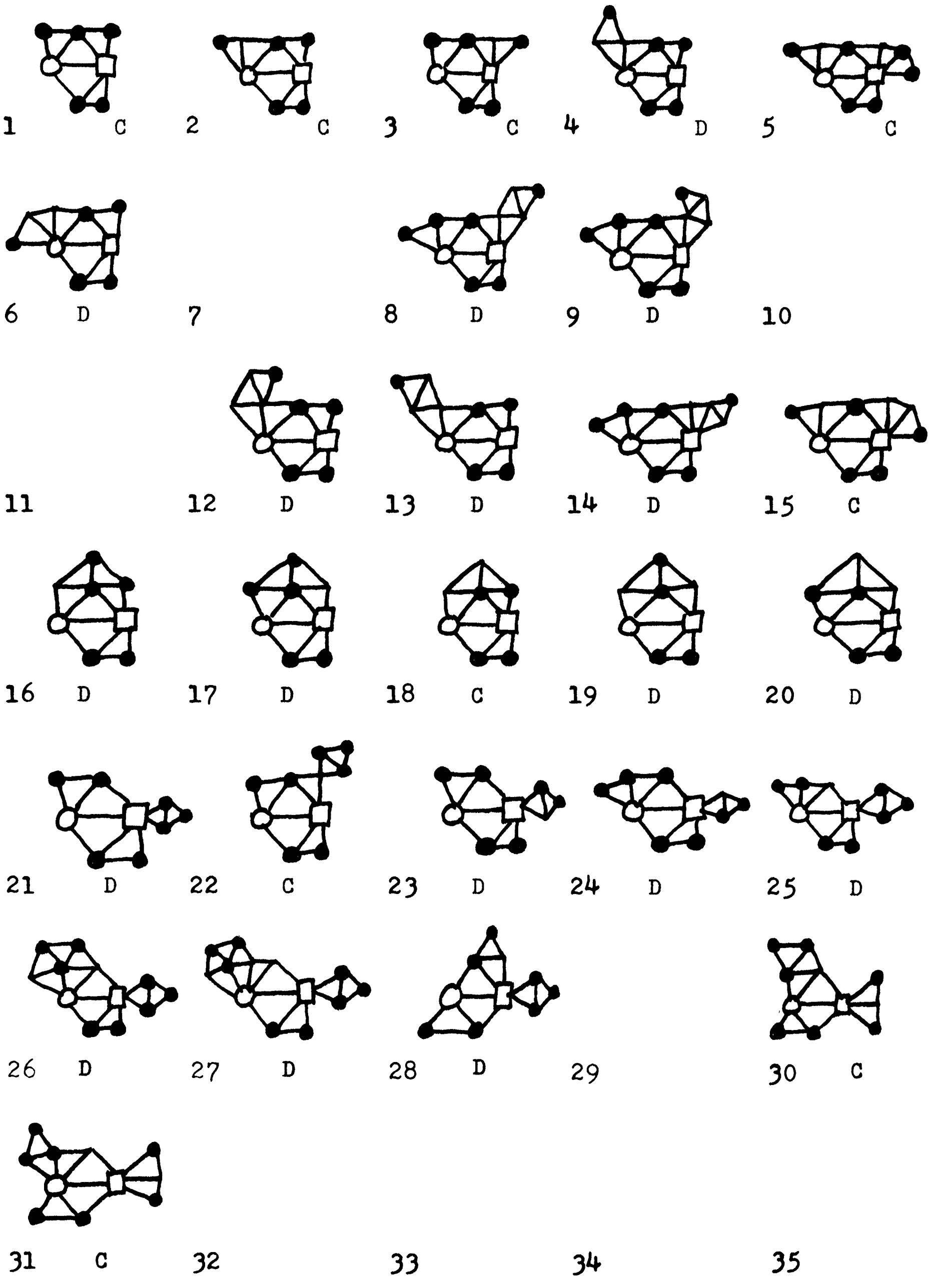
1 C
2 C
3 C
4 D
5 C
6 D
7
8 D
9 D
10
11
12 D
13 D
14 D
15 C
16 D
17 D
18 C
19 D
20 D
21 D
22 C
23 D
24 D
25 D
26 D
27 D
28 D
29
30 C
31 C
32
33
34
35

Figure 25

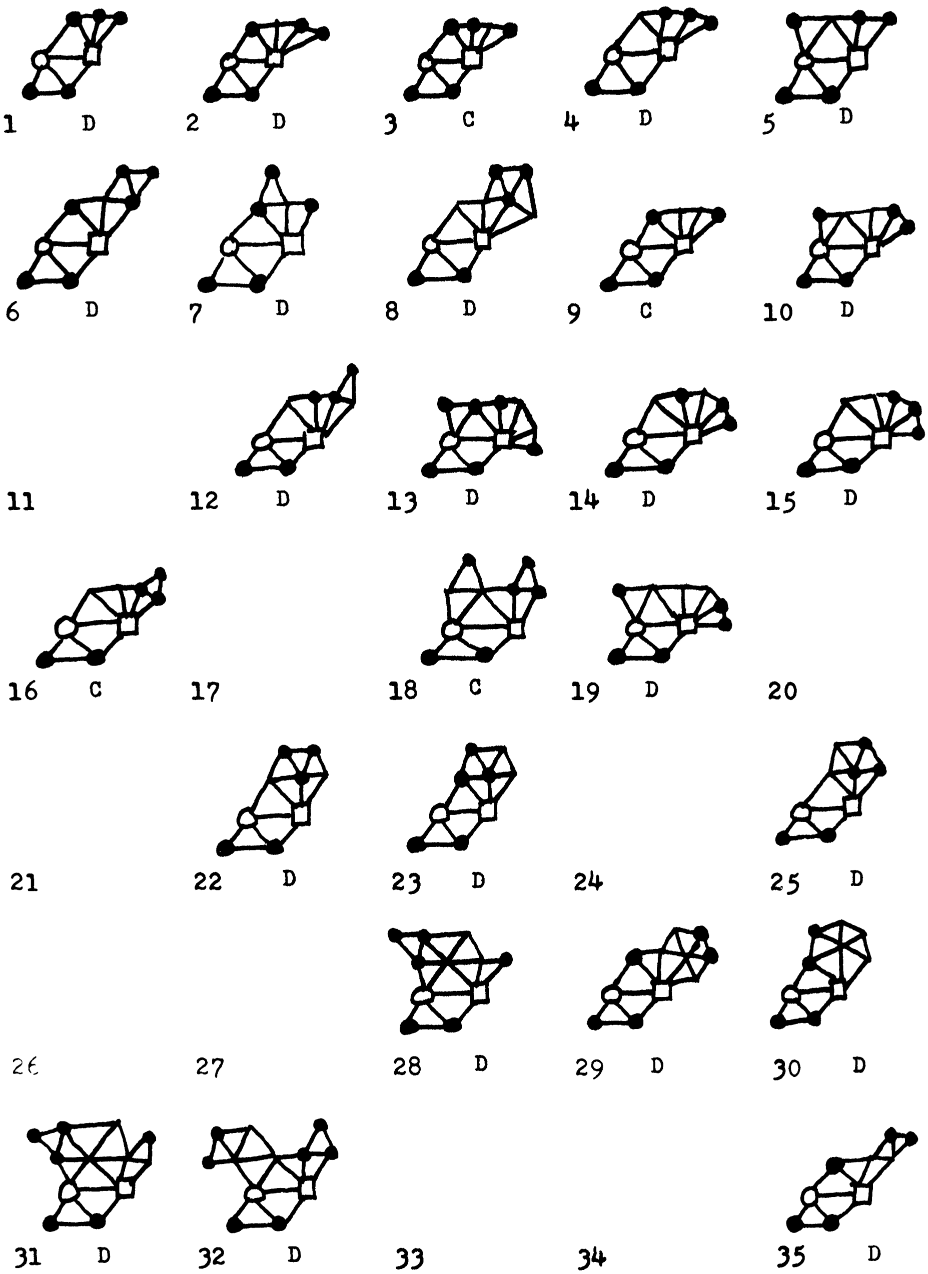
1 D
2 D
3 C
4 D
5 D
6 D
7 D
8 D
9 C
10 D
11
12 D
13 D
14 D
15 D
16 C
17
18 C
19 D
20
21
22 D
23 D
24
25 D
26
27
28 D
29 D
30 D
31 D
32 D
33
34
35 D

Figure 26

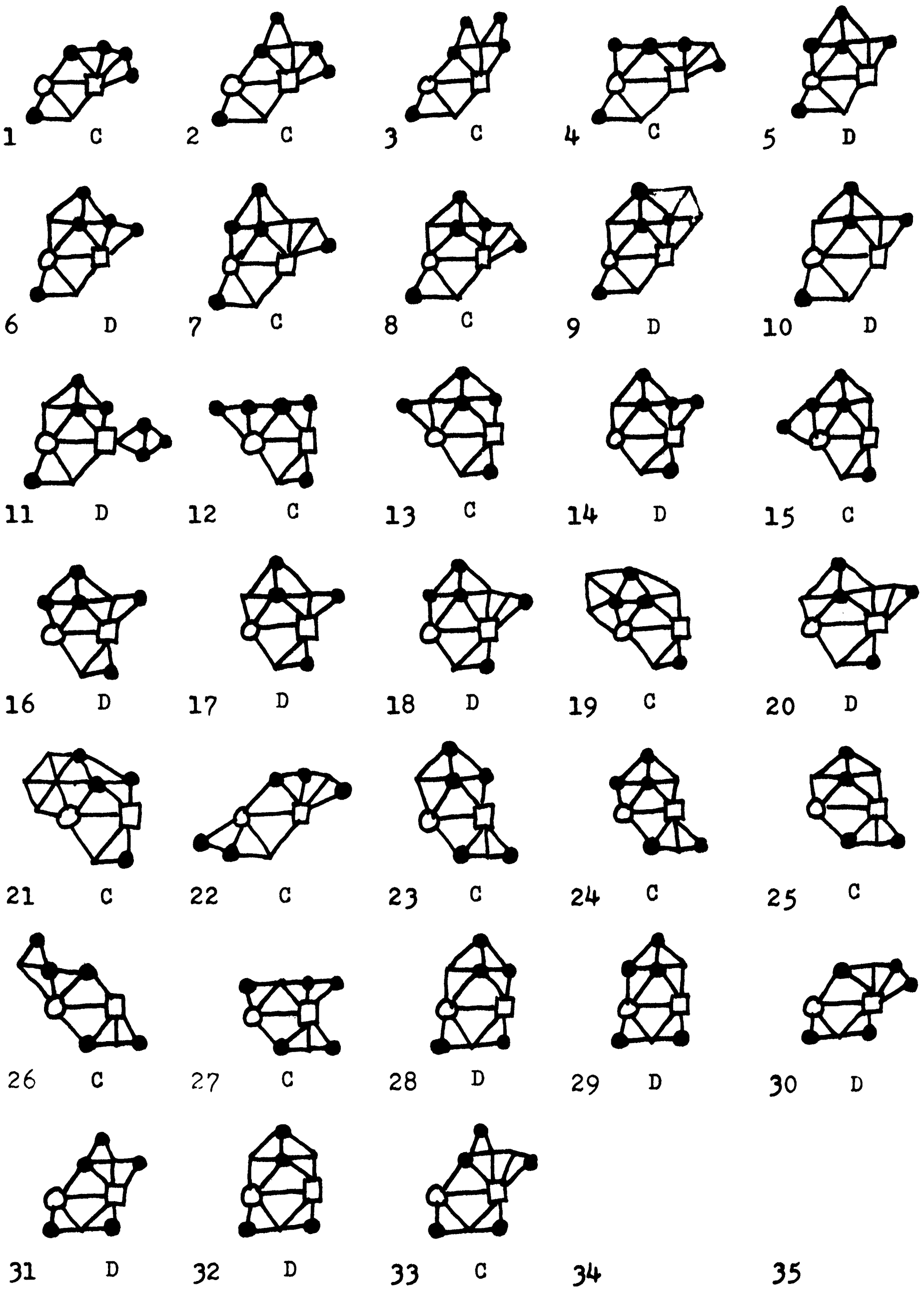
1 C
2 C
3 C
4 C
5 D
6 D
7 C
8 C
9 D
10 D
11 D
12 C
13 C
14 D
15 C
16 D
17 D
18 D
19 C
20 D
21 C
22 C
23 C
24 C
25 C
26 C
27 C
28 D
29 D
30 D
31 D
32 D
33 C
34
35

Figure 27

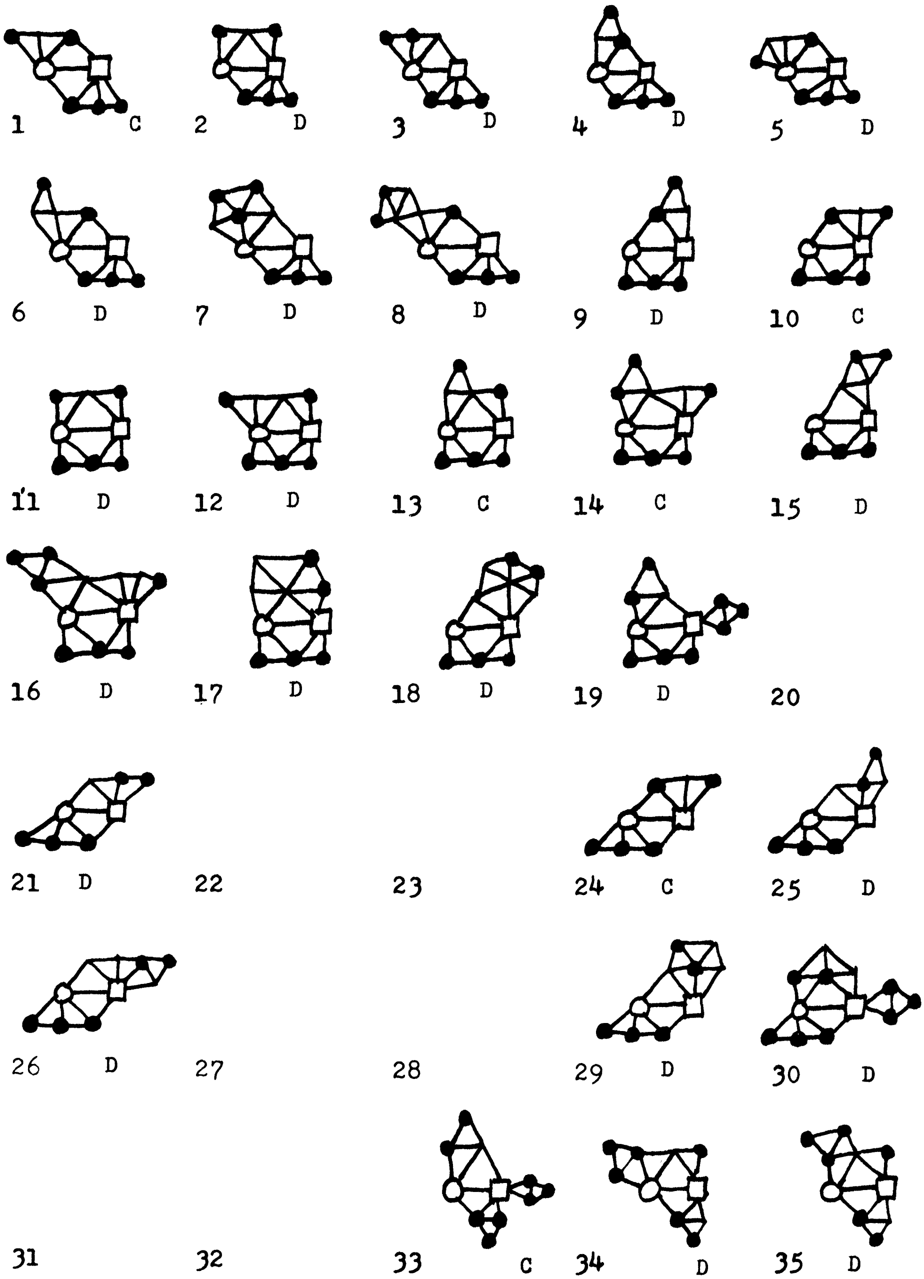
1 C
2 D
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 C
11 D
12 D
13 C
14 C
15 D
16 D
17 D
18 D
19 D
20
21 D
22
23
24 C
25 D
26 D
27
28
29 D
30 D
31
32
33 C
34 D
35 D

Figure 28

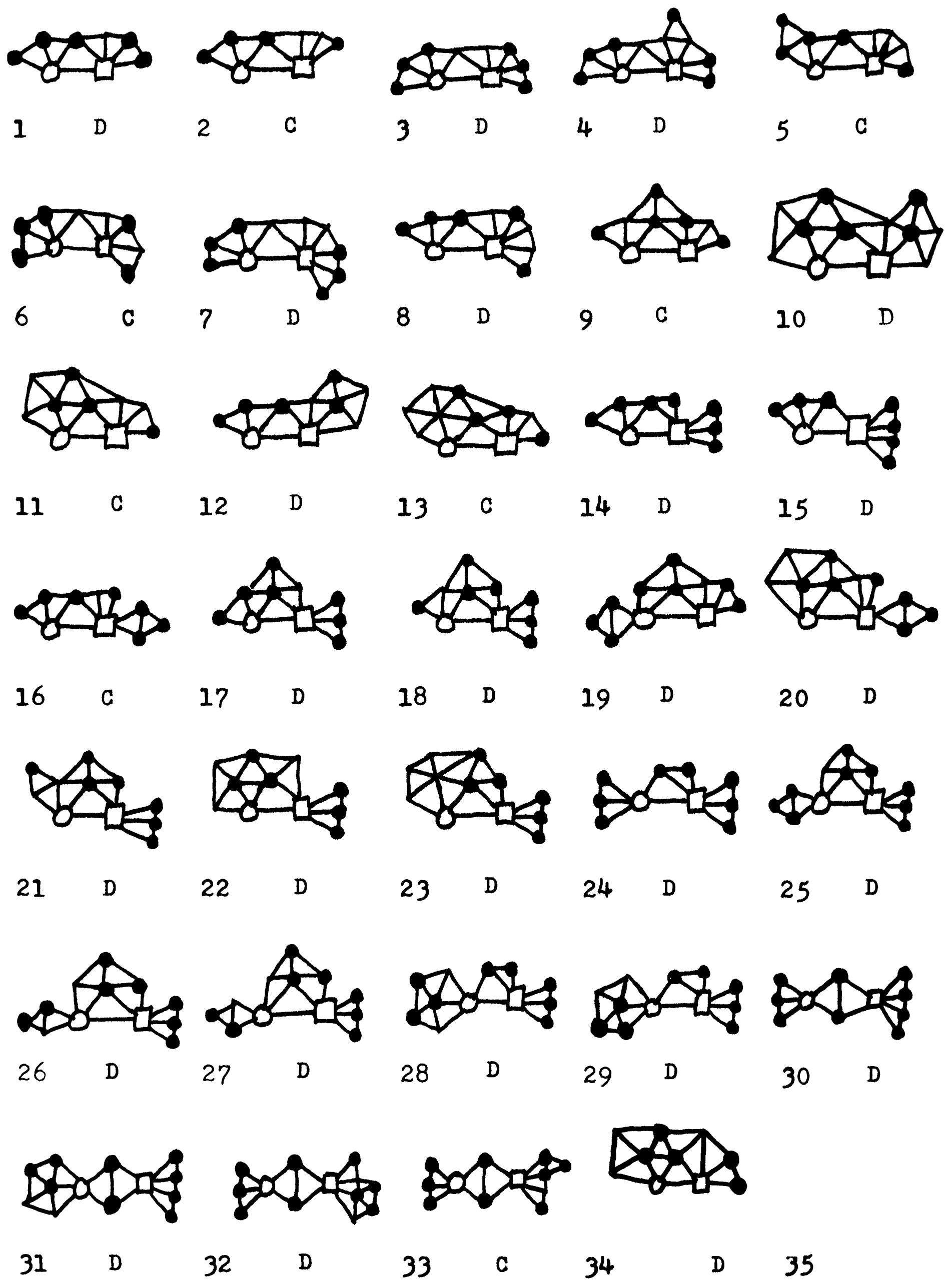

1 D
2 C
3 D
4 D
5 C
6 C
7 D
8 D
9 C
10 D
11 C
12 D
13 C
14 D
15 D
16 C
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 D
32 D
33 C
34 D
35

Figure 29

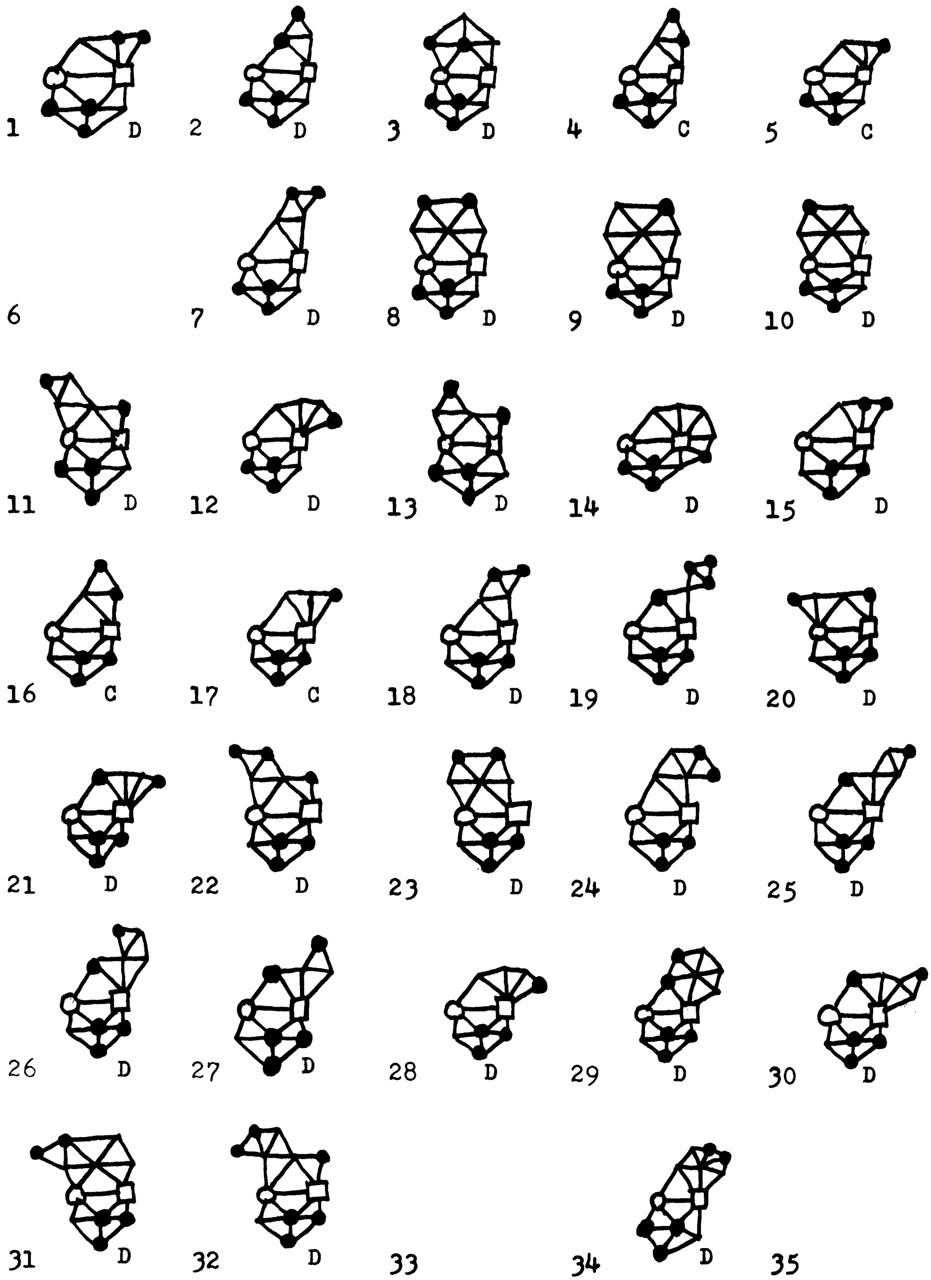
1 D
2 D
3 D
4 C
5 C
6
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 C
17 C
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 D
32 D
33
34 D
35

Figure 30

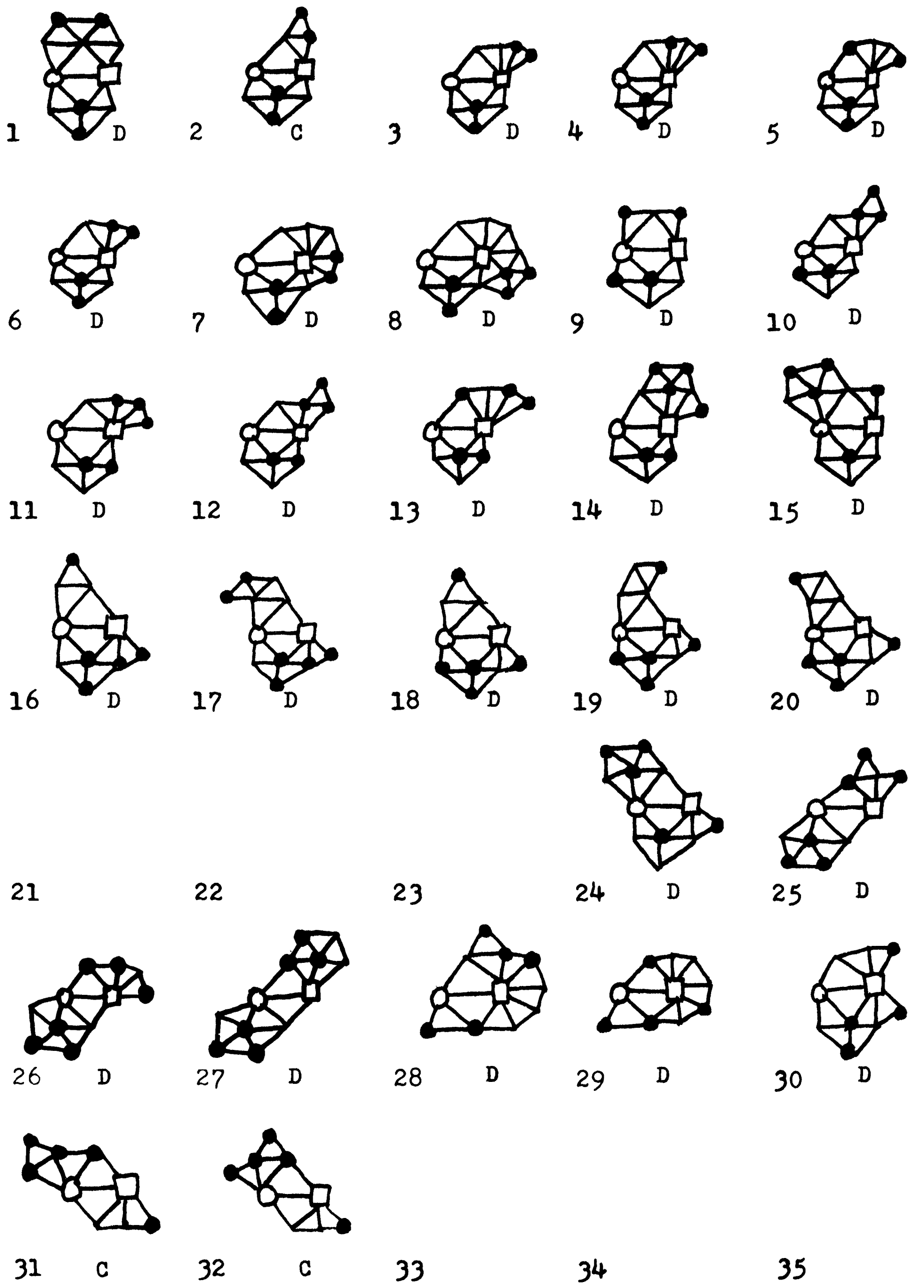
1 D
2 C
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21
22
23
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 C
32 C
33
34
35

Figure 31

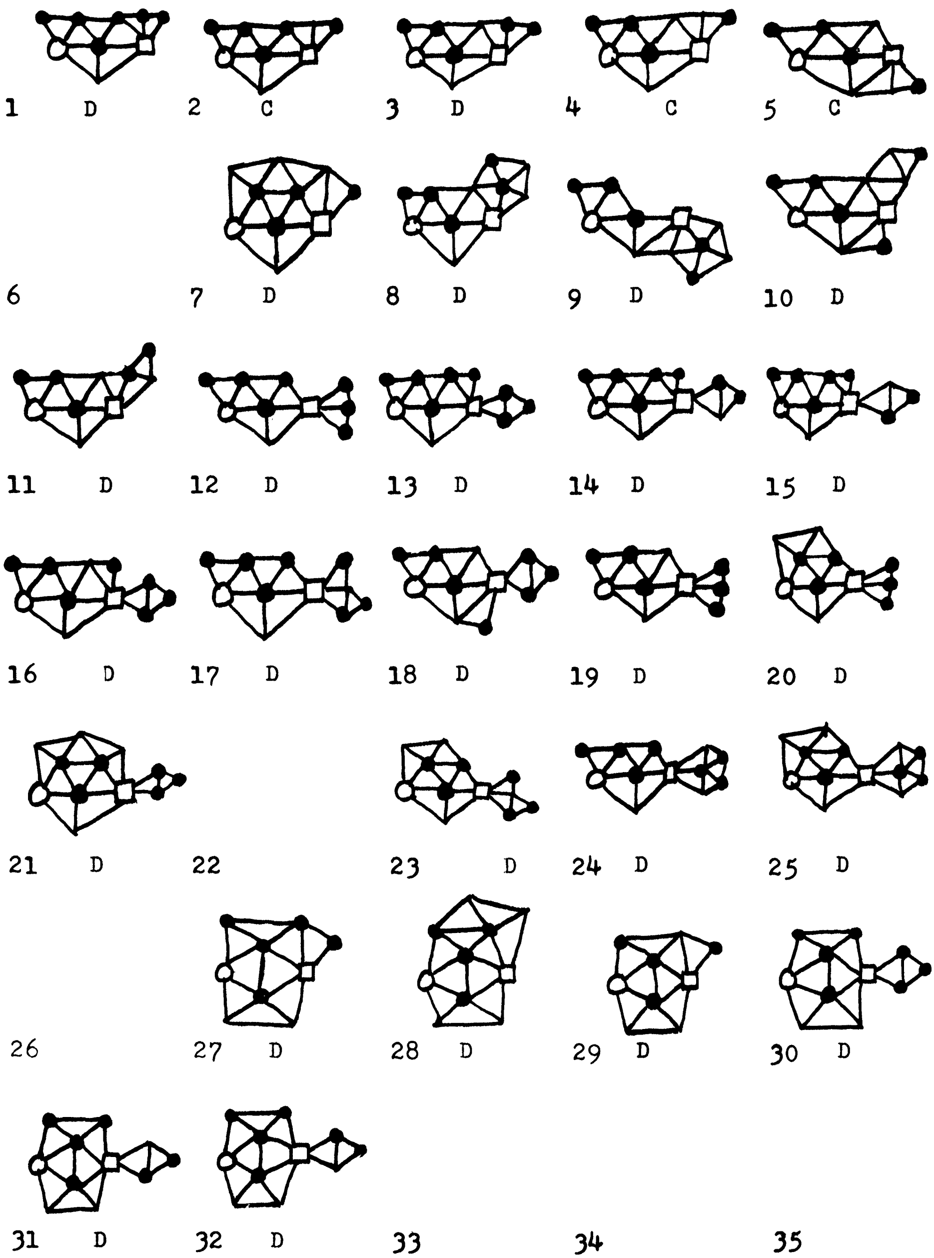
1 D
2 C
3 D
4 C
5 C
6
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22
23 D
24 D
25 D
26
27 D
28 D
29 D
30 D
31 D
32 D
33
34
35

Figure 32

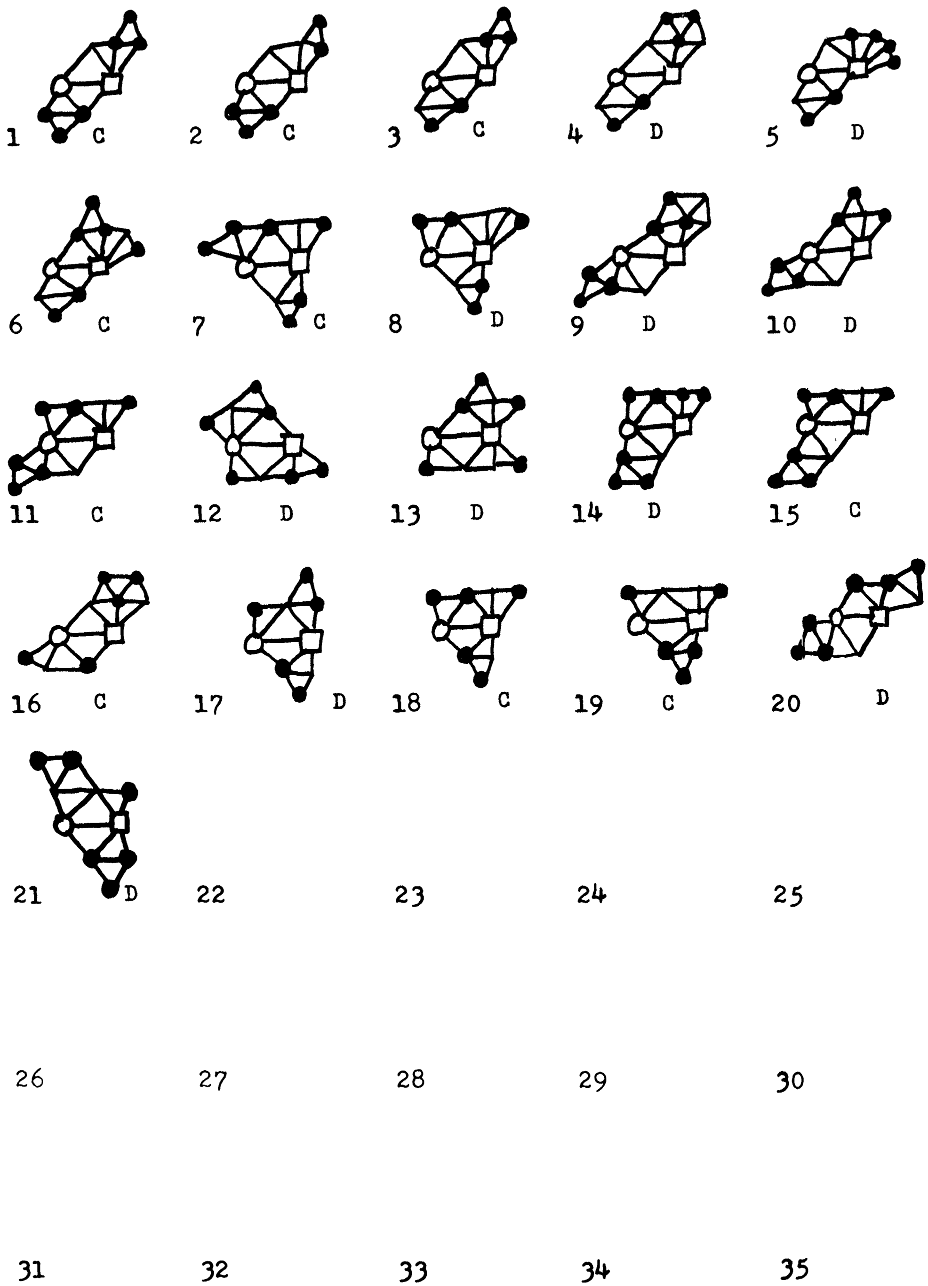
1 C
2 C
3 C
4 D
5 D
6 C
7 C
8 D
9 D
10 D
11 C
12 D
13 D
14 D
15 C
16 C
17 D
18 C
19 C
20 D
21 D
22
23
24
25
26
27
28
29
30
31
32
33
34
35

Figure 33

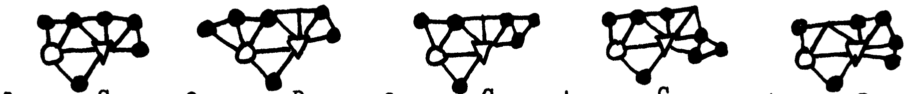

1 C 2 D 3 C 4 C 5 D

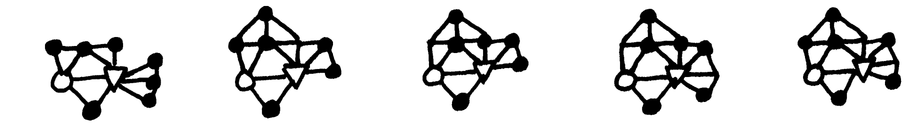

6 C 7 C 8 C 9 C 10 C

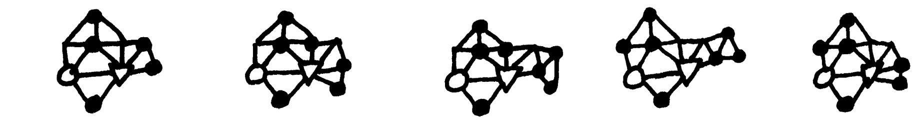

11 C 12 C 13 C 14 C 15 C

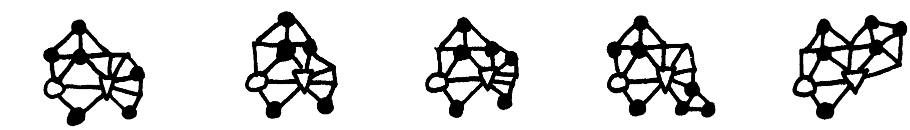

16 D 17 C 18 D 19 D 20 C

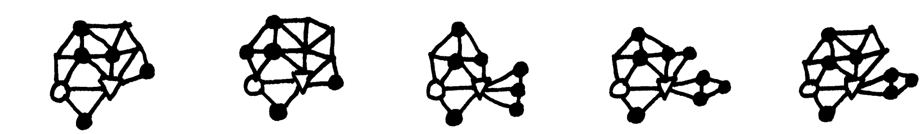

21 C 22 C 23 C 24 C 25 C

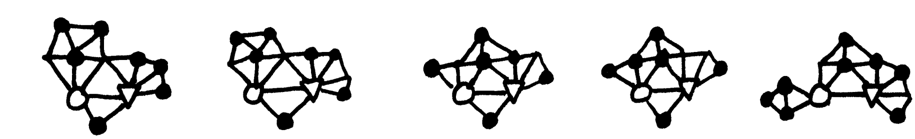

26 D 27 D 28 C 29 C 30 C

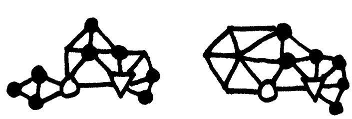

31 D 32 D 33 34 35

Figure 34

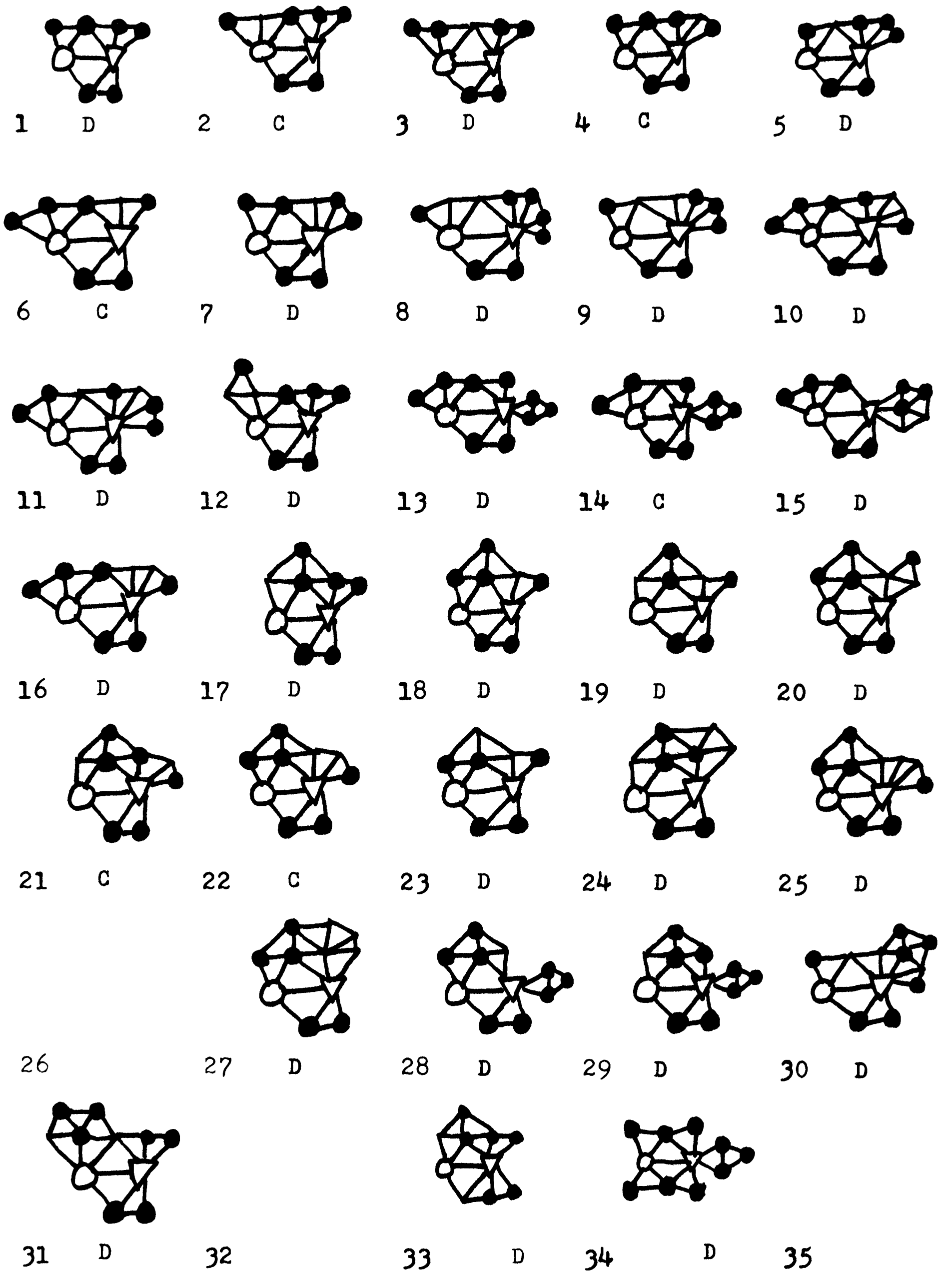
1 D
2 C
3 D
4 C
5 D
6 C
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 C
15 D
16 D
17 D
18 D
19 D
20 D
21 C
22 C
23 D
24 D
25 D
26
27 D
28 D
29 D
30 D
31 D
32
33 D
34 D
35

Figure 35

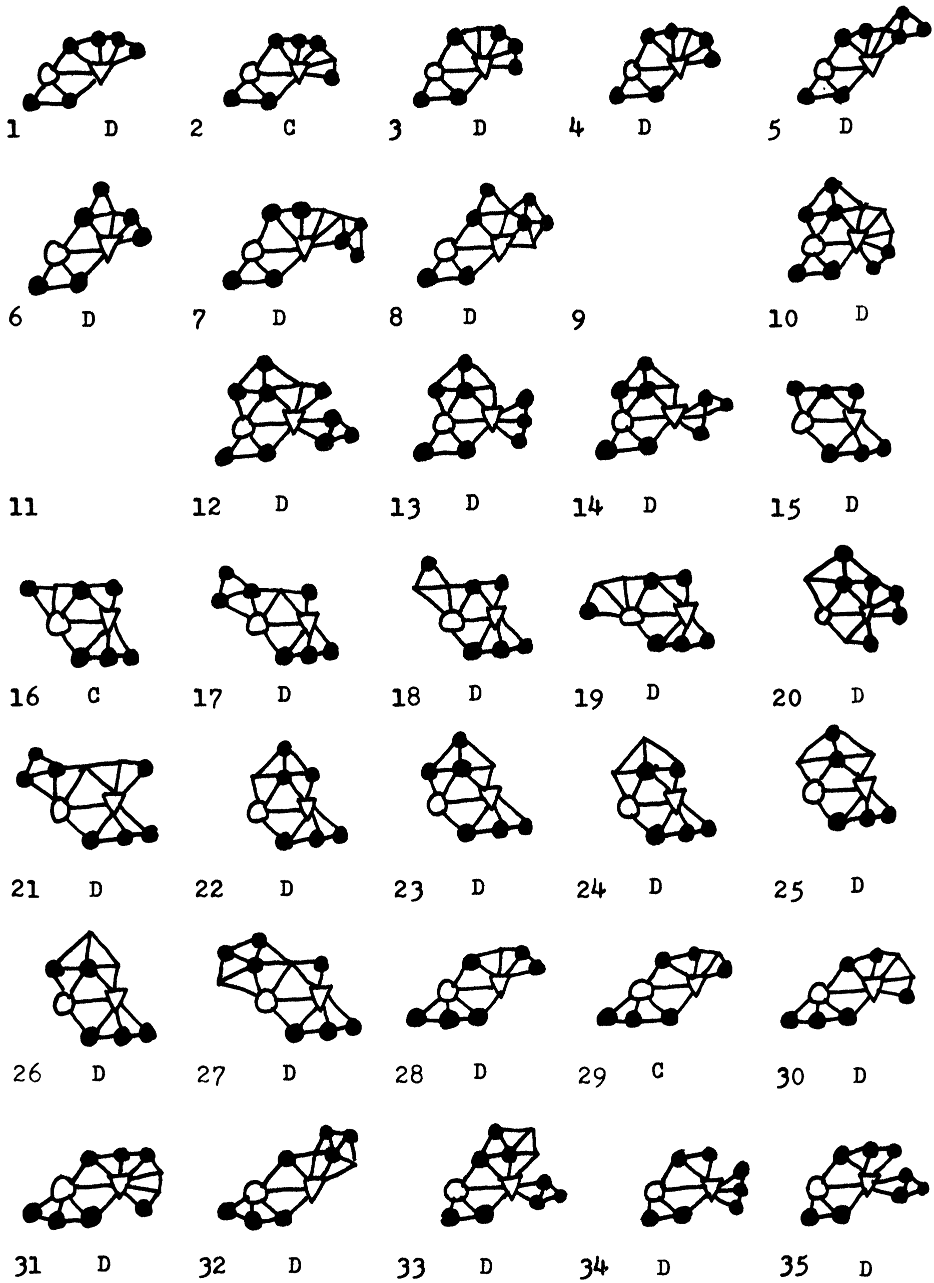
1 D
2 C
3 D
4 D
5 D
6 D
7 D
8 D
9
10 D
11
12 D
13 D
14 D
15 D
16 C
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 C
30 D
31 D
32 D
33 D
34 D
35 D

Figure 36

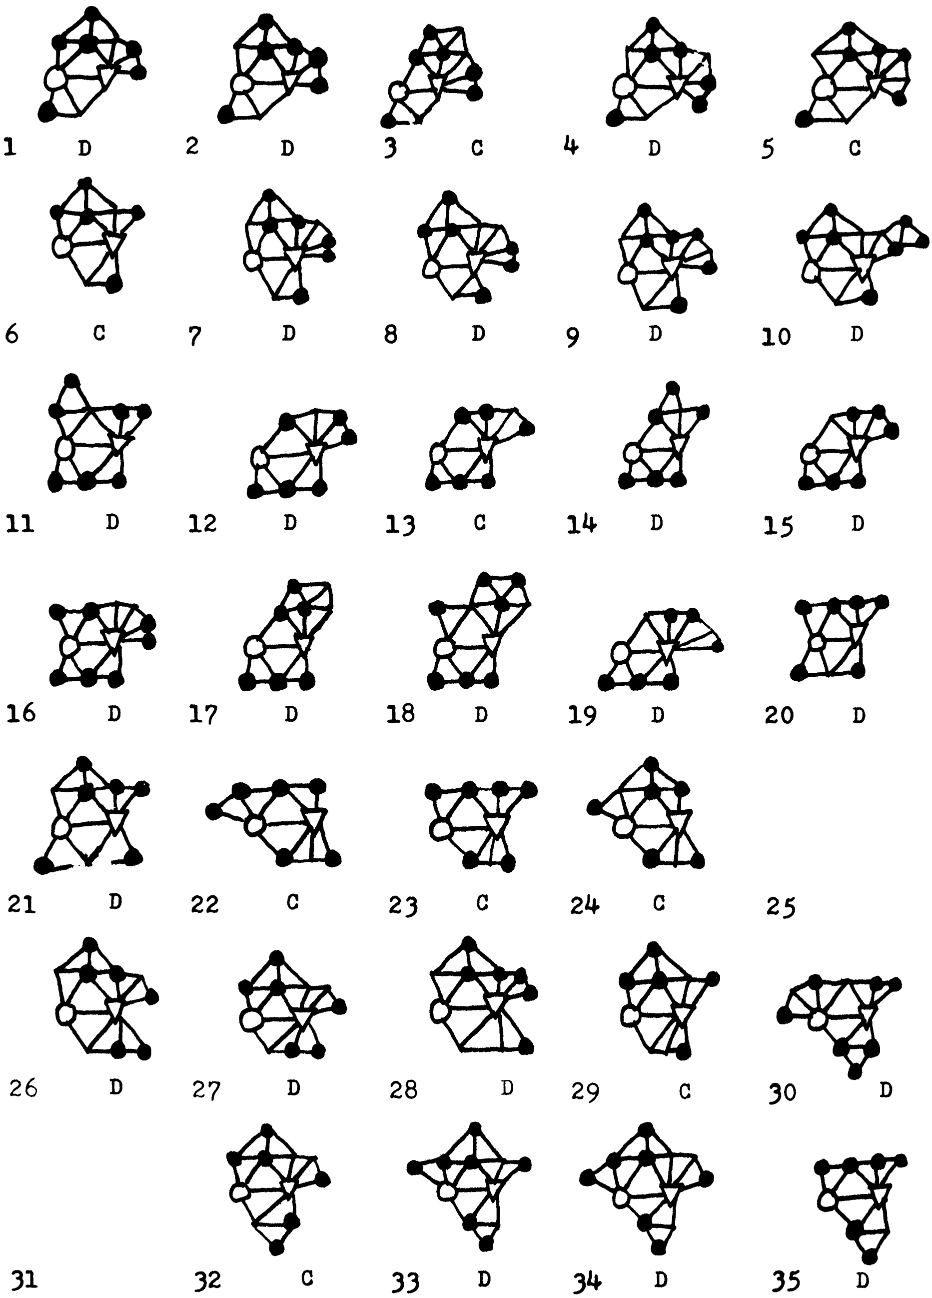
1 D
2 D
3 C
4 D
5 C
6 C
7 D
8 D
9 D
10 D
11 D
12 D
13 C
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 C
23 C
24 C
25
26 D
27 D
28 D
29 C
30 D
31
32 C
33 D
34 D
35 D

Figure 37

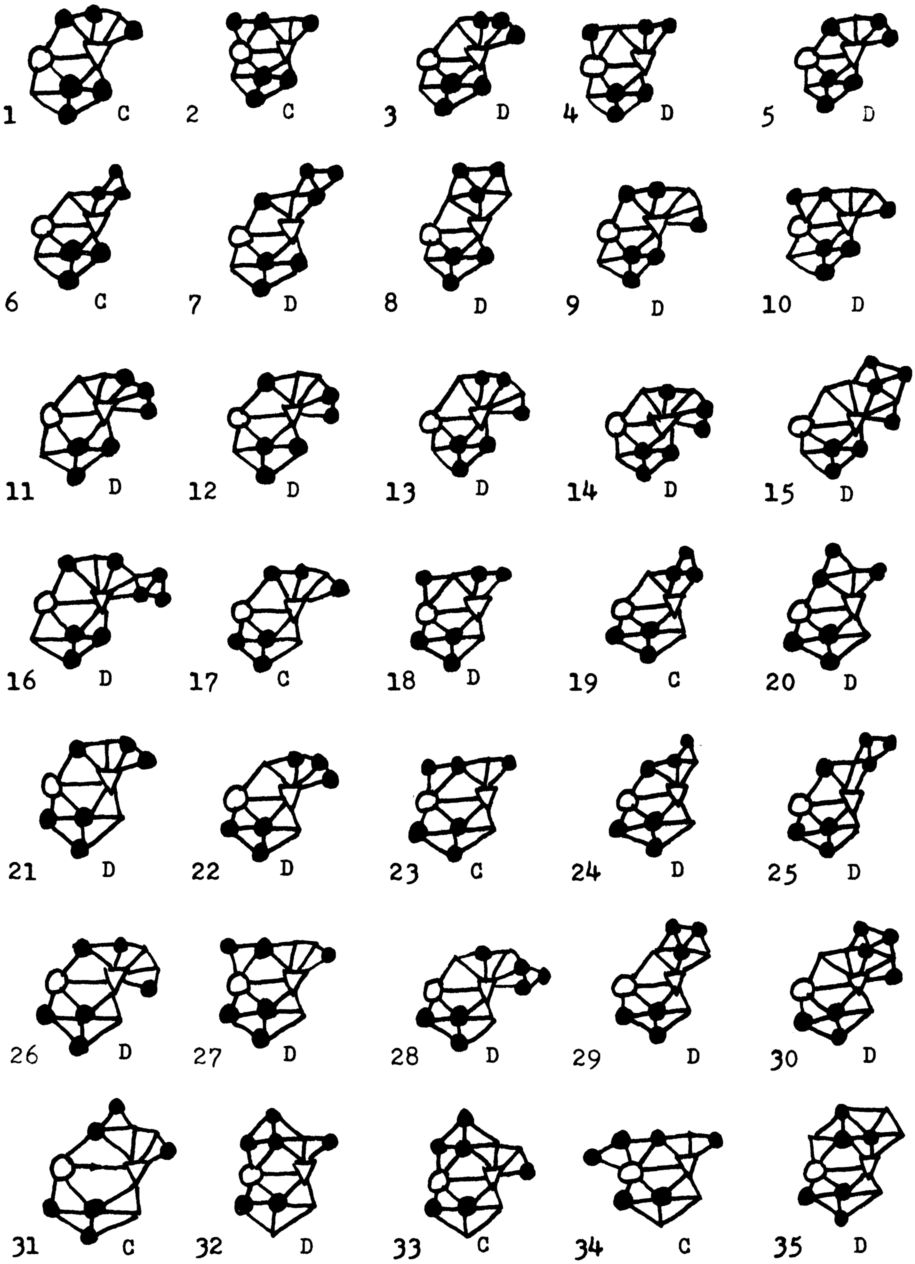

1 C
2 C
3 D
4 D
5 D
6 C
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17 C
18 D
19 C
20 D
21 D
22 D
23 C
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 C
32 D
33 C
34 C
35 D

Figure 38

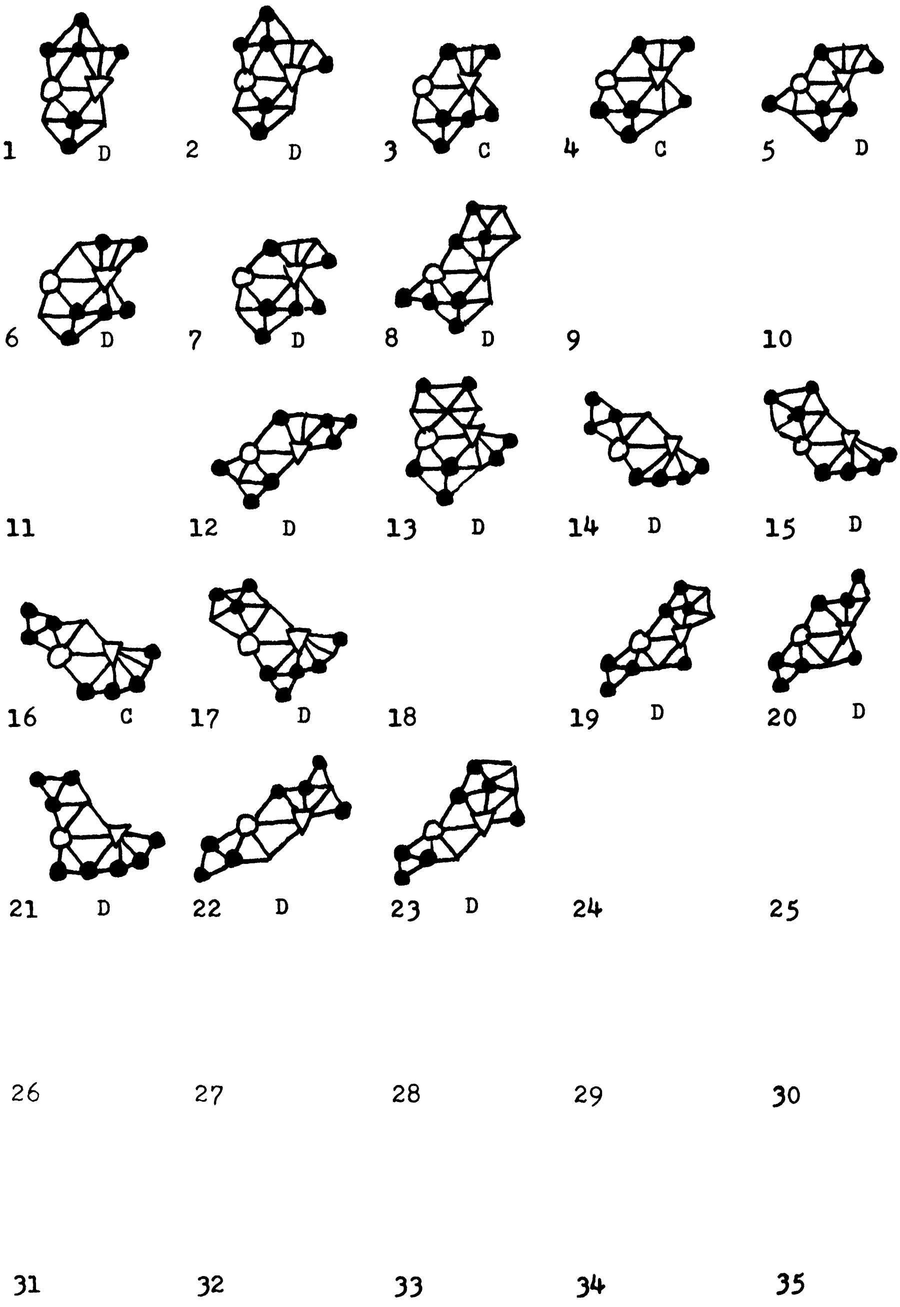
1 D
2 D
3 C
4 C
5 D
6 D
7 D
8 D
9
10
11
12 D
13 D
14 D
15 D
16 C
17 D
18
19 D
20 D
21 D
22 D
23 D
24
25
26
27
28
29
30
31
32
33
34
35

Figure 39

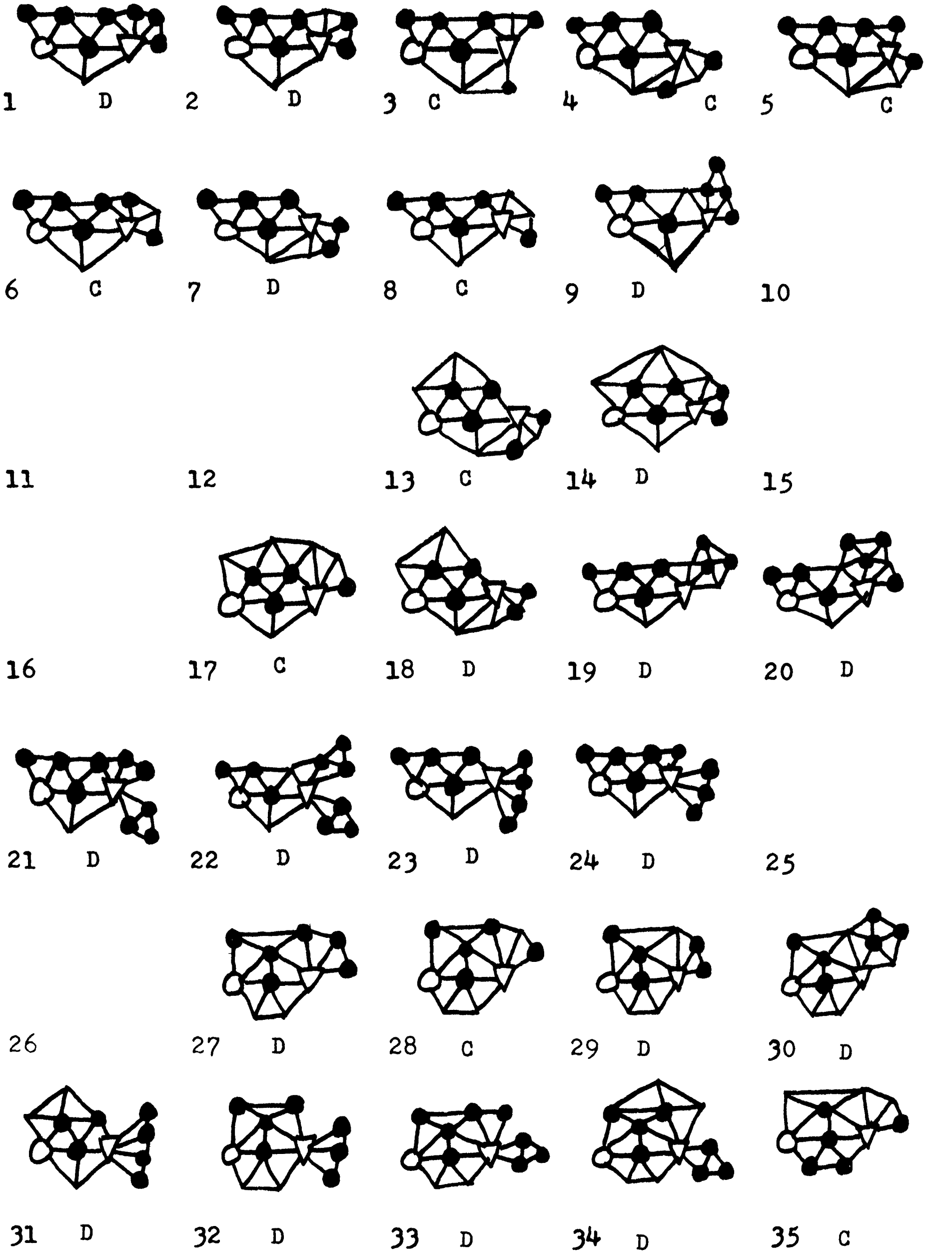
1 D
2 D
3 C
4 C
5 C
6 C
7 D
8 C
9 D
10
11
12
13 C
14 D
15
16
17 C
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25
26
27 D
28 C
29 D
30 D
31 D
32 D
33 D
34 D
35 C

Figure 40

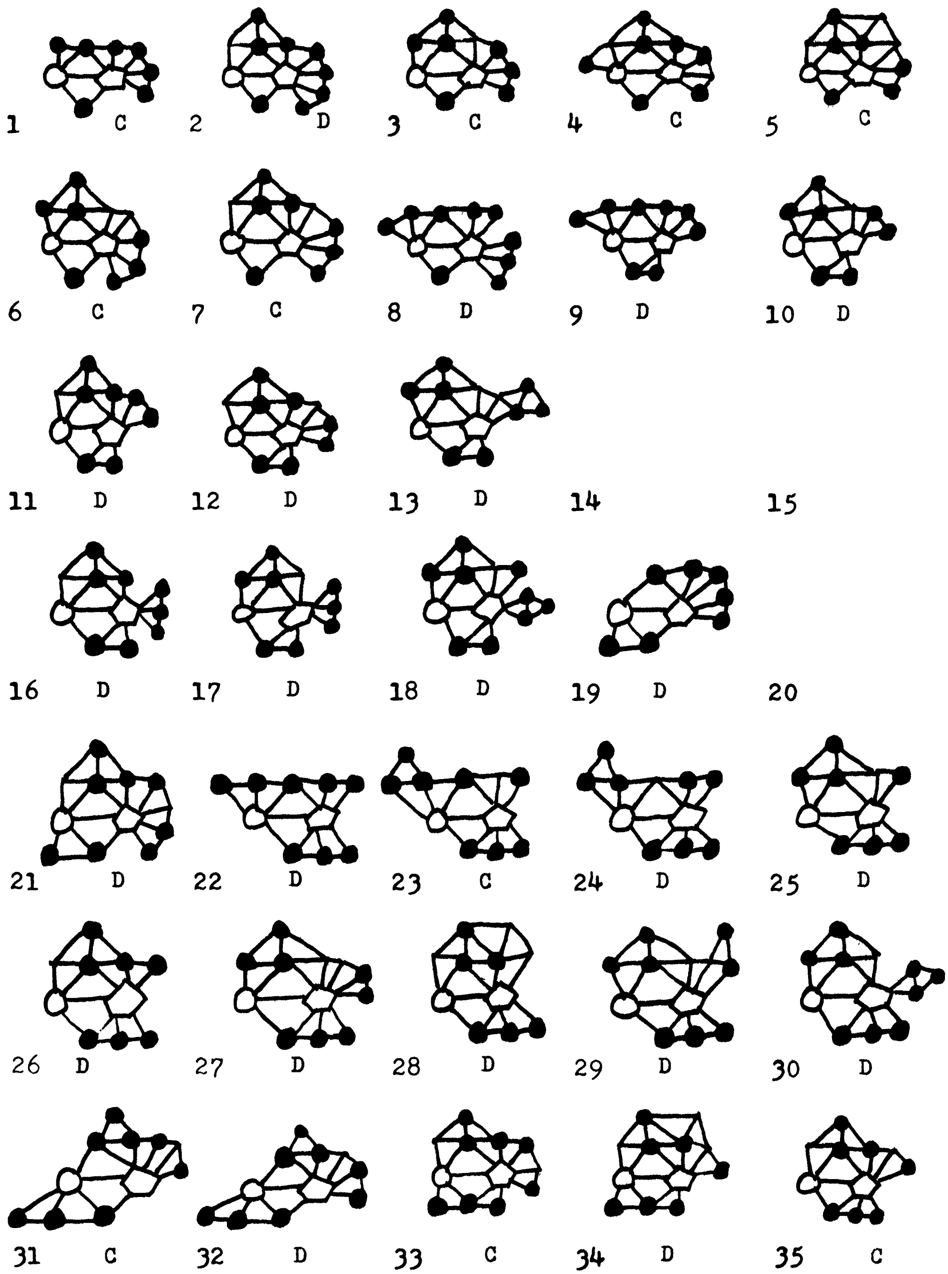
1 C
2 D
3 C
4 C
5 C
6 C
7 C
8 D
9 D
10 D
11 D
12 D
13 D
14
15
16 D
17 D
18 D
19 D
20
21 D
22 D
23 C
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 C
32 D
33 C
34 D
35 C

Figure 41

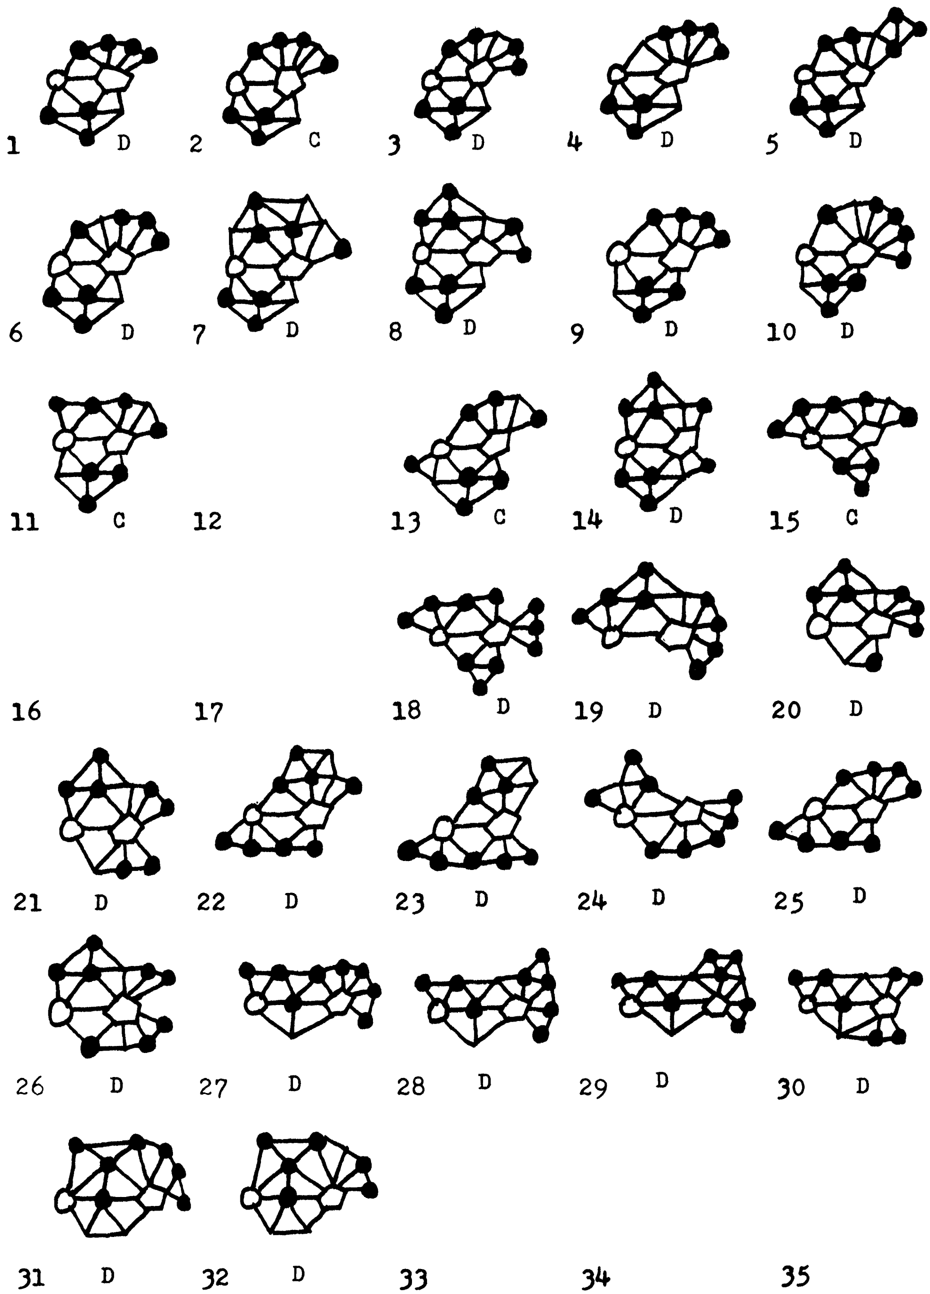
1 D
2 C
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 C
12
13 C
14 D
15 C
16
17
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 D
32 D
33
34
35

Figure 42

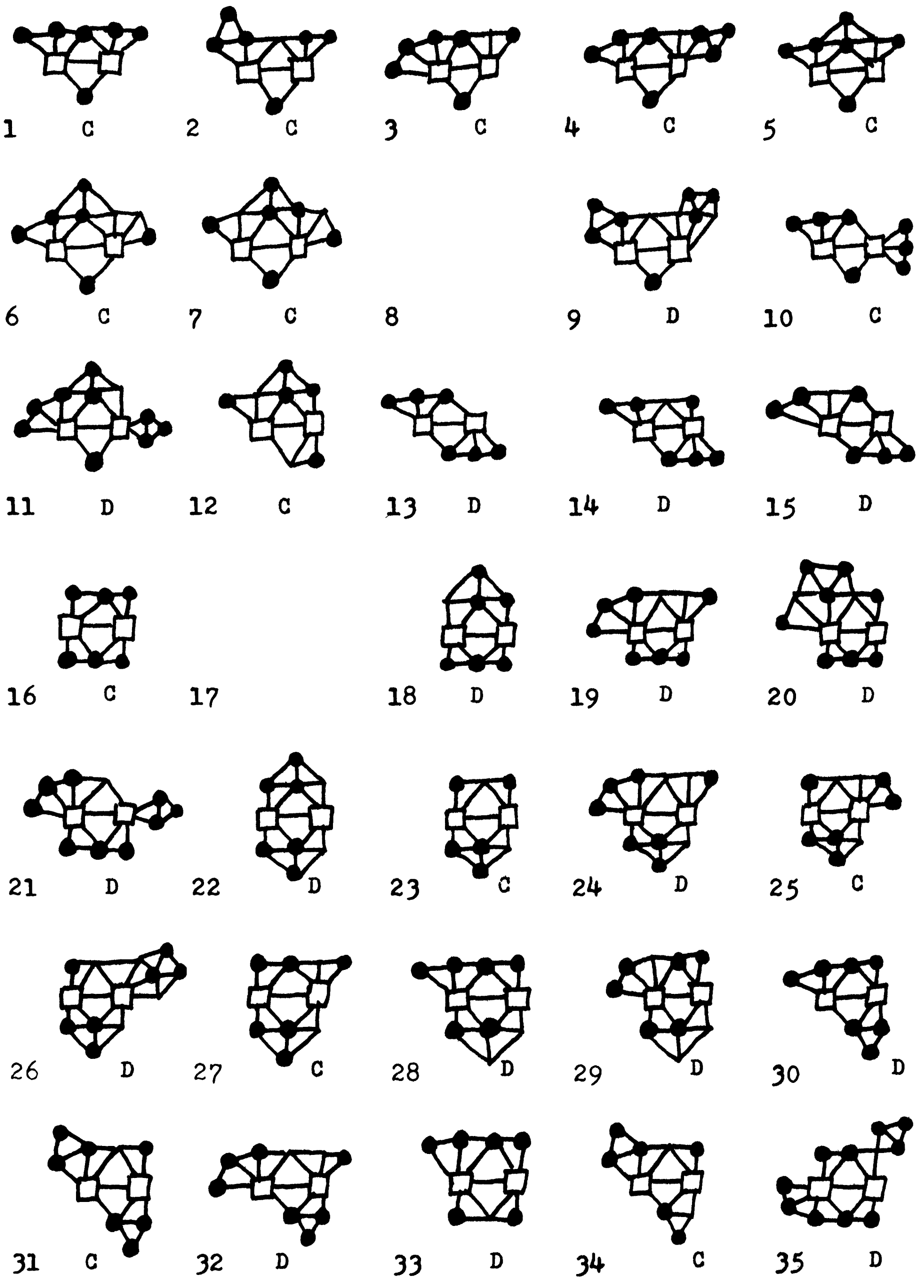
1 C
2 C
3 C
4 C
5 C
6 C
7 C
8
9 D
10 C
11 D
12 C
13 D
14 D
15 D
16 C
17
18 D
19 D
20 D
21 D
22 D
23 C
24 D
25 C
26 D
27 C
28 D
29 D
30 D
31 C
32 D
33 D
34 C
35 D

Figure 43

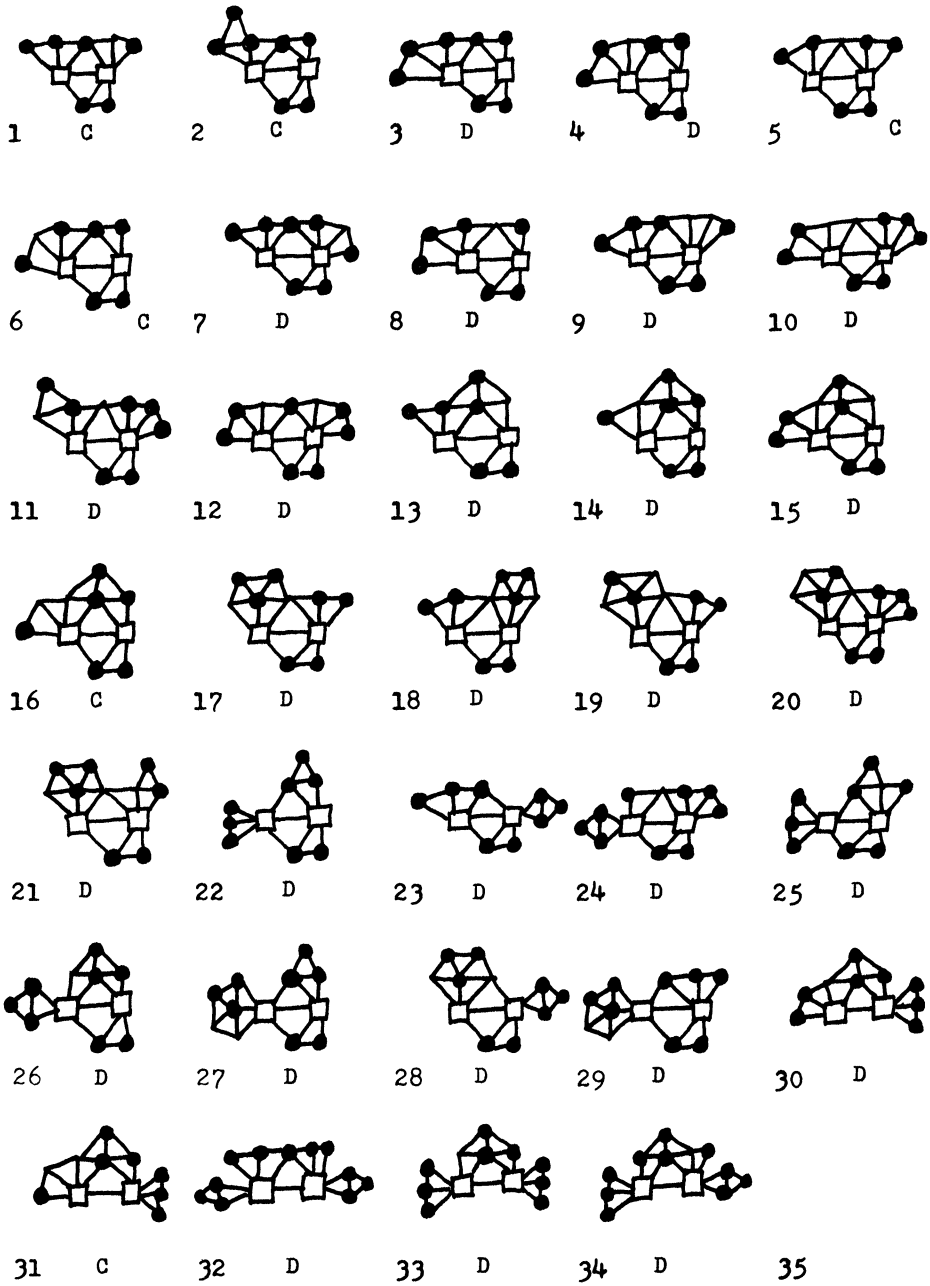
1 C
2 C
3 D
4 D
5 C
6 C
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 C
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 C
32 D
33 D
34 D
35

Figure 44

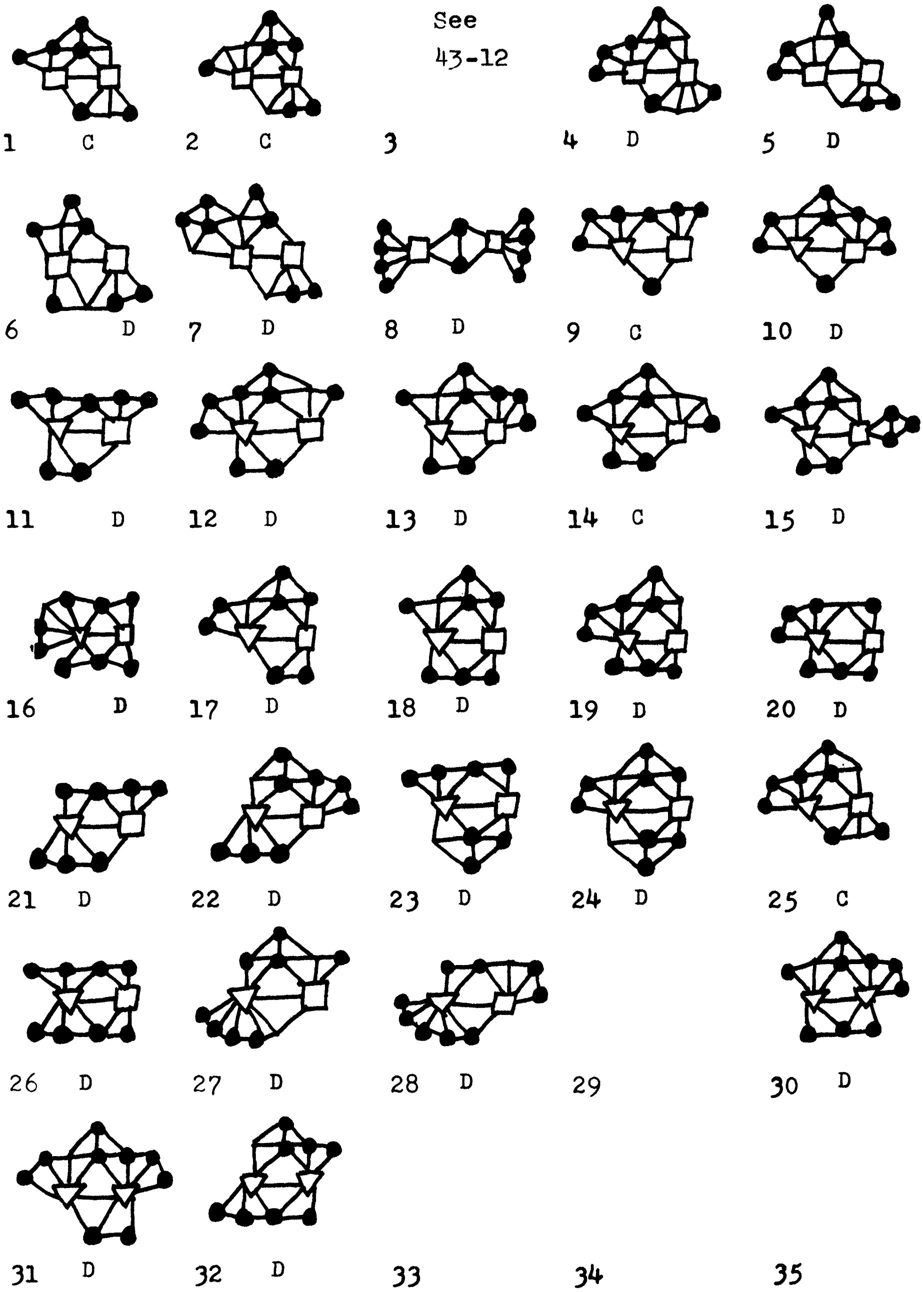
See
43-12
1 C
2 C
3
4 D
5 D
6 D
7 D
8 D
9 C
10 D
11 D
12 D
13 D
14 C
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 C
26 D
27 D
28 D
29
30 D
31 D
32 D
33
34
35

Figure 45

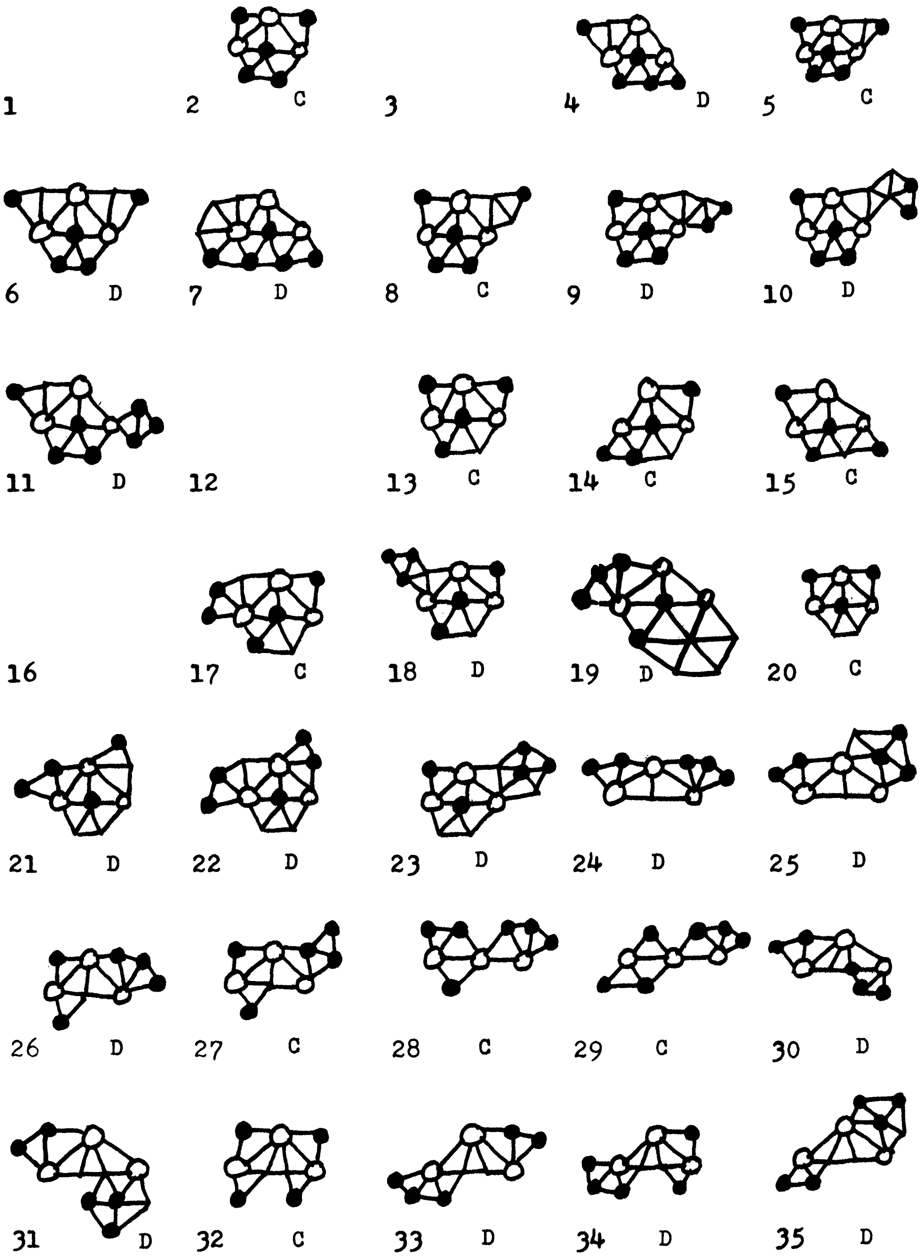
1
2 C
3
4 D
5 C
6 D
7 D
8 C
9 D
10 D
11 D
12
13 C
14 C
15 C
16
17 C
18 D
19 D
20 C
21 D
22 D
23 D
24 D
25 D
26 D
27 C
28 C
29 C
30 D
31 D
32 C
33 D
34 D
35 D

Figure 46

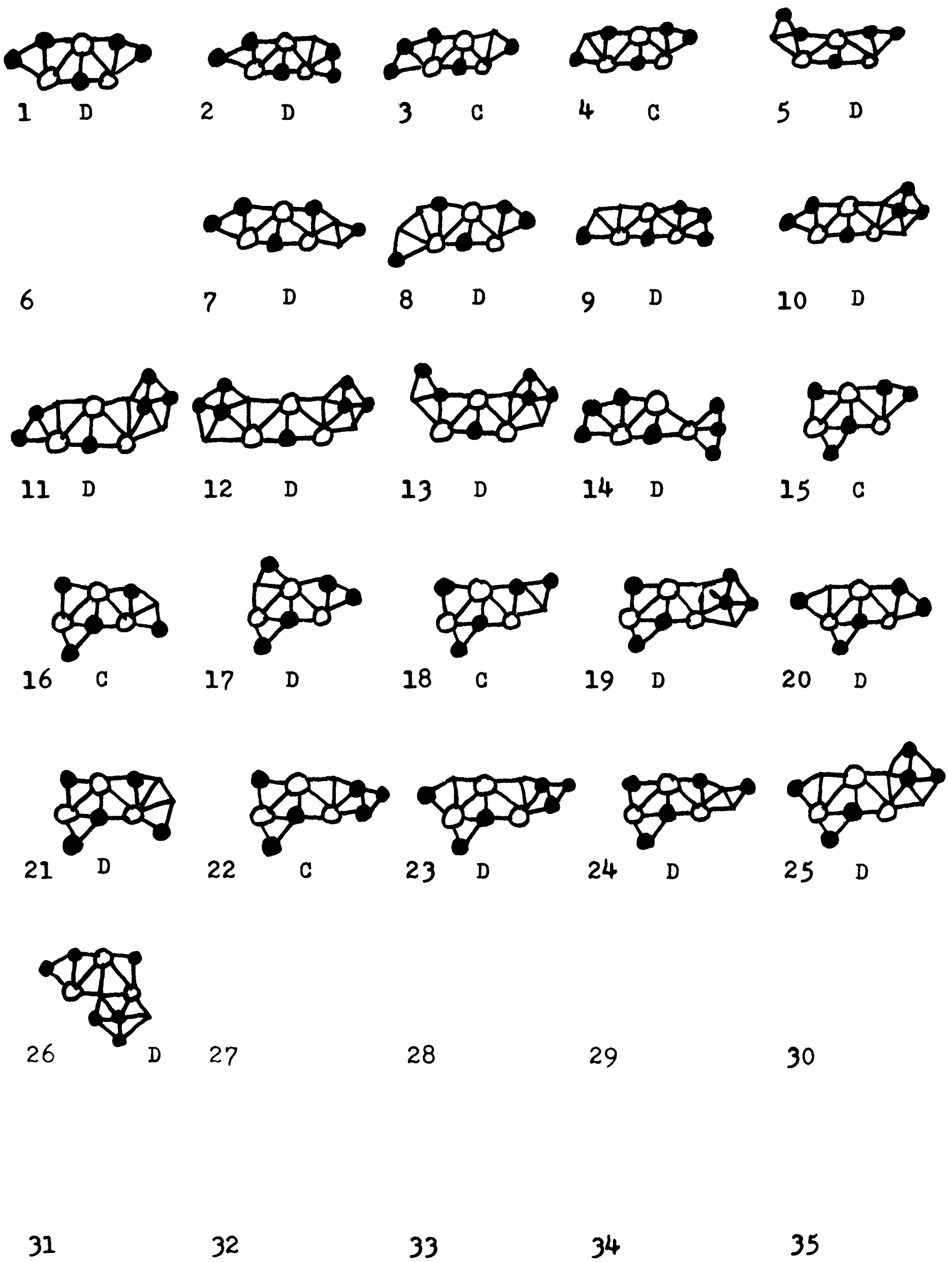
1 D
2 D
3 C
4 C
5 D
6
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 C
16 C
17 D
18 C
19 D
20 D
21 D
22 C
23 D
24 D
25 D
26 D
27
28
29
30
31
32
33
34
35

Figure 47

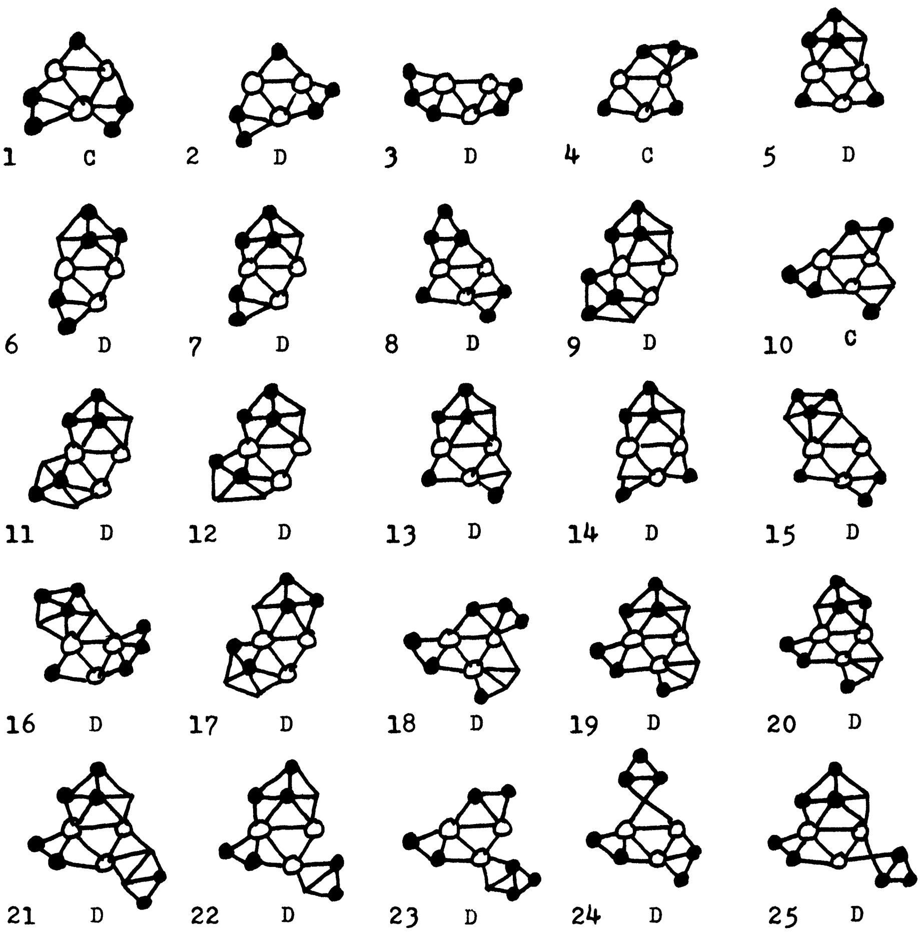

1 C 2 D 3 D 4 C 5 D

6 D 7 D 8 D 9 D 10 C

11 D 12 D 13 D 14 D 15 D

16 D 17 D 18 D 19 D 20 D

21 D 22 D 23 D 24 D 25 D

26 27 28 29 30

31 32 33 34 35

Figure 48

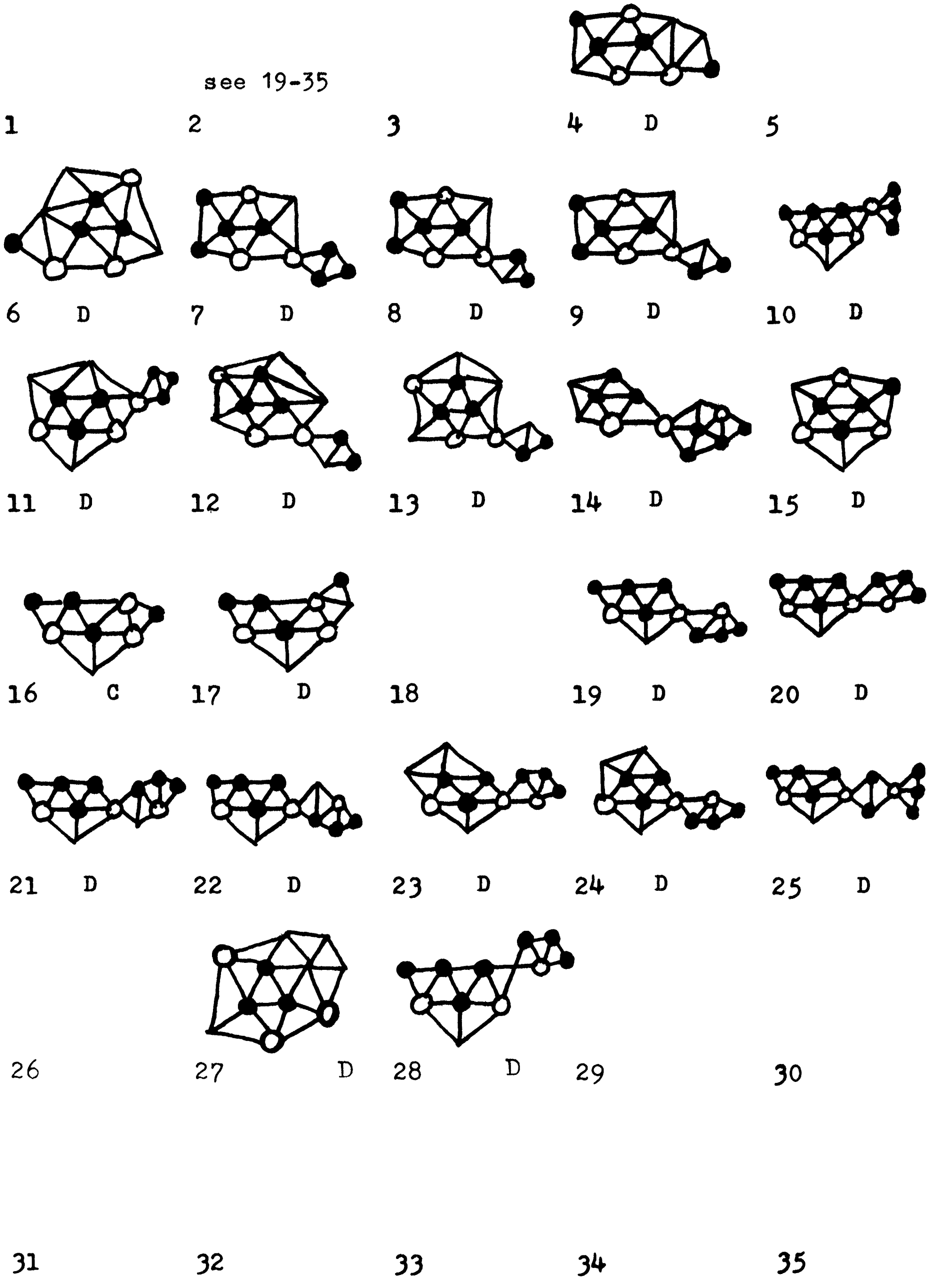
see 19-35
1
2
3
4 D
5
6 D
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 C
17 D
18
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26
27 D
28 D
29
30
31
32
33
34
35

Figure 49

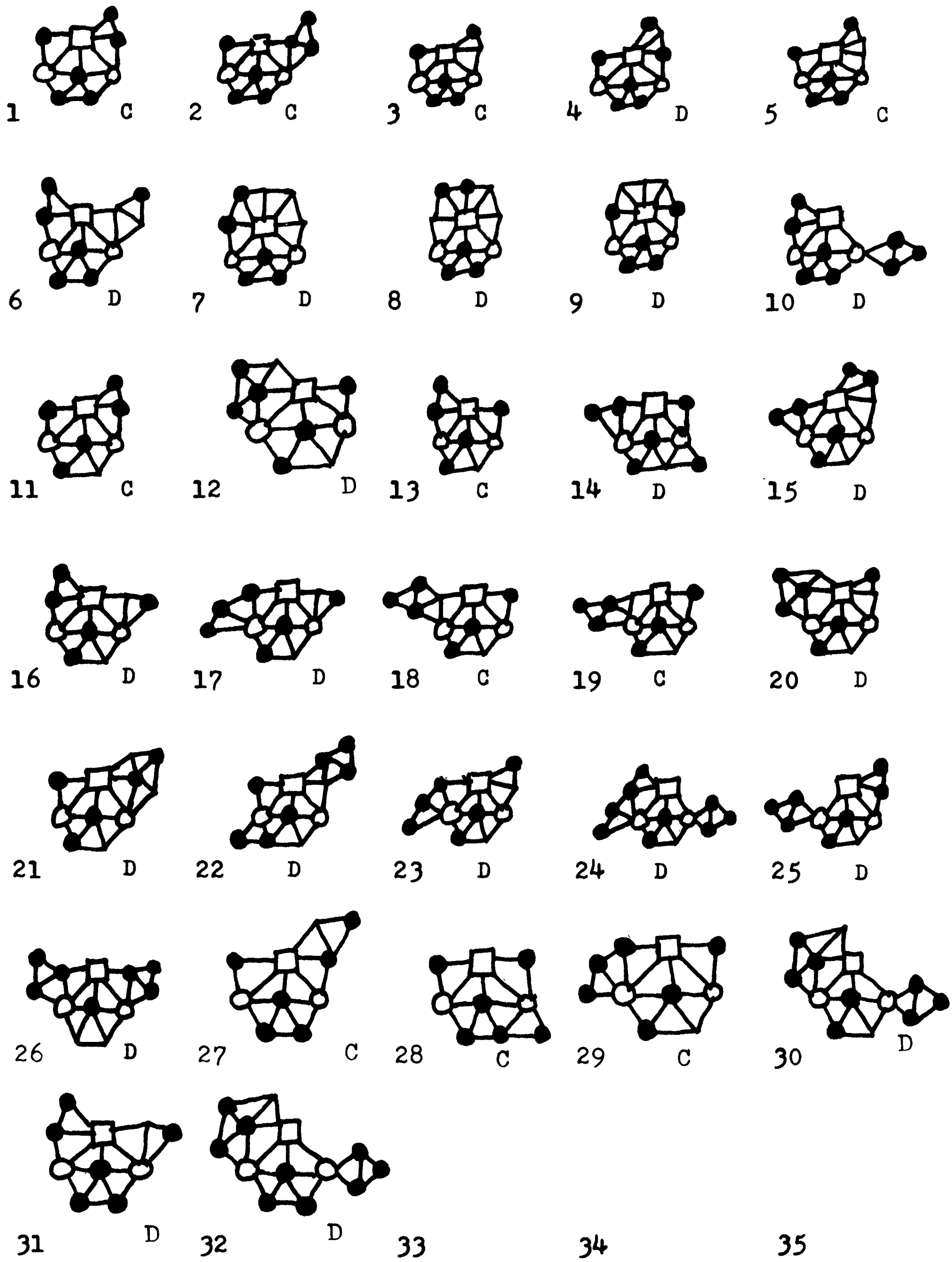
1 C
2 C
3 C
4 D
5 C
6 D
7 D
8 D
9 D
10 D
11 C
12 D
13 C
14 D
15 D
16 D
17 D
18 C
19 C
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 C
28 C
29 C
30 D
31 D
32 D
33
34
35

Figure 50

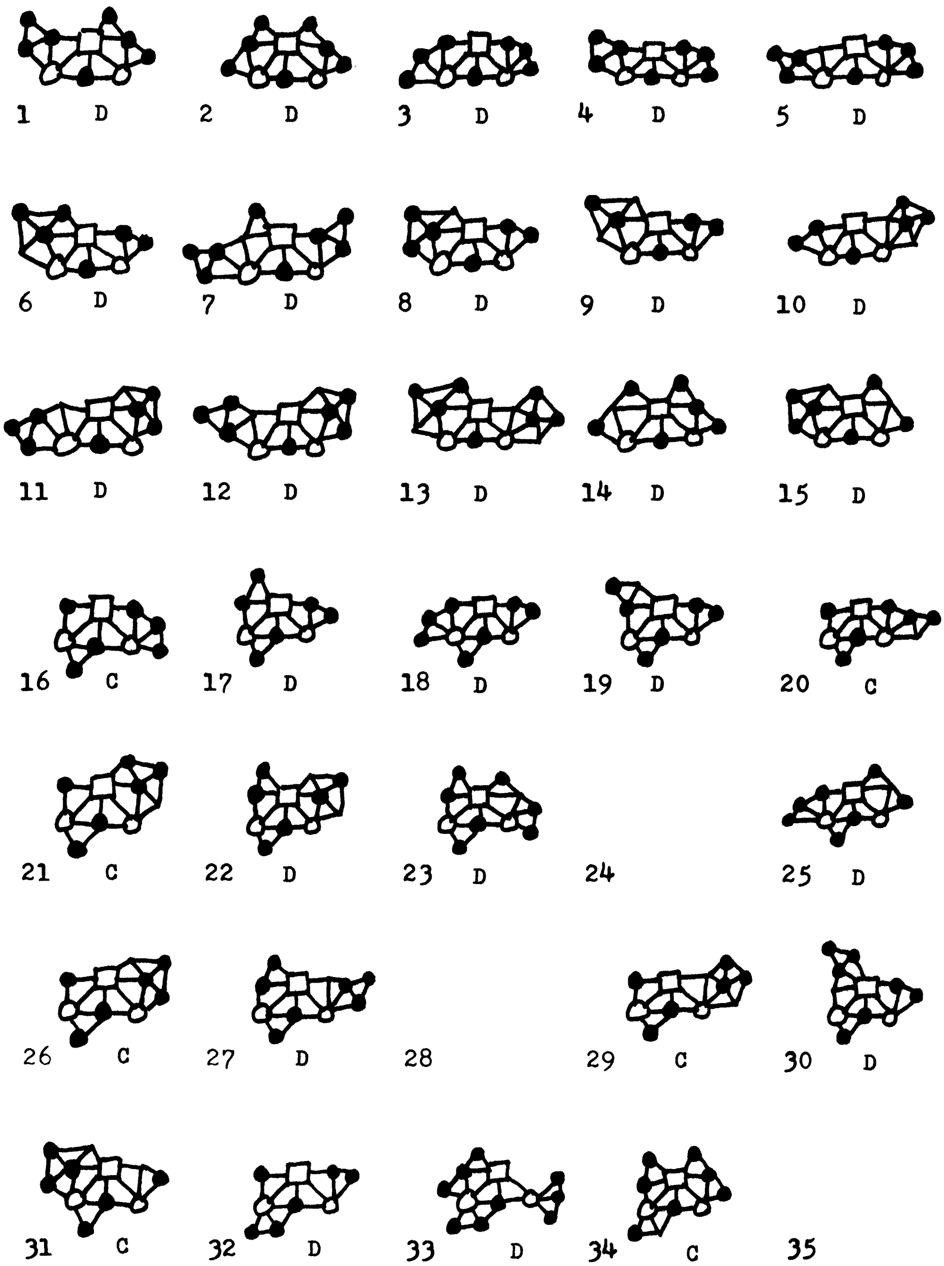
1 D
2 D
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 C
17 D
18 D
19 D
20 C
21 C
22 D
23 D
24
25 D
26 C
27 D
28
29 C
30 D
31 C
32 D
33 D
34 C
35

Figure 51

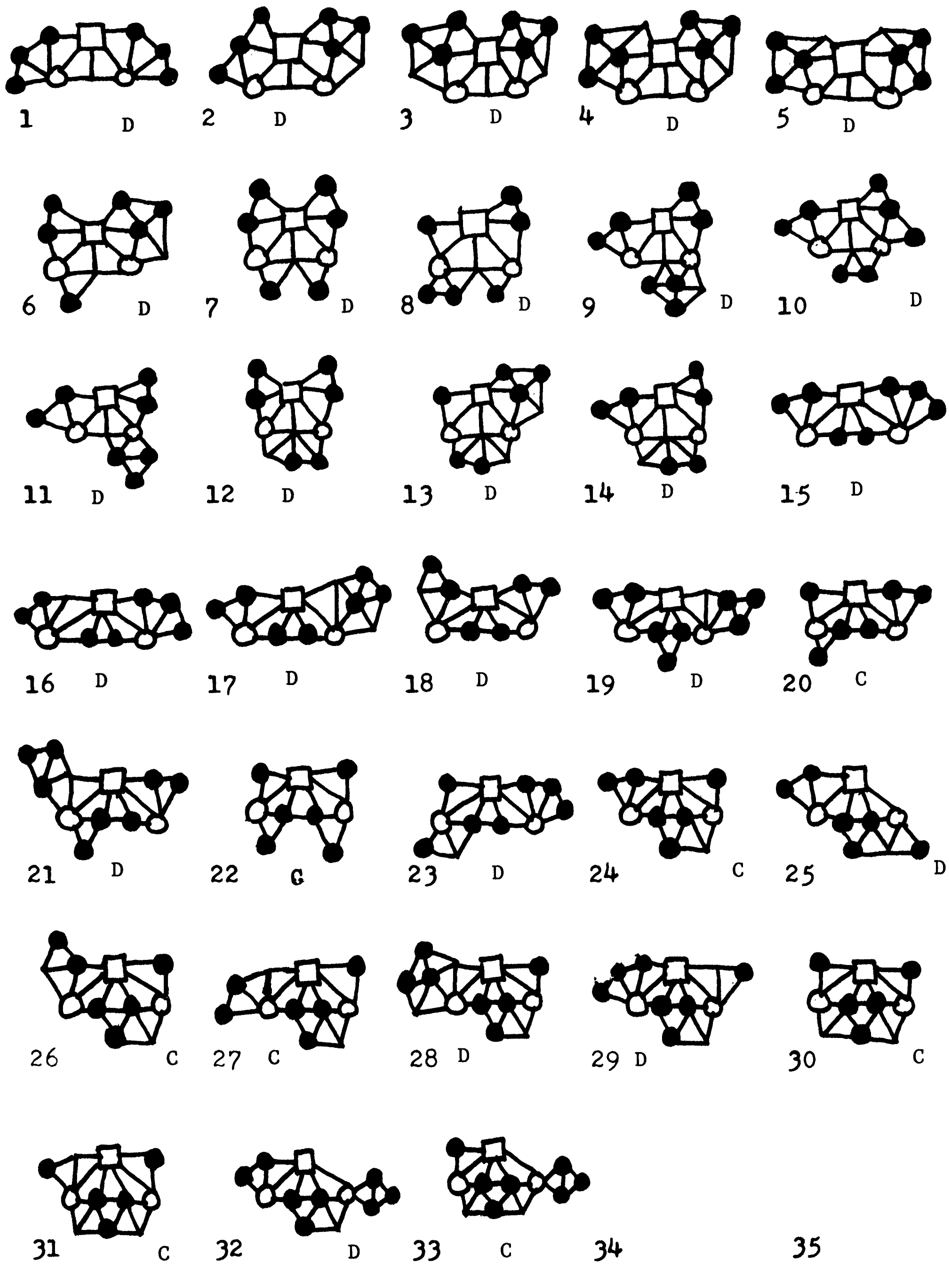
1 D
2 D
3 D
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17 D
18 D
19 D
20 C
21 D
22 C
23 D
24 C
25 D
26 C
27 C
28 D
29 D
30 C
31 C
32 D
33 C
34
35

Figure 52

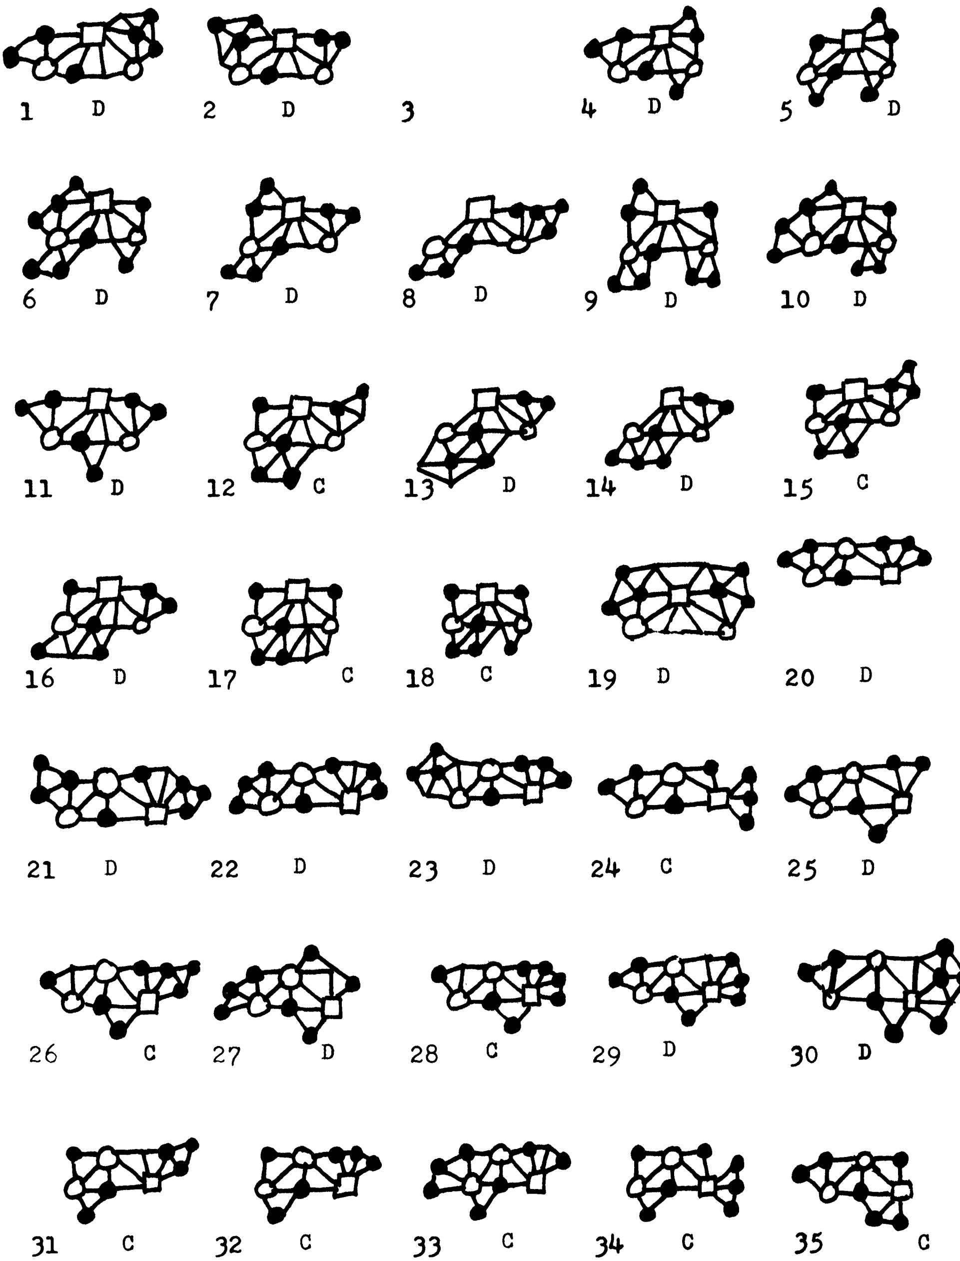
1 D
2 D
3
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 C
13 D
14 D
15 C
16 D
17 C
18 C
19 D
20 D
21 D
22 D
23 D
24 C
25 D
26 C
27 D
28 C
29 D
30 D
31 C
32 C
33 C
34 C
35 C

Figure 53

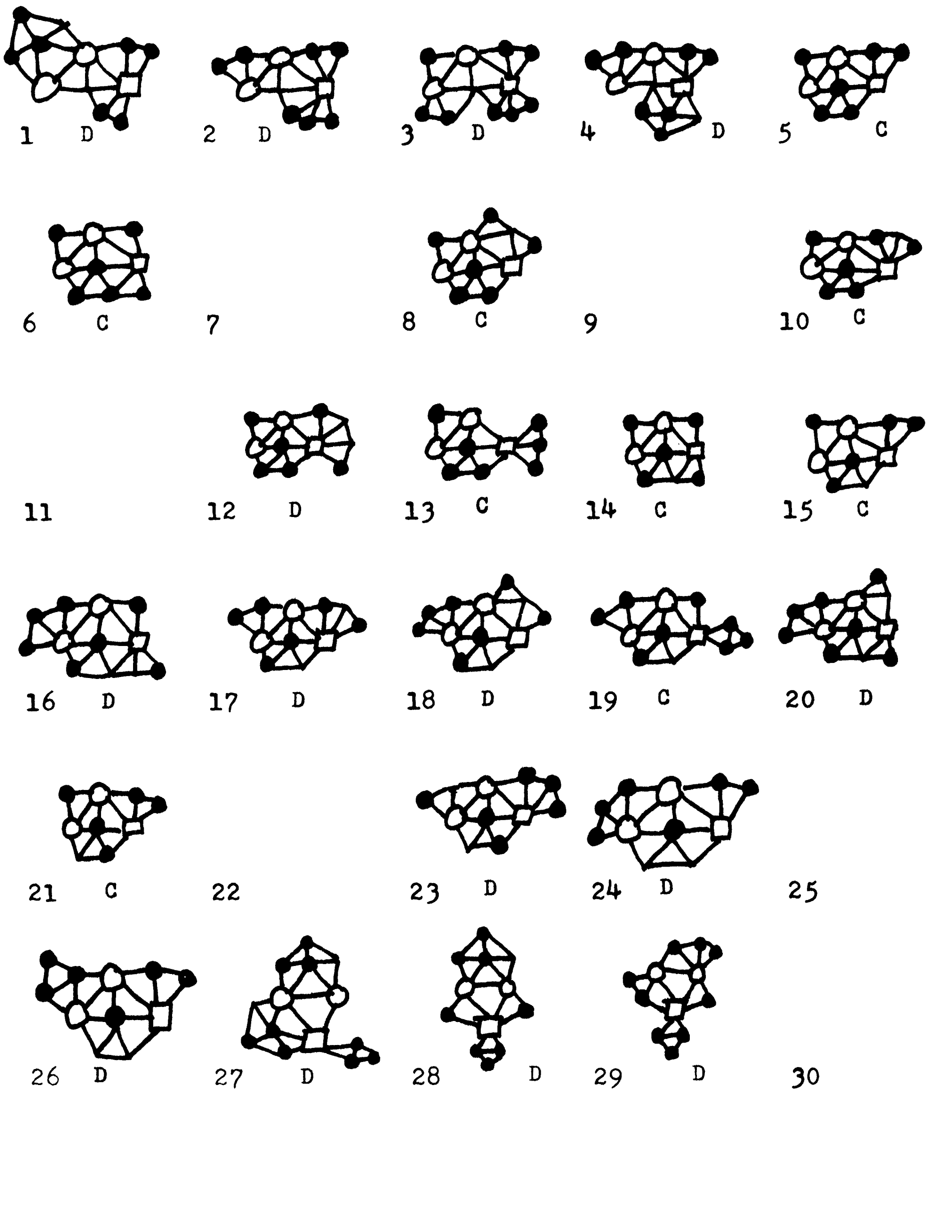
1 D
2 D
3 D
4 D
5 C
6 C
7
8 C
9
10 C
11
12 D
13 C
14 C
15 C
16 D
17 D
18 D
19 C
20 D
21 C
22
23 D
24 D
25
26 D
27 D
28 D
29 D
30
31
32
33
34
35

Figure 54

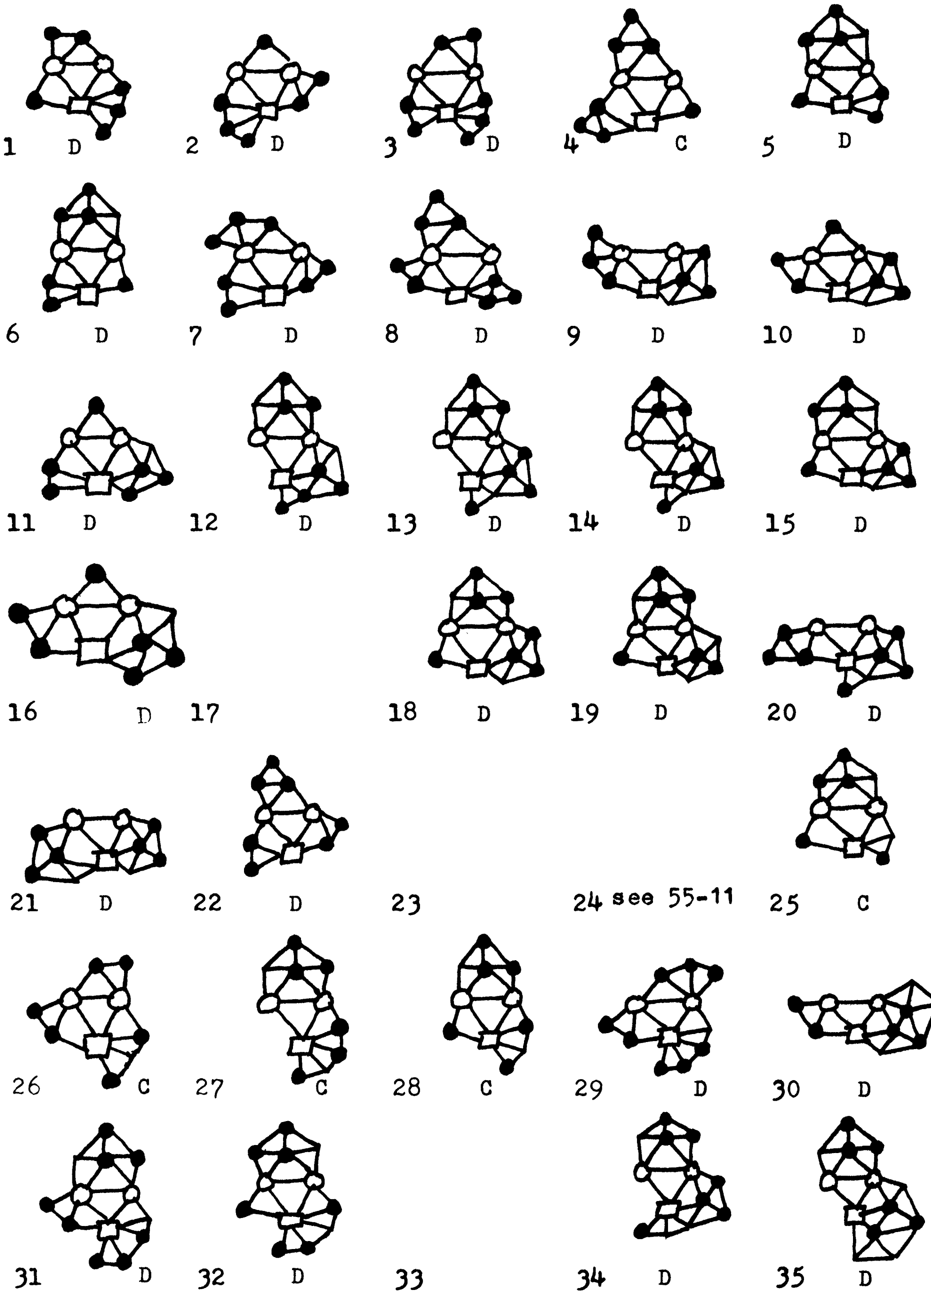
1 D
2 D
3 D
4 C
5 D
6 D
7 D
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17
18 D
19 D
20 D
21 D
22 D
23
24 see 55-11
25 C
26 C
27 C
28 C
29 D
30 D
31 D
32 D
33
34 D
35 D

Figure 55

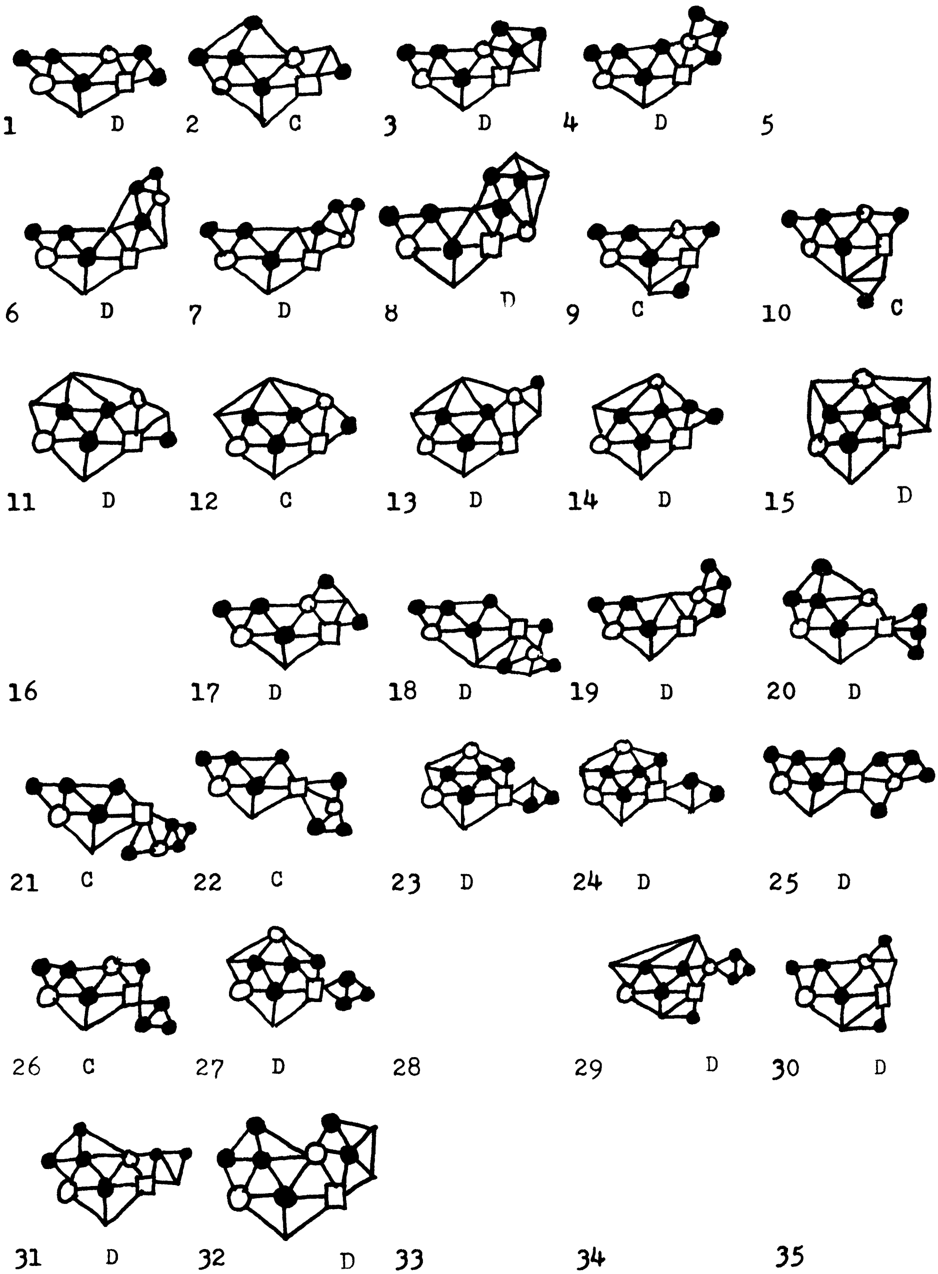
1 D
2 C
3 D
4 D
5
6 D
7 D
8 D
9 C
10 C
11 D
12 C
13 D
14 D
15 D
16
17 D
18 D
19 D
20 D
21 C
22 C
23 D
24 D
25 D
26 C
27 D
28
29 D
30 D
31 D
32 D
33
34
35

Figure 56

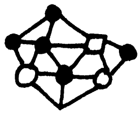

1 C

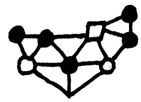

2 C

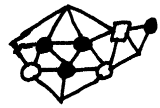

3 C

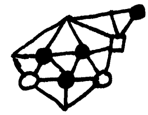

4 C

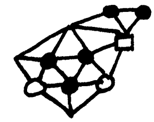

5 D

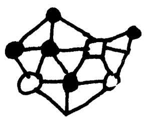

6 D

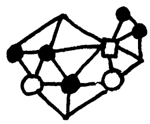

7 D

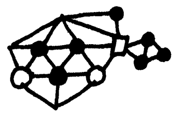

8 D

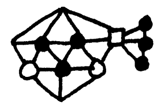

9 D

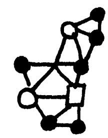

10 D

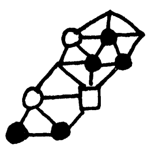

11 D

12

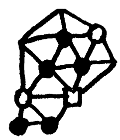

13 D

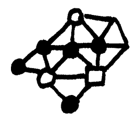

14 D

15

16 17 18 19 20

21 22 23 24 25

26 27 28 29 30

31 32 33 34 35

Figure 57

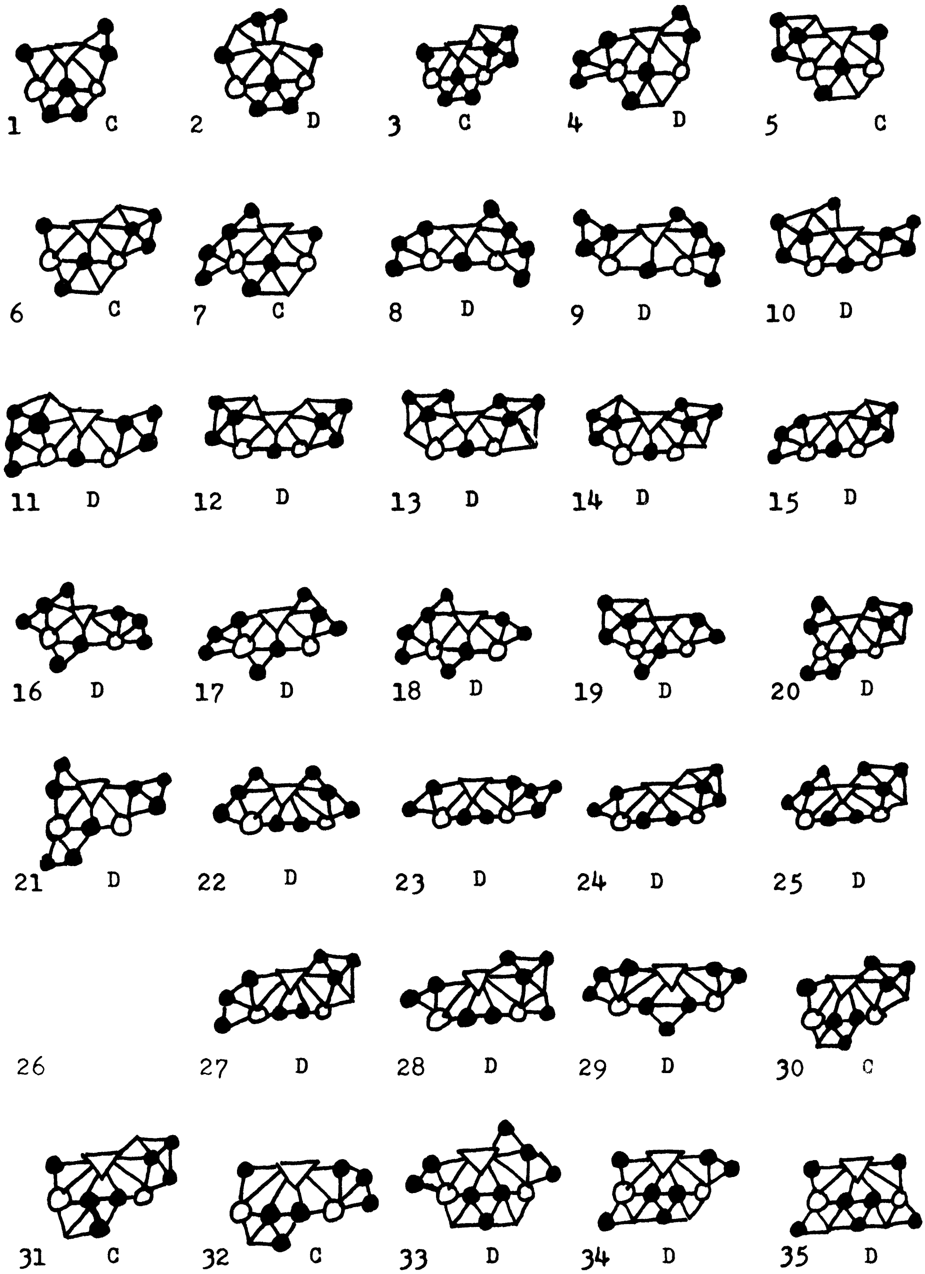
1 C
2 D
3 C
4 D
5 C
6 C
7 C
8 D
9 D
10 D
11 D
12 D
13 D
14 D
15 D
16 D
17 D
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26
27 D
28 D
29 D
30 C
31 C
32 C
33 D
34 D
35 D

Figure 58

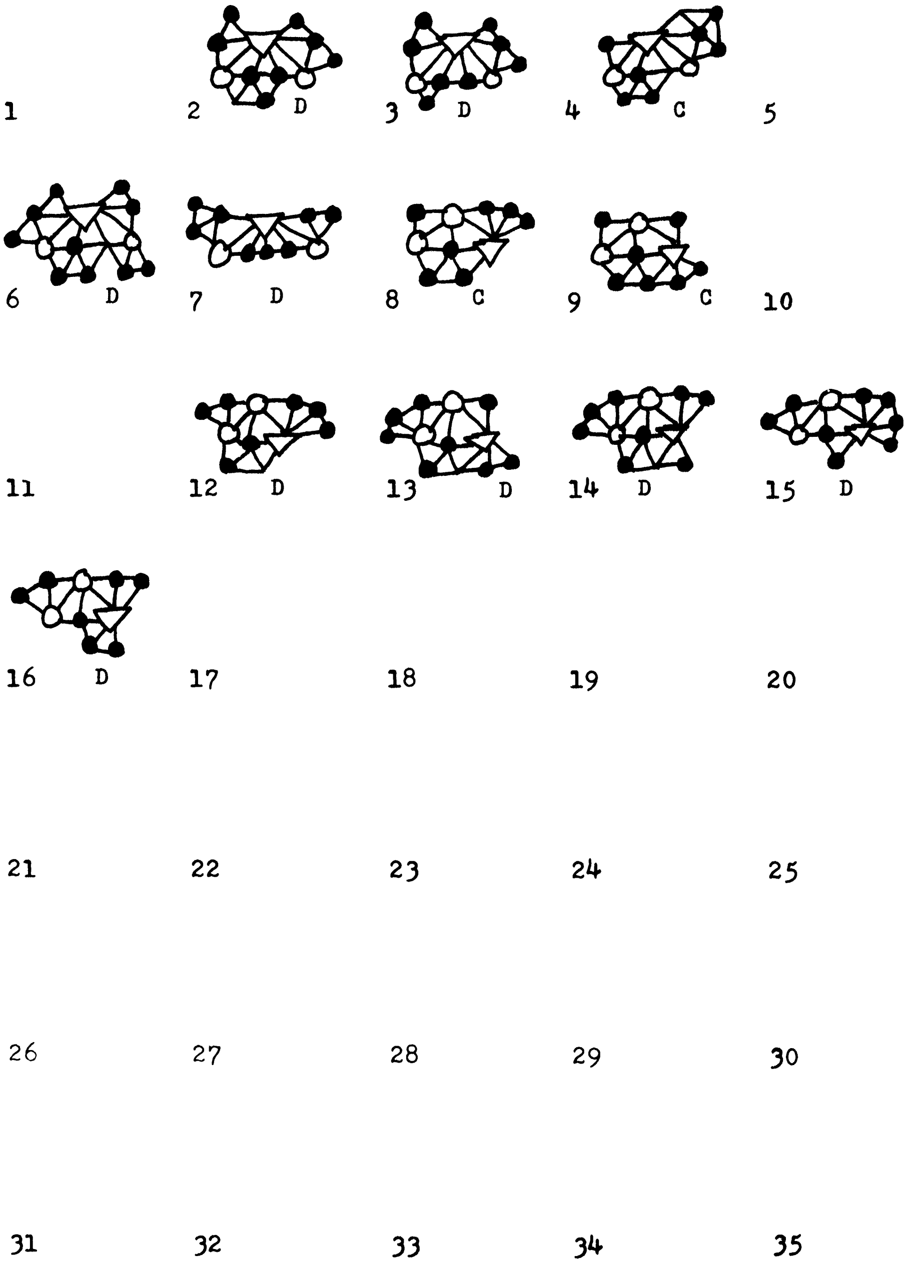
1
2 D
3 D
4 C
5
6 D
7 D
8 C
9 C
10
11
12 D
13 D
14 D
15 D
16 D
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35

Figure 59

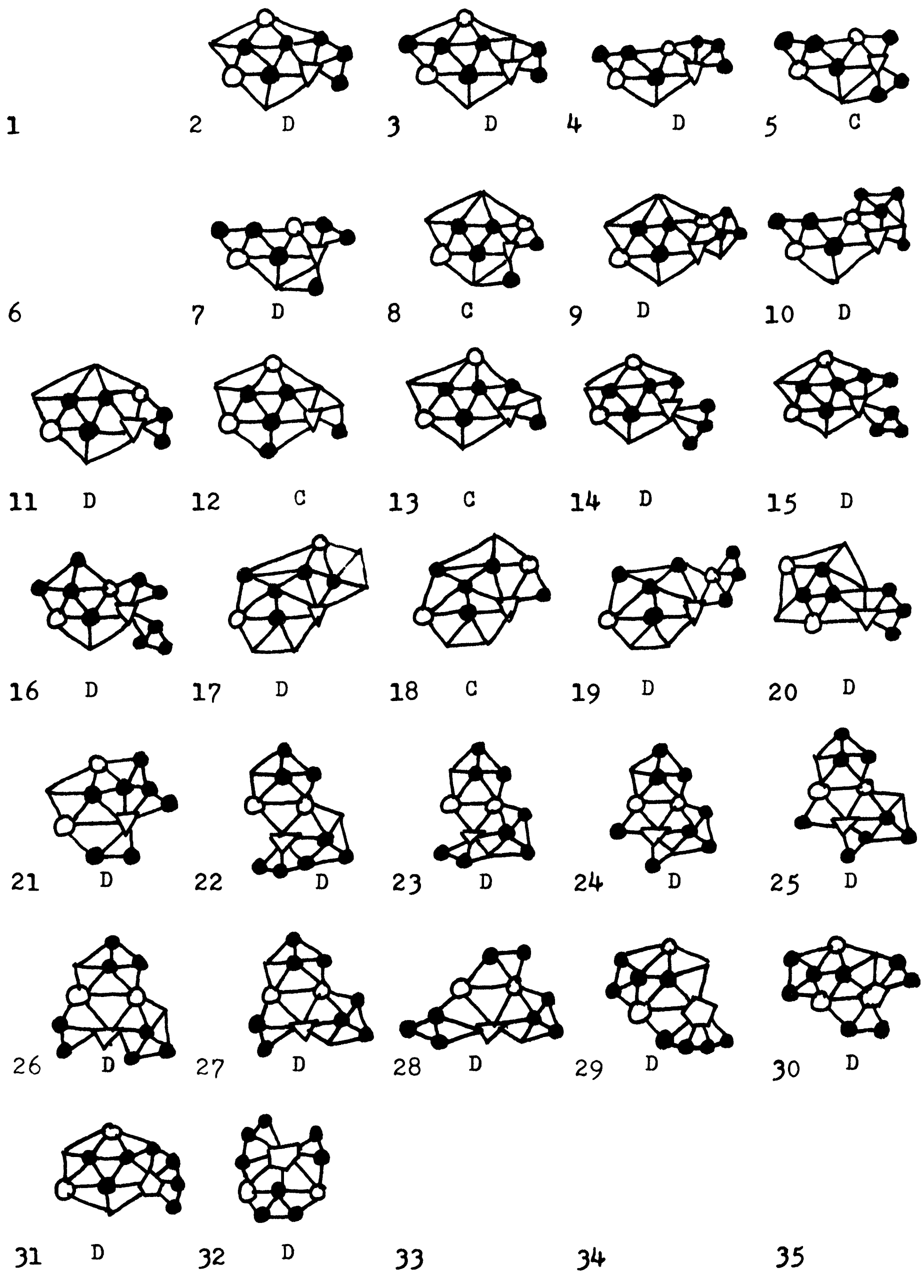
1
2 D
3 D
4 D
5 C
6
7 D
8 C
9 D
10 D
11 D
12 C
13 C
14 D
15 D
16 D
17 D
18 C
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 D
29 D
30 D
31 D
32 D
33
34
35

Figure 60

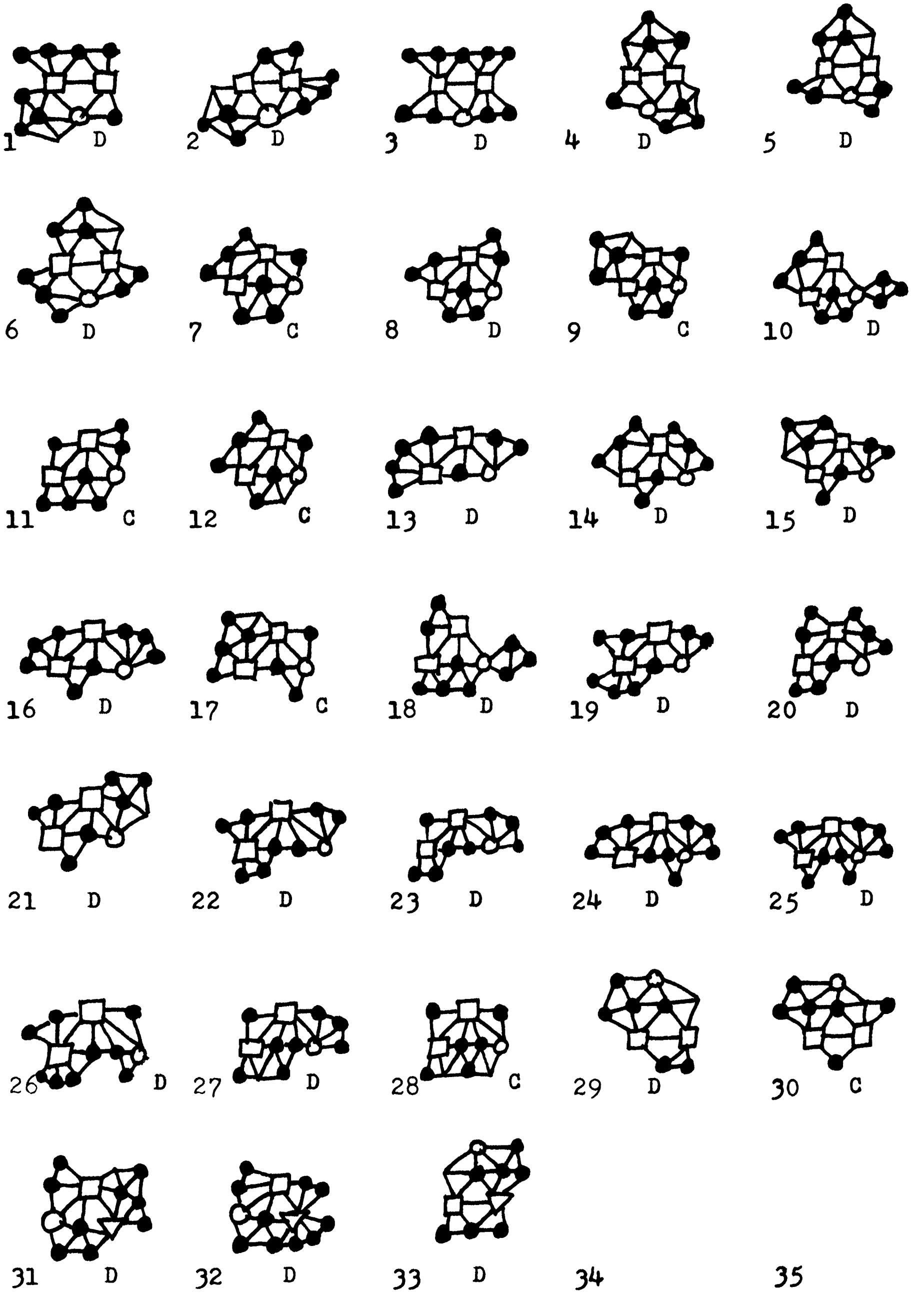
1 D
2 D
3 D
4 D
5 D
6 D
7 C
8 D
9 C
10 D
11 C
12 C
13 D
14 D
15 D
16 D
17 C
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25 D
26 D
27 D
28 C
29 D
30 C
31 D
32 D
33 D
34
35

Figure 61

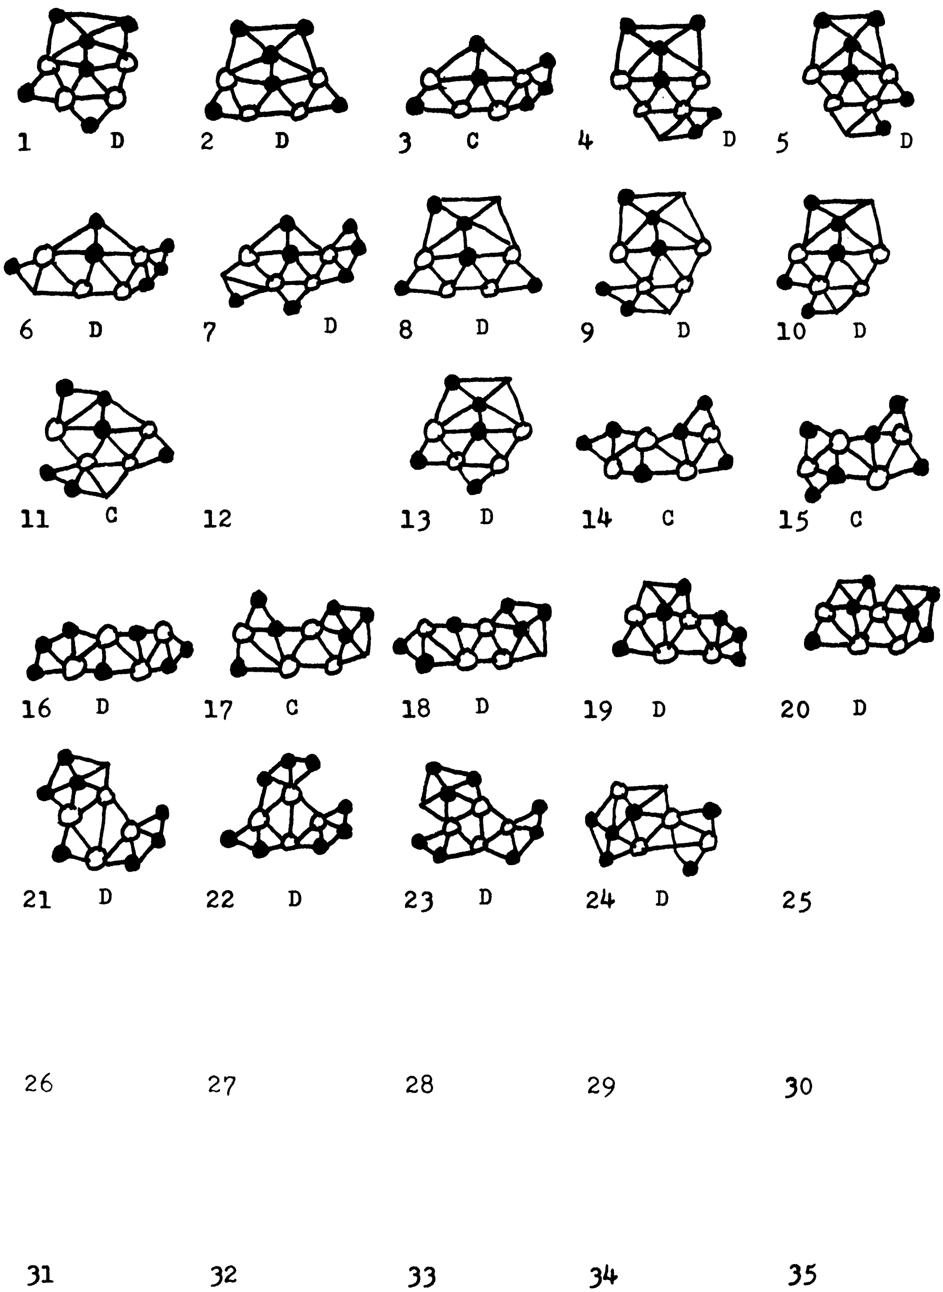
1 D
2 D
3 C
4 D
5 D
6 D
7 D
8 D
9 D
10 D
11 C
12
13 D
14 C
15 C
16 D
17 C
18 D
19 D
20 D
21 D
22 D
23 D
24 D
25
26
27
28
29
30
31
32
33
34
35

Figure 62

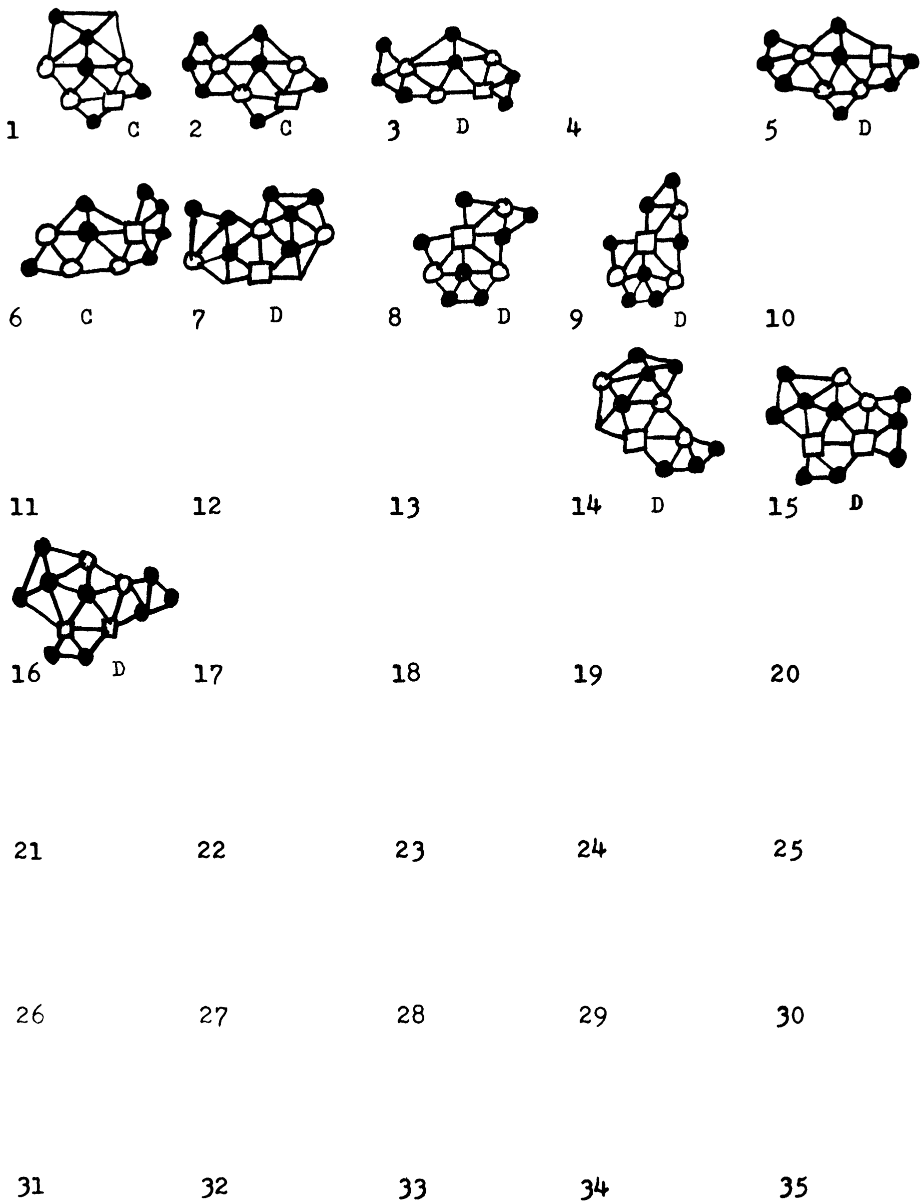

Figure 63

UNIVERSITY OF ILLINOIS
URBANA, ILLINOIS
WILKES COLLEGE
WILKES-BARRE, PENNSYLVANIA