#17

GORDON PASK
Systems Research, Ltd.

A PROPOSED EVOLUTIONARY MODEL

Ladies and gentlemen, I hope you will bear with me if I give an informal talk this evening, because I would like to compress four different things into as short a time as possible.

In the first place, I would like to point out why it is we may be impelled to regard some systems as though there were elements in them that made decisions, because I so believe that if, and only if, we do this can we call them self-organizing systems.

Second, I would like to outline very briefly indeed, a general application of the sort of model we get, if we do look at systems just like this.

Third, I would like to go into a more detailed model, though not in detail, and here you must pardon me completely, please, because I am not going to put the arithmetic of it on the board, and I hope you will take it that I can add numbers together, though in fact I cannot; my equations always make something different, but it works.

Finally, I should like to have the presumption to make a few odd comments intended to tie together some loose ends and to establish the community of the subject. I don't hope to do very much here, but I think there are one or two things which could be said in the light of models of this kind, which tend to unify the different approaches which we have heard about today, and no doubt will hear about in the next day.

The first issue then, is the question: why do we think about systems as though they contain decision-making elements?

Now, we are self-organizing systems and we wander around in a world which is full of wonderful black boxes, Dr. Ashby's black boxes. Some of them are turtles; some are turtledoves; some are mocking birds; some of them go "Poop!" and some go "Pop!"; some are computers; this sort of thing. Now these things we tend to categorize in odd ways.

Some black boxes, I go up to and say, "This thing is a chance machine". What do I mean by this? I mean precisely that I know just what sort of inquiry I want to make, being a self-organizing system, about this thing.

I know that it is a chance wheel, it is a roulette wheel, it has got certain positions where it can stop, and I know—I call it a chance machine because I know this—that if I observe it for a long time, I cannot tell where it will stop next. In this sense, I am uncertain about what will happen, but I am not at all uncertain about what sort of things I ought to look at.

Now, again we go up to the poor computer and we say, "Ah, you are a determinate machine", and so he is; we can take him to pieces and find out what happens inside him, and if we think about noise, this is a thing we agree to exclude because it doesn't refer to the kind of question we want to ask about computers at all.

But finally there is a nasty little class of systems that I think are the ones we call self-organizing systems, which includes you, gentlemen. I go up to you and I have a conversation with you. Now, of course, you are an awfully random thing, because you burble out words. On the other hand, if I can establish a conversation with you, this is no longer the case. Why is it no longer the case? Because, of course, I am uncertain about what you will say next, in the same sense that I am uncertain about where the chance wheel will arrive. But, my main uncertainty about you is of a different sort, it's an uncertainty about what sort of inquiries I should make.

Now, it may be the case that this defeats me altogether, and I cannot talk to you at all. On the other hand, it may be that I can so adapt myself as a self-organizing system, to put it mathematically, I can so change my representation, change the sort of inquiries I wish to make about you, that I can make sense of you.

In other words, it is a deliberate expedient, because, for some reason or other, it would be useful if I talked to you. I adopt this particular procedure of changing the representation in order that this consistency in the behavior of the system, which we can express in all kinds of manners by saying there is a group property in the transformatin of the system itself, shall come about.

Now, systems of this kind we tend to call self-organizing systems. When it happens that we *must* adopt this expedient, whether we like it or not, they are self-organizing systems. But, it is a confusing and blurred one, because there are many cases where we can adopt this expedient if we want to, but need not do so. I don't mind which case we have got. The absolute distinction can be made, I think, on the basis of the Gabor–McKay theory. You will notice that if I had been talking in terms of Shannon information and Shannon communication theory, I would not have made one of the distinctions I have made. I make it only because I separate the metrical and the logical aspects of information. I talk about logons and metrons separately. And I think that this is an important distinction in the present case. And it will be discussed much more ably than I can discuss it, I believe, later. It has already been discussed, incidentally, by Cowan here, and I believe he is going to make some more comments on this subject.

Now there are some funny things about these systems. For one thing, if we say of them that they learn, we cannot really distinguish that statement, because of the peculiar mixup of structural and metrical information, from a statement that they evolve, and this evening I am going to be particularly concerned with the evolution of an apparatus which, in a particular stationary condition of the system, we may then say is a learning mechanism.

In other words, taking the body of this afternoon's talk, the networks which McCulloch's group and Jerry Lettvin, at the physiological level, talk about as logical filters, I take as structures which we can understand *if* we find them. The sort of model I am going to discuss now, doesn't refer in the least bit to how they work, it only refers to how they shall come about in a system which is initially unstructured or moderately structured.

In doing this, I think I assume above all that we drink "Beer"—the pun is no worse than "Torus".

BEER: Touché!

PASK: Now, the kind of system we do have, when we do talk about it as a self-organizing system, is a system in which the elementary particles we are dealing with are not the elementary particles with which a physicist will commonly deal. These are replaced with unitary elements which may be considered to be automata, players, decision makers, "neurons" or the like. They can go "Poop!" and send a signal to another, the implication of the signal being that the state of some remote element is changed by the fact that this one goes "Poop!"

In order that they shall go "Poop!" they must feed. I do not mind how you represent the feeding, but it is important that they do feed. It is a condition on the measurability of the system rather than on its energetics, but it is convenient to present this for explanation as though it were food or energy coming into the system that is required in order that the signalling activity take place. It is a measure on the system in the sense that this is the way we are going to talk about and find out about the state of the system. It is a conservable quantity.

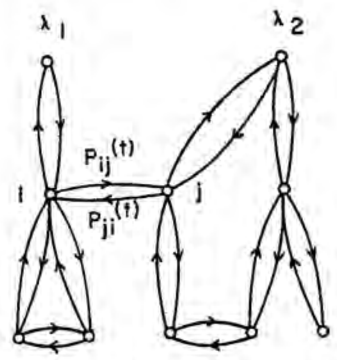

FIG. 1. Formal representation of food distribution network.

We consider that food enters a system and goes through a diffusion network. All I mean by a diffusion network is a system of tubes and basins, say, over which we can define food neighborhoods. A formal representation of a food diffusion network is shown in Fig. 1. It is a directed graph with nodes. The lines connecting nodes have quantities associated with them that represent the food impedance, the amount of resistance to the passage of food between nodes. The nodes at the top represent food sources; those at the bottom represent the nodes accessible to our elements, at which they feed. Such a formal representation is only to insure that we can define food neighborhoods, so that if one element feeds at a node we can say that it will deplete the food available to another element at a neighboring node, but will not affect the amount available to a more distant zone.

Now in addition to our elements and a food distribution network, we require a material in which the signalling of the elements builds and maintains signal pathways through the expenditure of food. Hence there is a signalwise connectivity among the elements; erected on the stage of every signalling activity is a structure which is determined, made, and maintained by the expenditure of food. This structure cannot exist and persist independently of the activity of the system, it comes into existence as a result of this activity and is maintained by it. Turn the system off, and it all disappears.

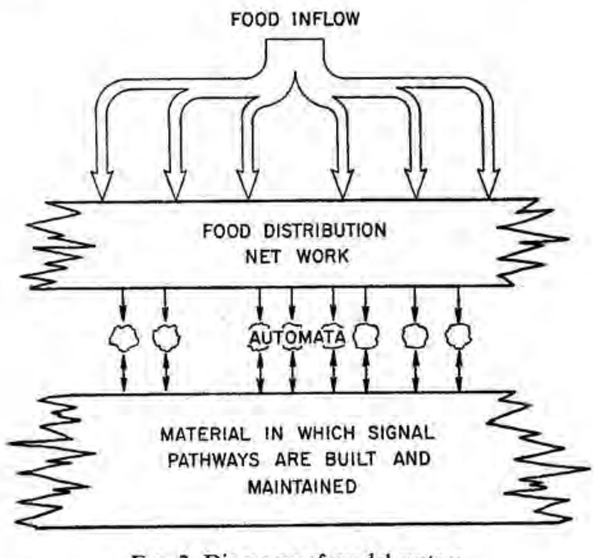

FIG. 2. Diagrams of model system.

Figure 2 is a diagram of the whole system. The important thing to note about such a system is that, as a result of the signalling activity of the elements, one builds up the patterns of structure by which the elements become connected. Those elements which become connected can then be in a position to correlate their strategies.

In particular, I am interested in those connections such that two elements can signal reciprocally. Such a connection would be represented by a cycle in the graph of the connection matrix. In a situation such as this, it is possible for two elements to correlate their strategies completely.

It is important to remember, because it is crucial to the next stage of my development, that we have decided, for reasons based only on our selfish considerations, only on efficiency or utility to our own observations, to call these elements "decision makers". We do not need to inquire what in fact is inside these elements. This is irrelevant to the matter.

Hence if I call such an element a player, and then go on to the alarming statement that such a cyclic connectivity can be a coalition, you will perhaps not take exception to me. The system is closely coupled, and if I regard the elements as players, I will call this a structure which can permit coalition. Please note that a coalition in this system is something which must have a structure associated with it, because it is nonsensical to talk about these things correlating their strategies if they cannot communicate, and they can communicate only if they establish and maintain a communication structure by the expenditure of food. There is a communication cost implicit in their bringing themselves into the same signal neighborhood.

Now, the point I want to make at this level, by considering this first very general model, and before we go on to the more accurate one, is the fact that we can derive some interesting conclusions about what coalition structures can exist for a given payoff function in the food network.

If the payoff function which we derive from the food network, the diffusion network, is the payoff function of a competitive game, so that no advantage is gained by the players cooperating, it is nonsensical to have the idea that this can form a self-organizing system. The self-organizing system is something which occurs when cooperative activity is favored or, to put it concisely, when the payoff function determines an essential game. When it does, there is still a restriction upon the coalition structures which can form; the restriction is introduced by any sensible assignment of the cost of maintaining these coalition structures.

For example, we have, in Fig. 3, examples of coalition structure involving the coalitions of α, β and γ. The second one involves a cycle, a single cycle, and it has the same maintenance cost as four independent elements. The remaining are also structures able to realize these coalitions, plus others, and they cost more.

Looking at Fig. 3 in a little more detail, we can see that there are different maintenance costs along here, assignable to different coalitions of linear elements, for example, there are the trivial coalitions of α and β alone, for, N, the number of players, equal to two. This is the only cost for the coalition of α and β, it is a unique affair.

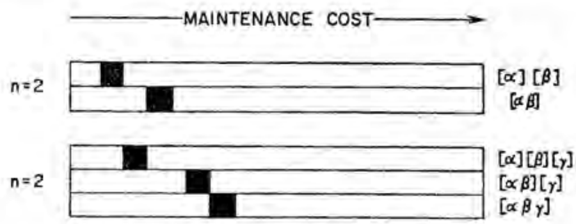

FIG. 3. Maintenance costs of coalitions.

In the case of coalitions of three, we have a cost for α, β and γ separately per player. We have a cost for α, β and γ separately, and we have a cost for the α, β, γ coalitions of which there are two, and so on for four and five players.

In other words, there are discrete maintenance cost levels which can be realized in this system. We may also add to this further constraints comparable to Luce's Psi function, which determine what possible coalition structures can emerge.

Now, a number of quite intriguing things occur when we consider how, when looking at a system like this, we might be able to make sharp, well-defined observations upon the state of the system. It is intriguing, for example, to plot different levels of maintenance costs at which coalition structures can occur (Fig. 4), and to consider what will happen if we make the surplus amount of food available to the system U^*, decrease, we can plot on the same time coordinate the probability P that a given observed structure, that is to say, a given connectivity F, equals a coalition structure, c; in other words, the probability that a given observed structural entity mediates a particular function.

Clearly if there is a level at which, say, coalition X can occur, then when we get down to the cost at which it just can occur, if it exists at all, then it is certainly being used, for otherwise it would collapse.

Hence, we have a local maximum in P. As we decrease U^*, we know that X cannot exist but Y perhaps can. Supposing Y does exist, we will have another maximum when we reach that point. In general, we will get a curve of this kind, as we decrease U^*, the local maxima of P will be the sharp-valued observations, the points at which we can make definite statements about the system.

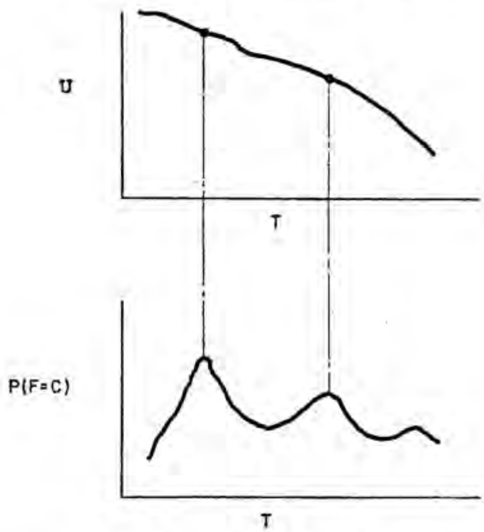

FIG. 4. Observation of coalitions.

This phenomenon is not unobservable in sociology. It is confirmed by anybody who plays around with the sort of network which, for example, Paul Weston makes with ionic resistors and neons; as you decrease the standing potential the thing locks into a stable state.

But this perhaps is a trivial case. The more important cases are in those systems, say, like your systems, Dr. Beurle, where you have critical values at which certain forms, certain modes can be maintained.

Now this is all I really want to say about the very general model, because I think that a certain clarity comes into the discussion when we consider the most primitive possible model we can conceive.

Now, for this purpose I have taken Edwin Abbott's "Flatland" as a universe. If you remember, Edwin Abbott wrote a little scientific fantasy about dimension. In this he was a two-dimensional figure, a square. There was a hierarchy of these figures, ranked according to the number of the sides they had, and they were supposed to have evolved, in some distant time, from one-dimensional creatures, and in fact he made a journey to "Lineland". Also it was possible that there were figures—indeed there was a missionary sphere—which lived in three-space. It is on this sort of format that I made my little model, and you must pardon its simplicity, which comes largely from having no computer and only a desk calculating machine, and a gentleman to work it for me.

What we did was to suppose a food distribution network in which the subset of nodes that are accessible to our automata for feeding lie in a plane. This accessible subset I call the field. It is obvious that we might apply other geometric constraints on the field, have it be the surface of a sphere or a torus, with consequent changes in the kinds of communication structures possible.

Now we will think of primitive little zero-dimensional automata, of which, since we are at least two-dimensional, we can distinguish two species: those that move longitudinally and those that move latitudinally.

But what does an automaton do? Because I drink "Beer", I have defined an automaton as that which is designed to survive in a specified universe. This means that the automaton itself is going to be subject to the same conditions that affect the coalition structures we spoke about. The automaton is not something that exists and persists in its own right; it has to pay for its creation and existence. I suppose its creation to occur by a process I call nucleation. If the food stored at any accessible node accumulates to a critical value, an automaton appears there.

An automaton, when it appears, sucks in food and builds this food into a communication structure, which structure is subject to degradation and must be maintained. The rate of feeding depends on the local concentration of food; the rate at which it feeds depends on how much food there is.

Here I want to point out an important thing about these automata, about any automaton which can be said to be designed to survive, designed to compete. To make my point, I will distinguish between two classes of automata. The first class includes those we most often come across. They are things which are able to make decisions, moves, signals or whatever. They do so on the basis of accumulating evidence about the activities of other automata and possibly about conditions in their environment engendered by changes other than the activities of their fellows. This evidence is signified by a vector of some sort, and the values of this vector are piped into a decision rule, the output of which is a move, a signal, or whatever.

It is conceivable, and note that it must be conceivable in the universe in which we define competition as the thing we are looking at, that such an automaton can encounter a situation which is undecidable. In such a case, automata of this class, the sort we encounter in "hill-climbing" devices, are given a fresh strategy from outside. They call for independent information from outside, so that they are no longer a closed system. They are invaded, as it were, by a wheel of chance or a table of independent numbers. They just ask for a number, and this decides their undecidable decision for them.

Now this is not the sort of automaton I have in my model. I have one of the other class, which, when presented with this same dilemma, either evolves, or dies. If it has enough substance, it evolves; if it does not, it's had it!

The manner of evolution can be expressed rather precisely in the language used either by Dr. Ashby or by Dr. Rashevsky, one of them in terms of states, the other in terms of biological functions and the graphs of these. The results of evolving will always be that the automaton which found a certain situation undecidable now becomes a larger, more complex automaton which can comprehend a larger world in which the situation may not be undecidable.

In our trivial little universe of these creatures moving around, we have not given a great many facilities to our evolving automata. In Fig. 5 I have shown the possible moves of the primitive automata.

As I mentioned, there are two species, a and b, capable of one-dimensional latitudinal and longitudinal moves, respectively. Each of these primitive automata is capable of just three moves, either up, down, or stay put, or left, right, or stay put. If he stays where he is he sucks up all the food and dies.

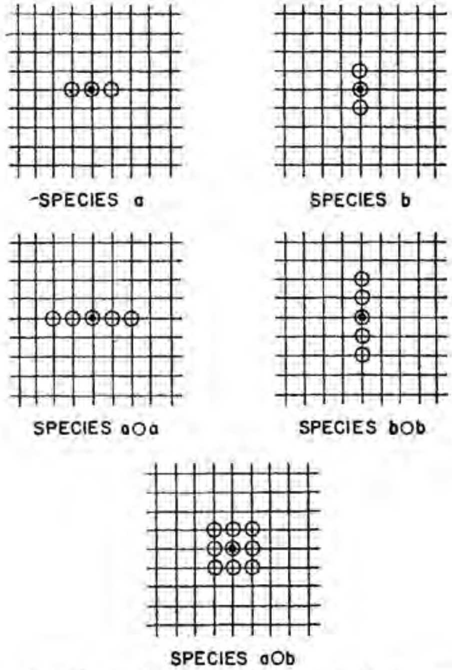

FIG. 5. Moves of primitive and evolved automata.

The evolved automata, of which a few are also shown in Fig. 5, are evolved by composing two or more of the primitive automata. If we compose two of the same species, we get beasts capable of five moves instead of three. More viable are the creatures obtained by composing one from each of the two primitive species. This creature is capable of nine moves, and has the advantage of having a two-dimensional move neighborhood. These compositions can be continued indefinitely, with each new species having all the moves of its predecessors. Although the move neighborhoods can never be other than squares or rectangles, the individual moves may be quite peculiar.

The evolutionary rule is exercised when automata get into difficulty, and when, having got into difficulty, they come together. The difficulties are engendered by the characteristics of the surroundings, the distribution of automata in the field and the food supply at the accessible nodes of the food distribution network.

I would now like to talk about the characteristics which we impose upon the food distribution network in order to be sure that out of this model there arises something which, when embodied in a fabric, but not on its own, would be a self-organizing system. The constraints are simply as follows; we say that when an automaton is given play on the accessible nodes of a food distribution network, the food impedance, that is, the impedance between the accessible node at which the automaton feeds and all inaccessible nodes with which that node is connected, is a function of its eating and of time. We make the weight of the connectivity of an accessible node to other nodes change as a function of feeding activity at that node.

Now, I would like to illustrate some of the tricks which these automata get up to when you actually run this model. Consider an automaton of this sort in a plane and consider just a line in this plane (Fig. 6). The food concentration at the nodes along this line are plotted on the left, the position of an automaton on the right.

Supposing we start out with a certain local perturbation of food concentration. Since one value is a little less than an adjacent value, this gives the thing a direction, so that he can sense that the food concentration here is higher than it is there. This determines the creature's move in that direction. Having determined this, it eats, and this determines its subsequent movement. If it is in an indefinitely extending plane, it just goes on until it encounters a boundary or another automaton, and it leaves behind it a wave of food depletion as shown in the lowest diagram of Fig. 6.

If you have rather more automata in the system, you get structures which are chain-like structures, due to the fact that automata tend to become nucleated and move into a region, where they approach each other reflecting each other, so that you get chains of oscillating automata which form coherent structures in the system. This is still in a plane.

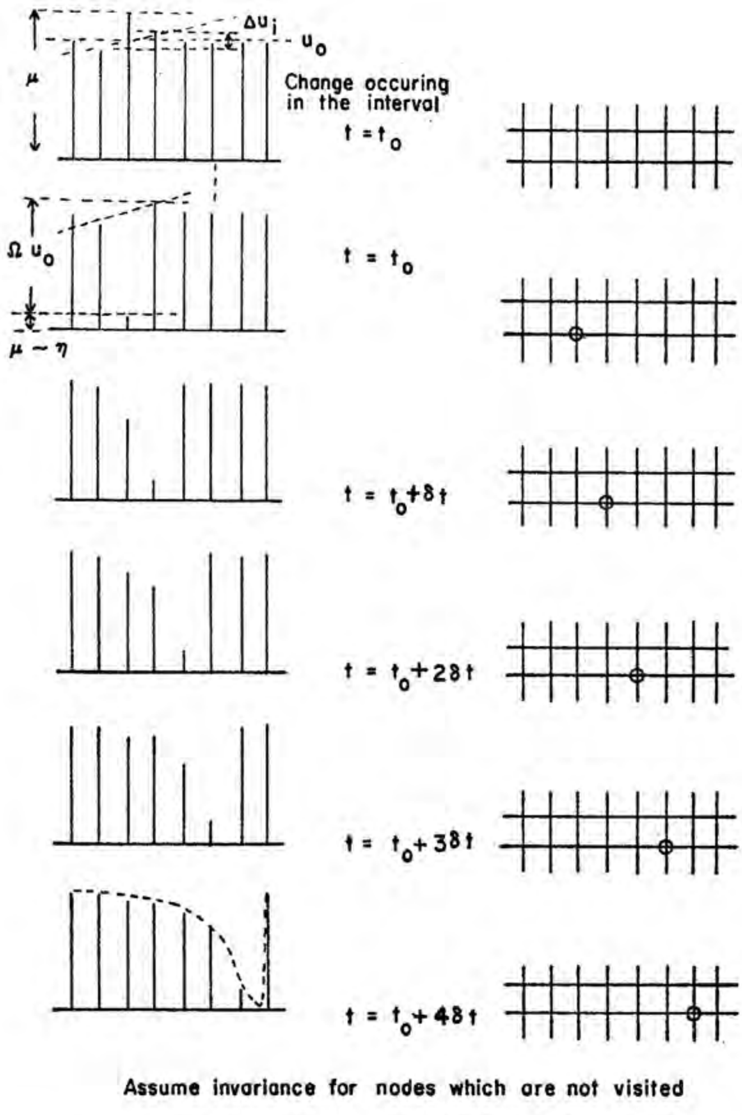

FIG. 6. Element movement and food depletion.

If we modify the topology of the thing a little and make, for example, a cylinder, we come across the possibility of cyclic action. Such cycles can act, rather obviously, as a sort of memory device, but they can also act as filtering devices. Supposing I establish a cycle in a cylindrical field (Fig. 7), and I establish a gradient down this supposedly indefinitely long cylinder, so automata tend to jump with the gradient, then the cycle of an automaton in the

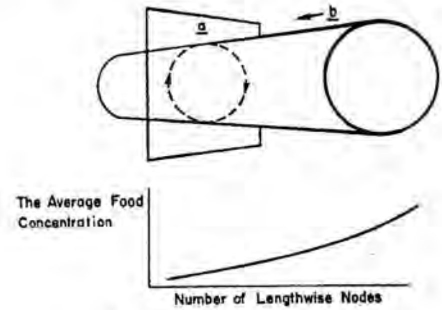

FIG. 7. Cylindrical field model.

transverse plane selectively prohibits or allows other automata to get through. Hence there is a certain sense in which filter-like structure, in terms of automata, are readily built up. For example, an automaton of composition *aoaobob* would be able to jump over this cycle altogether, it would not mind it, or at least a certain percentage of them would not, depending on whether they were on odd or even numbered nodes.

In sum, when you muck around a bit, and in particular when you define two neighborhoods with respect to the supply or source of food, it is possible to establish, in such a system, both cooperative and competitive activity. The dotted line in Fig. 8 is intended to show a source of food and, in a cylindrical field, you will see there are two alternative and energetically equivalent cycles, γ_1 and γ_2, which can be established in this system, and which are liable to change into the other. A particularly interesting system, which I will not discuss at all, but which I think is worth notice because of the structures, which may be discussed later, is that of a multiple-genus torus, where you have got the possibility of independent cycling where different species can come together (Fig. 9). I have a hunch it is no more than this, but so far we have not been able to realize such a field through lack of time or facilities, but it should be interesting to do so. I would just like to make one

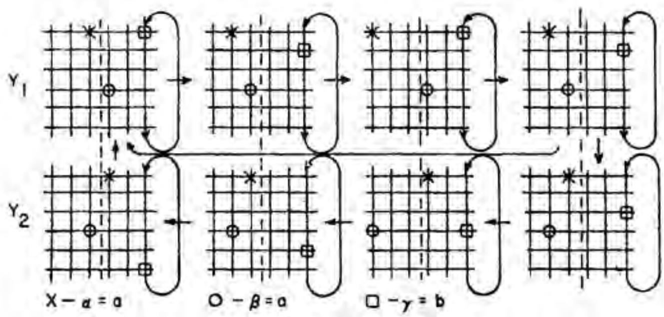

FIG. 8. Oscillatory movement in cylindrical field.

conjecture and please take this as such: that it seems that one way of introducing synthetic *a priori*'s into a system of this sort would be to produce topologies of this kind.

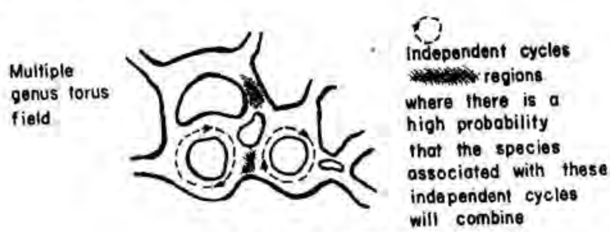

FIG. 9. High genus field.

Now I would like to pass on from this to considering what happens when we have a whole lot of automata interacting with the food supply network on which they live. Clearly in this case, there is a perfectly good sense in which the activity of an automaton, and in particular, a species of automaton, generated by the evolution permitted as a result of their previous feeding, will structure the world around them so that this particular species alone is favored.

Hence a system of this sort, and it can be argued quite rigorously, is a self-replicating system. Furthermore since these structures, being geometrically bounded, are constrained, there will be a finite size to the structures and things will tend to come apart when they reach this critical limit. What I would like to do is give a special name to this odd kind of structure, which is a close coupling between a lot of automata and the world they live in. I will call it a domain. I will suppose that the domain is an existent in this sort of system, for there is no chance to discuss it adequately at the moment. I am particularly interested in what happens to a domain when, for example, we give the elements a lot of food. I

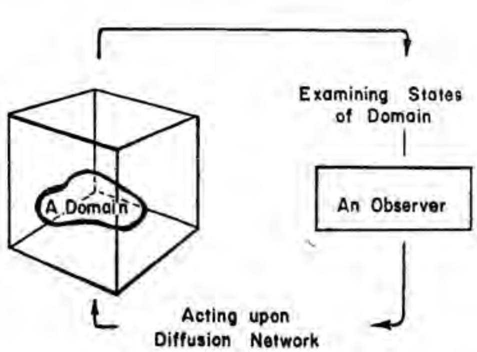

FIG. 10. Domain observation.

am particularly interested in this because it occurs to me this is one of the ways in which a game-theoretic explanation leads to clarity where others do not.

For example, suppose a three-space, in which we have got these creatures wandering around, and we have a domain (Fig. 10). Suppose I, as an observer, can by some miraculous means put my finger on this ephemeral entity and say that it is an organization which is wandering around, it is automata in equilibrium with their environment, which they modify, which is wandering around in a cube. Supposing I could put my finger on it in this way, I would be doing an operation, if I allowed it more food, if I favored this entity, equivalent to a linear transformation of the payoff function of the game. This is an entirely explicit thing to do. I simply add to each entry in the pay-off matrix, a small positive number. The result of this is that more automata can live within the domain, and they will be the sort of automata which happen to be playing in this region. I will be favoring just those players.

Or, if you like, you can regard the whole system as having a gigantic payoff function, and you can say I am favoring the activities of a given coalition, for indeed, there is a certain identity, a certain similarity rather, between the domain and the notion of coalition which we advanced earlier.

Now, supposing I do this—and I agree that it is not a thing I can easily do—what happens is that the density of automata increases in this region. Here I think we get a very interesting result from our model, which I think has been repeated often enough to assert it. At a certain critical stage it locks solid. And at this stage, we must change our whimsical, though not unnecessary description of what happens. Instead of talking about automata in the region of this highly rewarded domain, we must now talk about chains and structures which exhibit exactly the logical characteristics of a model nervous system.

We have a refractory interval, a partial refractory interval, an impulse which is transmitted with a wave of food depletion. If we "Poop!" at one end of the chain of neurons of the sort I illustrated on the board, we get a result at the other end which is transmitted by this local energy depletion process. It does not much matter where the automata move, because they are not allowed to move very far. They are constrained by their own kind. It does matter a great deal in what order they move when they move.

The domain always locks solid like this, and at the end of the chains there are link violations which quite obviously have temporal and spatial summative characteristics which, hopefully perhaps, I would like to identify with synaptic connections.

It seems to me of interest that the operation of this hypothetical "thumb-putting-on" procedure will lead us, with the choice of parameters which, perhaps, I have taken as a hunch, always to this result: that the center of the domain, regarded as a critter, walking around a world in which it feeds, becomes structured and acquires this rather net-like sort of nervous system.

And I think it is also interesting that I cannot really describe what is happening in terms of putting my thumb on it, but I can describe quite precisely what is happening in terms of the game-theoretic model.

There are one or two other things which I think we might point out. I did make a model in these terms. I ran a small program of these terms, to describe the development of a population of small social amoebae, cellular slime molds, which seemed to turn out quite successfully. I set up conditions whereby we had these

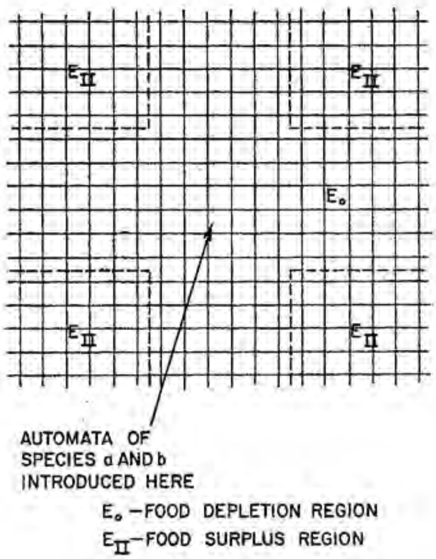

FIG. 11. Maze experiment.

creatures wandering around in a regionally depleted environment, and I added to the system the possibility of specific signalling. I gave them some "acrasin". And the result of this experiment will be perhaps interesting, although I cannot describe it in detail for lack of time. Figure 11 shows the field. We found that without any signal potentialities, given an experimental setup in which we have a region where there is food depletion and a region where there is plenty of food, and a structured network, then to get from one to the other you had to be a coalition, you had to combine. We do not know how many get out when you grant the possibility of signalling, that is to say, when you vary the signal possibility given to each individual one so that they can track their kind. This sort of model is interesting, and interesting perhaps in the same way that the structuring one is.

Finally I want to comment on the payoff, because we have gladly supposed it is food, but really in some senses this can be ridiculous. Alternatively, it is perfectly possible for a domain to feed on automata. There is no real reason why we should take, in the gigantic cube of Fig. 10, only the supply of food as being that

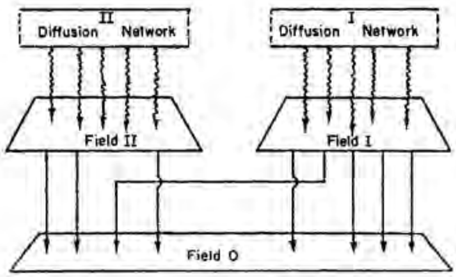

FIG. 12. Domain migration diagram.

which limits and causes competition. Equally, we may regard the migration of other domains of particular families of automata as being a source of food, or if we like, as initially a kind of catalytic action, since the specificity of these automata with respect to the domain into which they migrate can induce an amplifying action, which is, I think, rather easy to conceive.

The sort of thing I am thinking of in broad outline is shown in Fig. 12. Here, I suppose, a diffusion network and another diffusion network, two fields, and automata which can migrate down in a cubic lattice, for example, onto another field where there are other domains, where the migration of these would affect the survival or extinction of the domains that exist below.

So we have got into the realm of payoff functions which are defined in several ways. And I think this is again a very interesting region.

Now, finally, what is this all about, why should we play with these automata? I think the reason why we play with them is in order to find how the structures which have been described as the structures engendering logical stability, redundancy of potential command, and redundancy of computation, occur, in our mushes, our threads, our Dr. Beurle's networks, and anywhere else we please, all over the place.

How is it that systems like this occur, systems which are essentially characterized by a non-zero sum partially competitive game in which, in addition to the usual concepts, structure must be paid for?

How is it that these domains acquire, as indeed they do, as is obvious from the discussions, just these characteristics? I do not know how it is, but I do hope by means of a much more sophisticated and a much better presented model we may find out.

DISCUSSION

VON FOERSTER: Thank you, Gordon, for a most delightful paper, and I think if somebody was asleep at the beginning he is now wide awake. And I would like to ask first one question, Gordon, if you may, perhaps, give a little explanation of the particular experiment with the amoeba, so that everybody really can follow what kind of a setup this was, and so that we really see what happens in this particular instant.

PASK: The conditions are as follows: we wanted to simulate a cellular slime mold population; if you recall, the life cycle of a slime mold is a crazy mixed-up thing. The creatures start out as amoebae, and they live anywhere on the place and they just look for food. These amoebae have a rather remarkable signalling system. They produce a substance called acrasin, which is simulated by certain steroids. This diffuses away from them and the appearance of acrasin, at the boundary of another amoeba, causes this amoeba to move towards the source, in a chemotactic manner. So if you look at the culture, you will eventually see aggregates and there are streams of amoeba moving into them. When they have collected sufficiently, they form an organized whole, in which cells appear to take different functional parts. There is a definite hind and middle and fore end of the thing.

Ultimately, and this is the part of the model I haven't realized, because I would have to put in a duplication rule which I haven't mentioned, the thing forms a spore body on the end of some sort of structure, depending upon the species; they form all sorts of peculiar kinds of cellular structures. A spore body is formed which leads back to further amoeba, which are disseminated by wind and water.

Now, this is the life cycle I was anxious to induce, and I was anxious to induce it given simply the constraint that it was possible for a certain sort of structure to exist, in a food-depleted environment, where the original members of the population could not exist. And I wanted to see how their existence was determined by giving them the acrasin signalling system which gets them from here to there.

The experiment was set up in the following manner. I took a field as shown in Fig. 11. Food concentration at EII regions is high enough to sustain anything. Elsewhere we have a depleted region; the amount of food therein is dropping all the time. There are automata in the depleted region, they are eating and taking away the possibility of living as simple automata, but they can only survive if they are automata that can move to the EII regions.

So in order to get through the depleted regions, they have got to combine. The experiment was a very simple one indeed, namely, to plot the number that escaped as a function of the amount of acrasin given them.

RAPOPORT: Why did they have to combine in order to get through?

PASK: The automata introduced into the depleted region of the maze were only of species *a* and *b*. To reach the EII regions they had to combine as *aob* automata.

RAPOPORT: These are real flesh and blood amoeba?

PASK: I am simulating the amoeba's life cycle. The matter is simply that in order to get out of the central depleted region, it is necessary for an automaton to turn at least one corner. In order to be able to turn a corner, it must be a combined sort of an automaton.

YOVITS: Is this a geometrical combination?

PASK: It is a geometrical combination.

YOVITS: It is sufficiently large so that part of it remains in a food-giving region, is that what you are trying to say?

PASK: No, there is a food-depleted region in which it is just viable.

YOVITS: If the whole automaton is within the region which has no food, will it die?

PASK: Well, you start out with simple automata. Now they would all die, because they eat all the food. So that in order to get out they must get through this region where they would die before they got through it. So in order to get through the region they must turn at least one corner.

YOVITS: After they combine, how do they get through the region?

PASK: The boundary here actually is in terms of having to turn the corner. They have got to be combined automata in order to get out.

NOVIKOFF: It isn't that their food demands are less when combined?

PASK: No, no, definitely. As a matter of fact, I have another model in which this is the case, but that particular one I was describing here is one that has to turn a corner.

ROSENBLATT: In this problem, this model actually produces a depletion of food there around it?

PASK: I am modelling in terms of my stupid little automata. But why do I model in them? Because it is the simplest system I know of in which there is a definite cooperative action which increases, in other words, cooperation occurs in conditions of depletion, and a definite cooperating organism is formed under conditions of plenty of food.

YOVITS: This is a hypothesis that if they turn a corner they will get out?

MCCULLOCH: May I speak in answer to that? The point as I see it, there is such a beast as an amoeba, but he is not talking about that beast, he is talking about a model composed of straight lines.

ROSENBLATT: But how does the real beast get out itself? Does it have to organize to get out too?

PASK: Yes, it has to actually, it gets out by sticking together. The act of turning a corner in my model is equivalent to the act of sticking together.

VON FOERSTER: I have a particular question. Now we have talked at the moment of actual living systems, but what I was thinking, in terms of Pask's model where one could really assign a mathematical functional way of formulation, or, if you would like, to use the second law of thermodynamics for open systems, one could assign to a model the precise values which would be necessary to carry out these confirmations. From this point of view, you see, we could make, a real one-to-one correspondence.

In this simple model, I think it might not be too complicated to begin a simple correlation between these two fields, and then develop or try to apply it to the actual system.

PASK: In other words, do it without hope that it may be a testable model.

BOWMAN: Just a very brief word. I speak now not on the subject as presented, but from purely a biological standpoint. I have at first hand observed the movement of a slime mold* in the plasmodium stage across a dusty dirt road. Now what particular gradient could have induced that I have not yet found out. Yet the animal has no eyes and apparently no structural organization at all, and yet it was deliberately crossing that dirt road.

ADDISON: Does it ever go down the road?

BOWMAN: My research was incomplete.

SHERWOOD: Did you try to deviate it?

BOWMAN: I did, it went around an obstacle and continued on the path.

SHERWOOD: 180 degrees?

BOWMAN: No, that I did not do, I put a board in its path, a dry board which it did not go over, but went around, kept on the move.

YOVITS: What does it do after it gets across the road?

BOWMAN: I followed this for about ten days. It finally found an old, rotten log and formed its spore stalks.

ADDISON: Did you try digging the soil it was on and turn it?

BOWMAN: My research isn't that complete.

BEER: How big was that colony?

BOWMAN: The size of my hand, about that size and shape.

PASK: That is a big one.

ROSEN: Just a little word, in a very trivial sense, on the growth of domains in ferromagnetics and ferro-electrics. They have many of the properties that you outlined except it is quite simple, and there the equivalent of your food would be the energy of the whole system.

PASK: Yes.

ROSEN: What is real interesting in the terms of the geometry you described that is required for this to get out of the food depletion area. In the ferro-electric domain, if you have one domain neighboring another domain which is oriented ninety degrees to it, this is a very stable condition, nothing happens.

On the other hand, if there are two domains that are oriented 180 degrees from each other, it is quite possible under the action of a gradient, which in this context is a field, for the larger domain to swallow up the smaller domain, and it does this in a peculiar fashion, putting out wedges which grow.

YOVITS: But is this the same phenomenon, isn't the essential action of these domains one to form minimum energy? Now these domains you are talking about do not tend to form minimum energy.

ROSEN: There is one more remark I want to make, nothing happens until

---

* *Editor's note*—Mr. Pask has commented that Dr. Bowman appears to be describing one of the acellular slime molds. Mr. Pask's model has reference only to cellular slime molds with an "acrasin" signalling system.

you get a lot of cooperation among a bunch of dipoles. An individual does not do a thing. As I say, this is very simple.

YOVITS: Because it happens not to be a minimum configuration.

ROSEN: It happens that in physics you have a nice principle to go back on. Perhaps one can find such a nice simple unified principle here.

PASK: Yes, look, I think the point I will perhaps bring out, and I think that your comment, Dr. Rosen, is not at all without relevance. I think it is a highly relevant thing, but I think there is a distinction to be made between these two concepts, although they are related.

ROSEN: One is very much simpler than the other.

RAPOPORT: I would like to ask a couple of questions. First with regard to your calling this a game-theoretic model. If I understand it, it has to do with the game-theoretical model inasmuch as it applies to coalition formations, each coalition commanding a certain payoff. Therefore, we are reminded of the game in characteristic function form, in which each coalition commands a certain payoff, but when one asks questions about what coalition you actually form, one finds that immediately one gets into hot water, because this straightforward answer, that that coalition will form which commands the greatest payoff, is by no means the case. Because if such a coalition should form, which commands the greatest payoff, it means that the anti-coalition would also form because that's the best thing that can happen to the other side, that would command the smallest payoff. And then one is faced with the dilemma, if the coalition which commands the greatest forms, why does the other coalition form, since it commands the least?

The whole meat of the coalition formation theory is that there is competition for the members. So that it is not at all true that the coalition which commands the greatest payoff will form, and indeed, the theory of games in characteristic function form, as you well know, in its original formulation, didn't have any answer whatsoever. Except for von Neumann's solution of games, they are a laugh. They are ridiculous, they tell you absolutely nothing. In most cases every coalition is a solution. Every coalition is an imputation, is a solution. So that one cannot even say anything about what will happen.

But what I want to ask you is whether, aside from these general natural selections which you were talking about, there is anything else that can maybe shed light upon this respect? That is my first question.

My second question has to do with something that you started with, and I thought that you would elaborate it further, I would very much like to hear your opinion on this.

You said the self-organizing systems are characterized by this peculiar trait, that when one looks at them one doesn't know what questions to ask. Is that so?

PASK: I said that there is a class of systems where we do not know what questions to ask.

RAPOPORT: Right. And I thought for a moment you were talking about a situation similar to the following one, and wondered if it has any connection.

There was a man that came from Mars, and he began watching a chess game as they were playing it on earth, and he decided that having watched 10,000 chess games he had a mathematical theory of chess, and he went ahead and developed it, and that theory was simply wonderful. It predicted with great accuracy what the rate of depletion will be of the chess pieces, and it is replicable; if you take the first 10,000 and the next 10,000 games, the rate of depletion will be exactly the same.

And he even developed differential equations from which these rates of depletions were deduced. He developed equations which told him the distribution of the length of the game, of the probability of white winning over black in every kind of game played, and vice versa. In fact, he developed every possible statistics of chess you can think of, and he developed a very good set of fundamental axioms from which all this mathematics was developed.

And then he brought his theory to an earthly chess player, and he told him that it wasn't chess, he told him he had asked the wrong questions.

Now, does this have anything to do with what you are thinking of?

PASK: It does have something to do with it.

RAPOPORT: Would you please comment then further, and also my first question, please?

PASK: I agree that your comments about the von Neuman n-person game and the coalition formation in it were entirely valid, and I think that the best way to answer the question, and to expose the possible utility of our model, is to develop slightly the conditions which arise within it.

In the first place, we are not thinking so much of the von Neumann model, as of the kind of model proposed by Luce, in which there is a thing called a Ψ function which is absolutely central.

Let B represent the coalition structure, that is the aggregate of coalitions, of the whole set of coalitions, of the game as the time T equals T zero; let ρ represent the correlated strategies adopted, in equilibrium, which will be some solution by these numbers of these coalitions. Then $\Psi(B\rho)$ is the set of coalition structure, common strategy pairs, which are admissible at $T_0 + \delta T$.

Now, we are interested in this in the following sense, that in a sociological situation we commonly have to guess at Ψ, which represents the social inertia. In the present kind of automaton, we are in a much happier situation, because even if we make it out of threads or goo or semiconductors, we know the Ψ function; we know the admissible coalition structures which can occur. These arise from the inertial parameters of the system.

RAPOPORT: There is a permissible transition in the system.

PASK: It is a permissible transition, given those in a given stage. The same is true here, surely. In the sense that: given I have a certain coalition structure, I can, for example, take a K game and just add one, or whatever it may be, whatever rule exists for this possibility. But in the K it won't be quite as simple as just adding one or subtracting one or something of this kind. It will, in fact, be more complicated, and it will be determined by the characteristics of the funny material at the bottom, and by the time constants of this and so forth. And so a certain amount of structure can come in this way.

The second thing is that we surely can make some assertions, although I agree that they are rather poverty stricken, about the payoff functions which can be in equilibrium with a given common strategy, a coalition pair.

So that looking at the thing again a little crudely, we can talk about equilibria which consist of a sequence of payoff functions, induced by the existence of a coalition structure here, and another one which arises because of this, and then the arrival of a coalition structure which is induced by the payoff function.

If it happens that transformation T of G_1 and G_2 form a group and it is a cyclic group, we have a stable condition. And these kinds of stable conditions are analogous, perhaps, to resonance hybrids, because it always happens that food maintenance cost of one coalition structure must map into another if the food maintenance cost, the average payoff, is constant over the set of G's.

Question number two, the man from Mars, yes, I think he is looking at a self-organizing system when he looks at this chess game. I think he is doing precisely the opposite of what Lewis Carroll did when he looked at the game of chess. This man from Mars is a kind of statistician, and I was wrong in my original suggestion that he looked this over as a self-organizing system. He is the kind of statistician who makes a mathematical model, perhaps like a model of learning, or something of this sort. And he says this is learning, but he is asking an inappropriate question about this because he really has no idea about chess, and couldn't ask questions about chess. He has got no experience in this game in common with the creature that plays it.

Now, when we talk about self-organizing systems, we commonly do have this much knowledge about the sort of question we should ask, namely a chess-like question. Lewis Carroll when he was pondering on the symmetries of chess wrote *Through the Looking Glass*, and he was not, of course, doing anything more than taking the part of a chessman in this game. He was pretending he was a decision maker or pretending the chessmen were. Of course, he played bad chess. But he was essentially asking questions about the chess game, as conceived by the chess player.

I suggest that if we look at it in this way we are looking at the self-organizing system. If we look at it from a macroscopic point of view or any other point of view you choose, which is a formalism in which we do not infer similarity with the people who commonly play this game, then we are not looking at a self-organizing system, we are simply describing it in an inappropriate metalanguage.